



Designing for Cisco Network Service Architectures (ARCH)

Foundation Learning Guide

(CCDP ARCH 300-320)

Fourth Edition



ciscopress.com

Marwan Al-shawi, CCDE No. 20130066
André Laurent, CCIE No. 21840 | CCDE No. 20120024

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



Designing for Cisco Network Service Architectures (ARCH) Foundation Learning Guide, Fourth Edition CCDP ARCH 300-320

Marwan Al-shawi, CCDE No. 20130066

André Laurent, CCDE No. 20120024, CCIE No. 21840

Cisco Press

800 East 96th Street

Indianapolis, Indiana 46240 USA

Designing for Cisco Network Service Architectures (ARCH) Foundation Learning Guide, Fourth Edition

Marwan Al-shawi and André Laurent

Copyright © 2017 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 December 2016

Library of Congress Control Number: 2016958010

ISBN-13: 978-1-58714-462-2

ISBN-10: 1-58714-462-x

Warning and Disclaimer

This book is designed to provide information about designing Cisco Network Service Architectures. 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 intlcs@pearson.com.

Feedback Information

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

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

We greatly appreciate your assistance.

Editor-in-Chief: Mark Taub

Alliances Manager, Cisco Press: Ron Fligge

Product Line Manager: Brett Bartow

Acquisitions Editor: Michelle Newcomb

Managing Editor: Sandra Schroeder

Development Editor: Ginny Munroe

Senior Project Editor: Tonya Simpson

Copy Editor: Chuck Hutchinson

Technical Editors: Denise Fishburne, Orhan Ergun

Editorial Assistant: Vanessa Evans

Cover Designer: Chuti Prasertsith

Composition: codeMantra

Indexer: Lisa Stumpf

Proofreader: Deepa Ramesh



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, CCRP, 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. (0812R)

About the Authors

Marwan Al-shawi, CCDE No. 20130066, is a Cisco Press author whose titles include the top Cisco certification design books *CCDE Study Guide* and *Designing for Cisco Network Service Architectures (ARCH) Foundation Learning Guide*, Fourth Edition. He also is an experienced technical architect. Marwan has been in the networking industry for more than 12 years and has been involved in architecting, designing, and implementing various large-scale networks, some of which are global service provider-grade networks. Marwan holds a Master of Science degree in internetworking from the University of Technology, Sydney. He enjoys helping and assessing network designs and architectures; therefore, he was selected as a Cisco Designated VIP by the Cisco Support Community (CSC) (official Cisco Systems forums) in 2012 and by the Solutions and Architectures subcommunity in 2014. In addition, Marwan was selected as a member of the Cisco Champions program in 2015 and 2016. In his spare time, Marwan provides CCDP- and CCDE-related training and blogs at netdesignarena.com.

André Laurent, 3xCCIE No. 21840, CCDE No. 20120024, is the worldwide director of engineering for enterprise networking sales at Cisco Systems and a Cisco Press author. Outside his own personal development, André has an equal passion for helping others develop their systems and assisting them with the certification process. André is recognized in the industry as a subject matter expert in the areas of routing, switching, security, and design. Although he wears a Cisco badge, André takes a neutral approach in helping clients establish a long-term business and technology vision covering necessary strategy, execution, and metrics for measuring impact.

About the Technical Reviewers

Denise “Fish” Fishburne, CCDE No. 20090014, CCIE No. 2639 (R&S, SNA), is an engineer and team lead with the Customer Proof of Concept Lab (CPOC) in North Carolina. Fish is a geek who absolutely adores learning and passing it on. She works on many technologies in the CPOC, but her primary technical strength is troubleshooting. Fish has been with Cisco since 1996 and CPOC since 2001, and has been a regular speaker at Networkers/Cisco Live since 2006. Cisco Live is a huge passion for Fish! As such, in 2009, she got even more deeply involved with it by becoming a Cisco Live session group manager. Look for Fish swimming in the bits and bytes all around you, or just go to www.NetworkingWithFish.com.

Orhan Ergun, CCDE No. 2014:0017, CCIE No. 2014:0017 (CCNP, CCDP, JNCIS, and JNCIP), is a network architect who focuses on service providers, data centers, virtualization, cloud, and network security. He has more than 13 years of IT experience and has worked on many medium- and large-scale network design and deployment projects. He teaches Cisco network design concepts and writes exam questions for Cisco Systems.

Dedications

I would like to dedicate this book to my wonderful mother for her continued support, love, encouragement, guidance, and wisdom, as well as to the people in my life who always support and encourage me.

And most importantly, I would like to thank God for all blessings in my life.

—*Marwan*

I would like to dedicate this book to the women in my life. My mother, for her unconditional dedication and love. My sister, for rescuing me from the drifter life and setting me up with my first job in the industry. My beautiful wife, who continues to stand by my side while encouraging me through all the new challenges, opportunities, and experiences life brings.

—*André*

Acknowledgments

A special thank you goes to the Pearson Cisco Press team for their support in making this book possible.

A big thank you goes to André for being part of this publication and adding his expert perspective. It's always a pleasure to work with an experienced and extremely helpful person like André.

We would like to give special recognition to the wonderful technical reviewers Denise Fishburne and Orhan Ergun for their valuable contributions in editing the book. Both Denise and Orhan are very experienced network designers and CCDE certified; therefore, their suggestions and feedback helped shape and optimize the quality of the contents on multiple areas.

In addition, a special thank you to Maurizio Portolani (Cisco Press author and distinguished system engineer at Cisco Systems) and John Weston (systems engineer at Cisco) for their help and support with the technical review and optimization of the ACI chapter.

Also, we want to thank Adrian Arumugam (network engineer for a major content provider) for his technical review and valuable comments of certain chapters.

Contents at a Glance

Introduction xxix

Part I Designing Reliable and Resilient Enterprise Layer 2 and Layer 3 Networks

Chapter 1 Optimal Enterprise Campus Design 1

Chapter 2 EIGRP Design 49

Chapter 3 OSPF Design 75

Chapter 4 IS-IS Design 101

Chapter 5 Border Gateway Protocol Design 145

Part II Enterprise IPv6 Design Considerations and Challenges

Chapter 6 IPv6 Design Considerations in the Enterprise 193

Chapter 7 Challenges of the Transition to IPv6 219

Part III Modern Enterprise Wide-Area Networks Design

Chapter 8 Service Provider–Managed VPNs 229

Chapter 9 Enterprise-Managed WANs 271

Chapter 10 Enterprise WAN Resiliency Design 323

Part IV Enterprise Data Center Designs

Chapter 11 Multitier Enterprise Data Center Designs 375

Chapter 12 New Trends and Techniques to Design Modern Data Centers 397

Chapter 13 Cisco Application-Centric Infrastructure 431

Chapter 14 Data Center Connections 477

Part V Design QoS for Optimized User Experience

Chapter 15 QoS Overview 513

Chapter 16 QoS Design Principles and Best Practices 553

Chapter 17 Campus, WAN, and Data Center QoS Design **567**

Chapter 18 MPLS VPN QoS Design **605**

Chapter 19 IPsec VPN QoS Design **619**

Part VI IP Multicast Design

Chapter 20 Enterprise IP Multicast Design **633**

Chapter 21 Rendezvous Point Distribution Solutions **665**

Part VII Designing Optimum Enterprise Network Security

Chapter 22 Designing Security Services and Infrastructure Protection **689**

Chapter 23 Designing Firewall and IPS Solutions **709**

Chapter 24 IP Multicast Security **743**

Chapter 25 Designing Network Access Control Solutions **759**

Part VIII Design Scenarios

Chapter 26 Design Case Studies **777**

Appendix A Answers to Review Questions **843**

Appendix B References **855**

Index **857**

Contents

	Introduction	xxix
Part I	Designing Reliable and Resilient Enterprise Layer 2 and Layer 3 Networks	
Chapter 1	Optimal Enterprise Campus Design	1
	Enterprise Campus Design Principles	2
	Hierarchy	3
	Access Layer	4
	Distribution Layer	5
	Core Layer	6
	Enterprise Campus Two-Tier Layer Model	8
	Enterprise Campus Three-Tier Layer Model	9
	Modularity	10
	Modular Enterprise Campus Architecture and Modular Enterprise Campus with OSPF	10
	Access-Distribution Block	13
	Flexibility	15
	Campus Network Virtualization	16
	Campus Network Virtualization Technologies and Techniques	17
	VLAN Assignment	17
	Virtual Routing and Forwarding	18
	Path Isolation Techniques	19
	Resiliency	23
	Enterprise Campus High-Availability Design Considerations	23
	VLANs, Trunking, and Link Aggregation Design Recommendations	24
	VLAN Design	24
	Trunking	27
	Link Aggregation	28
	First-Hop Redundancy Protocol (FHRP)	31
	IP Gateway Redundancy Optimization with VSS	35
	Layer 2 to Layer 3 Boundary Design Options and Considerations	36
	Distribution-to-Distribution Link Design Considerations	36
	A Summary of Enterprise Campus HA Designs	44
	Summary	46
	Review Questions	46
	References	48

Chapter 2 EIGRP Design 49

Scalable EIGRP Design Overview	50
EIGRP with Multiple Autonomous Systems	50
EIGRP Queries	52
Multiple EIGRP Autonomous System Drivers	53
EIGRP Multilayer Architectures	53
EIGRP Two-Layer Hierarchy Architecture	56
EIGRP Three-Layer Hierarchy Architecture	57
EIGRP Hub-and-Spoke Design	60
Summarization Challenges	61
<i>Route Summarization Black Holes</i>	61
<i>Route Summarization and Suboptimal Routing</i>	63
EIGRP Hub-and-Spoke Scalability Optimization	65
<i>EIGRP Stub Leaking</i>	67
EIGRP DMVPN Scaling	69
EIGRP Fast Convergence Design Considerations	70
Bidirectional Forwarding Detection	70
EIGRP Graceful Restart/NSF Considerations	71
Summary	72
Review Questions	72

Chapter 3 OSPF Design 75

OSPF Scalability Design Considerations	76
Adjacent Neighbors	76
Routing Information in the Area and the Routed Domain	78
Numbers of Routers in an Area	80
Number of Areas per ABR	81
OSPF Area Design Considerations	82
OSPF Hierarchy	84
Area and Domain Summarization	85
OSPF Full-Mesh Design	87
OSPF Hub-and-Spoke Design	88
OSPF ABR Placement in Hub-and-Spoke Design	89
Number of Areas in OSPF Hub-and-Spoke Design	91
OSPF Network Types in Hub-and-Spoke Design	92

OSPF Convergence Design Considerations and Optimization Techniques	93
Event Detection	94
OSPF Event Propagation	94
OSPF Event Processing	96
OSPF Flooding Reduction	97
OSPF Database Overload Protection	97
Summary	98
Review Questions	99

Chapter 4 IS-IS Design 101

Protocol Overview	102
IS-IS Characteristics	103
Integrated IS-IS Routing	104
IS-IS Hierarchical Architecture Overview	105
IS-IS Router and Link Types	106
IS-IS Adjacencies	108
IS-IS Versus OSPF	110
Similarities Between IS-IS and OSPF	110
OSPF and IS-IS Characteristics	110
Integrated IS-IS and OSPF Area Designs	112
<i>OSPF Area Design</i>	112
<i>Integrated IS-IS Area Design</i>	113
IS-IS Technical Deep Dive	114
IS-IS Addressing	114
<i>IS-IS Packets</i>	117
<i>IS-IS Information Data Flow</i>	118
<i>IS-IS Network Types</i>	119
<i>IS-IS Protocol Operations</i>	119
<i>Level 1 and Level 2 LSPs and IIHs</i>	121
IS-IS Link-State Packets Flooding	122
IS-IS LSDB Synchronization	123
IS-IS Design Considerations	124
IS-IS Routing Logic Overview	125
<i>Advanced IS-IS Routing</i>	126
Route Leaking	126
Asymmetric Versus Symmetric IS-IS Routing	129

IS-IS Routing over NBMA Hub-and-Spoke	132
IS-IS Routing over a Full-Mesh Network	133
Flat IS-IS Routing Design	134
Hierarchal IS-IS Design	135
IS-IS Routes Summarization	136
Integrated IS-IS for IPv6	138
<i>IS-IS Single-Topology Restrictions</i>	138
<i>Multitopology IS-IS for IPv6</i>	140
Final Thoughts on IS-IS Routing Design	141
Summary	142
Review Questions	142
Chapter 5 Border Gateway Protocol Design	145
BGP Overview	146
BGP Speaker Types	147
BGP Loop Prevention and Split-Horizon Rule	148
BGP Path Attributes and Path Selection (Review)	149
<i>BGP Path Attributes</i>	150
<i>How BGP Selects Paths</i>	150
Designing Scalable iBGP Networks	152
iBGP Scalability Limitations	152
iBGP Scalability Solutions	152
<i>BGP Route Reflectors</i>	153
<i>BGP Confederations</i>	155
<i>BGP Confederations Versus BGP Route Reflectors</i>	157
BGP Route Reflector Design	158
Route Reflector Split-Horizon Rule	158
BGP Route Reflectors Redundancy Design Options and Considerations	159
<i>Route Reflector Clusters</i>	160
<i>Loop-Prevention Mechanisms</i>	162
<i>Congruence of Physical and Logical Networks</i>	165
<i>Hierarchical Route Reflector Design</i>	167
Route Reflector Potential Network Design Issues	169
Enhancing the Design of BGP Policies with BGP Communities	169
BGP Community Attribute Overview	169
Well-Known BGP Communities	170

BGP Named Community List	171
Planning for the Use of BGP Communities	171
Case Study: Designing Enterprise wide BGP Policies Using BGP Communities	172
Enterprise BGP Policy Requirements	173
BGP Community Solution Design	174
<i>Solution Detailed Design and Traffic Flow</i>	175
BGP Load-Sharing Design	177
Single-Homing Versus Multihoming	177
Dual-Homing and Multihoming Design Considerations	178
<i>Single-Homed, Multiple Links</i>	178
<i>Dual-Homed to One ISP Using a Single Local Edge Router</i>	180
<i>Dual-Homed to One ISP Using Multiple Edge Routers</i>	182
<i>Multihoming with Two ISPs Using a Single Local Edge Router</i>	183
<i>Multihoming with Two ISPs Using Multiple Local Edge Routers</i>	186
Summary	189
Review Questions	189

Part II Enterprise IPv6 Design Considerations and Challenges

Chapter 6 IPv6 Design Considerations in the Enterprise 193

IPv6 Deployment and Design Considerations	194
Business and Network Discovery Phase	196
Assessment Phase	196
Planning and Design Phase	196
Implementation and Optimization Phases	197
Considerations for Migration to IPv6 Design	197
Acquiring IPv6 Prefixes	197
<i>Provider Independent Versus Provider Assigned</i>	198
Where to Start the Migration	199
Migration Models and Design Considerations	200
<i>IPv6 Island</i>	200
<i>IPv6 WAN</i>	201
IPv6 Transition Mechanisms	203
Dual Stack	205
NAT64 and DNS64	206
Manual Tunnels	208
Tunnel Brokers	209

	6 Rapid Deployment	210
	Dual-Stack Lite (DS-Lite)	211
	Locator/ID Separation Protocol (LISP)	212
	<i>LISP Site Edge Devices</i>	213
	<i>LISP Infrastructure Devices</i>	213
	Final Thoughts on IPv6 Transition Mechanisms	216
	Summary	217
	Review Questions	217
Chapter 7	Challenges of the Transition to IPv6	219
	IPv6 Services	219
	Name Services	220
	<i>Implementation Recommendations</i>	220
	Addressing Services	220
	<i>Implementation Recommendations</i>	221
	Security Services	221
	Link Layer Security Considerations	221
	Application Support	222
	<i>Application Adaptation</i>	223
	<i>Application Workarounds</i>	223
	Control Plane Security	224
	Dual-Stack Security Considerations	225
	Tunneling Security Considerations	225
	Multihoming	226
	Summary	226
	Review Questions	227
Part III	Modern Enterprise Wide-Area Networks Design	
Chapter 8	Service Provider–Managed VPNs	229
	Choosing Your WAN Connection	230
	Layer 3 MPLS VPNs	233
	MPLS VPN Architecture	234
	Enterprise Routing Considerations	236
	Provider Edge (PE) Router Architecture	237
	<i>Route Distinguishers</i>	238
	<i>Route Target (RT)</i>	240
	PE-CE Routing Protocol	241
	<i>Using EIGRP as the PE-CE Routing Protocol</i>	241

	<i>Using OSPF as the PE-CE Routing Protocol</i>	247
	<i>Using BGP as the PE-CE Routing Protocol</i>	252
	Case Study: MPLS VPN Routing Propagation	255
	Forwarding in MPLS VPN	258
	Layer 2 MPLS VPN Services	259
	Virtual Private Wire Service (VPWS)	259
	Virtual Private LAN Service (VPLS)	261
	<i>VPLS Scalability Considerations</i>	263
	<i>VPLS Resiliency Considerations</i>	265
	VPLS Versus VPWS	266
	Summary	267
	Review Questions	268
Chapter 9	Enterprise-Managed WANs	271
	Enterprise-Managed VPN Overview	272
	GRE Overview	273
	Multipoint GRE Overview	275
	Point-to-Point and Multipoint GRE Comparison	276
	IPsec Overview	278
	IPsec and GRE	280
	IPsec and Virtual Tunnel Interface	281
	IPsec and Dynamic VTI	283
	DMVPN Overview	283
	DMVPN Phase 1	287
	DMVPN Phase 2	289
	DMVPN Phase 3	292
	Case Study: EIGRP DMVPN	295
	EIGRP over DMVPN Phase 1	295
	EIGRP over DMVPN Phase 2	297
	EIGRP over DMVPN Phase 3	299
	DMVPN Phase 1–3 Summary	302
	DMVPN and Redundancy	302
	Case Study: MPLS/VPN over GRE/DMVPN	304
	SSL VPN Overview	312

FlexVPN Overview	314
FlexVPN Architecture	315
FlexVPN Capabilities	315
FlexVPN Configuration Blocks	315
GETVPN	317
Summary	320
Review Questions	321

Chapter 10 Enterprise WAN Resiliency Design 323

WAN Remote-Site Overview	324
MPLS Layer 3 WAN Design Models	326
Common Layer 2 WAN Design Models	329
Common VPN WAN Design Models	331
3G/4G VPN Design Models	335
Remote Site Using Local Internet	337
Remote-Site LAN	339
Case Study: Redundancy and Connectivity	343
ATM WAN Design	344
Remote-Site (Branch Office) WAN Design	346
Regional Offices WAN Design	348
Basic Traffic Engineering Techniques	351
NGWAN, SDWAN, and IWAN Solution Overview	354
Transport-Independent Design	356
Intelligent Path Control	356
Application Optimization	356
Secure Connectivity	357
Management	357
IWAN Design Overview	358
IWAN Hybrid Design Model	359
Cisco PfR Overview	361
Cisco PfR Operations	362
Cisco IWAN and PfRv3	363
Cisco PfRv3 Design and Deployment Considerations	366
Enterprise WAN and Access Management	367
APIC-EM	368
Design of APIC-EM	370
Summary	371
Review Questions	372

Part IV Enterprise Data Center Designs

Chapter 11 Multitier Enterprise Data Center Designs 375

- Case Study 1: Small Data Centers (Connecting Servers to an Enterprise LAN) 376
- Case Study 2: Two-Tier Data Center Network Architecture 378
- Case Study 3: Three-Tier Data Center Network Architecture 380
 - Data Center Inter-VLAN Routing 381
 - End of Row Versus Top of Rack Design 383
 - Fabric Extenders 385
 - Data Center High Availability 388
 - Network Interface Controller Teaming 392
- Summary 394
- Review Questions 394

Chapter 12 New Trends and Techniques to Design Modern Data Centers 397

- The Need for a New Network Architecture 397
- Limitations of Current Networking Technology 398
- Modern Data Center Design Techniques and Architectures 400
 - Spine-Leaf Data Center Design 400
 - Network Overlays 402
 - Cisco Fabric Path* 402
 - Virtual Extensible LAN (VXLAN)* 407
 - VXLAN Tunnel Endpoint 408
 - Remote VTEP Discovery and Tenant Address Learning 411
 - VXLAN Control-Plane Optimization 413
 - Software-Defined Networking 414
 - How SDN Can Help* 416
 - Selection Criteria of SDN Solutions* 417
 - SDN Requirements* 419
 - SDN Challenges* 419
 - Direction of Nontraditional SDN* 421
- Multitenant Data Center 422
 - Secure Tenant Separation 422
 - Layer 3 Separation with VRF-Lite* 423
 - Device-Level Virtualization and Separation* 424

- Case Study: Multitenant Data Center 425
- Microsegmentation with Overlay Networks 427

- Summary 428
- Review Questions 429
- References 430

Chapter 13 Cisco Application-Centric Infrastructure 431

- ACI Characteristics 432
- How the Cisco ACI Addresses Current Networking Limitations 432
- Cisco ACI Architecture Components 434
 - Cisco Application Policy Infrastructure Controller (APIC) 434
 - APIC Approach Within the ACI Architecture* 436
 - Cisco ACI Fabric 437
- ACI Network Virtualization Overlays 441
- Application Design Principles with the Cisco ACI Policy Model 447
 - What Is an Endpoint Group in Cisco ACI? 450
 - Design EPGs* 451
 - ACI Fabric Access Polices 454
 - Building Blocks of a Tenant in the Cisco ACI 456
 - Crafting Applications Design with the Cisco ACI 459
 - ACI Interaction with External Layer 2 Connections and Networks 461
 - Connecting ACI to the Outside Layer 2 Domain* 462
 - ACI Integration with STP-Based Layer LAN* 464
- ACI Routing 465
 - First-Hop Layer 3 Default Gateway in ACI 465
 - Border Leaves 467
 - Route Propagation inside the ACI Fabric 468
 - Connecting the ACI Fabric to External Layer 3 Domains 470
 - Integration and Migration to ACI Connectivity Options 471
- Summary 473
- Review Questions 475
- References 476

Chapter 14 Data Center Connections 477

- Data Center Traffic Flows 478
 - Traffic Flow Directions 478
 - Traffic Flow Types 479

The Need for DCI	482
IP Address Mobility	484
Case Study: Dark Fiber DCI	490
Pseudowire DCI	495
Virtual Private LAN Service DCI	496
Customer-Managed Layer 2 DCI Deployment Models	497
<i>Any Transport over MPLS over GRE</i>	497
<i>Customer-Managed Layer 2 DCI Deployment</i>	498
<i>Layer 2 DCI Caveats</i>	501
<i>Overlay Transport Virtualization DCI</i>	501
Overlay Networking DCI	507
Layer 3 DCI	507
Summary	509
Review Questions	510

Part V Design QoS for Optimized User Experience

Chapter 15 QoS Overview 513

QoS Overview	514
IntServ versus DiffServ	514
Classification and Marking	516
Classifications and Marking Tools	516
Layer 2 Marking: IEEE 802.1Q/p Class of Service	517
Layer 3 Marking: IP Type of Service	519
Layer 3 Marking: DSCP Per-Hop Behaviors	520
Layer 2.5 Marking: MPLS Experimental Bits	524
Mapping QoS Markings between OSI Layers	524
Layer 7 Classification: NBAR/NBAR2	526
Policers and Shapers	527
Token Bucket Algorithms	529
Policing Tools: Single-Rate Three-Color Marker	532
Policing Tools: Two-Rate Three-Color Marker	533
Queuing Tools	535
Tx-Ring	536
Fair Queuing	537
CBWFQ	538

- Dropping Tools 541
 - DSCP-Based WRED 541
 - IP ECN 547
- Summary 550
- Review Questions 550

Chapter 16 QoS Design Principles and Best Practices 553

- QoS Overview 553
- Classification and Marking Design Principles 554
- Policing and Remarking Design Principles 556
- Queuing Design Principles 557
- Dropping Design Principles 557
- Per-Hop Behavior Queue Design Principles 558
- RFC 4594 QoS Recommendation 559
- QoS Strategy Models 560
 - 4-Class QoS Strategy 561
 - 8-Class QoS Strategy 562
 - 12-Class QoS Strategy 564
- Summary 565
- Review Questions 565

Chapter 17 Campus, WAN, and Data Center QoS Design 567

- Campus QoS Overview 568
 - VoIP and Video 568
 - Buffers and Bursts 569
 - Trust States and Boundaries 570
 - Trust States and Boundaries Example* 571
 - Dynamic Trust State* 572
 - Classification/Marking/Policing QoS Model 573
 - Queuing/Dropping Recommendations 574
 - Link Aggregation “EtherChannel” QoS Design 575
 - Practical Example of Campus QoS Design 576
- WAN QoS Overview 588
 - Platform Performance Considerations 589
 - Latency and Jitter Considerations 590
 - Queuing Considerations 591
 - Shaping Considerations 592
 - Practical Example of WAN and Branch QoS 593

	Data Center QoS Overview	594
	High-Performance Trading Architecture	595
	Big Data Architecture	596
	Case Study: Virtualized Multiservice Architectures	596
	Data Center Bridging Toolset	597
	Case Study: DC QoS Application	599
	Summary	601
	Review Questions	603
Chapter 18	MPLS VPN QoS Design	605
	The Need for QoS in MPLS VPN	605
	Layer 2 Private WAN QoS Administration	607
	Fully Meshed MPLS VPN QoS Administration	608
	MPLS DiffServ Tunneling Modes	609
	Uniform Tunneling Mode	612
	Short-Pipe Tunneling Mode	612
	Pipe Tunneling Mode	614
	Sample MPLS VPN QoS Roles	615
	Summary	617
	Review Questions	617
Chapter 19	IPsec VPN QoS Design	619
	The Need for QoS in IPsec VPN	619
	VPN Use Cases and Their QoS Models	621
	IPsec Refresher	621
	IOS Encryption and Classification: Order of Operations	623
	MTU Considerations	625
	DMVPN QoS Considerations	626
	GET VPN QoS Considerations	629
	Summary	630
	Review Questions	631
Part VI	IP Multicast Design	
Chapter 20	Enterprise IP Multicast Design	633
	How Does IP Multicast Work?	634
	Multicast Group	635
	IP Multicast Service Model	636
	Functions of a Multicast Network	638

Multicast Protocols	638
Multicast Forwarding and RPF Check	639
Case Study 1: RPF Check Fails and Succeeds	641
Multicast Protocol Basics	642
Multicast Distribution Trees Identification	644
PIM-SM Overview	645
Receiver Joins PIM-SM Shared Tree	646
Registered to RP	647
PIM-SM SPT Switchover	649
Multicast Routing Table	652
Basic SSM Concepts	654
SSM Scenario	655
Bidirectional PIM	657
PIM Modifications for Bidirectional Operation	658
<i>DF Election</i>	658
<i>DF Election Messages</i>	660
Case Study 2: DF Election	660
Summary	662
Review Questions	663

Chapter 21 Rendezvous Point Distribution Solutions 665

Rendezvous Point Discovery	665
Rendezvous Placement	667
Auto-RP	668
<i>Auto-RP Candidate RPs</i>	670
<i>Auto-RP Mapping Agents</i>	670
<i>Auto-RP and Other Routers</i>	670
<i>Case Study: Auto-RP Operation</i>	670
<i>Auto-RP Scope Problem</i>	674
PIMv2 BSR	676
<i>PIMv2 BSR: Candidate RPs</i>	677
<i>PIMv2 BSR: Bootstrap Router</i>	678
<i>PIMv2 BSR: All PIMv2 Routers</i>	678
<i>BSR Flooding Problem</i>	678
IPv6 Embedded Rendezvous Point	679
Anycast RP Features	681
Anycast RP Example	682

MSDP Protocol Overview	683
MSDP Neighbor Relationship	683
Case Study: MSDP Operation	684
Summary	686
Review Questions	687

Part VII Designing Optimum Enterprise Network Security

Chapter 22 Designing Security Services and Infrastructure Protection 689

Network Security Zoning	690
Cisco Modular Network Architecture	691
Cisco Next-Generation Security	696
Designing Infrastructure Protection	696
Infrastructure Device Access	698
Routing Infrastructure	699
Device Resiliency and Survivability	700
Network Policy Enforcement	701
Switching Infrastructure	702
SDN Security Considerations	703
Summary	705
Review Questions	705

Chapter 23 Designing Firewall and IPS Solutions 709

Firewall Architectures	709
Virtualized Firewalls	712
Case Study 1: Separation of Application Tiers	714
Securing East-West Traffic	716
Case Study 2: Implementing Firewalls in a Data Center	717
Case Study 3: Firewall High Availability	720
IPS Architectures	726
Case Study 4: Building a Secure Campus Edge Design (Internet and Extranet Connectivity)	729
Campus Edge	730
Connecting External Partners	737
<i>Challenges of Connecting External Partners</i>	737
<i>Extranet Topology: Remote LAN Model</i>	737
<i>Extranet Topology: Interconnect Model</i>	738
<i>Extranet: Security and Multitenant Segmentation</i>	739

Summary 740
Review Questions 741

Chapter 24 IP Multicast Security 743

Multicast Security Challenges 744
Problems in the Multicast Network 744
Multicast Network Security Considerations 745
 Network Element Security 746
 Security at the Network Edge 748
 Securing Auto-RP and BSR 749
 MSDP Security 751
 PIM and Internal Multicast Security 752
 Multicast Sender Control 753
 Multicast Receiver Controls 755
 Multicast Admission Controls 757
Summary 757
Review Questions 758

Chapter 25 Designing Network Access Control Solutions 759

IEEE 802.1X Overview 759
Extensible Authentication Protocol 763
802.1X Supplicants 765
IEEE 802.1X Phased Deployment 767
Cisco TrustSec 768
 Profiling Service 768
 Security Group Tag 769
Case Study: Authorization Options 772
Summary 775
Review Questions 775

Part VIII Design Scenarios

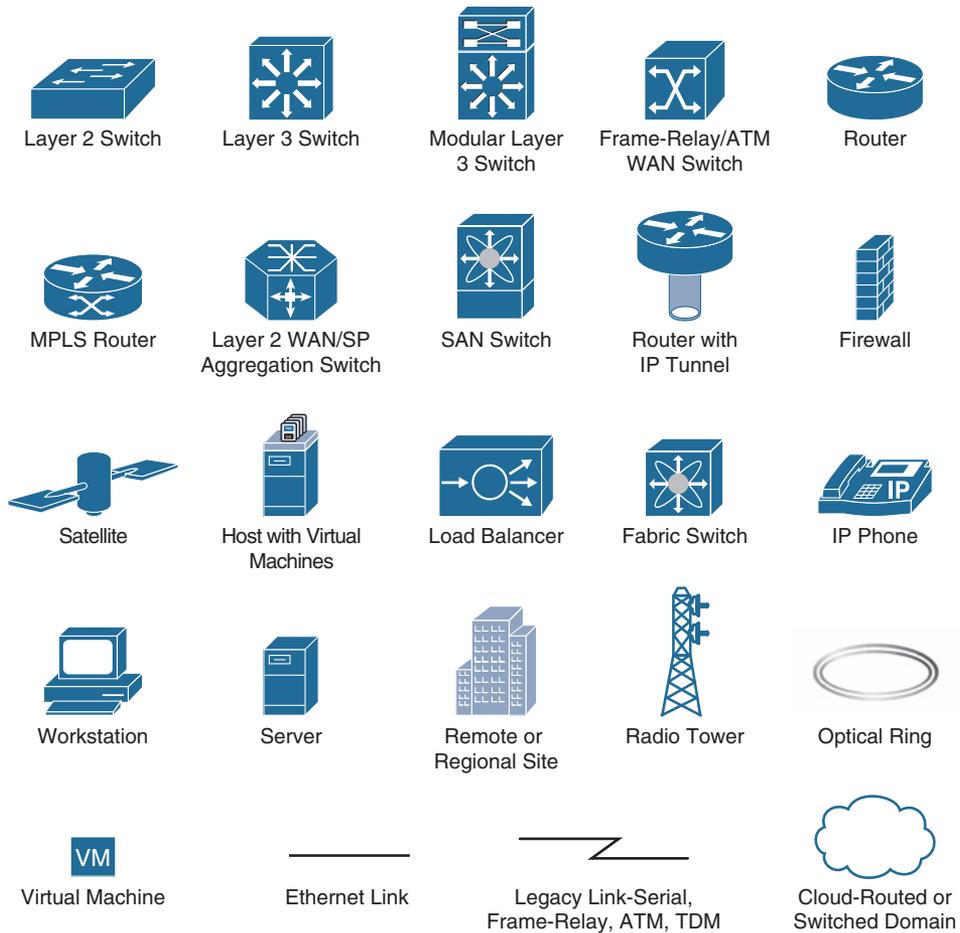
Chapter 26 Design Case Studies 777

Case Study 1: Design Enterprise Connectivity 778
 Detailed Requirements and Expectations 778
 Design Analysis and Task List 779
 Selecting a Replacement Routing Protocol 780
 Designing for the New Routing Protocol 780

OSPF Design Optimization	782
Planning and Designing the Migration from the Old to the New Routing	785
Scaling the Design	787
Case Study 2: Design Enterprise BGP Network with Internet Connectivity	788
Detailed Requirements and Expectations	788
Design Analysis and Task List	791
Choosing the Routing Protocol	792
Choosing the Autonomous System Numbers	792
BGP Connectivity	795
<i>BGP Sessions</i>	795
<i>BGP Communities</i>	796
Routing Policy	797
<i>Routing Policy in North American Sites</i>	797
<i>Routing Policy in European and Asian Sites</i>	799
Internet Routing	803
<i>Public IP Space Selection</i>	803
<i>Main HQ Multihoming</i>	804
<i>Default Routing</i>	805
Case Study 3: Design Enterprise IPv6 Network	807
Detailed Requirements and Expectations	808
Design Analysis and Task List	809
Choosing the IP Address Type for the HQ	809
Connecting the Branch Sites	810
Deployment Model	812
Addressing	813
<i>Address Provisioning</i>	814
Communication Between Branches	815
Application and Service Migration	815
Case Study 4: Design Enterprise Data Center Connectivity	816
Detailed Requirements and Expectations	817
Design Analysis and Task List	818
Selecting the Data Center Architecture and Connectivity Model	818
DCN Detailed Connectivity	819

Connecting Network Appliances	821
Data Center Interconnect	822
Data Center Network Virtualization Design	823
Case Study 5: Design Resilient Enterprise WAN	825
Detailed Requirements and Expectations	825
Design Analysis and Task List	826
Selecting WAN Links	828
WAN Overlay	828
Case Study 6: Design Secure Enterprise Network	830
Detailed Requirements and Expectations	831
Security Domains and Zone Design	832
Infrastructure and Network Access Security	833
Layer 2 Security Considerations	834
Main and Remote Location Firewalling	835
Case Study 7: Design QoS in the Enterprise Network	835
Detailed Requirements and Expectations	835
Traffic Discovery and Analysis	836
QoS Design Model	837
QoS Trust Boundary	838
Congestion Management	838
Scavenger Traffic Considerations	839
MPLS WAN DiffServ Tunneling	839
Appendix A Answers to Review Questions	843
Appendix B References	855
Index	857

Icons Used in This Book



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.

Reader Services

Register your copy at www.ciscopress.com/title/9781587144622 for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account*. Enter the product ISBN 9781587144622 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

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

Introduction

Enterprise environments require networks designed for performance, availability, and scalability to achieve outcomes. Seasoned IT professionals with progressive end-to-end network design expertise are crucial in ensuring networks deliver to meet today's requirements while future-proofing investments. For senior network design engineers, principal system engineers, network/solution architects, and CCDA professionals looking to build on your fundamental Cisco network design expertise, the Cisco CCDP certification program focuses on advanced addressing and routing protocols, WANs, service virtualization, and integration strategies for multilayered enterprise architectures.

This exam tests a candidate's knowledge and skills needed to design or help in designing an enterprise network. Successful candidates will be able to design and understand the inner workings of all elements within the common enterprise network, including internal routing, BGP routing, modern WAN connectivity, modern data center and data center interconnect, basic network security considerations, advanced quality-of-service design, transition to IPv6, and multicast routing design.

Goals of This Book

Designing Cisco Network Service Architectures (ARCH) enables network designers, engineers, architects, and CCDP candidates to perform the conceptual, intermediate, and detailed design of a network infrastructure that supports desired network solutions over intelligent network services to achieve effective performance, scalability, and availability. By applying solid Cisco network solution models and recommended design practices, ARCH enables learners to provide viable, stable enterprise internetworking solutions. This book presents concepts and examples necessary to design converged enterprise networks. Also, this new edition has content addressing software-defined networks (SDNs). You will learn additional aspects of modular campus design, advanced routing designs, WAN service designs, enterprise data center design, and security design.

Who Should Read This Book

Besides those who are planning or studying for the CCDP certification, this book is for

- Network designers, architects, consultants, or engineers seeking a thorough understanding of enterprise network design
- Network engineers or architects who are studying for the CCDE certification and need to improve their foundational knowledge of modern enterprise network design
- Anyone wanting to understand basic and advanced network design with an intermediate to advanced level of experience

How This Book Is Organized

This book is organized into eight distinct sections.

Part I of the book explains briefly the various design approaches, requirements, and principles required to design an optimum enterprise campus network. Also, it focuses on enterprise routing design, covering the different design options, considerations, and design implications with regard to business and other design requirements.

- **Chapter 1, “Optimal Enterprise Campus Design”:** This chapter discusses how to design a scalable and reliable enterprise campus taking into account applications and business requirements.
- **Chapter 2, “EIGRP Design”:** This chapter highlights, analyzes, and discusses different design options and considerations of EIGRP that any network designer must be aware of.
- **Chapter 3, “OSPF Design”:** This chapter looks at the different design options and considerations of OSPF that any network designer must be aware of, such as OSPF area design.
- **Chapter 4, “IS-IS Design”:** This chapter discusses IS-IS level design. It also compares the key functionalities of IS-IS and OSPF as link-state routing protocols.
- **Chapter 5, “Border Gateway Protocol Design”:** This chapter highlights, analyzes, and discusses different design options and considerations of BGP that any network designer must be aware of. It also provides some advanced BGP design approaches to address enterprise design needs.

Part II of the book focuses on IPv6 and how to plan and migrate your network to be IPv6 enabled along with the different design considerations and implications.

- **Chapter 6, “IPv6 Design Considerations in the Enterprise”:** This chapter highlights and explains the different design considerations and approaches of migrating IPv4 networks to IPv6.
- **Chapter 7, “Challenges of the Transition to IPv6”:** This chapter discusses the different challenges associated with migration to IPv6 that you need to take into account.

Part III of the book focuses on the different models of modern enterprise wide-area network design.

- **Chapter 8, “Service Provider–Managed VPNs”:** This chapter highlights and discusses the MPLS Layer 3 and Layer 2 VPN-based WAN modes along with the different design considerations and aspects that you need to be aware of.
- **Chapter 9, “Enterprise-Managed WAN”:** This chapter discusses the different enterprise-controlled VPN-based WAN models that can be used in today’s enterprise networks.

- **Chapter 10, “Enterprise WAN Resiliency Design”:** This chapter explains how to optimize the enterprise-managed WAN model to design a resilient overlay WAN model.

Part IV of the book focuses on the design options and technologies required to design an enterprise data center network.

- **Chapter 11, “Multitier Enterprise Data Center Designs”:** This chapter analyzes, explains, and compares the different data center design options and where each should be used.
- **Chapter 12, “New Trends and Techniques to Design Modern Data Centers”:** This chapter analyzes, explains, and compares the different modern data center design options and technologies and the drivers of each. It also introduces you to the data center overlay and SDN concepts.
- **Chapter 13, “Cisco Application-Centric Infrastructure”:** This chapter analyzes and explains the foundations of the Cisco ACI and the design concepts and terms that are ACI-specific, along with the different migration options from a traditional data center network to an ACI-based data center network.
- **Chapter 14, “Data Center Connections”:** This chapter analyzes, explains, and compares the different data center interconnect design options and considerations.

Part V of the book focuses on designing quality of service (QoS) for an optimized user experience and dives deeper, discussing QoS design for the different places in the network.

- **Chapter 15, “QoS Overview”:** This chapter explains the different QoS design concepts, techniques, and tools that any design engineer needs to be fully aware of its foundations.
- **Chapter 16, “QoS Design Principles and Best Practices”:** This chapter explains the different QoS design principles and strategies required to design a reliable QoS-enabled network.
- **Chapter 17, “Campus, WAN, and Data Center QoS Design”:** This chapter explains the best-practice design principles for enabling QoS in campus, WAN, and data center networks.
- **Chapter 18, “MPLS VPN QoS Design”:** This chapter covers the basics of designing QoS for MPLS VPN networks.
- **Chapter 19, “IPsec VPN QoS Design”:** This chapter reviews QoS-related considerations for IPsec VPNs.

Part VI of the book is an entry point to IP multicast services. It presents the functional model of IP multicast and gives an overview of technologies that are present in IP multicasting. The part is composed of an introduction to IP multicast concepts as well as a discussion of distribution trees and protocols.

- **Chapter 20, “Enterprise IP Multicast Design”:** This chapter reviews the foundations of IP multicast and how a multicast-enabled network delivers traffic from a source to a receiver. Also, it explains the most current scalable IP multicast routing protocol.
- **Chapter 21, “Rendezvous Point Distribution Solutions”:** This chapter offers an overview of RP distribution solutions. It explains the drawbacks of manual RP configuration and describes the Auto-RP and the BSR mechanisms. The chapter also introduces the concept of Anycast RP, which works in combination with the MSDP.

Part VII of the book focuses on how to design security services and what solutions are available today to implement network-level security.

- **Chapter 22, “Designing Security Services and Infrastructure Protection”:** This chapter explains how to secure the network infrastructure as it is a critical business asset.
- **Chapter 23, “Designing Firewall and IPS Solutions”:** This chapter explains the common firewall and IPS architectures, high-availability modes, and firewall virtualization along with design recommendations.
- **Chapter 24, “IP Multicast Security”:** This chapter describes the challenges with IP multicast security along with recommendations of how to secure a multicast network edge, Auto-RP, BSR, and MSDP.
- **Chapter 25, “Designing Network Access Control Solutions”:** This chapter discusses the different access control design approaches, including IEEE 802.1X–based access control and Cisco TrustSec technology.

Part VIII of the book offers some design scenarios that help you, as design engineer, practice designing technology solutions based on business and technical requirements.

- **Chapter 26, “Design Case Studies”:** This chapter provides different design scenarios that cover the design of IGP, BGP, WAN, data center networks, security, IPv6, and QoS.

QoS Design Principles and Best Practices

Upon completing this chapter, you will be able to

- Describe basic classification and marking design principles
- Describe basic policing and remarking design principles
- Explain queuing design principles
- Explain basic dropping design principles
- Explain what are per-hop behavior queue design principles
- Explain the role of RFC 4594 recommendation
- List and describe QoS strategy models
- Describe the 4-class QoS strategy model
- Describe the 8-class QoS strategy model
- Describe the 12-class QoS strategy model

Now that we have covered the various tools for enabling quality of service (QoS) in the network, it is possible to create a QoS strategy that best meets an organization's requirements. This chapter presents some best practice QoS design principles and QoS strategy models that are used to implement the numerous QoS tools we have at our disposal. Remember that usually more than one solution fits the given QoS requirements, so simplifying the models leveraged can significantly accelerate and ensure proper QoS deployment.

QoS Overview

Quality of service is critical to ensuring application performance consistency and optimized end-user experiences. As discussed in Chapter 15, "QoS Overview," the fundamental purpose of QoS is to manage contention for network resources while

addressing applications that require differentiated levels of service. Prior to developing a QoS strategy, you must perform the proper discovery to identify current and future applications and application characteristics within the environment. This information, coupled with an understanding of the end-to-end network design and traffic patterns, will drive the QoS design strategy model that is most appropriate for the business. Following are some common questions that you need to answer:

- What traffic needs to be classified and marked?
- Is it possible to leverage a 4-class, 8-class, or 12-class QoS strategy model from end to end?
- Will traffic-marking characteristics stay in place as data traverses the infrastructure?
- What traffic needs to be prioritized?
- What traffic requires bandwidth reservations?
- What traffic needs to be policed?
- Is shaping required at the WAN edge or at other places within the infrastructure such as the Data Center Interconnect (DCI)?
- How can congestion management and congestion avoidance techniques be leveraged to optimize TCP traffic?

Classification and Marking Design Principles

The first fundamental design principle is that QoS policies should always be enabled in hardware whenever possible. Some Cisco routers perform QoS in software, and such behavior can increase the load on the CPU. Cisco Catalyst switches have dedicated hardware called application-specific integrated circuits (ASIC), which are used to perform QoS operations. Switches can perform complex QoS policies under maximum traffic load without any marginal CPU spike. Some platforms, such as the Cisco ASR, can perform QoS operations (such as queuing) in dedicated hardware ASICs, but other functions (such as deep packet inspection) are still processed in software via the CPU.

Based on design recommendations, classification and marking should be done closest to the source of traffic as administratively and technically possible. This design principle promotes DiffServ and per-hop behaviors (PHB) as the recommended end-to-end design.

Note “As administratively close as possible” refers to an administrative domain, in scenarios in which you are not controlling the end-to-end traffic flow path of a packet; you need to classify/mark as close to the source as possible within your administrative domain.

As a rule, it is not recommended to trust markings set by end users leveraging PCs or other endpoint devices. End users can intentionally or unintentionally abuse QoS policies that trust markings of end devices. If users and unclassified applications take advantage of the configured QoS policy as a result of trusting end devices, this can result in easily starving priority queues with nonpriority traffic, ruining quality of service for real-time applications. However, if QoS markings for end devices and associated applications are administered centrally across the enterprise, this can be an acceptable design option. An additional area of exception might also include wireless devices that can leverage Wireless Multimedia (WMM) QoS provisioning in the upstream direction.

The next important recommendation is to use Differentiated Services Code Point (DSCP) marking whenever technically possible. DSCP markings are the recommended method for marking IP traffic for the following reasons:

- It has support for end-to-end Layer 3 marking.
- It is a more granular method of marking that supports 64 levels as compared to class of service (CoS) and MPLS Experimental EXP, which have 8 levels.
- It is more extensible than Layer 2 markings as these markings are lost when media changes.

To provide interoperability on the border between enterprise and service provider networks, you should use standard-based DSCP PHB markings because the use of such markings can streamline interoperability and compliance with service provider classes of service. Classification and marking design principles covered in this section are illustrated in Figure 16-1.

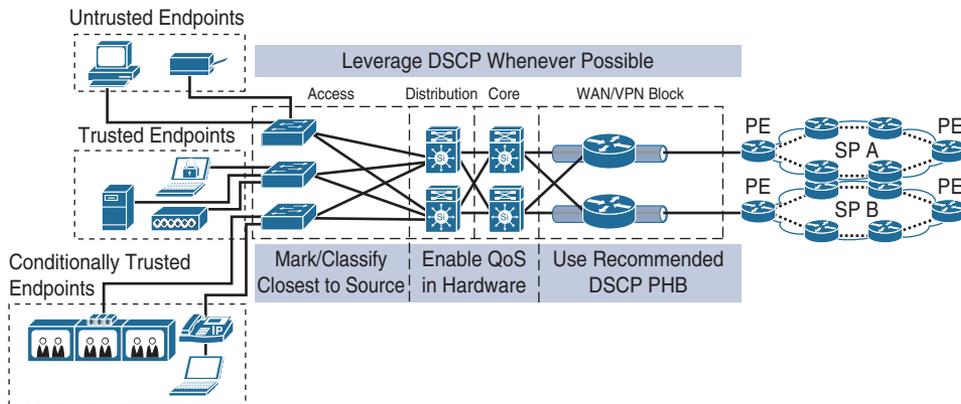


Figure 16-1 QoS Classification and Marking Architecture

Policing and Remarking Design Principles

Traffic that is unwanted should be discarded as soon as possible to preserve network resources from unnecessary consumption. Undesirable traffic can be the result of denial of service (DoS) or worm attacks. Furthermore, excessive unwanted traffic could cause a network outage as a result of high impact on the CPU and memory resources of network devices. Malicious traffic can mask under legitimate TCP/UDP ports that are used by well-known applications, and this traffic can create large amounts of unwanted traffic. Traffic behavior must be monitored and marked down as close as possible to the source under such circumstances.

Traffic should be marked down using RFC recommendations. Those recommendations ensure interoperability and end-to-end QoS network design. Examples of these recommendations are RFC 2597 and RFC 2698, where excess traffic with marking of AFx1 should be marked down to AFx2 or AFx3. Note that 2 or 3 in AFx2 and AFx3 represent drop probability. This markdown principle should be combined properly with other QoS tools. For example, with DSCP-based WRED, AFx2 should be dropped more aggressively than AFx1 but less aggressively than AFx3. Figure 16-2 illustrates the policing and remarking design principles covered in this section.

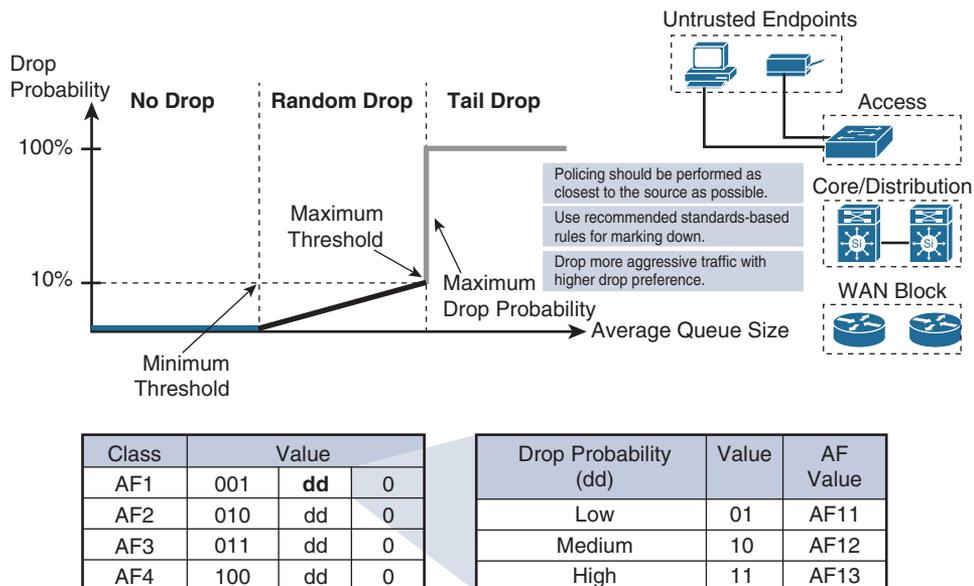


Figure 16-2 Policing and Remarking Concepts

Queuing Design Principles

The only way to provide QoS service guarantees to business-critical applications is to enable queuing to every node that has the potential for congestion. Queuing should be enabled regardless of whether congestion is occurring rarely or frequently. Although frequently deployed at the WAN edge, this principle must be applied not only to congested WAN links but also within the campus network. Speed mismatch, link aggregation, and link subscription ratios can create congestion in the network devices by filling up queuing buffers.

Because each distinctive application class requires unique QoS service requirements, it is recommended you provide a distinctive queue for each traffic class. One of the main justifications for leveraging distinctive queues is that each QoS service class can accept certain QoS-enabled behaviors such as bandwidth allocation and dropping ratios.

It is recommended you use a minimum of four standards-based queuing behaviors on all platforms and service provider links when deploying end-to-end QoS across the network infrastructure:

- RFC 3246 Expedited Forwarding PHB (used for real-time traffic)
- RFC 2597 Assured Forwarding PHB (used for guaranteed bandwidth queue)
- RFC 2474 Default Forwarding PHB (default nonprioritized queue, best effort)
- RFC 3662 Lower Effort Per-Domain Behavior (less than best-effort queue, bandwidth constrained)

Dropping Design Principles

As covered in Chapter 15, congestion avoidance mechanisms are used to selectively drop packets when a predefined limit is reached. As a review, by dropping packets early, congestion avoidance helps prevent bottlenecks downstream the network. Congestion avoidance mechanisms include RED and WRED. If WRED is designed per recommendations where every traffic class has its own queue, WRED should be used for only some types of queues (not necessarily all of them).

It is recommended that WRED not be used for the strict-priority queue, scavenger traffic queue, and control traffic queue. Traffic for the strict-priority queue and control traffic queue are highly sensitive to dropping. Scavenger traffic is often provisioned with a small amount of bandwidth, typically below 1 percent, and for this type of queue, WRED is not needed. Considering that the WRED feature is performed in software, enabling WRED for scavenger traffic class will consume additional CPU resources with no significant gain.

For AF-marked queues with DSCP-based WRED, typically traffic marked with AFx3 is more aggressively dropped than AFx2, which is in turn more aggressively dropped than AFx1.

All traffic types that are not explicitly defined in other queues fall into default (DF) traffic class. For this traffic class, it is recommended to enable WRED. WRED should be enabled in the default queue because, as explained in Chapter 15, it increases throughput by reducing the TCP synchronization effect. In the case of the default queue where all different traffic types are equally marked with a DSCP value of zero, there is no mechanism to fairly weight less aggressive applications when WRED is not enabled.

Per-Hop Behavior Queue Design Principles

The goal of convergence in the network is to enable voice, video, and data applications to seamlessly coexist in the network by providing each with appropriate QoS service expectations and guarantees.

When real-time applications are the only ones that consume link bandwidth, non-real-time applications' performance can be significantly degraded. Extensive testing results show that there is significant performance impact on non-real-time applications when more than one-third of the links is used by real-time applications as part of a strict-priority queue. Thus, it is recommended that no more than a third of link bandwidth be used for strict-priority queuing. This principle prevents non-real-time applications from being dropped out of their required QoS recommendations. In other words, it is recommended that no more than 33 percent of the bandwidth be used for the expedite forwarding (EF) queue. It is also important to note that this 33 percent design principle is simply a best practices design recommendation and not necessarily a mandatory rule.

It is recommended that a minimum of one queue be provisioned for assured forwarding per-hop behavior (AF PHB), but up to four subclasses can be defined within the AF class: AF1x, AF2x, AF3x, and AF4x. Each queue belonging to the specified AF subclass must have a bandwidth guarantee that corresponds to the application requirements of that traffic subclass.

The default forwarding (DF) class consists of all traffic that is not explicitly defined in other queues. If an enterprise is using many applications, it is important to have adequate space for those traffic types. It is recommended that typically 25 percent of link bandwidth be used for this service class. Figure 16-3 illustrates an example of bandwidth allocation leveraging these recommended best practices.

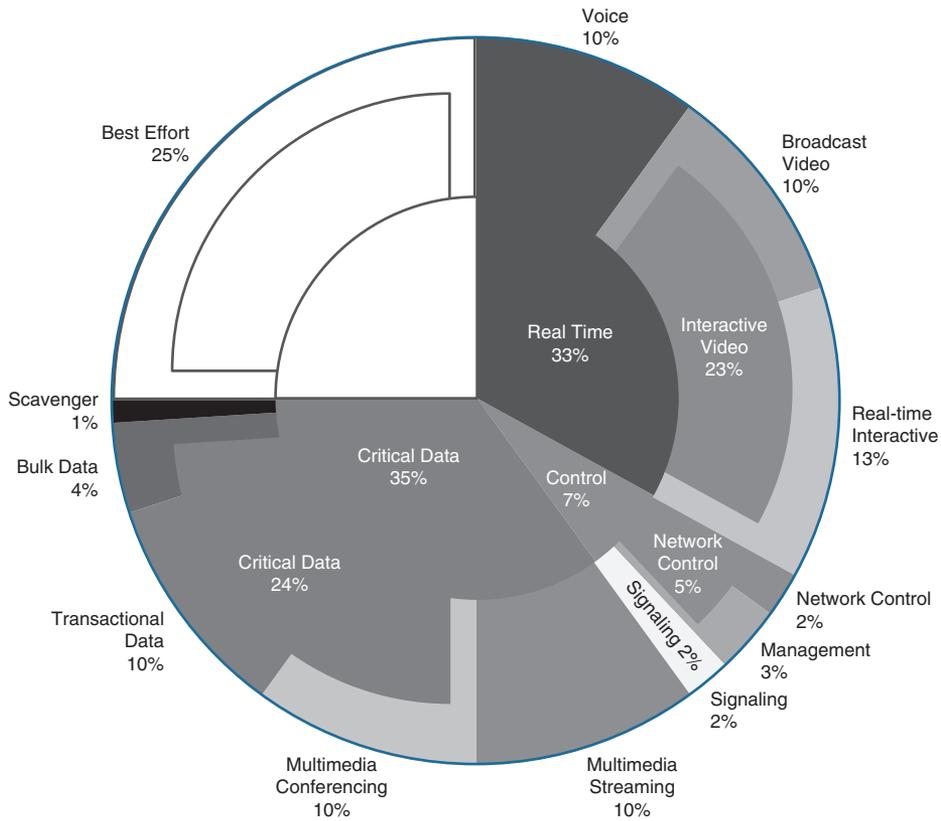


Figure 16-3 *Bandwidth Allocation Example*

RFC 4594 QoS Recommendation

RFC 4594 QoS provides guidelines for marking, queuing, and dropping principles for different types of traffic. Cisco has made a minor modification to its adoption of RFC 4594, namely the switching of Call-Signaling and Broadcast Video markings (to CS3 and CS5, respectively). A summary of Cisco's implementation of RFC 4594 is presented in Figure 16-4.

Cisco Implementation of RFC 4594-Based QoS				
Application Class	Per-Hop Behavior	Admission Control	Queuing and Dropping	Application Examples
VoIP Telephony	EF	Required	Priority Queue (PQ)	Cisco IP Phones (G.711, G.729)
Broadcast Video	CS5	Required	(Optional) PQ	Cisco IP Video Surveillance/Cisco Enterprise TV
Real-time Interactive	CS4	Required	(Optional) PQ	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Unified Personal Communicator, WebEx
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoDs)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Call Signaling	CS3		BW Queue	SCCP, SIP, H.323
Ops/Admin/Mgmt (OAM)	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	ERP Apps, CRM Apps, Database Apps
Bulk Data	AF1		BW Queue + DSCP WRED	E-mail, FTP, Backup Apps, Content Distribution
Best Effort	DF		Default Queue + RED	Default Class
Scavenger	CS1		Min BW Queue (Deferential)	YouTube, iTunes, BitTorrent, Xbox Live

Cisco Swapped CS5 and CS3

Figure 16-4 QoS Marking—RFC 4594

RFC 4594 is the recommendation but not the standard; it resides in the category of draft proposal RFCs. It recommends guidelines on how to configure 14 traffic classes that are associated with 28 different code-point marking values. Note that some of the PHBs shown in Figure 16-4 include multiple DSCP-associated values. For example, the AF class for multimedia streaming can have AF31, AF32, and AF33 DSCP values. RFC 4594 includes information on which PHBs should be used for certain traffic types and also what queuing and dropping mechanism should be used for that same traffic class.

Some sample recommendations highlighted in Figure 16-4 include

- Voice traffic should be marked to EF/DSCP 46.
- Voice should be queued using strict-priority queuing.
- Broadcast video traffic should be marked to CS5/DSCP 40.
- Multimedia conferencing should be treated with an AF PHB, provisioned with a guaranteed-bandwidth queue.

RFC 4594 is not a final RFC standard and will more than likely continue to be developed considering that needs and trends for QoS application requirements change over the time.

QoS Strategy Models

Before applying any QoS tools, organizations need to define the strategy and goals for different applications running in their network. This will result in defining a certain number of traffic classes to meet the end-to-end QoS objectives of an organization.

Three basic QoS strategy models can be deployed, depending on the granularity of applications running within an organization's network:

- 4-Class QoS Strategy Model
- 8-Class QoS Strategy Model
- 12-Class QoS Strategy Model

Although the more classes you define, the more specific and granular traffic treatment will be per application, the selection of a certain strategy model must be based on application requirements coupled with the WAN provider QoS model (if there is any WANs with QoS). The following sections provide a detailed view into each of these QoS strategy models.

4-Class QoS Strategy

The 4-class QoS strategy model is the simplest of the three models (in terms of QoS polices) and typically accounts for telephony, signaling, transactional/mission-critical, and best-effort data. When businesses deploy telephony applications in their network, three classes of traffic are typically required (telephony, signaling, and default/best effort).

Typically, the fourth class is the Assured Forwarding (AF) class. The AF class is used for transactional and mission-critical data applications such as SQL databases. The AF class can also be used for multimedia conferencing, multimedia streaming, and bulk data applications.

The 4-class QoS strategy model, as shown in Figure 16-5, is an example of where an organization has deployed IP telephony. In addition to separating telephony, signaling, and default/best-effort traffic, the organization has defined one mission-critical transactional data class.

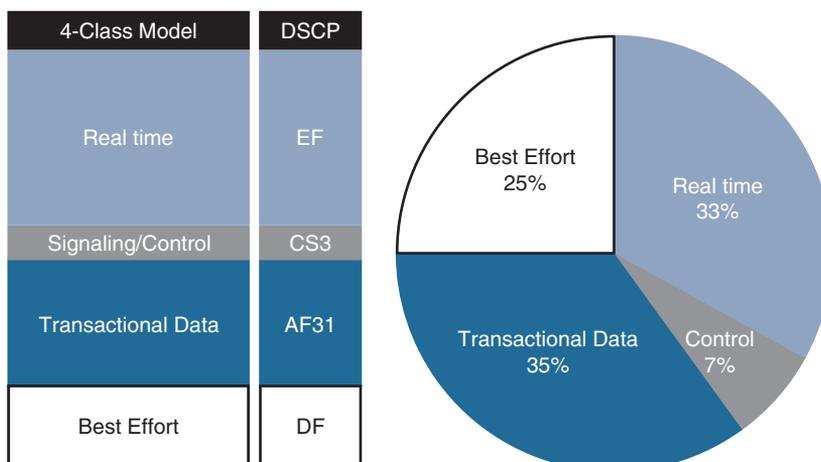


Figure 16-5 *The 4-Class QoS Strategy Model*

The four traffic classes of QoS markings and guarantees are as follows:

- **Voice (Real time):** Marked with EF and provisioned to leverage up to one-third of link bandwidth
- **Signaling:** Marked with CS3 and provisioned to leverage a minimum of 7 percent of link bandwidth
- **Mission-critical data (Transactional Data):** Marked with AF31 and provisioned to leverage 35 percent of link bandwidth
- **Default (best-effort data):** Marked with DF and provisioned to take advantage of 25 percent of link bandwidth

Voice and signaling guarantees must be selected based on the volume of voice calls and the VoIP codec that is used through the given link. Mission-critical data is selected based on the decision of the director of each company department who has given info about critical business application needs to the networking team.

8-Class QoS Strategy

The 8-class QoS strategy model builds upon the 4-class model and includes the following additional classes:

- Multimedia conferencing
- Multimedia streaming
- Network control
- Scavenger

The two additional multimedia traffic types in this model are multimedia conferencing and multimedia streaming. The explicitly defined network control traffic class is used for applications such as network routing protocol updates or network infrastructure control traffic such as OAM. The 8-class QoS strategy model is illustrated in Figure 16-6.

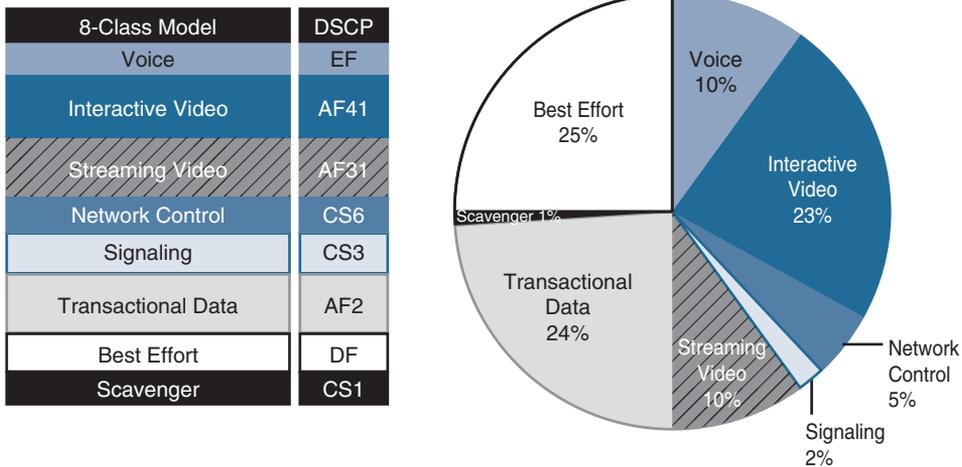


Figure 16-6 *The 8-Class QoS Strategy Model*

As can be seen from Figure 16-6, the recommendations for each traffic class in this model are as follows:

- **Voice:** Marked with EF and limited to 10 percent of link bandwidth in a strict-priority queue
- **Multimedia conferencing (Interactive video):** Marked with AF41 or sometimes as EF and limited to 23 percent of link bandwidth in a strict-priority queue
- **Multimedia streaming:** Marked with AF31 and guaranteed 10 percent of link bandwidth with WRED enabled
- **Network control:** Marked with CS6 and guaranteed 5 percent of link bandwidth
- **Signaling:** Marked with CS3 and provisioned with minimum of 2 percent of link bandwidth
- **Transactional data:** Marked with AF21 and provisioned with 24 percent of link bandwidth with WRED enabled
- **Default (best-effort data):** Marked with DF and provisioned with 25 percent of link bandwidth
- **Scavenger:** Marked with CS1 and provisioned with a maximum of 1 percent of link bandwidth

Note It is important to note the difference as some traffic types, such as voice traffic, are limited by bandwidth defined in a strict-priority queue, and other traffic types, such as multimedia streaming, have guaranteed provisioned bandwidth.

12-Class QoS Strategy

The 12-class QoS strategy model builds upon the 8-class model and includes the following additional classes:

- Real-time Interactive
- Broadcast Video
- Management/OAM
- Bulk Data

The 12-class QoS strategy model represents Cisco's interpretation of the RFC 4594 recommendation and, as previously noted, incorporates a slight modification by swapping the markings used for signaling and broadcast video. The 12-class QoS strategy model is illustrated in Figure 16-7.

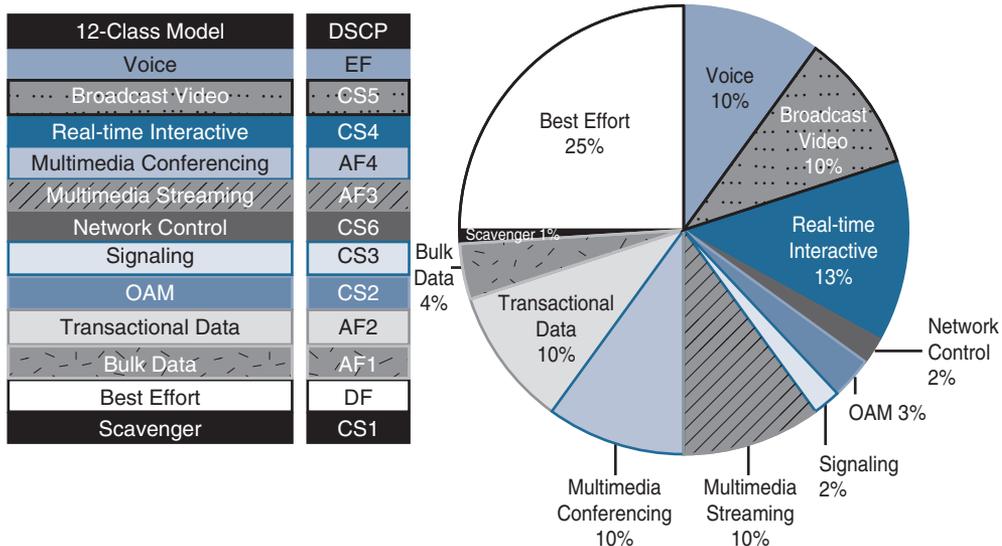


Figure 16-7 The 12-Class QoS Strategy Model

As can be seen from Figure 16-7, the recommendations for each traffic class in this model are as follows:

- **Voice:** Marked with EF and limited to 10 percent of link bandwidth in a strict-priority queue
- **Broadcast video:** Marked with CS5 or sometimes as EF and limited to 10 percent of link bandwidth in a strict-priority queue
- **Real-time interactive:** Marked with CS4 or sometimes as EF and limited to 13 percent of link bandwidth in a strict-priority queue

- **Multimedia conferencing:** Marked with AF41 or sometimes as EF and limited to 10 percent of link bandwidth in a strict-priority queue
- **Multimedia streaming:** Marked with AF31 and guaranteed 10 percent of link bandwidth with WRED enabled
- **Network control:** Marked with CS6 and provisioned as guaranteed bandwidth 2 percent of link bandwidth
- **Signaling:** Marked with CS3 and provisioned with a minimum of 2 percent of link bandwidth
- **Management/OAM:** Marked with CS2 and provisioned with a minimum of 3 percent of link bandwidth
- **Transactional data:** Marked with AF21 and provisioned with 10 percent of link bandwidth with WRED enabled
- **Bulk data:** Marked with AF11 and provisioned with 4 percent of link bandwidth with WRED enabled
- **Default (best-effort data):** Marked with DF and provisioned with 25 percent of link bandwidth
- **Scavenger:** Marked with CS1 and provisioned with a maximum of 1 percent of link bandwidth

Summary

- Use QoS policies in hardware rather than in software whenever possible.
- Classify, mark, and police applications as close to the source as possible.
- Use DSCP marking whenever possible.
- Define a queue for the traffic class and enable queuing on each node that has potential congestion.
- Limit the strict-priority queue to one-third of the link bandwidth.
- Do not use WRED for priority or scavenger traffic classes.
- Use one of the three QoS strategy models to govern end-to-end QoS design.

Review Questions

After answering the following questions, please refer to Appendix A, “Answers to Review Questions,” for the answers.

1. Which of the following is recommended for a QoS queuing design?
 - a. You should implement queuing policy very selectively.
 - b. Classes should share queues in order to save resources.

- c. You should use at minimum 4 classes of queuing behavior.
 - d. You should use at minimum 11 classes of queuing behavior.
- 2.** Match the application classes with their PHBs as per RFC 4594.
- | | |
|-----------------------|------|
| VoIP Telephony | EF |
| Transactional Data | CS1 |
| Network Control | CS6 |
| Call Signaling | CS4 |
| Real-time Interactive | AF21 |
- 3.** Select the four classes of the 4-class QoS model.
- a. Voice, signaling, mission-critical data, and best effort
 - b. Video, signaling, mission-critical data, and best effort
 - c. Voice, signaling, mission-critical data, and scavenger
 - d. Real-time interactive, signaling, mission-critical data, and best effort
- 4.** Why is it recommended to leverage DSCP markings wherever possible?
- a. Support for end-to-end Layer 3 marking.
 - b. It is a more granular method of marking that supports 64 levels as compared to CoS and MPLS EXP, which have 8 levels.
 - c. It is more extensible than Layer 2 markings because these markings are lost when media change.
 - d. All the above.
 - e. None of the above.
- 5.** Traffic should be marked down using which RFC recommendations? (Select two.)
- a. RFC 2957
 - b. RFC 2597
 - c. RFC 2698
 - d. RFC 2968

Index

Numbers

3G/4G VPN design models, 335
4-class QoS strategy model, 561–562
6RD (6 Rapid Deployment), IPv6,
210–211
6RD border relay, 210
6RD prefix, 211
6RD-capable router, 210
8-class 1P1Q3T egress queueing,
581–588
8-class 1P1Q3T ingress queueing,
580–581
8-class QoS strategy model, 562–563
12-class QoS strategy model,
564–565
/40 prefix, 197
/48 prefix, 198
/56 prefix, 198
802.1p, QoS (quality of service),
517–519
802.1Q, 27
 QoS (quality of service), 517–519
802.1X, 759–763

message flow, 763
phased deployment, 767
supplicants, 765–766

Symbols

*, G (star comma G), 644, 645
PIM-SM (Protocol-Independent
Multicast—Sparse Mode), 653

A

ABR placement, hub-and-spoke
design, OSPF (Open Shortest Path
First), 89–90
access control lists (ACLs), 702
access coverage, WAN connections,
232
access layer, enterprise campus
design, 4–5
access management, enterprise WAN,
367–368
access restrictions, 740
access-distribution block, enterprise
campus design, 13–15

- ACI (Application-Centric Infrastructure), 431**
 - ANP (application network profile), 449, 459–460
 - application design, 459–460
 - architecture, 434
 - APIC (Application Policy Infrastructure Controller), 434–437*
 - fabric, 437–440*
 - characteristics, 432
 - EPG (endpoint groups), 450–453
 - external Layer 2 connections and networks, 461–465
 - fabric access policies, 454–455
 - integration and migration connectivity options, 471–473
 - network virtualization overlays, 441–446
 - networking limitations, 432–434
 - route propagation inside ACI fabric, 468–470
 - routing, 465
 - border leaves, 467–468*
 - first-hop layer 3 default gateway, 465–466*
 - STP-based layer LANs, 464–465
 - tenants, 456–459
- ACI APIC cluster, 440**
- ACI fabric**
 - connecting to external Layer 3 domains, 470–471
 - route propagation, 468–470
- ACI policy model, application design, 447–450**
- ACLs (access control lists), 702**
- acquiring IPv6 prefixes, 197–198**
- active passive failover mode, ASA firewall, 722**
- active/active mode, firewalls, 722**
- adaptive security appliance (ASA), 696**
- Adaptive Security Virtual Appliance (ASAv), 713–714**
- additive keyword, 177**
- Address Family Translation (AFT), 206**
- address provisioning, 814**
- addressing**
 - enterprise IPv6 networks case study, 813–814
 - IS-IS (Intermediate System-to-Intermediate System), 114–116
- addressing services, IPv6, 220–221**
- adjacencies, IS-IS (Intermediate System-to-Intermediate System), 108–109, 120**
- adjacent neighbors, OSPF (Open Shortest Path First), 76–77**
- AF (Assured Forwarding), 561**
- AF drop probability, 521**
- AF PHB, 521**
- AF profiles, 546**
- AFT (Address Family Translation), 206**
- aggregation layer deployment model, DCI (Data Center Interconnect), 499**
- aggressive mode, IKE (Internet Key Exchange), 279**
- AH (Authentication Header), 278**
- algorithms, token bucket algorithms, 529–531**
- analysis and task list**
 - enterprise BGP network with Internet connectivity case study, 791
 - enterprise data center connectivity case study, 818

- enterprise IPv6 networks case study, 809
- resilient enterprise WANs case study, 826–827
- analyzing enterprise connectivity, 779–780
- ANP (application network profile), 449
 - ACI (Application-Centric Infrastructure), 459–460
- anti-replay window sizing, 630
- Any Transport over MPLS over GRE (AToMoGRE), DCI (Data Center Interconnect), 497–498
- Anycast RP, 681
 - examples, 682–683
 - MSDP (Multicast Source Discovery Protocol), 683
- AnyConnect Secure Mobility Client, 623, 765–766
- APIC (Application Policy Infrastructure Controller), 357–358, 434–437, 439
- APIC-EM (Application Policy Infrastructure Controller Enterprise Module), 357–358, 368–370
 - design, 370–371
- application adaptation, IPv6, 223
- application design
 - ACI (Application-Centric Infrastructure), 459–460
 - ACI policy model, 447–450
- application migration, enterprise IPv6 networks case study, 815–816
- application network profile (ANP), 449
- application optimization, WAN, 356–357
- Application Policy Infrastructure (APIC), 357–358
- application support, IPv6, 222–223
 - application adaptation, 223
 - application workarounds, 223–224
- application tiers, separating, 714–716
- Application Visibility Control (AVC), 357
- application workarounds, IPv6, 223–224
- Application-Centric Infrastructure. *See* ACI (Application-Centric Infrastructure)
- application-specific integrated circuits (ASIC), 554
- architecture
 - ACI (Application-Centric Infrastructure), 434
 - APIC (*Application Policy Infrastructure Controller*), 434–437
 - fabric*, 437–440
 - big data architecture, data center QoS, 596
 - EAP (Extensible Authentication Protocol), 763–764
 - firewalls, 709–712
 - FlexVPN, 315
 - hierarchical architecture, IS-IS (Intermediate System-to-Intermediate System), 105–106
 - HPT (high-performance trading), data center QoS, 595
 - IPS (intrusion prevention system), 726–729
 - modular network architecture, 691–695
 - zones*, 695
 - MPLS VPNs, 234–236
 - multilayer architectures, EIGRP (Enhanced Interior Gateway Routing Protocol), 53–56

- new network architecture, 397–398
- ONE (Open Network Environment) architecture, 435
- provider edge (PE) routers, 237–238
 - route distinguishers*, 238–239
 - route target (RT)*, 240–241
- three-layer hierarchy architecture,
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 57–59
- three-tier data center network architecture, 380–381
- two-layer hierarchy architecture,
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 56–57
- two-tier data center network architecture, 378–380
- virtualized multiservice architectures, 596–597
- area, OSPF (Open Shortest Path First)**
 - number of areas per ABR, 81–82
 - numbers of routers in an area, 80–81
 - routing information, 78–80
- area design**
 - IS-IS (Intermediate System-to-Intermediate System), 113
 - OSPF (Open Shortest Path First), 82–83, 112–113
- ARP inspection, 702**
- AS (autonomous systems), EIGRP (Enhanced Interior Gateway Routing Protocol), 50–52**
 - multiple autonomous system drivers, 53
- AS (autonomous systems) number**
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 243–244
 - PE-CE routing protocol, 242–243
- ASA (adaptive security appliance), 696, 712**
 - FirePOWER services, 727
- ASA 1000V, 714**
- ASA clustering, 723**
- ASA firewall active/passive failover mode, 722**
- ASA SFR, 726–727**
- ASAv (Adaptive Security Virtual Appliance), 713–714**
- ASBRs (autonomous system border routers), 79**
- Asian sites, routing policies, 799–802**
- ASIC (application-specific integrated circuits), 554**
- as-override, 254**
- assessment phase, IPv6, 196**
- asymmetric routing versus symmetric routing, IS-IS (Intermediate System-to-Intermediate System), 129–132**
- asymmetrical routing issues, GLBP (Gateway Load Balancing Protocol), 34**
- ATM WAN design, 344–346**
- AToMoGRE (Any Transport over MPLS over GRE), 497–498**
- attacks**
 - multicast traffic, 753
 - preventing, 703
- attributes, BGP (Border Gateway Protocol)**
 - extended community attributes, 241–242
 - path attributes, 150
- authentication, 740**
- Authentication Header (AH), 278**
- authentication servers, 760**
- authenticators, 760**

authorization, 740
 authorization options case study,
 772–775
 autonomous system border routers
 (ASBRs), 79
 autonomous system numbers,
 choosing, 792–794
 autonomous systems. *See* AS
 (autonomous systems)
 Auto-RP, 667, 668–669
 candidate RPs, 670
 case studies, 670–674
 mapping agents, 670
 multicast network edge security,
 749–751
 operations, 671–674
 routers, 670
 scope problems, 674–676
 AVC (Application Visibility Control),
 357
 A-VPLS (Advanced VPLS), 496

B

backdoor links between customer
 sites, PE-CE routing protocol
 BGP (Border Gateway Protocol),
 254–255
 EIGRP (Enhanced Interior Gateway
 Routing Protocol), 245–247
 OSPF (Open Shortest Path First),
 250–251
 backoff messages, DF election
 messages, 660
 backoff timers, 94
 bandwidth allocation, 558–559
 bandwidth keyword, 539
 baseline network policy enforcement,
 701–702
 baseline switching security, 702
 bestpath as-path multipath-relax, 183
 BFD (bidirectional forwarding
 detection), EIGRP (Enhanced
 Interior Gateway Routing
 Protocol), 70–71
 BFD echo, 71
 BGP (Border Gateway Protocol), 146
 case studies, 172–177
 communities, 169–170
 named communities, 171
 planning for, 171–172
 well-known BGP communities,
 170–171
 confederations, 155–156
 versus route reflectors, 157
 dual-homing, 178
 extended community attributes,
 241–242
 load-sharing design, 177
 single-homing versus multi-
 homing, 177–178
 loop prevention, 148–149
 multihoming, 178
 overview, 146–147
 path attributes, 150
 path selection, 150–151
 PE-CE routing protocol, 252–254
 backdoor links between
 customer sites, 254–255
 peer-forwarding rules, 158
 route reflectors, 153–155
 congruence of physical and
 logical networks, 165–167
 hierarchical route reflector
 design, 167–168
 loop prevention, 162–165
 network design issues, 169

- redundancy*, 159–160
 - route reflector cluster-ID*, 161–162
 - route reflector clusters*, 160–161
 - split-horizon rule*, 158–159
 - single-homed, multiple links, 178–180
 - speaker types, 147–148
 - split-horizon rule, 148–149
 - traffic engineering techniques, 352–353
 - TTL Security Check, 700
 - bgp always-compare-med**, 151
 - BGP ASN design**, 792–794
 - bgp bestpath med missing-as-worst**, 151
 - BGP communities**, 796–797
 - BGP connectivity**
 - BGP communities, 796–797
 - BGP sessions, 795–796
 - BGP Originator-ID attribute**, 162
 - BGP sessions**, 795–796
 - bidirectional forwarding detection (BFD)**, EIGRP (Enhanced Interior Gateway Routing Protocol), 70–71
 - BIDIR-PIM (bidirectional PIM)**, 657, 754
 - DF election, 658–659
 - DF election messages, 660
 - PIM modifications, 658
 - big data architecture**, data center QoS, 596
 - black holes**, route summarization, EIGRP (Enhanced Interior Gateway Routing Protocol), 61–63
 - bootstrap router (BSR)**, 667
 - Border Gateway Protocol**. *See* BGP (Border Gateway Protocol)
 - border leaf devices**, 439
 - border leaves**, ACI (Application-Centric Infrastructure), 467–468
 - boundaries**, trust states and, 570–573
 - branch border routers**, 366
 - branch master controller**, 366
 - branch offices**, remote-site WAN design, 346–348
 - branch sites**, connecting, 810–812
 - bridge domains**, tenants, ACI (Application-Centric Infrastructure), 456–457
 - broadcast links**, IS-IS (Intermediate System-to-Intermediate System), 119
 - BSR (bootstrap router)**, 667
 - multicast network edge security, 749–751
 - PIMv2, 676–677
 - PIMv2 BSR, 678
 - securing, 751
 - buffering**, 535
 - buffers**, QoS (quality of service), 569–570
 - building a secure campus edge design (Internet and extranet connectivity) case study**, 729–740
 - bursts**, QoS (quality of service), 569–570
-
- ## C
-
- Campus Edge network**, 730–736
 - characteristics, 730–731
 - DMZs (demilitarized zones), 732–733
 - firewalls, 731–735

- internal networks, connecting, 733–734
- Internet, connecting, 731
- campus network virtualization, 16–23**
 - path isolation, 19–23
 - VLAN assignment, 17–18
 - VRF (virtual routing and forwarding), 18
- campus QoS, 568**
 - design examples, 576–588
- candidate RPs, 676–677**
 - Auto-RP, 670
 - PIMv2 BSR, 677–678
- candidate-RP announce packets, 750**
- candidate-RP discovery packets, 750**
- capabilities, FlexVPN, 315**
- case studies**
 - authorization options, 772–775
 - Auto-RP operation, 670–674
 - building a secure campus edge design (Internet and extranet connectivity), 729–740
 - dark fiber DCI, 490–494
 - DC QoS application, 599–601
 - design enterprise BGP network with Internet connectivity, 788
 - analysis and task list, 791*
 - BGP connectivity, 795–797*
 - choosing autonomous system numbers, 792–794*
 - choosing routing protocols, 792*
 - Internet routing, 803–807*
 - requirements and expectations, 788–791*
 - routing policies, 797–802*
 - design enterprise connectivity, 778
 - analysis and task list, 779–780*
 - designing for new routing protocols, 780–782*
 - migrating from old to new routing, 785–787*
 - OSPF design optimization, 782–785*
 - requirements and expectations, 778–779*
 - scaling, 787–788*
 - selecting replacement routing protocols, 780*
 - design enterprise data center connectivity, 816–817
 - analysis and task list, 818*
 - connecting network appliances, 821–822*
 - data center interconnect, 822–823*
 - data center network virtualization design, 823–825*
 - DCN detailed connectivity, 819–821*
 - requirements and expectations, 817–818*
 - selecting architecture and connectivity model, 818–819*
 - design enterprise IPv6 network, 807
 - addressing, 813–814*
 - analysis and task list, 809*
 - application and service migration, 815–816*
 - choosing IP address types for HQ, 809–810*
 - communication between branches, 815*
 - connecting branch sites, 810–812*

- deployment models*, 812
- requirements and expectations*, 808–809
- design QoS in the enterprise network, 835
 - congestion management*, 838–839
 - MPLS WAN DiffServ tunneling*, 839–841
 - QoS design model*, 837–838
 - QoS trust boundary*, 838
 - requirements and expectations*, 835–836
 - scavenger traffic*, 839
 - traffic discovery and analysis*, 836–837
- design resilient enterprise WANs, 825
 - analysis and task list*, 826–827
 - requirements and expectations*, 825–826
 - selecting WAN links*, 828
 - WAN overlays*, 828–830
- design secure enterprise networks, 830
 - firewalls*, 835
 - infrastructure and network access security*, 833–834
 - Layer 2*, 834–835
 - requirements and expectations*, 831
 - security domains and zone design*, 832
- designing enterprisewide BGP policies using BGP communities, 172–177
- DF election, 660–662
- EIGRP DMVPN, 295–302
- firewall high availability, 720–725
- implementing firewalls in a data center, 717–720
- MPLS VPN routing propagation, 255–258
- MPLS/VPN over GRE/DMVPN, 304–312
- MSDP operations, 684–686
- multitenant data centers, 425–426
- redundancy and connectivity, 343–354
- RPF check fails and succeeds, 641–642
- separation of application tiers, 714–716
- small data centers (connecting servers to an enterprise LAN), 376–378
- three-tier data center network architecture, 380–381
- two-tier data center network architecture, 378–380
- virtualized multiservice architectures, 596–597
- Catalyst switches**, 554, 571, 574
- CBWFQ (class-based weighted fair queuing)**, 536, 538–541, 591
 - WAN/branch edge, 592
- cellular connectivity**, 335
- CGA (cryptographically generated access)**, 222
- challenges of SDN (software-defined networking)**, 419–421
- characteristics**
 - ACI (Application-Centric Infrastructure), 432
 - Campus Edge network, 730–731
 - DiffServ, 516
 - ECN (explicit congestion notification), 550
 - IntServ (Integrated Services), 516

- IS-IS (Intermediate System-to-Intermediate System), 103–104, 110–112
- OSPF (Open Shortest Path First), 110–112
- PIM-SM (Protocol-Independent Multicast—Sparse Mode), 645
- SDN controller characteristics, 418
- SSM (source-specific multicast), 654
- traffic policing, 529
- traffic shaping, 529
- choke points**
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 54
 - summarization and, 55–56
- choosing**
 - autonomous system numbers, 792–794
 - WAN connections, 230–233
- CIR (committed information rate), 530**
- Cisco AnyConnect Secure Mobility client, 765–766**
- Cisco Application-Centric Infrastructure. *See* ACI (Application-Centric Infrastructure)**
- Cisco ASA 5500-X Series Next-Generation Firewall, 696**
- Cisco ASA 5506-X, 696**
- Cisco ASA 5512-X, 696**
- Cisco ASA 5555-X, 696**
- Cisco FabricPath, 402–407**
- Cisco FirePOWER, NGIPS (next-generation IPS), 696, 726–727**
- Cisco Identity Services Engine (ISE), 768**
- Cisco IOS, encryption, 623–625**
- Cisco IOS XR software, 750**
- Cisco modular network architecture, 691–695**
- Cisco next-generation security, 696**
- Cisco Security Group Tag (SGT), 769–772**
- Cisco TrustSec, 768**
 - Profiling Service, 768–769
 - SGT (Security Group Tag), 769–772
- Cisco Web Security Appliance (WSA), 735–736**
- cLACP (Cluster Link Aggregation Control Protocol), 724**
- class-based weighted fair queueing (CBWFQ), 536, 538–541**
- classification, QoS (quality of service), order of operations, 623–625**
- classification and marking, QoS (quality of service)**
 - design principles, 554–555
 - Layer 2 marking, 517–519
 - Layer 2.5 marking: MPLS experimental bits, 524
 - Layer 3 marking: DSCP per-hop behaviors, 520–523
 - Layer 3 marking: IP type of service, 519–520
 - Layer 7: NBAR/NBAR2, 526–527
 - mapping markings between OSI layers, 524–525
 - traffic policing and shaping, 527–529, 532
- classification/marketing/policing QoS model, 573–574**
- classifications and marking tools, QoS (quality of service), 516–517**
- client-server traffic, 479**

- CLNP (Connectionless Network Protocol), 102
- CLNS (Connectionless Network Service), 102
- Cluster ID, 164–165
- Cluster Link Aggregation Control Protocol (cLACP), 724
- Cluster-List attribute, 163
- committed information rate (CIR), 530
- communication between branches, enterprise IPv6 networks case study, 815
- communities, BGP (Border Gateway Protocol), 169–170, 796–797
 - named communities, 171
 - planning for, 171–172
 - well-known BGP communities, 170–171
- comparing
 - 802.1X deployment modes, 767
 - control planes and data planes, 414–415
 - DMVPN (Dynamic Multipoint VPN) *and GET VPN*, 629
 - phases*, 302
 - EF and AF profiles, 546
 - enterprise campus access-distribution design models, 45
 - IntServ and DiffServ, 514–516
 - MSDP and BGP features, 752
 - point-to-point GRE and multipoint GRE, 276–277
 - QoS design drivers and considerations based on the PIN, 602
 - RP deployments, 667
 - traffic shaping and traffic policing, 529
 - virtual firewall models, 714
 - VPLS and VPWS, 266–267
- complete sequence numbers (CSNP), 123–124
- confederations, BGP (Border Gateway Protocol), 155–156
 - versus BGP route reflectors, 157
- configuration blocks, FlexVPN, 315–316
- congestion avoidance, 541, 575
- congestion management, QoS in the enterprise network case study, 838–839
- congruence of physical and logical networks, route reflectors, BGP (Border Gateway Protocol), 165–167
- connecting
 - ACI fabric to external Layer 3 domains, 470–471
 - ACI to outside Layer 2 domains, 462–465
 - branch sites, 810–812
 - external partners, 737
 - internal networks, Campus Edge network, 733–734
 - Internet, Campus Edge network, 731
 - network appliances, 821–822
 - servers to enterprise LANs, 376–378
- Connectionless Network Protocol (CLNP), 102
- Connectionless Network Service. *See* CLNS (Connectionless Network Service)
- connectivity, case studies, redundancy and connectivity, 343–354
- connectivity model, MPLS VPNs, 606
- content and application security, 695

contracts, 449

control plane, 697

control plane optimization, VXLAN (virtual extensible LAN), 413–414

control plane policing (CoPP), 747

control plane protection, 697

control plane security, 414–415
IPv6, 224

convergence
EtherChannel convergence, 28
OSPF (Open Shortest Path First), 93
event detection, 94
event processing, 96–97
event propagation, 94–96
WAN connections, 231

CoPP (control plane policing), 747

core layer, enterprise campus design, 6–7

core layer deployment model, DCI (Data Center Interconnect), 499

CQ (custom queueing), 536

critical VLANs, 773

cryptographically generated access (CGA), 222

CSNP (complete sequence number), 123–124

custom queueing (CQ), 536

customer edge (CE) routers, 235

customer-managed Layer 2 DCI deployment models, 497
aggregation layer deployment model, 499
Any Transport over MPLS over GRE (AToMoGRE), 497–498
core layer deployment model, 499
limitations of, 501

overlay transport virtualization DCI, 501–506
separate DCI layer deployment model, 500

CWDM, 490

D

dark fiber DCI, 490–494

data center bridging toolset, 597–598

Data Center Interconnect. *See* DCI (Data Center Interconnect)

data center network virtualization design, 823–825

data center QoS
big data architecture, 596
data center bridging toolset, 597–598
DC QoS application case study, 599–601
HPT (high-performance trading), 595
overview, 594
virtualized multiservice architectures, 596–597

data center traffic flows
DCI (Data Center Interconnect). *See* DCI (Data Center Interconnect)
traffic flow directions, 478–479
traffic flow types, 479–482

data centers
case studies, implementing firewalls in a data center, 717–720
end of row versus top of rack design, 383–384
fabric extenders, 385–388
high availability, 388–392
interconnecting, 822–823
inter-VLAN routing, 381–383
modern data centers. *See* modern data centers

- new network architecture, 397–398
- NIC teaming, 392–393
- small data centers (connecting servers to an enterprise LAN), 376–378
- three-tier data center network architecture, 380–381
- two-tier data center network architecture, 378–380
- data flow, IS-IS (Intermediate System-to-Intermediate System), 118–119**
- data plane, 414–415, 697**
- data plane protection, 697**
- Database Overload Protection, OSPF (Open Shortest Path First), 97–98**
- DC QoS application, 599–601**
- DCB (Data Center Bridging) toolset, 597–598**
- DCI (Data Center Interconnect), 482–483**
 - customer-managed Layer 2 DCI deployment models, 497
 - aggregation layer deployment model, 499*
 - Any Transport over MPLS over GRE (AToMoGRE), 497–498*
 - core layer deployment model, 499*
 - limitations of, 501*
 - overlay transport virtualization DCI, 501–506*
 - separate DCI layer deployment model, 500*
 - dark fiber DCI, 490–494
 - IP address mobility, 484–490
 - Layer 3, 507–509
 - LISP (locator/ID separation protocol), 487–489
 - overlay networks, 507
 - pseudowire DCI, 495
 - virtual private LAN service DCI, 496
- DCN connectivity, enterprise data center connectivity, 819–821**
- DCN connectivity model, 820**
- decision process, IS-IS (Intermediate System-to-Intermediate System), 119**
- default forwarding (DF), 558**
- default routing, 805–807**
- default VLANs, 773**
- delays, jitter and latency, WAN QoS, 590–591**
- demilitarized zones (DMZs), 710**
- dense mode protocols, 642**
- deployment**
 - IPv6, 194–195
 - assessment phase, 196*
 - discovery phase, 196*
 - implementation and optimization phases, 197*
 - planning and design phase, 196–197*
 - PfRv3, 366–367
 - phased deployment, 802.1X, 767
- deployment models**
 - DHCPv6 deployment model, 814
 - DMVPN (Dynamic Multipoint VPN), 285
 - enterprise IPv6 networks, case study, 812
- design**
 - APIC-EM (Application Policy Infrastructure Controller Enterprise Module), 370–371
 - campus QoS, examples, 576–588
 - IPv6, 194–195

- assessment phase, 196*
- discovery phase, 196*
- implementation and optimization phases, 197*
- planning and design phase, 196–197*
- link aggregation of EtherChannel interface, 575–576
- designated forwarder (DF), BIDIR-PIM (bidirectional PIM), 658**
- designing**
 - enterprise BGP network with Internet connectivity, 788
 - analysis and task list, 791*
 - BGP connectivity, 795–797*
 - choosing autonomous system numbers, 792–794*
 - choosing routing protocols, 792*
 - Internet routing, 803–807*
 - requirements and expectations, 788–791*
 - routing policies, 797–802*
 - enterprise connectivity, 778
 - analysis and task list, 779–780*
 - designing for new routing protocols, 780–782*
 - migrating from old to new routing, 785–787*
 - OSPF design optimization, 782–785*
 - requirements and expectations, 778–779*
 - scaling, 787–788*
 - selecting replacement routing protocols, 780*
 - enterprise data center connectivity, 816–817
 - analysis and task list, 818*
 - connecting network appliances, 821–822*
 - data center interconnect, 822–823*
 - data center network virtualization design, 823–825*
 - DCN detailed connectivity, 819–821*
 - requirements and expectations, 817–818*
 - selecting architecture and connectivity model, 818–819*
 - enterprise IPv6 networks, 807
 - addressing, 813–814*
 - analysis and task list, 809*
 - application and service migration, 815–816*
 - choosing IP address types for HQ, 809–810*
 - communication between branches, 815*
 - connecting branch sites, 810–812*
 - deployment models, 812*
 - requirements and expectations, 808–809*
 - infrastructure protection, 696–697
 - for new routing protocols, 780–782
 - QoS in the enterprise network case study, 835
 - congestion management, 838–839*
 - MPLS WAN DiffServ tunneling, 839–841*
 - QoS design model, 837–838*
 - QoS trust boundary, 838*
 - requirements and expectations, 835–836*
 - scavenger traffic, 839*
 - traffic discovery and analysis, 836–837*

- resilient enterprise WANs, 825
 - analysis and task list*, 826–827
 - requirements and expectations*, 825–826
 - selecting WAN links*, 828
 - WAN overlays*, 828–830
- secure enterprise networks, 830
 - firewalls*, 835
 - infrastructure and network access security*, 833–834
 - Layer 2 security considerations*, 834–835
 - requirements and expectations*, 831
 - security domains and zone design*, 832
- device profiling, 769
- device resiliency, 24
- device-level virtualization, separation, 424–425
- DF (default forwarding), 558
- DF (designated forwarder),
 - BIDIR-PIM (bidirectional PIM), 658
 - DF election, 658–659
 - DF election messages, 660
- DF election, case studies, 660–662
- DF election messages, BIDIR-PIM (bidirectional PIM), 660
- DHCP snooping, 702
- DHCPv6, 220
- DHCPv6 deployment model, 814
- DiffServ (Differentiated Services), 515–516
- discovery phase, IPv6, deployment and design, 196
- Distance Vector Multicast Routing Protocol (DVMRP), 756
- distribution layer, enterprise campus design, 5–6
- distribution-to-distribution interconnect
 - multitier access model, 37–41
 - routed access model, 41–42
 - virtual switch model, 43–44
- distribution-to-distribution link design, 36–37
- DMVPN (Dynamic Multipoint VPN), 621
 - benefits of, 286
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 69
 - limitations of, 287
 - overview, 283–287
 - Phase 1, 287–289
 - EIGRP*, 295–297
 - Phase 2, 289–292
 - EIGRP*, 297–299
 - Phase 3, 292–295
 - EIGRP*, 299–301
 - QoS (quality of service), 626–628
 - redundancy, 302–304
 - VPN WAN design models, 331–333
- DMZs (demilitarized zones), 710
 - Campus Edge network, 732–733
- DNS64, IPv6, 206–208
- domains, IS-IS (Intermediate System-to-Intermediate System), 104
- drop probability, 543
 - DSCP, 522
- dropping design principles, QoS (quality of service), 557–558
- dropping modes, RED (random early detection), 543–544
- dropping recommendations, QoS (quality of service), 574–575

dropping tools, DSCP-based WRED, 541–546

DSCP (Differentiated Services Code Point)

- drop probability, 522
- IP precedence mapping, 523
- markings, 555

DSCP MPLS EXP bits, 611

DSCP-based WRED, 541–546

DS-Lite, IPv6, 211–212

dual domains, 104

dual IS-IS, 104–105

dual stack, IPv6, 205–206

dual-bucket policing, 532–533

dual-homed to one ISP using a single local edge router, 180–181

dual-homed to one ISP using multiple edge routers, 182–183

dual-homing, 178

dual-rate metering. *See* policing tools

Dual-Stack Lite, IPv6, 211–212

dual-stack security, IPv6, 225

DVMRP (Distance Vector Multicast Routing Protocol), 756

DVTI (Dynamic VTI), IPsec and, 283

DWDM, 490

Dynamic Multipoint VPN. *See* DMVPN (Dynamic Multipoint VPN)

dynamic trust states, 572–573

dynamic VLAN assignments, 772–774

Dynamic VTI (DVTI), IPsec and, 283

E

EAP (Extensible Authentication Protocol), 762, 763–765

- types of, 764–765

EAP chaining, 765

EAP method, 762

EAP over LAN (EAPOL), 762

EAP-Chaining, 766

EAP-FASTv2 (Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling), 765

EAPOL (EAP over LAN), 762, 763

east-west traffic, 478

- securing, 716–717

eBGP (external BGP), 151

ebgp multihop, 179

ECN (explicit congestion notification), 520, 547–550

- characteristics, 550
- operations, 549
- WRED, 548–549

e-commerce, 693

edge routers

- dual-homed to one ISP using a single local edge router, 180–181
- dual-homed to one ISP using multiple edge routers, 182–183
- multihoming with two ISPs using a single local edge router, 183–186
- multihoming with two ISPs using multiple local edge routers, 186–188

EF PHB, 521

EF profiles, 546

EF traffic, 546

EFC (Ethernet Flow Control), 598

EGP (Exterior Gateway Protocol), 146

egress tunnel router (ETR), 213

EIGRP (Enhanced Interior Gateway Routing Protocol), 49–50

- AS (autonomous systems), 50–52
- BFD (bidirectional forwarding detection), 70–71

- DMVPN (Dynamic Multipoint VPN)
 - Phase 1*, 295–297
 - Phase 2*, 297–299
 - Phase 3*, 299–301
 - scaling*, 69
- fast convergence design, 70
- GR (graceful restart), 71–72
- hub-and-spoke design, 60–61
 - scalability optimization*, 65–68
 - summarization challenges*, 61–65
- multilayer architectures, 53–56
- multiple autonomous system drivers, 53
- with multiple autonomous systems, 50–52
- PE-CE routing protocol, 241–242
 - backdoor links between customer sites*, 245–247
 - different AS number*, 243–244
 - same AS number*, 242–243
 - some sites only*, 244–245
- queries, 52–53
- scalable EIGRP design, 50
- stub leaking, 67–68
- three-layer hierarchy architecture, 57–59
- two-layer hierarchy architecture, 56–57
- EIGRP DMVPN, case study, 295–302
- election**
 - DF election, BIDIR-PIM (bidirectional PIM), 658–659
 - DF election case study, 660–662
 - DF election messages, BIDIR-PIM (bidirectional PIM), 660
- encapsulating security payload (ESP)**, 278
- end of row versus top of rack design, 383–384
- endpoint groups (EPG), 449
- Enhanced Interior Gateway Routing Protocol. *See* EIGRP (Enhanced Interior Gateway Routing Protocol)
- enhanced VXLAN (eVXLAN), 443–444
- enterprise BGP network with Internet connectivity, designing, 788
 - analysis and task list, 791
 - BGP connectivity, 795–797
 - choosing autonomous system numbers, 792–794
 - choosing routing protocols, 792
 - Internet routing, 803–807
 - requirements and expectations, 788–791
 - routing policies, 797–802
- enterprise branch**, 692
- enterprise campus**, 692
- enterprise campus access-distribution design models**, comparing, 45
- enterprise campus design**, 2–3
 - distribution-to-distribution link design, 36–37
 - flexibility, 15–16
 - campus network virtualization*, 16–23
- hierarchies, 3
 - access layer*, 4–5
 - core layer*, 6–7
 - distribution layer*, 5–6
 - three-tier layer model*, 9–10
 - two-tier layer model*, 8–9
- high-availability enterprise campus. *See* high-availability enterprise campus
- modularity, 10

- access-distribution block*, 13–15
 - OSPF (Open Shortest Path First)*, 10–12
- resiliency, 23
- enterprise connectivity, designing**, 778
 - analysis and task list, 779–780
 - designing for new routing protocols, 780–782
 - migrating from old to new routing, 785–787
 - OSPF design optimization, 782–785
 - requirements and expectations, 778–779
 - scaling, 787–788
 - selecting replacement routing protocols, 780
- enterprise core**, 692
- enterprise data center connectivity, designing**, 816–817
 - analysis and task list, 818
 - connecting network appliances, 821–822
 - data center interconnect, 822–823
 - data center network virtualization design, 823–825
 - DCN detailed connectivity, 819–821
 - requirements and expectations, 817–818
 - selecting architecture and connectivity model, 818–819
- enterprise Internet edge**, 692
- enterprise IPv6 networks, designing**, 807
 - addressing, 813–814
 - analysis and task list, 809
 - application and service migration, 815–816
 - choosing IP address types for HQ, 809–810
 - communication between branches, 815
 - connecting branch sites, 810–812
 - deployment models, 812
 - requirements and expectations, 808–809
- enterprise LANs, connecting servers to**, 376–378
- enterprise routing, WAN**, 236–237
- enterprise WAN, access management**, 367–368
- enterprise WAN edge**, 692
- enterprise-managed VPNs**, 272
 - case studies
 - EIGRP DMVPN*, 295–302
 - MPLS/VPN over GRE/DMVPN*, 304–312
 - DMVPN (Dynamic Multipoint VPN)
 - overview*, 283–287
 - Phase 1*, 287–289
 - Phase 2*, 289–292
 - Phase 3*, 292–295
 - GRE (generic routing encapsulation), 273–275
 - IPsec, 278–280
 - overview, 272–273
- EoMPLS**, 497–498
- EoR (End of Row) design**, 383–384
- EPG (endpoint groups)**, 449
 - ACI (Application-Centric Infrastructure), 450–453
 - extending, 462–463
- equal-cost multipath routing**, 724
- ESP (encapsulating security payload)**, 278
- EtherChannel, link aggregation of EtherChannel interface**, 575–576
- EtherChannel convergence**, 28

- Ethernet, 480–481, 721
 - Ethernet Flow Control (EFC), 598
 - ETR (egress tunnel router), 213
 - European sites, routing policies, 799–802
 - event detection, OSPF (Open Shortest Path First), 94
 - event processing, OSPF (Open Shortest Path First), 96–97
 - event propagation, OSPF (Open Shortest Path First), 94–96
 - eVXLAN (enhanced VXLAN), 443–444
 - explicitly congestion notification (ECN), 520
 - extended community attributes, BGP (Border Gateway Protocol), 241–242
 - Extensible Authentication Protocol (EAP), 762, 763–765
 - Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FASTv2), 765
 - Exterior Gateway Protocol (EGP), 146
 - external Layer 2 connections and networks, ACI (Application-Centric Infrastructure), 461–465
 - external Layer 3 domains, connecting, ACI fabric, 470–471
 - external partners, connecting, 737
 - extranet topology
 - interconnect model, 738–739
 - remote LAN model, 737–738
 - extranets, security, 739–740
- F**
-
- fabric
 - ACI (Application-Centric Infrastructure), 437–440
 - ACI fabric
 - connecting to external Layer 3 domains, 470–471
 - route propagation, 468–469
 - fabric access policies, ACI (Application-Centric Infrastructure), 454–455
 - fabric extenders, 385–388
 - FabricPath, 402–407
 - fair-queue keyword, 539
 - fair-queueing, 537–538
 - fast convergence design, EIGRP (Enhanced Interior Gateway Routing Protocol), 70
 - FCoE (Fibre Channel over Ethernet), 597–598
 - FCoE Initialization Protocol (FIP), 388
 - FEX (fabric extenders), 385–388
 - FHR (first-hop routers), 637, 644
 - FHRP (First-Hop Redundancy Protocol), 31–35
 - remote-site LANs, 342–343
 - Fibre Channel over Ethernet (FCoE), 597–598
 - FIP (FCoE Initialization Protocol), 388
 - FirePOWER, 726–727
 - FirePOWER IPS appliance, 728
 - FirePOWER IPS deployment modes, 728–729
 - FirePOWER IPS module deployment modes, 728
 - FireSIGHT Management Center, 727
 - firewall clustering, 722–723
 - firewall modes, 719–720
 - firewall permissions, 740
 - firewall placement, in DC networks, 718
 - firewall virtualization, 712–714

firewalls, 695

- architecture, 709–712
- ASA (adaptive security appliance), 712
- Campus Edge network, 731–735
- case studies, separation of application tiers, 714–716
- DMZs (demilitarized zones), 710
- high availability, 720–725
- implementing in data centers, case studies, 717–720
- IOS firewalls, 712
- nonredundant firewall connectivity design, 721
- routed mode, 719
- secure enterprise networks, 835
- single-tier firewalls, 710
- transparent mode, 719
- two-tier firewall, 710
- virtualization, 712–714

first-hop layer 3 default gateway, ACI (Application-Centric Infrastructure), 465–466**First-Hop Redundancy Protocol. *See* FHRP (First-Hop Redundancy Protocol)****first-hop router (FHR), 644****first-hop routers (FHR), 637****first-in, first-out queueing, 535****flat IS-IS routing design, 134–135****flexibility, enterprise campus design, 15–16**

- campus network virtualization, 16–23

FlexVPN, 314

- architecture, 315
- capabilities, 315
- configuration blocks, 315–316

flooding, LSPs (link state packets), IS-IS (Intermediate System-to-Intermediate System), 122–123**flooding problems, PIMv2 BSR, 678–679****flooding reduction, OSPF (Open Shortest Path First), 97****forward process, IS-IS (Intermediate System-to-Intermediate System), 119****forwarding, MPLS VPNs, 258–259****front door virtual routing and forwarding (fVRF), 338****full drop (tail drop), 544****full-mesh design**

- IS-IS (Intermediate System-to-Intermediate System), 133–134
- OSPF (Open Shortest Path First), 87–88

fully meshed MPLS VPN QoS, 608–609**fVRF (front door virtual routing and forwarding), 338**

- IWAN Hybrid design model, 360

G

Gateway Load Balancing Protocol.

- See* GLBP (Gateway Load Balancing Protocol)

generic routing encapsulation (GRE), 208**GET VPN, 317–320, 621**

- QoS (quality of service), 629–630

GLBP (Gateway Load Balancing Protocol), 31–35**global synchronization, 541****GM (group member) router, 629****GR (graceful restart), EIGRP, 71–72**

GRE (generic routing encapsulation), 208

comparing point-to-point GRE and multipoint GRE, 276–277

DMVPN (Dynamic Multipoint VPN), case studies, 304–312

IPsec, 280–281, 622–623

multipoint GRE (mGRE), 275–276

overview, 273–275

group member (GM) router, 629

group-to-RP m mapping, 670–674

guest VLANs, 773

H

Head of Line (HOL), 598

hierarchal IS-IS design, 135–136

hierarchical architecture, IS-IS (Intermediate System-to-Intermediate System), 105–106

hierarchical route reflector design, BGP (Border Gateway Protocol), 167–168

hierarchies

enterprise campus design, 3

access layer, 4–5

core layer, 6–7

distribution layer, 5–6

three-tier layer model, 9–10

two-tier layer model, 8–9

OSPF (Open Shortest Path First), 84–85

high availability

data centers, 388–392

firewalls, 720–725

high-availability enterprise campus, 23–24

distribution-to-distribution interconnect

multitier access model, 37–41

routed access model, 41–42

with virtual swi, 43–44

FHRP (First-Hop Redundancy Protocol), 31–35

link aggregation, 28–31

overview, 44–46

trunking, 27

VLAN design, 24–26

high-performance trading (HPT), data center QoS, 595

HOL (Head of Line), 598

hold-interval, 95

hop-by-hop easy virtual network (EVN) based, 20

hop-by-hop VRF-lite based, 19

Hot Standby Router Protocol. *See* HSRP (Hot Standby Router Protocol)

HPT (high-performance trading), data center QoS, 595

HQ, choosing IP address types for, 809–810

HSRP (Hot Standby Router Protocol), 31

hub border router, 365

hub mast controller (MC), 365

hub-and-spoke design

DMVPN (Dynamic Multipoint VPN), 285

EIGRP (Enhanced Interior Gateway Routing Protocol), 60–61

scalability optimization, 65–68

summarization challenges, 61–65

NBMA hub-and-spoke, IS-IS (Intermediate System-to-Intermediate System), 132–133

OSPF (Open Shortest Path First), 88

ABR placement, 89–90
network types, 92–93
number of areas, 91

H-VPLS, 263–264

I

iBGP, 148

scalability limitations, 152
 scalability solutions, 152–153
 confederations, 155–156
 route reflectors, 153–155

Identity Services Engine (ISE), 768

IEEE 802.1X, 759–763

message flow, 763
 phased deployment, 767
 supplicants, 765–766

IETF (Internet Engineering Task Force), 31

IGMP, multicast receiver controls, 755–757

IGMP membership report, 646–647

IGMPv3, SSM (source-specific multicast), 655

IIF (incoming interface), 653

IIH PDUs, 117–118

IIHs, IS-IS (Intermediate System-to-Intermediate System), 121–122

IKE (Internet Key Exchange), 278
 phases of, 278–279

IKE GDOI (Group Domain of Interpretation), 317–318

IKEv2, FlexVPN, 316

implementation and optimization phases, IPv6, 197

implementing, firewalls in a data center, 717–720

incoming interface (IIF), 653

information data flow, IS-IS (Intermediate System-to-Intermediate System), 118–119

infrastructure

network infrastructure devices, resiliency and survivability, 700–701

routing infrastructure, security, 699–700

secure enterprise networks, designing, 833–834

switching infrastructure, 702–703

infrastructure device access, 698–699

infrastructure devices, LISP (locator/ID separation protocol), 213–216

infrastructure protection, 695
 designing, 696–697

ingress traffic filtering, 702

ingress tunnel router (ITR), 213

inline mode, 727

inside zone, IPS (intrusion prevention system), 726

integrated IS-IS, 104–105
 for IPv6, 138–141

Integrated Services. *See* IntServ

integration options, ACI (Application-Centric Infrastructure), 471–473

intelligent path control, WAN, 356

Intelligent WAN. *See* IWAN (Intelligent WAN)

Intelligent WAN (IWAN), 354–355

inter-AS MPLS VPN, WAN connections, 232

interconnect model, 738–739

interconnecting, data centers, 822–823

- inter-DC traffic, 478
- interdomain, 639
- interface-based PIM neighbor filtering, 752
- Intermediate System-to-Intermediate System. *See* IS-IS (Intermediate System-to-Intermediate System)
- internal multicast security
 - multicast admission controls, 757
 - multicast receiver controls, 755–757
 - multicast sender control, 753–755
 - PIM (Protocol-Independent Multicast), 752
- internal networks, connecting, Campus Edge network, 733–734
- Internet
 - connecting, Campus Edge network, 731
 - remote sites, using local Internet, 337–339
- Internet Engineering Task Force (IETF), 31
- Internet Key Exchange (IKE), 278
 - phases of, 278–279
- internet keyword, 171
- Internet routing
 - default routing, 805–807
 - multihoming, 804–805
 - public IP space selection, 803–804
- Inter-Switch Link (ISL), 27
- inter-VLAN routing, 381–383
- intradomain, 639
- intranet data center, 692
- intrusion prevention system. *See* IPS (intrusion prevention system)
- IntServ (Integrated Services), 514–515, 516
- IOS encryption, order of operations, 623–625
- IOS firewalls, 712
- IOS XR software, 750
- IP address mobility, 484–490
- IP address types, choosing for HQ, 809–810
- IP ECN, 547–550
- IP gateway redundancy, VSS (virtual switching system), 35–36
- ip msdp sa-filter, 755
- IP multicast, 633–634
 - how it works, 634–635
 - multicast forwarding and RPF check, 639–641
 - multicast groups, 635–636
 - multicast networks, 638
 - multicast protocols, 638–639, 642–644
 - PIM-SM (Protocol-Independent Multicast—Sparse Mode). *See* PIM-SM (Protocol-Independent Multicast—Sparse Mode)
 - security, 743
 - challenges of*, 744
 - SSM (source-specific multicast). *See* SSM (source-specific multicast)
- ip multicast boundary, 754
- IP multicast service model, 636–637
- IP packet DiffServ DS field, 522
- ip pim accept-register, 755
- ip pim register-rate-limit, 755
- ip pim rp-announce-filter rp-list, 746
- IP precedence mapping, DSCP, 523
- IP RTP priority queueing, 536
- IP source guard, 702
- IP spoofing protection, 702
- ip tcp adjust mss [size]626
- IP type of service, QoS (quality of service), 519–520

- IP-in-IP (IPIP), 208**
- IPIP (IP-in-IP), 208**
- IPS (intrusion prevention system), 696**
 - architecture, 726–729
 - security, 695
- IPsec, 278–280, 284**
 - DVTI (Dynamic VTI), 283
 - GRE (generic routing encapsulation), 622–623
 - GRE (generic routing encapsulation) and, 280–281
 - VTI (virtual tunnel interface) and, 281–282
- IPsec SA anti-replay, 630**
- IPsec VPNs**
 - modes, 621–623
 - QoS (quality of service), 619–620
 - MTU (maximum transmission unit), 625–626*
 - use cases, 621*
- IPv4 addresses, 194**
- IPv6, 194**
 - 6RD (6 Rapid Deployment), 210–211
 - application support, 222–223
 - application adaptation, 223*
 - application workarounds, 223–224*
 - control plane security, 224
 - deployment and design, 194–195
 - assessment phase, 196*
 - discovery phase, 196*
 - implementation and optimization phases, 197*
 - planning and design phase, 196–197*
 - DNS64, 206–208
 - dual stack, 205–206
 - Dual-Stack Lite, 211–212
 - dual-stack security, 225
 - integrated IS-IS, 138–141
 - link layer security, 221–222
 - manual tunnels, 208–209
 - migration
 - acquiring IPv6 prefixes, 197–198*
 - transition mechanisms, 203–205*
 - where to start, 199–200*
 - migration models
 - IPv6 islands, 200–201*
 - IPv6 WAN, 201–203*
 - multihoming, 226
 - NAT64, 206–208
 - transition mechanisms, 216–217
 - tunnel brokers, 209
 - tunneling security, 225–226
- IPv6 embedded RP, 679–681**
- IPv6 islands, 200–201**
- IPv6 services, 219–220**
 - addressing services, 220–221
 - name services, 220
 - security services, 221
- IPv6 WAN, 201–203**
- ISE (Identity Services Engine), 768, 771**
- IS-IS (Intermediate System-to-Intermediate System), 87, 102, 141–142**
 - addressing, 114–116
 - adjacencies, 108–109, 120
 - characteristics, 103–104, 110–112
 - domains, 104
 - flat routing design, 134–135
 - hierarchal IS-IS design, 135–136

hierarchical architecture, 105–106
 information data flow, 118–119
 integrated IS-IS, 104–105
 for IPv6, 138–141
 level 1/level 2 LSPs, 121–122
 link state packets flooding, 122–123
 LSDB synchronization, 123–124
 network types, 119
 OSPF versus, 110–112
 area design, 112–113
 overview, 102–103
 packets, 117
 protocol operations, 119–121
 route summarization, 136–138
 router and link types, 106–108
 routing, 125–126
 asymmetric versus symmetric, 129–132
 full-mesh design, 133–134
 NBMA hub-and-spoke, 132–133
 route leaking, 126–129
 single topology restrictions, 138–139
 IS-IS PDUs, 117
 ISL (Inter-Switch Link), 27
 ITR (ingress tunnel router), 213
 IWAN (Intelligent WAN), 354–355
 AVC (Application Visibility Control), 357
 PfR (Performance Routing), 356
 PfRv3, 363–366
 secure connectivity, 357
 IWAN design, 358–359
 IWAN Hybrid design model, 361
 IWAN Hybrid WAN design model, 359
 IWAN WAN aggregation (hub) designs, 359

J

jitter, WAN QoS, 590–591

K

keywords

additive, 177
 bandwidth, 539
 fair-queue, 539
 internet, 171

KS (key server), 629

L

L3Out, connecting ACI fabric to external Layer 3 domains, 470–471

LAN segments, 703

LANs, remote-site LANs, 339–343

latency, WAN QoS, 590–591

Layer 2 attacks, 753

Layer 2 connections and networks, ACI (Application-Centric Infrastructure), 461–465

Layer 2 DCI:LISP based, 488

Layer 2 hub-and-spoke WAN QoS design model, 607

Layer 2 marking, QoS (quality of service), 517–519

Layer 2 MPLS VPN, 259

Layer 2 outside connections, ACI (Application-Centric Infrastructure), 463–464

Layer 2 private WAN QoS, 607

Layer 2 switch networks with STP, 703

Layer 2 VPN provisioning models, 497

- Layer 2 WAN design models, 329–331
- Layer 2.5 marking: MPLS
 - experimental bits, QoS (quality of service), 524
- Layer 3 DCI, 507–509
- Layer 3 marking: DSCP per-hop behaviors, QoS (quality of service), 520–523
- Layer 3 marking: IP type of service, QoS (quality of service), 519–520
- Layer 3 MPLS VPNs, 233–234
- Layer 3 separation with VRF-Lite, 423–424
- Layer 7: NBAR/NBAR2, QoS (quality of service), 526–527
- leaf nodes, ACI (Application-Centric Infrastructure), 467
- leaf switches, 401, 439
- level 1 router, IS-IS (Intermediate System-to-Intermediate System), 107
- level 1/level 2 LSPs, IS-IS (Intermediate System-to-Intermediate System), 121–122
- level 1/level 2 router, IS-IS (Intermediate System-to-Intermediate System), 107
- level 2 router, IS-IS (Intermediate System-to-Intermediate System), 107
- limitations of
 - ACI (Application-Centric Infrastructure), networking limitations, 432–434
 - current networking technology, 398–399
 - customer-managed Layer 2 DCI deployment models, 501
- link aggregation, high-availability enterprise campus, 28–31
 - link aggregation of EtherChannel interface, QoS (quality of service), 575–576
- Link Layer Discovery Protocol (LLDP), 464
- link layer security, IPv6, 221–222
- link types, IS-IS (Intermediate System-to-Intermediate System), 106–108
- Link-State Database Overload Protection, OSPF (Open Shortest Path First), 97–98
- link-state routing protocols, designing, 781
- LISP (Locator/ID Separation Protocol), 212–216
 - LISP (locator/ID separation protocol), DCI (Data Center Interconnect), 487–489
 - LISP infrastructure devices, 213–216
 - LISP site edge devices, 213
- LLDP (Link Layer Discovery Protocol), 464
- LLQ (low-latency queueing), 536, 540, 591
- load balancing
 - enterprise routing, WAN, 237
 - EtherChannel, 575
- load-sharing design, BGP (Border Gateway Protocol), 177
 - single-homing versus multihoming, 177–178
- Locator/ID Separation Protocol (LISP), 212–216
- loop prevention
 - BGP (Border Gateway Protocol), 148–149
 - route reflectors, BGP (Border Gateway Protocol), 162–165
- low-latency queueing (LLQ), 536, 540

- LSA throttling timers, 96
- LSDB synchronization, IS-IS (Intermediate System-to-Intermediate System), 123–124
- LSPs (link state packets), IS-IS, 121–123
 - flooding, IS-IS (Intermediate System-to-Intermediate System), 122–123

M

- MAB (MAC Authentication and Bypass), 769
- main HQ multihoming, Internet routing, 803–804
- main mode, IKE (Internet Key Exchange), 279
- managed CE service, WAN connections, 232
- managed VPNs, 230
- management, WAN, 357–358
- management access, securing, to infrastructure devices, 698–699
- management network, 693
- management plane, 697
- management plane protection, 697
- management restricted zones, 695
- manual tunnels, IPv6, 208–209
- mapping QoS markings between OSI layers, 524–525
- mapping agents, Auto-RP, 670
- Map-Resolver (MR), 214
- Map-Server (MS), 213–214
- mark probability denominator, 543
- markings
 - DSCP (Differentiated Services Code Point), 555
 - mapping QoS markings between OSI layers, 524–525
- masquerading, 754
- maximum threshold, 543
- maximum transmission unit (MTU), 80
- max-interval, 95
- MCP (Mis-Cabling Protocol), 464
- MEC (Multichassis EtherChannel), 30
- message flow, 802.1X, 763
- messages, DF election messages, 660
- mGRE (multipoint GRE), 275–276, 284
 - versus point-to-point GRE, 276–277
- microsegmentation, overlay networks, 427–428
- migrating
 - from old to new routing, designing enterprise connectivity, 785–787
 - from RIPv2 to OSPF, 785
- migration, IPv6
 - acquiring IPv6 prefixes, 197–198
 - transition mechanisms, 203–205
 - where to start, 199–200
- migration models
 - IPv6 islands, 200–201
 - IPv6 WAN, 201–203
- migration options, ACI (Application-Centric Infrastructure), 471–473
- minimum threshold, 543
- Mis-Cabling Protocol (MCP), 464
- mobility, IP address mobility, 484–490
- models
 - 3G/4G VPN design models, 335
 - ACI policy model, 447–450
 - classification/marketing/policing QoS model, 573–574
 - customer-managed Layer 2 DCI deployment models. *See* customer-managed Layer 2 DCI deployment models
 - DCN connectivity model, 820

- deployment models
 - DMVPN (Dynamic Multipoint VPN)*, 285
 - enterprise IPv6 networks case study*, 812
- enterprise campus access-distribution design models, comparing, 45
- interconnect model, 738–739
- IP multicast service model, 636–637
- IWAN Hybrid design model, 361
- Layer 2 hub-and-spoke WAN QoS design model, 607
- migration models
 - IPv6 islands*, 200–201
 - IPv6 WAN*, 201–203
- MPLS VPNs connectivity model, 606
- network-centric security model, 715
- QoS (quality of service), 12-class QoS strategy model, 564–565
- QoS design model, 837–838
- QoS strategy models, 560–561
 - 4-class QoS strategy model*, 561–562
 - 8-class QoS strategy model*, 562–563
- remote LAN model, 737–738
- three-tier layer model, enterprise campus design, 9–10
- three-tiered e-commerce application functional model, 714
- two-tier layer model, enterprise campus design, 8–9
- modern data centers, 400**
 - microsegmentation, with overlay networks, 427–428
 - multitenant data centers, 422
 - secure tenant separation*, 422–425
- network overlays, 402
 - Cisco FabricPath*, 402–407
 - VXLAN (virtual extensible LAN)*, 407–408
- SDN (software-defined networking), 414–416
 - benefits of*, 416–417
 - challenges of*, 419–421
 - nontraditional SDN*, 421
 - requirements*, 419
 - selection criteria*, 417–418
- spine-leaf topologies, 400–401
- VTEP (VXLAN tunnel endpoint), 408–411
- modes**
 - active/active mode, 722
 - ASA firewall active/passive failover mode, 722
 - FirePOWER IPS deployment modes, 728–729
 - firewall modes, 719–720
 - inline mode, 727
 - IPsec VPNs, 621–623
 - monitor-only mode, 727
- modular enterprise campus with OSPF, 10–12**
- modular network architecture, 691–695**
 - security zones, 695
- modularity, enterprise campus design, 10**
 - access-distribution block, 13–15
 - OSPF (Open Shortest Path First), 10–12
- modules, 692–693**
- monitor-only mode, 727**
- MP-BGP (Multiprotocol BGP), 468–470, 639**

- MP-BGP EVPN (Multiprotocol Border Gateway Protocol Ethernet Virtual Private Network), 413–414
- MPLS (Multiprotocol Label Switching), 230
 - Layer 3 MPLS VPNs, 233–234
 - MPLS VPNs, architecture, 234–236
- MPLS DiffServ tunneling modes, 609–611
- MPLS EXP, 612–613
- MPLS headers, 524
- MPLS Layer 3 WAN design models, 326–329
- MPLS uniform DiffServ tunneling mode, 612
- MPLS VPNs
 - architecture, 234–236
 - connectivity model, 606
 - forwarding, 258–259
 - fully meshed MPLS VPN QoS, 608–609
 - Layer 2 MPLS VPN, 259
 - QoS (quality of service), 605–607
 - MPLS DiffServ tunneling modes*, 609–611
 - pipe tunneling mode*, 614–615
 - role mapping*, 616
 - sample roles*, 615–617
 - short-pipe tunneling mode*, 612–614
 - uniform tunneling mode*, 612
 - routing propagation, 255–258
- MPLS WAN DiffServ tunneling, QoS in the enterprise network case study, 839–841
- MQC, 536
- MR (Map-Resolver), 214
- mrinfo, 756
- MS (Map-Server), 213–214
- MSDP (Multicast Source Discovery Protocol), 639, 654
 - multicast network edge security, 751–752
 - neighbor relationships, 683
 - operations, 684–686
 - RP (Rendezvous Point), 683
- mtrace, 756
- MTU (maximum transmission unit), 80
 - QoS (quality of service), 625–626
 - WAN connections, 232
- multicast. *See also* IP multicast, security challenges, 744
- multicast admission controls, 757
- multicast boundary, 749, 752
- multicast distribution trees, 642
- multicast distribution trees identification, 644–645
- multicast forwarding, 645
 - RPF check, 639–641
- multicast groups, 635–636
- Multicast Information Protocol, 748
- multicast network edge, security, 748–749
 - Auto-RP and BSR, 749–751
 - MSDP (Multicast Source Discovery Protocol), 751–752
- multicast networks, 638
 - network element security, 746–748
 - problems in, 744–745
 - security considerations, 745–746
- multicast protocols, 638–639, 642–644
- multicast receiver controls, 755–757
- multicast rekeying, 318–319
- multicast routing protocols, 642

- multicast routing tables, PIM-SM (Protocol-Independent Multicast—Sparse Mode), 652–653
 - multicast sender control, 753–755
 - Multichassis EtherChannel (MEC), 30
 - multicontext mode, firewall virtualization, 712
 - multihoming, 178
 - Internet routing, 804–805
 - IPv6, 226
 - versus single-homing, BGP (Border Gateway Protocol), 177–178
 - with two ISPs using a single local edge router, 183–186
 - with two ISPs using multiple local edge routers, 186–188
 - multihop GRE tunneling based, 21
 - multihop MPLS core based, 22–23
 - multi-hypervisor-ready fabric, 445
 - multilayer architectures, EIGRP (Enhanced Interior Gateway Routing Protocol), 53–56
 - multiple autonomous system drivers, EIGRP (Enhanced Interior Gateway Routing Protocol), 53
 - multiple autonomous systems, EIGRP. *See* AS (autonomous systems)
 - multipoint GRE (mGRE), 275–276, 284
 - versus point-to-point GRE, 276–277
 - multipology IS-IS, for IPv6, 140–141
 - Multiprotocol BGP (MP-BGP), 468–469
 - Multiprotocol Border Gateway Protocol Ethernet Virtual Private Network (MP-BGP EVPN), 413–414
 - Multiprotocol Label Switching. *See* MPLS (Multiprotocol Label Switching)
 - Multiprotocol Label Switching Virtual Private Networks. *See* MPLS VPNs
 - multitenant data centers, 422
 - case studies, 425–426
 - secure tenant separation, 422–425
 - multitenant segmentation, extranets, 739–740
 - multitier, access-distribution block, 13
 - multitier access model, distribution-to-distribution interconnect, 37–41
 - multitier data center designs
 - data center high availability, 388–392
 - end of row versus top of rack design, 383–384
 - fabric extenders, 385–388
 - inter-VLAN routing, 381–383
 - NIC teaming, 392–393
 - small data centers (connecting servers to an enterprise LAN), 376–378
 - two-tier data center network architecture, 378–380
-
- N
- name services, IPv6, 220
 - named communities, BGP (Border Gateway Protocol), 171
 - NAT64, IPv6, 206–208
 - NBAR (Network-Based Application Recognition), 526–527
 - NBAR2 (next-generation NBAR), 526–527, 837
 - NBMA hub-and-spoke, IS-IS (Intermediate System-to-Intermediate System), 132–133
 - neighbor relationships, MSDP (Multicast Source Discovery Protocol), 683

- NetFlow, 837
- network access control
 - authorization options case study, 772–775
 - Cisco TrustSec, 768
 - Profiling Service*, 768–769
 - SGT (Security Group Tag)*, 769–772
 - EAP (Extensible Authentication Protocol), 763–765
 - IEEE 802.1X, 759–763
 - secure enterprise networks, 833–834
- Network Access Manager, 766
- network and security management, 695
- network appliances, connecting, 821–822
- network bgp router, 151
- network design issues, route reflectors, BGP (Border Gateway Protocol), 169
- network element security, 746–748
- network infrastructure devices, resiliency and survivability, 700–701
- network interface controller teaming, 392–393
- Network Layer 2 separation, 423
- Network Layer 3 separation, 422
- network overlays, modern data centers, 402
 - Cisco FabricPath, 402–407
 - VXLAN (virtual extensible LAN), 407–408
- network policy enforcement, 701–702
- network resiliency, 24
- network security zoning, 690–691
- network separation, multitenant data centers, 422–423
- network service access points (NSAPs), 102
- network services separation, 423
- network targeted attacks, security, 747
- network types
 - hub-and-spoke design, OSPF (Open Shortest Path First), 92–93
 - IS-IS (Intermediate System-to-Intermediate System), 119
- network virtualization overlays, ACI (Application-Centric Infrastructure), 441–446
- Network-Based Application Recognition (NBAR), 526–527
- network-centric security model, 715
- networking limitations, ACI (Application-Centric Infrastructure), 432–434
- networking technology, limitations of, 398–399
- networks
 - multicast networks
 - problems in*, 744–745
 - security considerations*, 745–746
 - overlay networks, microsegmentation, 427–428
- new network architecture, data centers, 397–398
- next-generation IPS (NGIPS), Cisco FirePOWER, 726–727
- next-generation NBAR (NBAR2), 526–527
- next-generation security, 696
- next-generation WAN (NGWAN), 354–355

Nexus ACI fabric software, 440
 Nexus switches, ACI fabric mode, 439
 NGIPS (next-generation IPS), Cisco FirePOWER, 726–727
 NGWAN (next-generation WAN), 354–355
 NHRP, 284
 DMVPN (Dynamic Multipoint VPN), Phase 2, 290
 NIC teaming, 392–393
 no drop, 543
 no next-hop-self, 298
 no-advertise, 170
 no-export, 170
 no-export-subconfed, 170
 nonclients, 155
 nonredundant firewall connectivity design, 721
 non-RR clients, 155
 nonstop forwarding (NSF), EIGRP (Enhanced Interior Gateway Routing Protocol), 71–72
 nontraditional SDN, 421
 nontunnel EAP, 763
 no-peer, 171
 North American sites, routing policies, 797–799
 north-south traffic, 478
 NSAPs (network service access points), 102
 NSF (nonstop forwarding), EIGRP (Enhanced Interior Gateway Routing Protocol), 71–72
 number of areas, hub-and-spoke design, OSPF (Open Shortest Path First), 91
 number of areas per ABR, OSPF (Open Shortest Path First), 81–82
 numbers of routers in an area, OSPF (Open Shortest Path First), 80–81

O

offer message, DF election messages, 660
 OILs (outgoing interface lists), 639
 ONE (Open Network Environment) architecture, 435
 ONF (Open Networking Foundation), 398
 OpenFlow, 415–416
 Open Network Environment (ONE) architecture, 435
 Open Networking Foundation (ONF), 398
 Open Shortest Path First. *See* OSPF (Open Shortest Path First)
 OpenFlow, ONF (Open Networking Foundation), 415–416
 open-source sniffing solutions, 837
 operational resiliency, 24
 operations
 Auto-RP, 671–674
 MSDP (Multicast Source Discovery Protocol), 684–686
 PfR (Performance Routing), 362–363
 operations zone, 694
 order of operations, QoS (quality of service), 623–625
 OSI layers, mapping QoS markings, 524–525
 OSPF (Open Shortest Path First), 75
 adjacent neighbors, 76–77
 area design, 82–83
 characteristics, 110–112
 convergence, 93
 event detection, 94
 event processing, 96–97
 event propagation, 94–96
 design optimization, 782–785

- DMVPN (Dynamic Multipoint VPN), 289
 - flooding reduction, 97
 - full-mesh design, 87–88
 - hierarchies, 84–85
 - hub-and-spoke design, 88
 - ABR placement*, 89–90
 - network types*, 92–93
 - number of areas*, 91
 - IS-IS versus, 110–112
 - area design*, 112–113
 - Link-State Database Overload Protection, 97–98
 - migrating from RIPv2, 785
 - modularity, enterprise campus design, 10–12
 - number of areas per ABR, 81–82
 - numbers of routers in an area, 80–81
 - PE-CE routing protocol, 247–250
 - backdoor links between customer sites*, 250–251
 - route summarization*, 251–252
 - routing information in the area and routed domain, 78–80
 - scalability design, 76
 - sham links, 250–251
 - summarization, 85–86
 - OSPF backbone area design, 781–782
 - OTV (overlay transport virtualization), DCI (Data Center Interconnect), 501–506
 - outgoing interface lists (OILs), 639
 - outside zone, IPS (intrusion prevention system), 726
 - overlay networks
 - ACI network virtualization overlays, 441–446
 - DCI (Data Center Interconnect), 507
 - microsegmentation, 427–428
 - overlay transport virtualization DCI, 501–506
 - overlay transport virtualization (OTV), DCI (Data Center Interconnect), 501–506
 - oversubscription, 380
-
- ## P
- P routers, 235
 - PA (Provider-Assigned) prefixes, 197–198
 - PaaS (platform as a service), 596
 - packet dropping, 547
 - packets, IS-IS (Intermediate System-to-Intermediate System), 117
 - PAgP (Port Aggregation Protocol), 28, 30
 - partial SNP. *See* PSNPs
 - partner and extranet modules, 693
 - pass messages, DF election messages, 660
 - passive interfaces, 700
 - path attributes, BGP (Border Gateway Protocol), 150
 - path isolation, campus network virtualization, 19–23
 - path selection, BGP (Border Gateway Protocol), 150–151
 - PAUSE frame, 598
 - PBR (policy-based routing), 724
 - PDUs, 117
 - PE (provider edge) routers,
 - architecture, 237–238
 - route distinguishers, 238–239
 - route target (RT), 240–241
 - peak information rate (PIR), 533

- PE-CE routing protocol, 241**
 - BGP (Border Gateway Protocol), 252–254
 - backdoor links between customer sites, 254–255*
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 241–242
 - backdoor links between customer sites, 245–247*
 - different AS number, 243–244*
 - same AS number, 242–243*
 - some sites only, 244–245*
 - OSPF (Open Shortest Path First), 247–250
 - backdoor links between customer sites, 250–251*
 - route summarization, 251–252*
- peer-forwarding rules, BGP (Border Gateway Protocol), 158**
- performance, platform performance, WAN QoS, 589–590**
- per-hop behavior (PHB)**
 - Layer 3 marking, 520–523
 - queue design principles, 558–559
- Pervasive SVI, 465–466**
- PFC (Priority-based Flow Control), 598**
- PfR (Performance Routing), 361–362**
 - IWAN (Intelligent WAN), 356
 - operations, 362–363
- PfRv3**
 - design and deployment, 366–367
 - IWAN (Intelligent WAN), 363–366
- Phase 1, DMVPN (Dynamic Multipoint VPN), 287–289**
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 295–297
- Phase 2, DMVPN (Dynamic Multipoint VPN), 289–292**
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 297–299
- Phase 3, DMVPN (Dynamic Multipoint VPN), 292–295**
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 299–301
- phased deployment, 802.1X, 767**
- phases of IKE (Internet Key Exchange), 278–279**
- PHB (per-hop behavior), Layer 3 marking, 520–523**
- PHB-DSCP bit mapping, 520**
- PI (Provider-Independent) prefixes, 197–198**
- PILE Forensic Accounting, enterprise BGP network with Internet connectivity case study. *See* enterprise BGP network with Internet connectivity**
- PIM (Protocol-Independent Multicast), 637**
 - BIDIR-PIM (bidirectional PIM), 658
 - internal multicast security, 752
 - multicast admission controls, 757*
 - multicast receiver controls, 755–757*
- PIM source-specific multicast (PIM-SSM), 646**
- PIM-DM (PIM Dense Mode), 750**
- PIM-SM (Protocol-Independent Multicast—Sparse Mode), 645–646**
 - (S, G), 653–654
 - **, G (star comma G), 653
 - bidirectional PIM (BIDIR-PIM), 657
 - characteristics, 645
 - IP multicast, 645–646
 - multicast routing tables, 652–653
 - receiver joins PIM-SM shared tree, 646–647

- RP registration, 647–648
- SPT switchover, 649–652
- SSM (source-specific multicast). *See* SSM (source-specific multicast)
- PIM-SM SPT switchover, 668
- PIM-SSM (PIM source-specific multicast), 646
- PIMv1, 756
- PIMv2 BSR, 676–677
 - BSR (bootstrap router), 678
 - candidate RPs, 677–678
 - flooding problems, 678–679
 - routers, 678
- PIMv6, 646
- PIN (Places-in-the-Network), 568
 - internal multicast security, multicast sender control, 753–755
- pipe mode, MPLS DiffServ tunneling modes, 610
- pipe tunneling mode, MPLS VPNs, 614–615
- PIR (peak information rate), 533
- placement of, RP (Rendezvous Point), 667–668
- Places-in-the-Network (PIN), 568
- planning and design phase, IPv6, 196–197
- platform performance, WAN QoS, 589–590
- PoE (Power over Ethernet), 378
- point-to-point GRE versus mGRE, 276–277
- point-to-point links, IS-IS (Intermediate System-to-Intermediate System), 119
- policies, ACI fabric access policies, 454–455
- policing and remarking design principles, QoS (quality of service), 556
- policing tools
 - single-rate three-color marker, 532–533
 - two-rate three-color marker, 533–535
- policing traffic, 527–529, 532
- policy-based centralized control, 418
- policy-based routing (PBR), 724
- Port Aggregation Protocol (PAgP), 28
- port extenders, 385–388
- Power over Ethernet (PoE), 378
- PQ (priority queueing), 535
- PQ-WFQ, 536
- prefixes
 - 6RD prefix, 211
 - acquiring IPv6 prefixes, 197–198
- prefix-suppression, 79
- preventing, attacks, 703
- priority command, 540
- priority queueing (PQ), 535
- Priority-based Flow Control (PFC), 598
- problems, in multicast networks, 744–745
- Profiling Service, 768–769
- protocol operations, IS-IS (Intermediate System-to-Intermediate System), 119–121
- Protocol-Independent Multicast (PIM), 637
- Protocol-Independent Multicast—Sparse Mode. *See* PIM-SM (Protocol-Independent Multicast—Sparse Mode)
- protocols
 - BGP. *See* BGP (Border Gateway Protocol)
 - EAP (Extensible Authentication Protocol), 762, 763–765

- EGP (Exterior Gateway Protocol), 146
 - EIGRP. *See* EIGRP (Enhanced Interior Gateway Routing Protocol)
 - FHRP (First-Hop Redundancy Protocol), 31–35
 - FIP (FCoE Initialization Protocol), 388
 - GLBP (Gateway Load Balancing Protocol), 31–35
 - HSRP (Hot Standby Router Protocol), 31
 - IS-IS. *See* IS-IS (Intermediate System-to-Intermediate System)
 - LISP (Locator/ID Separation Protocol), 212–216
 - LLDP (Link Layer Discovery Protocol), 464
 - MCP (Mis-Cabling Protocol), 464
 - MPLS (Multiprotocol Label Switching), 230
 - Multicast Information Protocol, 748
 - multicast protocols, 638–639, 642–644
 - PAgP (Port Aggregation Protocol), 28, 30
 - PE-CE routing protocol, 241
 - BGP (Border Gateway Protocol), 252–254*
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 241–242*
 - OSPF (Open Shortest Path First), 247–250*
 - routing protocol authentication mechanisms, 699
 - SAP (Session Announcement Protocol), 748
 - SXP (Security Group Tag Exchange Protocol), 770
 - VRRP (Virtual Router Redundancy Protocol), 31
 - provider (P) networks, 235
 - provider edge (PE) routers, 235
 - architecture, 237–238
 - route distinguishers, 238–239*
 - route target (RT), 240–241*
 - Provider-Assigned (PA) prefixes, 197–198
 - provider-assigned approach, IPv6 WAN, 201
 - Provider-Independent (PI) prefixes, 197–198
 - provider-independent approach, IPv6 WAN, 201–202
 - Proxy Tunnel Router (PxTR), 214
 - pseudowire DCI, 495
 - PSNPs (partial number packets), 123–124
 - public access zones, 690, 694
 - public IP space selection, Internet routing, 803–804
 - public zones, 690, 694
 - pure IP domain, 104
 - pure ISO domain, 104
 - PxTR (Proxy Tunnel Router), 214
- ## Q
-
- QoS (quality of service), 514, 745
 - buffers, 569–570
 - bursts, 569–570
 - campus QoS
 - design examples, 576–588*
 - overview, 568*
 - classification, order of operations, 623–625
 - classification and marking

- classification and marking tools*, 516–517
- Layer 2 marking*, 517–519
- Layer 2.5 marking: MPLS experimental bits*, 524
- Layer 3 marking: DSCP per-hop behaviors*, 520–523
- Layer 3 marking: IP type of service*, 519–520
- Layer 7: NBAR/NBAR2*, 526–527
- mapping markings between OSI layers*, 524–525
- traffic policing and shaping*, 527–529, 532
- classification and marking design principles, 554–555
- classification/marking/policing QoS model, 573–574
- classifications and marking tools, 516–517
- data center QoS, 594
 - big data architecture*, 596
 - DC QoS application case study*, 599–601
 - HPT (high-performance trading)*, 595
- DMVPN (Dynamic Multipoint VPN), 626–628
- dropping design principles, 557–558
- dropping tools, DSCP-based WRED, 541–546
- GETVPN, 629–630
- IP ECN, 547–550
- IPsec VPN, 619–620
 - MTU (maximum transmission unit)*, 625–626
 - use cases*, 621
- Layer 2 private WAN QoS, 607
- link aggregation of EtherChannel interface, 575–576
- MPLS VPNs, 605–607
 - fully meshed MPLS VPN QoS*, 608–609
 - MPLS DiffServ tunneling models*, 609–611
 - pipe tunneling mode*, 614–615
 - sample roles*, 615–617
 - short-pipe tunneling mode*, 612–614
 - uniform tunneling mode*, 612
- overview, 553–554
- per-hop behavior queue design principles, 558–559
- policing and remarking design principles, 556
- policing tools, 532–533
- queueing
 - CBWFQ (class-based weighted fair queueing)*, 538–541
 - fair-queueing*, 537–538
 - Tx-Ring*, 536–537
- queueing design principles, 557
- queueing tools, 535–536
- queueing/dropping recommendations, 574–575
- RFC 4594, 559–560
- token bucket algorithms, 529–531
- traffic descriptors, 516–517
- traffic policing, 527–529
- traffic shaping, 527–529
- trust boundary, QoS in the enterprise network case study, 838
- trust states, boundaries and, 570–573
- video, 568–569
- VoIP (voice over IP), 568–569

WAN connections, 231
 WAN QoS. *See* WAN QoS

QoS design model, 837–838

QoS in the enterprise network case study, 835

designing

- congestion management, 838–839*
- MPLS WAN DiffServ tunneling, 839–841*
- QoS design model, 837–838*
- QoS trust boundary, 838*
- requirements and expectations, 835–836*
- scavenger traffic, 839*
- traffic discovery and analysis, 836–837*

QoS strategy models, 560–561

- 4-class QoS strategy model, 561–562
- 8-class QoS strategy model, 562–563
- 12-class QoS strategy model, 564–565

quality of service (QoS). *See* QoS (quality of service), WAN connections, 231

queries, EIGRP (Enhanced Interior Gateway Routing Protocol), 52–53

queueing, 535

- 8-class 1P1Q3T egress queueing, 581–588
- 8-class 1P1Q3T ingress queueing, 580–581
- CBWFQ (class-based weighted fair queueing), 538–541
- fair-queueing, 537–538
- Tx-Ring, 536–537
- WAN QoS, 591–592

queueing design principles, QoS (quality of service), 557

queueing recommendations, QoS (quality of service), 574–575

queueing tools, 535–536

R

RA spoofing, 222

rACLs (receive access control lists), 747

RADIUS (Remote Authentication Dial-In User Service), 762, 763

random drop, 544

random early detection (RED), 542

- dropping modes, 543–544

rate-limiting PIM register messages, 752

receive access control lists (rACLs), 747

receive process, IS-IS (Intermediate System-to-Intermediate System), 118

receiver joins PIM-SM shared tree, 646–647

Recovery Point Objective (RPO), 482

Recovery Time Objective (RTO), 482

RED (random early detection), 542

- dropping modes, 543–544

redundancy

- case studies, redundancy and connectivity, 343–354
- DMVPN (Dynamic Multipoint VPN), 302–304

Regional Internet Registries (RIR), 809

regional offices WAN design, 348–351

rekeying options, 318–319

- Remote Authentication Dial-In User Service (RADIUS), 762, 763
- remote LAN model, 737–738
- remote sites
 - local Internet, 337–339
 - WAN, 324–326
- remote VPN solutions, 272
- remote VTEP discovery, 411–413
 - tenant address learning, 411–413
- remote-site LANs, 339–343
- remote-site WAN design, 346–348
- Rendezvous Point. *See* RP (Rendezvous Point)
- replacement routing protocols, selecting, 780
- requirements
 - enterprise BGP network with Internet connectivity case study, 788–791
 - for enterprise connectivity, 778–779
 - enterprise data center connectivity design, 817–818
 - enterprise IPv6 networks case study, 808–809
 - QoS in the enterprise network case study, 835–836
 - resilient enterprise WANs case study, 825–826
 - for SDN, 419
 - secure enterprise networks case study, 831
- resiliency
 - enterprise campus design, 23
 - high-availability enterprise campus, 23–24
 - network infrastructure devices, 700–701
 - VPLS (Virtual Private LAN Service), 265–266
- resilient enterprise WANs, designing, 825
 - analysis and task list, 826–827
 - requirements and expectations, 825–826
 - selecting WAN links, 828
 - WAN overlays, 828–830
- REST, 422
- restricted VLANs, 773
- restricted zones, 690, 694
- reverse path forwarding (RPF), 635
- RFC 791, 523
- RFC 2474, 523
- RFC 2597, 556
- RFC 3168, 547
- RFC 3171, 636
- RFC 3956, 679
- RFC 4594, 559–560
- RIPv2, migrating to OSPF, 785
- RIR (Regional Internet Registries), 809
- role mapping, MPLS VPNs, 616
- route distinguishers, provider edge (PE) routers, 238–239
- route filtering, 224
- route leaking, IS-IS (Intermediate System-to-Intermediate System), 126–129
- route reflector clients, 155
- route reflector cluster-ID, BGP (Border Gateway Protocol), 161–162
- route reflector clusters, BGP (Border Gateway Protocol), 160–161
- route reflectors, BGP (Border Gateway Protocol), 153–155
 - versus confederations, 157

- congruence of physical and logical networks, 165–167
- hierarchical route reflector design, 167–168
- loop prevention, 162–165
- network design issues, 169
- redundancy, 159–160
- route reflector cluster-ID, 161–162
- route reflector clusters, 160–161
- split-horizon rule, 158–159
- route summarization**
 - black holes, EIGRP (Enhanced Interior Gateway Routing Protocol), 61–63
 - IS-IS (Intermediate System-to-Intermediate System), 136–138
 - OSPF (Open Shortest Path First), PE-CE routing protocol, 251–252
 - suboptimal routing, EIGRP (Enhanced Interior Gateway Routing Protocol), 63–65
- route target (RT), provider edge (PE) routers, 240–241
- routed access, access-distribution block, 14–15
- routed access model, distribution-to-distribution interconnect, 41–42
- routed domains, OSPF (Open Shortest Path First), 78–80
- routed mode, firewalls, 719
- router hardening, 745
- router types, IS-IS (Intermediate System-to-Intermediate System), 106–108
- routers**
 - Auto-RP, 670
 - customer edge (CE) routers, 235
 - P routers, 235
 - PIMv2 BSR, 678
 - provider edge (PE) routers, 235
- routing**
 - ACI (Application-Centric Infrastructure), 465
 - border leaves*, 467–468
 - first-hop layer 3 default gateway*, 465–466
 - default routing, 805–807
 - enterprise routing, WAN, 236–237
 - Internet routing, 803–807
 - inter-VLAN routing, 381–383
 - IS-IS (Intermediate System-to-Intermediate System), 125–126
 - asymmetric versus symmetric*, 129–132
 - flat IS-IS routing design*, 134–135
 - full-mesh design*, 133–134
 - NBMA hub-and-spoke*, 132–133
 - route leaking*, 126–129
- routing information, area and routed domain, OSPF (Open Shortest Path First), 78–80**
- routing infrastructure, security, 699–700**
- routing policies**
 - Asian sites, 799–802
 - enterprise BGP network with Internet connectivity, case study, 797–802
 - European sites, 799–802
 - North American sites, 797–799
- routing policy language (RPL), 169**
- routing propagation, MPLS VPNs, 255–258**

routing protocol authentication mechanisms, 699

routing protocols, choosing, for enterprise BGP network with Internet connectivity design, 792

RP (Rendezvous Point), 665

- Anycast RP, 681
 - examples, 682–683*
- Auto-RP, 668–669
 - candidate RPs, 670*
 - case studies, 670–674*
 - mapping agents, 670*
 - routers, 670*
 - scope problems, 674–676*
- candidate RPs, 676–677
- IPv6 embedded RP, 679–681
- MSDP (Multicast Source Discovery Protocol), 683
 - neighbor relationships, 683*
 - operations case study, 684–686*
- PIMv2 BSR, 676–677
 - BSR (bootstrap router), 678*
 - candidate RPs, 677–678*
 - flooding problems, 678–679*
 - routers, 678*
- placement of, 667–668

RP (Rendezvous Point) discovery, 665–667

RP deployments, comparing, 667

RP registration, PIM-SM (Protocol-Independent Multicast—Sparse Mode), 647–648

RPF (reverse path forwarding), 635

RPF check

- case studies, 641–642
- multicast forwarding, 639–641

RPL (routing policy language), 169

RPO (Recovery Point Objective), 482

RT (route target), provider edge (PE) routers, 240–241

RTO (Recovery Time Objective), 482

S

(S, G)

PIM-SM (Protocol-Independent Multicast—Sparse Mode), 653–654

PIM-SM SPT switchover, 649–652

SA (Security Association), 278

sandbox infrastructures, 740

SAP (Session Announcement Protocol), 748

scalability

iBGP, 152–153

confederations, 155–156

VPLS (Virtual Private LAN Service), 263–265

WAN connections, 231

scalability design, OSPF (Open Shortest Path First), 76

scalability optimization

DMVPN (Dynamic Multipoint VPN), EIGRP (Enhanced Interior Gateway Routing Protocol), 69

hub-and-spoke design, EIGRP (Enhanced Interior Gateway Routing Protocol), 65–68

scalable EIGRP design, 50

scalable passive monitoring, PFRv3, 364

scaling, enterprise connectivity design, 787–788

scavenger traffic, QoS in the enterprise network case study, 839

scheduling, 535

WFQ (weighted fair queueing), 537–538

- scope problems, Auto-RP, 674–676
- SDN (software-defined networking), 398, 414–416
 - benefits of, 416–417
 - challenges of, 419–421
 - nontraditional SDN, 421
 - requirements, 419
 - security, 703–704
 - selection criteria, 417–418
- SDN controller characteristics, 418
- SDWAN (software-defined WAN), 354–355
- secure connectivity, WAN, 357
- secure enterprise networks, designing, 830
 - firewalls, 835
 - infrastructure and network access security, 833–834
 - Layer 2 security, 834–835
 - requirements and expectations, 831
 - security domains and zone design, 832
- secure neighbor discovery (SeND), 222
- secure network access, 695
- secure network design, 695
- Secure Sockets Layer (SSL) VPN, 312–313
- Secure Sockets Layer virtual private network (SLL VPN), 221
- secure tenant separation, multitenant data centers, 422–425
- securing
 - BSR (bootstrap router), 751
 - east-west traffic, 716–717
 - management access, to infrastructure devices, 698–699
- security
 - control plane security, IPv6, 224
 - dual-stack security, IPv6, 225
 - extranets, 739–740
 - firewalls. *See* firewalls
 - infrastructure device access, 698–699
 - internal multicast security, 752
 - IP multicast, 743
 - challenges of*, 744
 - link layer security, IPv6, 221–222
 - multicast network edge, 748–749
 - Auto-RP and BSR*, 749–751
 - MSDP (Multicast Source Discovery Protocol)*, 751–752
 - multicast networks, 745–746
 - network element security, 746–748
 - network infrastructure devices, resiliency and survivability, 700–701
 - network policy enforcement, 701–702
 - network security zoning, 690–691
 - next-generation security, 696
 - routing infrastructure, 699–700
 - SDN (software-defined networking), 703–704
 - switching infrastructure, 702–703
 - tunneling security, IPv6, 225–226
- Security Association (SA), 278
- security domains, designing, 832
- security group access control lists (SGACL), 770
- Security Group Tag Exchange Protocol (SXP), 770
- Security Group Tag (SGT), 769–772
- security services, IPv6, 221

- security zones, modular network architecture, 695
- segmentation, multitenant segmentation, extranets, 739–740
- selecting
 - data center architecture and connectivity model, 818–819
 - replacement routing protocols, 780
 - WAN links, 828
- selection criteria, SDN (software-defined networking), 417–418
- SeND (secure neighbor discovery), 222
- send-community, 169–170
- separate DCI layer deployment model, 500
- separating, application tiers, 714–716
- sequence number packets (SNPs), 123
- server-server traffic, 480
- service graphs, 459
- service migration, enterprise IPv6 networks case study, 815–816
- service provider-managed VPNs 230
- service-level agreement (SLA), WAN connections, 231
- Session Announcement Protocol (SAP), 748
- SGACL (security group access control lists), 770
- SGT (Security Group Tag), 769–772
- sham links, OSPF (Open Shortest Path First), 250–251
- shaping traffic, 527–529, 532
 - WAN QoS, 592–593
- shared distribution trees, 643–644
- shared trees, 642, 643–644
- shortest path trees (SPT), 637
- short-pipe mode, MPLS DiffServ tunneling modes, 610
- short-pipe tunneling mode, MPLS VPNs, 612–614
- show ip community-list, 171
- show ip pim rp mapping, 671
- SIA (stuck in active), 52
- simple demarcation, 329
- single topology restrictions, IS-IS (Intermediate System-to-Intermediate System), 138–139
- single-homed, multiple links, BGP (Border Gateway Protocol), 178–180
- single-homing, versus multihoming, BGP (Border Gateway Protocol), 177–178
- single-rate three-color marker, 532–533
- single-tier firewalls, architecture, 710
- site-to-site VPN solutions, 272–273
- SLA (service-level agreement), WAN connections, 231
- SLAAC (Stateless Address Autoconfiguration), 221
- SLL VPN (Secure Sockets Layer virtual private network), 221
- small data centers (connecting servers to an enterprise LAN), connecting servers to an enterprise LAN, 376–378
- smart probing, 364
- SNPs (sequence number packets), 123
- software-defined networking (SDN), 398, 414–416
 - benefits of, 416–417
 - challenges of, 419–421
 - nontraditional SDN, 421
 - requirements, 419
 - selection criteria, 417–418

- software-defined WAN (SDWAN), 354–355
- solution manageability, 355
- source distribution trees, 643
- source-rooted trees, 642
- source-specific multicast. *See* SSM (source-specific multicast)
- source-specific multicast mode, 655
- spanned EtherChannel, 724
- sparse mode protocols, 642
- speaker types, BGP (Border Gateway Protocol), 147–148
- SPF-Hold, 96
- SPF-Max, 96
- SPF-Start, 96
- spince switches, 439
- spine switches, 401
- spine-leaf topologies, modern data centers, 400–401
- split brain, 485
- split-horizon rule, BGP (Border Gateway Protocol), 148–149
 - route reflectors, 158–159
- spoke-to-spoke, DMVPN (Dynamic Multipoint VPN), 285
- SP-provided VPN services, 230
- SPT (shortest path trees), 637
- SPT switchover, PIM-SM (Protocol-Independent Multicast—Sparse Mode), 649–652
- SSL (Secure Sockets Layer) VPN, 312–313
- SSM (source-specific multicast), 654–656
 - characteristics, 654
- SSM out-of-band source directory, 656
- stages of Pfrv2, 363
- start-interval, 94
- Stateless Address Autoconfiguration (SLAAC), 221
- storage traffic, 480–482
- STP blocking links, GLBP (Gateway Load Balancing Protocol), 35
- STP-based layer LANs, ACI (Application-Centric Infrastructure), 464–465
- stub leaking, EIGRP (Enhanced Interior Gateway Routing Protocol), 67–68
- stuck in active (SIA), 52
- suboptimal bandwidth utilization, 541–542
- suboptimal routing, route
 - summarization, EIGRP (Enhanced Interior Gateway Routing Protocol), 63–65
- summarization
 - choke points and, 55–56
 - hub-and-spoke design, EIGRP (Enhanced Interior Gateway Routing Protocol), 61–65
 - OSPF (Open Shortest Path First), 85–86
 - route summarization, IS-IS (Intermediate System-to-Intermediate System), 136–138
- supplicants, 759
 - 802.1X, 765–766
- supported traffic, WAN connections, 232
- survivability, network infrastructure devices, 700–701
- SVI (switched virtual interface), 468
- switched virtual interface (SVI), 468
- switching infrastructure, 702–703
- SXP (Security Group Tag Exchange Protocol), 770

symmetric routing versus asymmetric routing, IS-IS (Intermediate System-to-Intermediate System), 129–132

synchronization, LSDB
synchronization, IS-IS (Intermediate System-to-Intermediate System), 123–124

T

TACACS+ 833

tail drop, 544

task lists, enterprise connectivity, 779–780

TCP windowing, 547

TDM (time-division multiplexing), 530

TEAP (Tunnel Extensible Authentication Protocol), 765

teleworker, 693

tenant address learning, remote VTEP discovery, 411–413

tenant separation

- device-level virtualization, 424–425
- multitenant data centers, 422–425

tenants

- ACI (Application-Centric Infrastructure), 456–459
- multitenant data centers, 422

TEP (tunnel endpoint), 441

theft of service, 754

three-layer hierarchy architecture, EIGRP (Enhanced Interior Gateway Routing Protocol), 57–59

three-tier data center network architecture, 380–381

three-tier layer model, enterprise campus design, 9–10

three-tiered e-commerce application functional model, 714

TID (Transport-Independent Design), 356

time-division multiplexing (TDM), 530

TLVs (type, length, value), 103

token bucket algorithms, 529–531

tools

- dropping tools, DSCP-based WRED, 541–546

- policing tools. *See* policing tools

- queueing tools, 535–536

topology depths, 54

ToR (Top of Rack) design, 383–384

traffic

- east-west traffic, 716–717

- scavenger traffic, 839

traffic descriptors, QoS (quality of service), 516–517

traffic discovery, QoS in the enterprise network case study, 836–837

traffic engineering techniques, 351–354

traffic filtering, Layer 2 segments, 703

traffic flow directions, 478–479

traffic flow types, 479–482

traffic policing

- ECN (explicit congestion notification), 547–550

- QoS (quality of service), 527–529

traffic shaping, QoS (quality of service), 527–529

traffic trombone, 487

trail drop, 547

transit border router, 366

transit link, remote-site LANs, 343
 transit master controller, 365
 transition mechanisms
 IPv6, 216–217
 IPv6 migration, 203–205
 transparent mode, firewalls, 719
 transport mode, IPsec VPN, 621
 transport options
 for remote sites using local Internet, 338–339
 remote sites using local Internet, 350–351
 Transport-Independent Design (TID), 356
 trunked demarcation, 329
 trunking, high-availability enterprise campus, 27
 trust CoS, 571
 trust DSCP, 571
 trust states
 boundaries and, 570–573
 dynamic trust states, 572–573
 TrustSec, 768
 Profiling Service, 768–769
 SGT (Security Group Tag), 769–772
 TTL Security Check, 7006
 tunnel broker approach, 202
 tunnel brokers, IPv6, 209
 Tunnel Extensible Authentication Protocol (TEAP), 765
 tunnel mode, IPsec VPN, 621
 tunneled EAP, 764
 tunneling modes, MPLS DiffServ tunneling modes, 609–611
 tunneling security, IPv6, 225–226
 tunnels, manual tunnels, IPv6, 208–209
 tuples, 103

two-layer hierarchy architecture, EIGRP (Enhanced Interior Gateway Routing Protocol), 56–57
 two-rate three-color marker, 533–535
 two-tier data center network architecture, 378–380
 two-tier firewall, architecture, 710
 two-tier layer model, enterprise campus design, 8–9
 Tx-Ring, 536–537, 591

U

unicast, 635
 unicast rekeying, 318
 unicast reverse pack forwarding (uRPF), 702
 uniform mode, MPLS DiffServ tunneling modes, 610
 uniform tunneling mode, MPLS VPNs, 612
 untrusted, 571
 update process, IS-IS (Intermediate System-to-Intermediate System), 118
 uRPF (unicast reverse pack forwarding), 702

V

VDCs (virtual device contexts), 424–425
 video, QoS (quality of service), 568–569
 virtual device contexts (VDCs), 424–425
 virtual extensible LAN (VXLAN), 407–408
 virtual firewalls, 712

- Virtual MAC (vMAC), 489
- virtual machines (VMs), 716–717
- virtual network interface cards (vNICs), 715–716
- Virtual Network Management Center (VNMC), 713
- virtual private LAN service DCI, 496
- Virtual Private LAN Service (VPLS), 259, 261–263, 265–266
 - scalability, 263–265
- Virtual Private Wire Service (VPWS), 259–261
- Virtual Router Redundancy Protocol (VRRP), 31
- virtual routing and forwarding. *See* VRF (virtual routing and forwarding)
- Virtual Security Gateway (VSG), 713
- virtual switch model, distribution-to-distribution interconnect, 43–44
- virtual switch (switch clustering), access-distribution block, 13–14
- virtual switching system (VSS), IP gateway redundancy, 35–36
- virtual tunnel interface (VTI), IPsec and, 281–282
- virtualization
 - campus network virtualization, 16–23
 - device-level virtualization, 424–425
- virtualized firewalls, 712–714
- virtualized multiservice architectures, 596–597
- Virtualized Multiservice Data Centers (VMDC), 596–597
- VLAN assignment
 - campus network virtualization, 17–18
 - dynamic VLAN assignments, 772–774
 - VLAN design, high-availability enterprise campus, 24–26
- vMAC (Virtual MAC), 489
- VMDC (Virtualized Multiservice Data Centers), 596–597
- VMs (virtual machines), 716–717
- vNICs (virtual network interface cards), 715–716
- VNMC (Virtual Network Management Center), 713
- voice traffic, QoS (quality of service), 568–569
- VoIP (voice over IP), QoS (quality of service), 568–569
- vPC, 388–392
 - firewall routing, 725
- VPLS (Virtual Private LAN Service), 259, 261–263
 - DCI (Data Center Interconnect), 496
 - resiliency, 265–266
 - scalability, 263–265
 - versus VPWS, 266–267
- VPN use cases, QoS (quality of service), 621
- VPN WAN design models, 331–335
- VPNs (virtual private networks)
 - enterprise-managed VPNs. *See* enterprise-managed VPNs
 - FlexVPN, 314
 - architecture*, 315
 - capabilities*, 315
 - configuration blocks*, 315–316
 - GETVPN, 317–320
 - Layer 3 MPLS VPNs, 233–234
 - managed VPNs, 230
 - MPLS VPNs, architecture, 234–236
 - security, 695

- service provider-managed VPNs 230
 - SSL (Secure Sockets Layer) VPN, 312–313
 - VPWS (Virtual Private Wire Service), 259–261**
 - versus VPLS, 266–267
 - VRF (virtual routing and forwarding)**
 - campus network virtualization, 18
 - firewalls, 712
 - VRF-Lite, Layer 3 separation, 423–424**
 - VRRP (Virtual Router Redundancy Protocol), 31**
 - VSG (Virtual Security Gateway), 713, 714**
 - VSS (virtual switching system), IP gateway redundancy, 35–36**
 - VTEP (VXLAN tunnel endpoint), 408–411, 441**
 - VTI (virtual tunnel interface), IPsec and, 281–282**
 - VXLAN (virtual extensible LAN), 407–408**
 - control plane optimization, 413–414
 - overlay networks, microsegmentation, 427–428
 - remote VTEP discovery, 411–413
 - VTEP (VXLAN tunnel endpoint), 408–411
 - VXLAN tunnel endpoint (VTEP), 408–411**
-
- W**
-
- WAN (Wide Area Network)**
 - 3G/4G VPN design models, 335
 - application optimization, 356–357
 - case studies, redundancy and connectivity, 343–354
 - enterprise routing, 236–237
 - intelligent path control, 356
 - IWAN (Intelligent WAN), 354–355
 - Layer 2 WAN design models, 329–331
 - management, 357–358
 - MPLS Layer 3 WAN design models, 326–329
 - NGWAN (next-generation WAN), 354–355
 - regional offices WAN design, 348–351
 - remote sites, local Internet, 337–339
 - remote-site LANs, 339–343
 - remote-site WAN design, 346–348
 - SDWAN (software-defined WAN), 354–355
 - secure connectivity, 357
 - TID (Transport-Independent Design), 356
 - traffic engineering techniques, 351–354
 - VPN WAN design models, 331–335
 - WAN aggregation, 325–326, 327**
 - WAN connections, choosing, 230–233**
 - WAN links, selecting, 828**
 - WAN overlays, resilient enterprise WANs case study, 828–830**
 - WAN QoS**
 - examples, 593–594
 - latency and jitter, 590–591
 - overview, 588–589
 - platform performance, 589–590
 - queueing, 591–592
 - shaping traffic, 592–593
 - WAN remote sites, overview, 324–326**

- WAN remote-site design models, 328
- WAN remote-site transport options, 325–326
- WAN/branch edge, 588–589
 - CBWFQ (class-based weighted fair queueing), 592
- WAN/VPN QoS design, 593
- WDM, 490
- web proxy, 740
- Web Security Appliance (WSA), 735–736
- weighted fair queueing (WFQ), 536
- well-known BGP communities, 170–171
- WFQ (weighted fair queueing), 536, 537–538
- winner messages, DF election messages, 660
- wired networks, 802.1X, 760

- wireless LAN controller (WLC), 771
- WLC (wireless LAN controller), 771
- WRED, 544–546, 547, 591
 - dropping design principles, 557–558
 - ECN (explicit congestion notification), 548–549
- WSA (Web Security Appliance), 735–736

X-Y-Z

- zone interface points, 690
- zones
 - designing, 832
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 54
 - modular network architecture, 695
- zoning, 690–691