



Study
Planner



Video
Training



Flash
Cards



Practice
tests



Review
Exercises

Official Cert Guide

Advance your IT career with hands-on learning

CCNP and CCIE Data Center Core DCCOR 350-601

ciscopress.com

SOMIT MALOO, CCIE NO. 28603, CCDE NO. 20170002

FIRAS AHMED, CCIE NO.14967

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



CCNP and CCIE Data Center Core DCCOR 350-601 Official Cert Guide

SOMIT MALOO, CCIE No. 28603, CCDE No. 20170002

FIRAS AHMED, CCIE No. 14967

Cisco Press

CCNP and CCIE Data Center Core DCCOR 350-601 Official Cert Guide

Somit Maloo & Firas Ahmed

Copyright © 2020 Cisco Systems, Inc.

Published by:
Cisco Press

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.

ScoutAutomatedPrintCode

Library of Congress Control Number: 2020930073

ISBN-13: 978-0-13-644962-1

ISBN-10: 0-13-644962-X

Warning and Disclaimer

This book discusses the content and skills needed to pass the 350-601 CCNP Data Center Core certification exam, which is the prerequisite for CCNP as well as CCIE certification. 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 authors 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: Arezou Gol

Director, ITP Product Management: Brett Bartow

Senior Editor: James Manly

Managing Editor: Sandra Schroeder

Development Editor: Ellie Bru

Senior Project Editor: Tonya Simpson

Copy Editor: Chuck Hutchinson

Technical Editor: Ozden Karakok

Editorial Assistant: Cindy Teeters

Cover Designer: Chuti Prasertsith

Composition: codeMantra

Indexer: Ken Johnson

Proofreader: Charlotte Kughen



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

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

About the Authors

Somit Maloo, CCIE No. 28603, CCDE No. 20170002, is a content lead from the data center team in the Learning@Cisco organization. He holds a master's degree in telecommunication networks and a bachelor's degree in electronics and telecommunication engineering. He is also a penta CCIE in routing and switching, service provider, wireless, security, and data center technologies. Somit holds various industry-leading certifications, including CCDE, PMP, RHCSA, and VMware VCIX6 in Data Center and Network Virtualization. Somit has extensive experience in designing and developing various data center courses for the official Cisco curriculum. He started his career as a Cisco TAC engineer. Somit has more than 10 years of experience in the networking industry, working mostly with data center networks. You can reach Somit on Twitter: @somitmaloo.

Firas Ahmed, CCIE No. 14967, is a solution architect from the enterprise data center team at Cisco Customer Experience (CX). He completed a master's degree in systems and control engineering following a bachelor's degree in computer engineering. Firas holds CCIE certificates in routing and switching, collaboration, wireless, security, and data center technologies in addition to industry-based certifications, including CISSP, PMP, VMware VCP6.5-DCV, ITIL, and GICSP. Firas has more than 15 years of experience in designing, developing, and supporting various data centers for enterprise and IoT customers. Firas has additional experience as a seasonal instructor in a number of community colleges in Toronto, where he taught various computer networking courses. You can reach Firas on Twitter: @dccor_firas.

About the Technical Reviewer

Ozden Karakok, CCIE No. 6331, is a technical consultant on data center technologies and solutions at Flint Consulting. She worked at Cisco for 19 years as a technical leader supporting data center solutions. Prior to joining Cisco, Ozden spent five years working for a number of Cisco's large customers in various telecommunication roles. She is a Cisco Certified Internetwork Expert in routing and switching, SNA/IP, and storage. She co-authored three Cisco Press books: *CCNA Data Center DCICN 200-150*, *CCNA Data Center DCICT 200-155*, and *Data Center Technologies DCICT 640-916*. Ozden holds a degree in computer engineering from Istanbul Bogazici University. You can reach Ozden on Twitter: @okarakok.

Dedications

Somit:

To my loving wife, Renuka, for her unending love and support.

To my wonderful parents, who supported me in every phase of my life.

To Navya and Namit, who agreed not to fight while Papa was working on the book.

To my aunt, Tara, for being the guiding angel in my life.

Firas:

To my amazing wife, Nora, who has been extremely supportive throughout this process. Thanks for letting me spend long hours on my computer once again!

To Ibrahim and Maryam, you are growing so fast. Never give up on what you want. If at first you don't succeed, try and try again. I love you more than anything!

To my parents, you are still the guiding light that keeps me on the right path.

Acknowledgments

Somit Maloo:

I would like to thank my co-author, Firas Ahmed, for working as a team to complete this book. Without your support, this book would not have been possible. I am thankful to all our professional editors, especially Mary Beth Ray, James Manly, and Ellie Bru, for their patience and guidance at every step of the book process. I would also like to thank our technical editor, Ozden Karakok, for her keen attention to detail and for agreeing to review the book, taking time out of her busy schedule.

Firas Ahmed:

I would like to thank my co-author, Somit Maloo, for taking the initiative to form this partnership and for his dedication in putting together the outline of this book. Thank you for your valuable input and continuous support throughout the process.

A special recognition to the technical editor, Ozden Karakok, for her valuable contributions in reviewing the book. Her suggestions and feedback helped improve the quality of the content in multiple areas.

Thanks to the Cisco Press team, especially Mary Beth Ray and James Manly, for believing in us, and Ellie Bru, for her guidance and extreme patience while editing and amending the chapters of the book.

A special credit to Hazim Dahir, distinguished engineer at Cisco Systems, for his help and support with the technical review of the book.

In addition, I want to thank my colleague Naveen Chapa for reviewing and providing constructive feedback that helped enhance the ACI chapter.

Contents at a Glance

Introduction xxx

Part I Networking

- Chapter 1 Implementing Routing in the Data Center 2
- Chapter 2 Implementing Data Center Switching Protocols 90
- Chapter 3 Implementing Data Center Overlay Protocols 148
- Chapter 4 Describe Cisco Application Centric Infrastructure 194
- Chapter 5 Cisco Cloud Services and Deployment Models 264
- Chapter 6 Data Center Network Management and Monitoring 276

Part II Storage

- Chapter 7 Implement Fibre Channel 346
- Chapter 8 Implement FCoE Unified Fabric 418
- Chapter 9 Describe NFS and NAS Concepts 460
- Chapter 10 Describe Software Management and Infrastructure Monitoring 470

Part III Compute

- Chapter 11 Cisco Unified Computing Systems Overview 514
- Chapter 12 Cisco Unified Computing Infrastructure Monitoring 610
- Chapter 13 Cisco Unified Compute Software and Configuration Management 640
- Chapter 14 Cisco HyperFlex Overview 684

Part IV Automation

- Chapter 15 Automation and Scripting Tools 712
- Chapter 16 Evaluate Automation and Orchestration Technologies 744

Part V Security

- Chapter 17 Network Security 792
- Chapter 18 Compute Security 864
- Chapter 19 Storage Security 892
- Chapter 20 Final Preparation 928

Appendix A: Answers to the “Do I Know This Already?” Quizzes	938
Appendix B: <i>CCNP and CCIE Data Center Core DCCOR 350-601 Official Cert Guide Exam Updates</i>	952
Glossary	954
Index	978

Online Elements

Glossary

Appendix C: Memory Tables

Appendix D: Memory Tables Answer Key

Appendix E: Study Planner

Reader Services

Other Features

In addition to the features in each of the core chapters, this book has additional study resources on the companion website, including the following:

Practice exams: The companion website contains an exam engine that enables you to review practice exam questions. Use these to prepare with a sample exam and to pinpoint topics where you need more study.

An online interactive Flash Cards application to help you drill on Key Terms by chapter.

Glossary quizzes: The companion website contains interactive quizzes that enable you to test yourself on every glossary term in the book.

More than 2 hours of video training: The companion website contains multiple hours of unique test-prep videos.

To access this additional content, simply register your product. To start the registration process, go to www.ciscopress.com/register and log in or create an account*. Enter the product ISBN 9780136449621 and click Submit. After 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.

Contents

Introduction xxx

Part I Networking

Chapter 1 Implementing Routing in the Data Center 2

“Do I Know This Already?” Quiz 2

Foundation Topics 5

OSPF 5

OSPF Link-State Advertisements 6

OSPF Areas 9

Designated Routers and Backup Designated Routers 11

OSPF Authentication 11

OSPF Configurations and Verifications 12

Border Gateway Protocol 23

BGP Peering 24

BGP Path Selection 25

Step 1: Comparing Pairs of Paths 25

Step 2: Determining the Order of Comparisons 27

Step 3: Determining the Best-Path Change Suppression 27

Multiprotocol BGP 28

BGP Configurations and Verifications 28

Bidirectional Forwarding Detection 36

Rapid Detection of Failures 37

BFD Configurations and Verifications 37

Multicast 41

Internet Group Management Protocol 41

Switch IGMP Snooping 44

Multicast Listener Discovery 44

Multicast Distribution Trees 45

Protocol Independent Multicast 48

PIM Rendezvous Points 52

PIM Designated Routers/Forwarders 53

Multicast Forwarding 53

Multicast Configurations and Verifications 54

Hot Standby Router Protocol 68

Virtual Router Redundancy Protocol 72

VRRP Operation 72

	VRRP Groups	74
	VRRP Router Priority and Preemption	74
	VRRP Authentication	75
	VRRP Tracking	75
	IPv6 First Hop Redundancy	76
	HSRP/VRRP Configurations and Verifications	77
	Exam Preparation Tasks	86
	Review All Key Topics	86
	Memory Tables	87
	Define Key Terms	87
	References	87
Chapter 2	Implementing Data Center Switching Protocols	90
	“Do I Know This Already?” Quiz	90
	Foundation Topics	93
	Spanning Tree Protocols	93
	STP Topology	93
	STP Port Types	94
	STP Extensions	94
	<i>STP Bridge Assurance</i>	95
	<i>BPDU Guard</i>	96
	<i>BPDU Filter</i>	96
	<i>Loop Guard</i>	96
	<i>Root Guard</i>	97
	Unidirectional Link Detection	97
	Rapid PVST+	98
	<i>Rapid PVST+ Ports</i>	100
	Spanning Tree Configurations and Verifications	102
	Port Channels	117
	<i>Port Channel Load Balance</i>	120
	Virtual Port Channel	121
	vPC Traffic Flows	124
	vPC Dual-Control Plane	125
	vPC Primary and Secondary Roles	126
	vPC Configuration Consistency	127
	vPC Duplicate Frames Prevention Mechanism	128
	vPC HSRP Gateway Considerations	130
	vPC ARP Synchronization	130

vPC Peer Gateway	130
Port Channel Configurations and Verifications	131
Exam Preparation Tasks	145
Review All Key Topics	145
Memory Tables	146
Define Key Terms	146
References	147

Chapter 3 Implementing Data Center Overlay Protocols 148

“Do I Know This Already?” Quiz	148
Foundation Topics	150
Overlay Transport Virtualization (OTV)	150
OTV Terminology	150
OTV Control Plane Function	151
<i>Multicast-Enabled Transport Infrastructure</i>	151
<i>Unicast-Only Transport Infrastructure (Adjacency-Server Mode)</i>	152
OTV Data Plane Function	154
<i>Unicast Traffic over OTV</i>	154
<i>Multicast Traffic over OTV</i>	156
<i>Broadcast Traffic over OTV</i>	156
Failure Isolation	157
STP Isolation	157
Unknown Unicast Handling	157
ARP Optimization	158
Broadcast Policy Control	159
Multihoming OTV	159
FHRP Isolation	162
OTV Configurations and Verifications	163
Virtual Extensible LAN (VXLAN) Overview	173
VXLAN Encapsulation and Packet Format	173
VXLAN Tunnel Endpoint	174
Virtual Network Identifier	175
VXLAN Control Plane	176
<i>VXLAN Flood and Learn Multicast-Based Control Plane</i>	176
<i>VXLAN MPBGP EVPN Control Plane</i>	178
VXLAN Gateways	178
VXLAN High Availability	179
VXLAN Tenant Routed Multicast	180
VXLAN Configurations and Verifications	182

Exam Preparation Tasks	191
Review All Key Topics	192
Define Key Terms	192
References	192
Chapter 4 Describe Cisco Application Centric Infrastructure	194
“Do I Know This Already?” Quiz	194
Foundation Topics	196
Cisco Application Centric Infrastructure (ACI) Overview	196
Cisco Application Policy Infrastructure Controller	198
Cisco Nexus 9000 Series Spine and Leaf Switches for Cisco ACI	201
Cisco ACI Initial Setup, Fabric Discovery, Access Policy, and VMM Domains	204
Cisco ACI Initial Setup	204
Cisco ACI Fabric Discovery	209
Startup with Cisco ACI Fabric Discovery and Configuration	210
Fabric Upgrade	212
ACI Policy Model	212
Tenants	214
Virtual Routing and Forwarding Objects	214
Bridge Domains and Subnets	214
Endpoint Groups	215
Application Profiles	215
Microsegmentation	215
Attachable Entity Profile	216
Cisco ACI Fabric Policies	216
Cisco ACI Virtual Machine Manager Domains	222
Cisco ACI integration with Microsoft SCVMM	223
Cisco ACI Integration with VMware vCenter	224
Integrating VMware Overlays with the Cisco ACI	225
Cisco ACI Virtual Edge	225
Cisco ACI Fabric: Tenants and Packet Flow	225
Cisco ACI Tenants	227
Virtual Routing and Forwarding	228
Bridge Domain and Subnets	229
Application Profile	230
Endpoint Group	231
Microsegmentations	231
ACI Contract	231

Taboo Contracts	233
vzAny Rule	233
Filters and Subjects	236
Management Tenant	237
<i>In-Band Management Access</i>	237
<i>Out-of-Band Management Access</i>	238
ACI VXLAN	239
ACI Intersubnet Tenant Traffic	241
Policy Identification and Enforcement	242
ACI Fabric Traffic Storm Control	243
ACI Fabric Traffic Load Balance	243
ACI Fabric Loop Detection	244
ACI Design Best Practices	245
ACI LAB Configurations Example	245
Building ACI Fabric	248
Creating Tenant	251
Creating Contract and Filter	254
Deploying a Three-Tier Application	257
Integrating with vCenter	259
Exam Preparation Tasks	262
Review All Key Topics	262
Define Key Terms	262
References	263

Chapter 5 Cisco Cloud Services and Deployment Models 264

“Do I Know This Already?” Quiz	264
Foundation Topics	266
What Is Cloud Computing?	266
Cloud Service Models	269
Software as a Service	269
Platform as a Service	270
Infrastructure as a Service	270
Cloud Deployment Models	272
Private Cloud	272
Public Cloud	272
Hybrid Cloud	273
Community Cloud	274
Exam Preparation Tasks	274
Review All Key Topics	274

Define Key Terms 275

References 275

Chapter 6 Data Center Network Management and Monitoring 276

“Do I Know This Already?” Quiz 276

Foundation Topics 278

Cisco Nexus NX-OS Software Installation, Updates, and Their Impacts 278

PowerOn Auto Provisioning (POAP) 283

Data Center Infrastructure Software Lifecycle Management 287

Nexus Software Maintenance Upgrade 287

Programmable Logical Devices Upgrade 289

Graceful Insertion and Removal 291

Nexus Nondisruptive In-Service Software Upgrade 295

Nexus Disruptive and Nondisruptive Upgrade/Downgrade Procedure 299

Nexus Configuration Management 303

NX-OS Configuration Save and Backup 303

Nexus Config Rollback and Checkpoint 303

Network Infrastructure Monitoring 306

NX-OS System Message Logging 306

Network Time Management 307

Network Time Protocol 307

Precision Time Protocol 313

NX-OS Simple Network Management Protocol 317

Nexus Smart Call Home 324

Nexus NetFlow 325

Switched Port Analyzer 330

Streaming Telemetry 337

Network Assurance Concept 341

Exam Preparation Tasks 344

Review All Key Topics 344

Memory Tables 344

Define Key Terms 345

References 345

Part II Storage

Chapter 7 Implement Fibre Channel 346

“Do I Know This Already?” Quiz 346

Foundation Topics 349

Fibre Channel Basics 349

Fibre Channel Topologies	350
Fibre Channel Port Types	353
<i>E Port</i>	353
<i>F Port</i>	354
<i>NP Ports</i>	354
<i>TE Port</i>	354
<i>TF Port</i>	354
<i>TNP Port</i>	354
<i>Fx Port</i>	354
<i>Auto Mode</i>	354
Fibre Channel Addressing	355
Flow Control	356
Switched Fabric Initialization	358
<i>Principal Switch Selection</i>	358
<i>Domain ID Distribution</i>	360
<i>FCID Allocation</i>	362
<i>Fabric Reconfiguration</i>	362
Device Registration: FLOGI, PLOGI, PRLI	362
FLOGI and FCNS Databases	363
CFS	364
CFS Features	365
CFS Fabric Lock	366
CFS <i>o</i> IP and CFS <i>o</i> FC	367
CFS Merge	368
CFS Regions	369
VSAN	370
VSAN Features	370
VSAN Attributes	372
VSAN Advantages	373
Dynamic Port VSAN Membership (DPVM)	373
VSAN Trunking	374
SAN Port Channels	381
Types of SAN Port Channels	381
Port Channel Load Balancing	383
Port Channel Modes	384
Zoning	389
Zoning Features	389
Zone Enforcement	391

	Full and Active Zone Set	392
	Autozone	395
	Zone Merge	395
	Smart Zoning	396
	Enhanced Zoning	397
	Device Alias	403
	Device Alias Features	403
	Device Alias Modes	404
	Device Alias Distribution	405
	Zone Aliases (FC Aliases) Versus Device Aliases	406
	NPIV and NPV	409
	Exam Preparation Tasks	416
	Review All Key Topics	416
	Memory Tables	417
	Define Key Terms	417
	References	417
Chapter 8	Implement FCoE Unified Fabric	418
	“Do I Know This Already?” Quiz	418
	Foundation Topics	420
	FCoE Overview	420
	Ethernet Enhancements	422
	<i>Priority-Based Flow Control (PFC)</i>	422
	<i>Enhanced Transmission Selection (ETS)</i>	423
	<i>Data Center Bridging Exchange (DCBX)</i>	424
	FCoE Frame Format	426
	Virtual Fibre Channel (VFC)	428
	FCoE Elements and Port Types	429
	FCoE Addressing and Forwarding	431
	FCoE Initialization Protocol (FIP)	432
	Benefits of FCoE	435
	FCoE Topology Options	435
	FCoE Single-Hop Topology	435
	<i>FCoE Direct-Attached Topology</i>	436
	<i>FCoE FEX Topology</i>	437
	<i>FCoE Remote-Attached Topology</i>	438
	FCoE Multi-Hop Topology	438

FCoE Implementations	439
FCoE Configuration on Cisco Nexus 7000 Series Switches	440
<i>Miscellaneous FCoE Configuration</i>	441
FCoE Configuration on Cisco Nexus 5000 Series Switches	442
FCoE over FEX	444
FCoE NPV	445
FCoE Verification	448
Exam Preparation Tasks	457
Review All Key Topics	457
Memory Tables	458
Define Key Terms	458
References	458

Chapter 9 Describe NFS and NAS Concepts 460

“Do I Know This Already?” Quiz	460
Foundation Topics	461
Describe NFS Concepts	461
Describe NAS Concepts	463
NAS Benefits	465
Cisco NSS3000 Series Network Storage System	465
Exam Preparation Tasks	467
Review All Key Topics	467
Memory Tables and Lists	468
Define Key Terms	468
References	468

Chapter 10 Describe Software Management and Infrastructure Monitoring 470

“Do I Know This Already?” Quiz	470
Foundation Topics	472
Cisco MDS NX-OS Setup Utility	472
Cisco MDS NX-OS Software Upgrade and Downgrade	480
Nondisruptive Upgrade on a Cisco MDS Fabric Switch	482
Disruptive Upgrade on a Cisco MDS Fabric Switch	487
Nondisruptive Downgrade on a Cisco MDS Fabric Switch	490
Disruptive Downgrade on a Cisco MDS Fabric Switch	495
EPLD Upgrade on Cisco MDS 9000 Series Switches	498
Infrastructure Monitoring	503
System Messages	503
Call Home	504

	Embedded Event Manager	505
	RMON	505
	SPAN	505
	<i>SPAN Configuration Example</i>	508
	<i>Remote SPAN</i>	509
	Exam Preparation Tasks	511
	Review All Key Topics	511
	Define Key Terms	511
	References	512
Part III	Compute	
Chapter 11	Cisco Unified Computing Systems Overview	514
	“Do I Know This Already?” Quiz	514
	Foundation Topics	516
	Cisco UCS Architecture	516
	Cisco UCS Components and Connectivity	518
	Cisco UCS 5108 Blade Server Chassis	520
	UCS Blade Servers	520
	Cisco UCS Rack Servers	521
	Cisco UCS Storage Servers	521
	Cisco UCS Mini	523
	Cisco UCS Fabric Infrastructure	524
	Cisco UCS 6454 Fabric Interconnect	524
	Cisco UCS 6300 Series Fabric Interconnects	526
	Fabric Interconnect and Fabric Extender Connectivity	527
	Cisco UCS Virtualization Infrastructure	533
	Cisco UCS Initial Setup and Management	536
	Fabric Interconnect Connectivity and Configurations	544
	<i>Uplink Connectivity</i>	546
	<i>Downlink Connectivity</i>	546
	Fabric Interconnect Port Modes	547
	Fabric Failover for Ethernet: High-Availability vNIC	549
	Ethernet Switching Mode	550
	UCS Device Discovery	556
	Chassis /FEX Discovery	556
	Rack Server Discovery Policy	557
	Initial Server Setup for Standalone UCS C-Series	557

Network Management	563
UCS Virtual LAN	563
<i>Named VLANs</i>	566
<i>Private VLANs</i>	570
UCS Identity Pools	571
<i>Universally Unique Identifier Suffix Pools</i>	572
<i>MAC Pools</i>	573
<i>IP Pools</i>	574
<i>Server Pools</i>	576
Service Profiles	577
UCS Server Policies	580
UCS Service Profile Templates	583
Quality of Service	589
<i>QoS System Classes</i>	589
<i>QoS System Classes Configurations</i>	590
<i>Configuring Quality of Service Policies</i>	591
UCS Storage	592
UCS SAN Connectivity	592
UCS SAN Configuration	596
Virtual Storage-Area Networks	597
<i>Named VSANs Configurations</i>	597
<i>Zones and Zone Sets</i>	599
World Wide Name Pool	603
SAN Connectivity Policies	605
Exam Preparation Tasks	606
Review All Key Topics	606
Define Key Terms	607
References	608

Chapter 12 Cisco Unified Computing Infrastructure Monitoring 610

“Do I Know This Already?” Quiz	610
Foundation Topics	612
Cisco UCS System Monitoring	612
Data Management Engine	612
Application Gateway	613
Northbound Interfaces	614
Cisco UCS Monitoring Events and Logs	614
Cisco UCS Monitoring Policies	616

	<i>Cisco UCS Simple Network Management Protocol</i>	618
	<i>Cisco UCS Call Home and Smart Call Home</i>	619
	<i>Cisco UCS Manager Database Health and Hardware Monitoring</i>	620
	<i>Cisco UCS NetFlow Monitoring</i>	620
	Traffic Monitoring	622
	<i>Traffic Monitoring Across Ethernet</i>	623
	<i>Traffic Monitoring Across Fibre Channel</i>	624
Cisco Intersight		629
	Intersight Management as a Service	630
	Intersight as a Telemetry Data Collection	632
	Cisco Intersight Supported Software	632
	Cisco Intersight Licensing	632
	Exam Preparation Tasks	637
	Review All Key Topics	637
	Define Key Terms	637
	References	638
Chapter 13	Cisco Unified Compute Software and Configuration Management	640
	“Do I Know This Already?” Quiz	640
	Foundation Topics	642
	Cisco UCS Configuration Management	642
	Creating and Running a Backup Operation	643
	Backup Policies	648
	Backup Policy Configuration	648
	Import Backups	650
	Enable the Import Operation	651
	System Restore	652
	Restoring the Configuration for a Fabric Interconnect	653
	UCS Firmware and Software Updates	654
	Firmware Version Terminology	661
	Firmware Upgrades Through Auto Install	662
	Direct Upgrade After Auto Install Procedure	666
	Install Infrastructure Firmware Procedure	670
	Upgrading the Server Firmware with Auto Install	673
	Standalone Cisco UCS C-Series Server Firmware Upgrade Using the Host Upgrade Utility (HUU)	675
	Downloading and Preparing the ISO for an Upgrade	676

Exam Preparation Tasks 682

Review All Key Topics 682

Define Key Terms 682

References 682

Chapter 14 Cisco HyperFlex Overview 684

“Do I Know This Already?” Quiz 684

Foundation Topics 686

Cisco HyperFlex Solution and Benefits 686

HyperFlex Benefits 689

Intelligent End-to-End Automation 690

Unified Management for All Workloads 691

Independent Resource Scaling 692

Superior Virtual Machine Density with Lower and Consistent Latency 693

HyperFlex as an Edge, Hybrid, and All-Flash Nodes 694

HyperFlex as an Edge Device 694

HyperFlex Hyperconverged Multicloud Platform (Hybrid or All-Flash) 696

HyperFlex All NVMe 697

Cisco HyperFlex Data Platform 698

HX Storage Cluster Physical Components 699

HX Data Platform High Availability 700

HX Data Platform Cluster Tolerated Failures 701

HX Data Platform Ready Clones 701

HX Data Platform Native Snapshots 701

HX Cluster Interfaces 702

HX Self-Encrypting Drives 702

Configuring a Local Encryption Key 703

Managing HX Disks in the Cluster 703

Managing HX Datastores 706

Expand Cisco HX System Clusters 707

Enabling HX Logical Availability Zones 708

Exam Preparation Tasks 710

Review All Key Topics 710

Define Key Terms 710

References 711

Part IV Automation**Chapter 15 Automation and Scripting Tools 712**

- “Do I Know This Already?” Quiz 712
- Foundation Topics 715
- EEM Overview 715
 - Policies 715
 - Event Statements 716
 - Action Statements 716
 - Configuring EEM 717
 - Verifying the EEM Configuration 718
- Scheduler 718
 - Configuring Scheduler 719
 - Verifying Scheduler Configuration 721
- Bash Shell for Cisco NX-OS 722
 - Managing Feature RPMs 724
 - Managing Patch RPMs 724
- Guest Shell for Cisco NX-OS 725
 - Accessing the Guest Shell 725
 - Resources Used for the Guest Shell 726
 - Capabilities in the Guest Shell 726
 - Managing the Guest Shell 728
- XML 730
 - Example 731
 - XML Syntax 732
- JSON 733
- Rest API 734
 - Authentication 735
 - Response 736
 - NX-API 737
 - NX-API Request and Response Elements 739
 - NX-API Developer Sandbox 741
- Exam Preparation Tasks 742
- Review All Key Topics 742
- Memory Tables 743
- Define Key Terms 743
- References 743

Chapter 16 Evaluate Automation and Orchestration Technologies 744

“Do I Know This Already?” Quiz	745
Foundation Topics	747
Ansible	747
Ansible Components	748
Important Ansible Concepts	749
Ansible CLI Tools	750
Cisco NX-OS and Ansible Example	750
Puppet	751
Puppet Workflow	752
Puppet and NX-OS Environment Integration	753
Puppet Master Installation	754
Puppet Agent Installation	754
Resource Types	756
Sample Manifest: OSPF	756
Puppet and Cisco UCS Manager Integration	757
Python	758
Python Package for Cisco	758
Using the CLI Command APIs	760
Python in Interactive Mode	761
Python in Noninteractive Mode	762
UCS Manager Python SDK	764
Convert to UCS Python	766
PowerOn Auto Provisioning (POAP)	767
Limitations of POAP	767
Network Requirements for POAP	767
POAP Configuration Script	768
POAP Process	768
<i>Power-Up Phase</i>	770
<i>USB Discovery Phase</i>	770
<i>DHCP Discovery Phase</i>	770
<i>Script Execution Phase</i>	772
<i>Post-Installation Reload Phase</i>	772
Configuring a Switch Using POAP	772
Cisco DCNM	772
Feature Details and Benefits	774
Cisco DCNM Web User Interface	779

Cisco UCS Director	782
Automation and Orchestration with Cisco UCS Director	783
Features and Benefits	784
Cisco UCS Director System Setup	785
PowerShell	787
Installing the Cisco UCS Director PowerShell Agent	787
Executing PowerShell Agent Commands	788
Exam Preparation Tasks	789
Review All Key Topics	789
Memory Tables	790
Define Key Terms	790
References	791

Part V Security

Chapter 17 Network Security 792

“Do I Know This Already?” Quiz	792
Foundation Topics	794
Authentication, Authorization, and Accounting	794
AAA Service Configuration Options	796
Authentication and Authorization User Login Process	797
AAA NX-OS Configurations	798
Role-Based Access Control	801
NX-OS User Roles and Rules	803
NX-OS RBAC Configurations	805
Nexus First-Hop Security	809
Nexus Dynamic ARP Inspection	810
NX-OS DAI Configurations	813
NX-OS DHCP Snooping	821
<i>DHCP Snooping Trusted and Untrusted Sources</i>	821
<i>DHCP Snooping Packet Validation</i>	822
<i>DHCP Snooping Option 82 Data Insertion</i>	823
<i>NX-OS DHCP Snooping Configuration</i>	823
Port Security	826
Nexus Port Secure MAC Address Maximum and Dynamic Address Aging	827
Port Security Violations and Actions	828
Nexus Port Types and Port Security	829
NX-OS Port Security Configuration	829

Nexus Control Plane Policing	831
Control Plane Packet	833
Classification for CoPP	834
<i>Rate-Controlling Mechanisms</i>	834
<i>Modular QoS Command-Line Interface</i>	836
NX-OS CoPP Configuration	838
Cisco ACI Contracts	845
Cisco ACI Contract Configuration Parameters	847
Create, Modify, or Remove Regular Contracts	848
Apply or Remove VRF Contracts	850
Inter-Tenant Contracts	851
Inter-Private Network Contracts Communication	852
Single Contract Bidirectional Reverse Filter	853
Single Contract Unidirectional with Multiple Filters	853
Multiple Contracts Unidirectional Single Filter	854
ACI Microsegmentation	854
Example: ACI Microsegmentation with VMs from a Single Application EPG	856
Example: ACI Microsegmentation with VMs in Different Application EPGs	857
ACI Microsegmentation Configurations	858
Exam Preparation Tasks	862
Review All Key Topics	862
Define Key Terms	862
References	863
Chapter 18 Compute Security	864
“Do I Know This Already?” Quiz	864
Foundation Topics	865
Securing UCS Management Using Authentication, Authorization, and Accounting	865
User RADIUS and TACACS+ Attributes	866
Two-Factor Authentication	869
UCS Web Session Refresh and Session Timeout Period	869
UCS LDAP Providers and Groups	869
<i>LDAP Group Mapping</i>	875
RADIUS and TACACS+ Authentication Configurations	878
UCS Remote Users Role Policy	882

Multiple Authentication Services Configuration	884
Keychains Authentication	884
NX-OS Keychain Configurations	885
Key Selection	888
Exam Preparation Tasks	889
Review All Key Topics	889
Define Key Terms	890
References	890

Chapter 19 Storage Security 892

“Do I Know This Already?” Quiz	892
Foundation Topics	894
Authentication, Authorization, and Accounting	894
Authentication	895
Authorization	895
Accounting	896
Server Groups	896
AAA Service Configuration Options	896
AAA Server Monitoring	896
Remote AAA Services	897
RADIUS	898
TACACS+	900
LDAP	903
Local AAA Services	907
AAA Authentication and Authorization Process	908
AAA Server Distribution	909
Merging RADIUS and TACACS+ Configurations	910
User Accounts and RBAC	910
User Roles	911
Rules	911
User Role Policies	913
RBAC Sample Configuration	914
Port Security	915
Port Security Configuration	917
<i>Method 1: Manual Database Configuration</i>	917
<i>Method 2: Auto-Learning Without CFS Distribution</i>	918
<i>Method 3: Auto-Learning with CFS Distribution</i>	919
Verification of Port Security	920

Fabric Binding	922
Fabric Binding Configuration	922
Port Security Versus Fabric Binding	924
Exam Preparation Tasks	925
Review All Key Topics	925
Memory Tables and Lists	926
Define Key Terms	926
References	926

Chapter 20 Final Preparation 928

Getting Ready	928
Tools for Final Preparation	929
Pearson Test Prep Practice Test Software and Questions on the Website	929
<i>Accessing the Pearson Test Prep Software Online</i>	929
<i>Accessing the Pearson Test Prep Software Offline</i>	929
Customizing Your Exams	930
Updating Your Exams	931
<i>Premium Edition</i>	931
Chapter-Ending Review Tools	932
Learn the Question Types Using the Cisco Certification Exam Tutorial	932
Suggested Plan for Final Review/Study	936
Summary	936

Appendix A Answers to the “Do I Know This Already?” Quizzes 938

Appendix B CCNP and CCIE Data Center Core DCCOR 350-601 Official Cert Guide Exam Updates 952

Glossary 954

Index 978

Online Elements

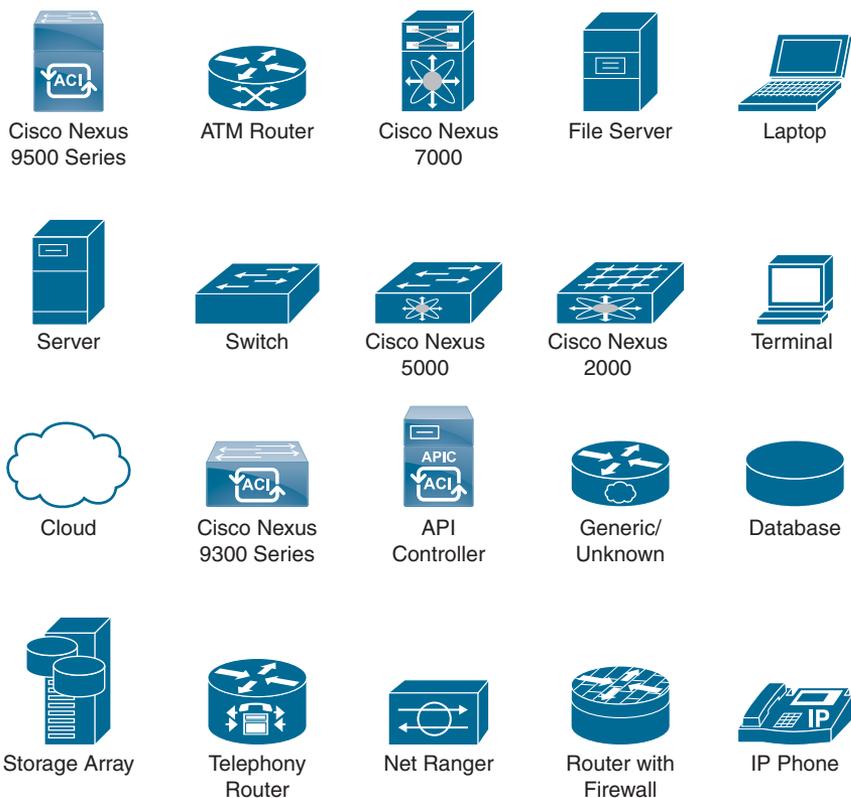
Glossary

Appendix C Memory Tables

Appendix D Memory Tables Answer Key

Appendix E Study Planner

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.

Introduction

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

Goals and Methods

The most important and somewhat obvious goal of this book is to help you pass the 350-601 CCNP Data Center Core Exam. In fact, if the primary objective of this book were different, the book's title would be misleading; however, the methods used in this book to help you pass the 350-601 CCNP Data Center Core Exam are designed to also make you much more knowledgeable about how to do your job. Although this book and the companion website together have more than enough questions to help you prepare for the actual exam, the method in which they are used is not simply to make you memorize as many questions and answers as you possibly can.

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

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

Who Should Read This Book?

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

So why should you want to pass the CCNP Data Center Core Exam? Because it's one of the milestones toward getting the CCNP and CCIE certification—no small feat in itself. What would getting the CCNP or CCIE mean to you? A raise, a promotion, recognition?

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

Strategies for Exam Preparation

The strategy you use for the CCNP Data Center Core Exam might be slightly different from strategies used by other readers, mainly based on the skills, knowledge, and experience you already have obtained. For instance, if you have attended the DCICN and DCICT course, you might take a different approach than someone who learned data center technologies via on-the-job training.

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

The Companion Website for Online Content Review

All the electronic review elements, as well as other electronic components of the book, exist on this book's companion website.

How to Access the Companion Website

To access the companion website, which gives you access to the electronic content with this book, start by establishing a login at www.ciscopress.com and register your book. To do so, simply go to www.ciscopress.com/register and enter the ISBN of the print book: 9780136449621. After you have registered your book, go to your account page and click the **Registered Products** tab. From there, click the **Access Bonus Content** link to get access to the book's companion website.

Note that if you buy the Premium Edition eBook and Practice Test version of this book from Cisco Press, your book will automatically be registered on your account page.

Simply go to your account page, click the Registered Products tab, and select Access Bonus Content to access the book's companion website.

How to Access the Pearson Test Prep Practice Test Software

You have two options for installing and using the Pearson Test Prep practice test software: a web app and a desktop app. To use the Pearson Test Prep practice test software, start by finding the registration code that comes with the book. You can find the code in these ways:

- **Print book:** Look in the cardboard sleeve in the back of the book for a piece of paper with your book's unique PTP code.

- **Premium Edition:** If you purchase the Premium Edition eBook and Practice Test directly from the Cisco Press website, the code will be populated on your account page after purchase. Just log in at www.ciscopress.com, click **Account** to see details of your account, and click the **Digital Purchases** tab.
- **Amazon Kindle:** For those who purchase a Kindle edition from Amazon, the access code will be supplied directly from Amazon.
- **Other Bookseller e-books:** Note that if you purchase an e-book version from any other source, the practice test is not included because other vendors to date have not chosen to vend the required unique access code.

NOTE Do not lose the activation code because it is the only means with which you can access the QA content with the book.

When you have the access code, to find instructions about both the PTP web app and the desktop app, follow these steps:

- Step 1.** Open this book's companion website, as shown earlier in this Introduction under the heading "How to Access the Companion Website."
- Step 2.** Click the Practice Exams button.
- Step 3.** Follow the instructions listed there both for installing the desktop app and for using the web app.

Note that if you want to use the web app only at this point, just navigate to www.pearsontestprep.com, establish a free login if you do not already have one, and register this book's practice tests using the registration code you just found. The process should take only a couple of minutes.

NOTE Amazon e-book (Kindle) customers: It is easy to miss Amazon's email that lists your PTP access code. Soon after you purchase the Kindle eBook, Amazon should send an email. However, the email uses very generic text and makes no specific mention of PTP or practice exams. To find your code, read every email from Amazon after you purchase the book. Also do the usual checks for ensuring your email arrives, like checking your spam folder.

NOTE Other e-book customers: As of the time of publication, only the publisher and Amazon supply PTP access codes when you purchase their e-book editions of this book.

How This Book Is Organized

Although this book could be read cover-to-cover, it is designed to be flexible and allow you to easily move between chapters and sections of chapters to cover just the material that you need more work with.

The core chapters, Chapters 1 through 20, cover the following topics:

- **Chapter 1, “Implementing Routing in the Data Center”:** This chapter discusses data center Layer 3 routing protocols, focusing on OSPF and BGP routing protocols. It also discusses multicast and First Hop Redundancy Protocols such as HSRP and VRRP.
- **Chapter 2, “Implementing Data Center Switching Protocols”:** This chapter discusses data center Layer 2 switching protocols, focusing on spanning tree and multiport aggregation. It also discusses virtual port channels (multichassis port channels).
- **Chapter 3, “Implementing Data Center Overlay Protocols”:** This chapter discusses various data center Overlay protocols, including Overlay Transport Virtualization (OTV) and Virtual Extensible LAN (VXLAN).
- **Chapter 4, “Describe Cisco Application Centric Infrastructure”:** This chapter discusses various aspects of Cisco ACI, including but not limited to fabric discovery, fabric access policies, fabric packet flow, tenants, and VMM domains.
- **Chapter 5, “Cisco Cloud Services and Deployment Models”:** This chapter discusses an overview of what cloud computing is along with cloud service models per the NIST 800-145 definition, such as Infrastructure as a Service (IaaS), Software as a Service (SaaS), and Platform as a Service (PaaS). It also discusses various cloud deployment models per the NIST 800-145 definition, such as public, private, community, and hybrid cloud.
- **Chapter 6, “Data Center Network Management and Monitoring”:** This chapter discusses data center network disruptive/nondisruptive upgrade procedures, network configurations, and infrastructure monitoring aspects in detail. It also discusses data center network assurance and data telemetry.
- **Chapter 7, “Implement Fibre Channel”:** This chapter discusses the Fibre Channel protocol in detail. It discusses Fibre Channel topologies, port types, switched fabric initialization, CFS distribution, VSAN, zoning, device alias, FLOGI, and FCNS databases. It also discusses NPV and NPIV features in detail.
- **Chapter 8, “Implement FCoE Unified Fabric”:** This chapter discusses the FCoE Unified Fabric Protocol in detail. It discusses various Ethernet enhancements that enable FCoE support on Ethernet interfaces. It also discusses FCoE topology options and various FCoE implementations—for example, FCoE over FEX and FCoE NPV.
- **Chapter 9, “Describe NFS and NAS Concepts”:** This chapter discusses NFS basics along with various NFS versions. It also discusses NAS basics with an overview of the Cisco NSS 3000 Series NAS product.
- **Chapter 10, “Describe Software Management and Infrastructure Monitoring”:** This chapter discusses how the Cisco MDS NX-OS Setup Utility helps to build an initial configuration file using the System Configuration dialog. It also discusses Cisco MDS NX-OS software upgrade and downgrade procedures, along with infrastructure monitoring features such as SPAN, RSPAN, RMON, and Call Home.

- **Chapter 11, “Cisco Unified Computing Systems Overview”:** This chapter discusses the Cisco Unified Computing System (UCS) architecture. It also discusses in detail UCS initial setup, along with network management aspects of Cisco UCS such as identity pools, policies, QoS, and templates.
- **Chapter 12, “Cisco Unified Computing Infrastructure Monitoring”:** This chapter discusses Cisco Unified Compute traffic monitoring and Intersight cloud management.
- **Chapter 13, “Cisco Unified Compute Software and Configuration Management”:** This chapter discusses Cisco UCS configuration management such as backup and restore. It also discusses aspects of firmware and software updates on Cisco UCS.
- **Chapter 14, “Cisco HyperFlex Overview”:** This chapter discusses the Cisco Hyperflex solution and benefits. It also discusses edge solutions that enable any application to be deployed, monitored, and managed anywhere.
- **Chapter 15, “Automation and Scripting Tools”:** This chapter discusses various automation and scripting tools. It discusses the Embedded Event Manager (EEM), Scheduler, Bash Shell, and Guest Shell for Cisco NX-OS software, and various data formats such as XML and JSON. It also discusses how the REST API can be used to configure Cisco NX-OS devices.
- **Chapter 16, “Evaluate Automation and Orchestration Technologies”:** This chapter discusses various automation and orchestration technologies. It discusses how Ansible, Puppet, and Python can be used to automate Cisco Data Center products. It also discusses the PowerOn Auto Provisioning (POAP) process, Cisco Data Center Network Manager (DCNM) tool, Cisco UCS Director (UCSD) tool, along with how the PowerShell Agent executes various tasks on UCSD.
- **Chapter 17, “Network Security”:** This chapter discusses network authentication, authorization, and accounting (AAA) and user role-based access control (RBAC). It also discusses various network security protocols in detail, including control plan policing, dynamic ARP inspection, DHCP snooping, and port security.
- **Chapter 18, “Compute Security”:** This chapter discusses Cisco UCS authentication and user role-based access control. It also discusses the keychain authentication method.
- **Chapter 19, “Storage Security”:** This chapter discusses various storage security features in detail. It discusses authentication, authorization, and accounting (AAA), user accounts, and RBAC. It also discusses configuration and verification of port security and fabric binding features on the Cisco MDS 9000 Series switches.
- **Chapter 20, “Final Preparation”:** This chapter suggests a plan for final preparation after you have finished the core parts of the book, in particular explaining the many study options available in the book.

Certification Exam Topics and This Book

The questions for each certification exam are a closely guarded secret. However, we do know which topics you must know to *successfully* complete this exam. Cisco publishes them as an exam blueprint for the Implementing Cisco Data Center Core Technologies (DCCOR 350-601) Exam. Table I-1 lists each exam topic listed in the blueprint along with a reference to the book chapter that covers the topic. These are the same topics you should be proficient in when working with Cisco data center technologies in the real world.

Table I-1 DCCOR Exam 350-601 Topics and Chapter References

DCCOR 350-601 Exam Topic	Chapter(s) in Which Topic Is Covered
1.0 Network	
1.1 Apply routing protocols	1
1.1.a OSPFv2, OSPFv3	1
1.1.b MP-BGP	1
1.1.c PIM	1
1.1.d FHRP	1
1.2 Apply switching protocols such as RSTP+, LACP and vPC	2
1.3 Apply overlay protocols such as VXLAN EVPN and OTV	3
1.4 Apply ACI concepts	4
1.4.a Fabric setup	4
1.4.b Access policies	4
1.4.c VMM	4
1.4.d Tenant policies	4
1.5 Analyze packet flow (unicast, multicast, and broadcast)	4
1.6 Analyze Cloud service and deployment models (NIST 800-145)	5
1.7 Describe software updates and their impacts	6
1.7.a Disruptive/nondisruptive	6
1.7.b EPLD	6
1.7.c Patches	6
1.8 Implement network configuration management	6
1.9 Implement infrastructure monitoring such as NetFlow and SPAN	6
1.10 Explain network assurance concepts such as streaming telemetry	6
2.0 Compute	
2.1 Implement Cisco Unified Compute System Rack Servers	11
2.2 Implement Cisco Unified Compute System Blade Chassis	11
2.2.a Initial setup	11
2.2.b Infrastructure management	11
2.2.c Network management (VLANs, pools and policies, templates, QoS)	11

DCCOR 350-601 Exam Topic	Chapter(s) in Which Topic Is Covered
2.2.d Storage management (SAN connectivity, Fibre Channel zoning, VSANs, WWN pools, SAN policies, templates)	11
2.2.e Server management (Server pools and boot policies)	11
2.3 Explain HyperFlex Infrastructure Concepts and benefits (Edge and Hybrid Architecture versus all-flash)	14
2.4 Describe firmware and software updates and their impacts on B-Series and C-Series servers	13
2.5 Implement compute configuration management (Backup and restore)	13
2.6 Implement infrastructure monitoring such as SPAN and Intersight	12
3.0 Storage Network	
3.1 Implement Fibre Channel	7
3.1.a Switch fabric initialization	7
3.1.b Port channels	7
3.1.c FCID	7
3.1.d CFS	7
3.1.e Zoning	7
3.1.f FCNS	7
3.1.g Device alias	7
3.1.h NPV and NPIV	7
3.1.i VSAN	7
3.2 Implement FCoE Unified Fabric (FIP and DCB)	8
3.3 Describe NFS and NAS concepts	9
3.4 Describe software updates and their impacts (Disruptive/ nondisruptive and EPLD)	10
3.5 Implement infrastructure monitoring	10
4.0 Automation	
4.1 Implement automation and scripting tools	15
4.1.a EEM	15
4.1.b Scheduler	15
4.1.c Bash Shell and Guest Shell for NX-OS	15
4.1.d REST API	15
4.1.e JSON and XML encodings	15
4.2 Evaluate automation and orchestration technologies	16
4.2.a Ansible	16
4.2.b Puppet	16
4.2.c Python	16
4.2.d POAP	16

DCCOR 350-601 Exam Topic	Chapter(s) in Which Topic Is Covered
4.2.e DCNM	16
4.2.f UCSD	16
4.2.g PowerShell	16
5.0 Security	
5.1 Apply network security	17
5.1.a AAA and RBAC	17
5.1.b ACI contracts and microsegmentation	17
5.1.c First-hop security features such as dynamic ARP inspection (DAI), DHCP snooping, and port security	17
5.1.d CoPP	17
5.2 Apply compute security	18
5.2.a AAA and RBAC	18
5.2.b Keychain authentication	18
5.3 Apply storage security	19
5.3.a AAA and RBAC	19
5.3.b Port security	19
5.3.c Fabric binding	19

Each version of the exam can have topics that emphasize different functions or features, and some topics can be rather broad and generalized. The goal of this book is to provide the most comprehensive coverage to ensure that you are well prepared for the exam. Although some chapters might not address specific exam topics, they provide a foundation that is necessary for a clear understanding of important topics. Your short-term goal might be to pass this exam, but your long-term goal should be to become a qualified data center professional.

It is also important to understand that this book is a “static” reference, whereas the exam topics are dynamic. Cisco can and does change the topics covered on certification exams often.

This exam guide should not be your only reference when preparing for the certification exam. You can find a wealth of information available at Cisco.com that covers each topic in great detail. If you think that you need more detailed information on a specific topic, read the Cisco documentation that focuses on that topic.

Note that as data center technologies continue to develop, Cisco reserves the right to change the exam topics without notice. Although you can refer to the list of exam topics in Table I-1, always check Cisco.com to verify the actual list of topics to ensure that you are prepared before taking the exam. You can view the current exam topics on any current Cisco certification exam by visiting the Cisco.com website, choosing **Menu**, and **Training & Events**, then selecting from the Certifications list. Note also that, if needed, Cisco Press might post additional preparatory content on the web page associated

with this book at <http://www.ciscopress.com/title/9780136449621>. It's a good idea to check the website a couple of weeks before taking your exam to be sure that you have up-to-date content.

Taking the CCNP Data Center Core Exam

As with any Cisco certification exam, you should strive to be thoroughly prepared before taking the exam. There is no way to determine exactly what questions are on the exam, so the best way to prepare is to have a good working knowledge of all subjects covered on the exam. Schedule yourself for the exam and be sure to be rested and ready to focus when taking the exam.

The best place to find out the latest available Cisco training and certifications is under the Training & Events section at Cisco.com.

Tracking Your Status

You can track your certification progress by checking <http://www.cisco.com/go/certifications/login>. You must create an account the first time you log in to the site.

How to Prepare for an Exam

The best way to prepare for any certification exam is to use a combination of the preparation resources, labs, and practice tests. This guide has integrated some practice questions and sample scenarios to help you better prepare. If possible, get some hands-on experience with ACI, Nexus, and UCS equipment. There is no substitute for real-world experience; it is much easier to understand the designs, configurations, and concepts when you can actually work with a live data center network.

Cisco.com provides a wealth of information about Application Centric Infrastructure (ACI), Nexus switches, and Unified Computing System—Blade and Rack servers, and data center LAN technologies and features.

Assessing Exam Readiness

Exam candidates never really know whether they are adequately prepared for the exam until they have completed about 30 percent of the questions. At that point, if you are not prepared, it is too late. The best way to determine your readiness is to work through the “Do I Know This Already?” quizzes at the beginning of each chapter and review the foundation and key topics presented in each chapter. It is best to work your way through the entire book unless you can complete each subject without having to do any research or look up any answers.

Cisco Data Center Certifications in the Real World

Cisco is one of the most recognized names on the Internet. Cisco Certified data center specialists can bring quite a bit of knowledge to the table because of their deep understanding of data center technologies, standards, and networking devices. This is why the Cisco certification carries such high respect in the marketplace. Cisco certifications demonstrate to potential employers and contract holders a certain professionalism, expertise, and dedication required to complete a difficult goal. If Cisco certifications were easy to obtain, everyone would have them.

Exam Registration

The 350-601 CCNP Data Center Core Exam is a computer-based exam, with around 100 to 110 multiple-choice, fill-in-the-blank, list-in-order, and simulation-based questions. You can take the exam at any Pearson VUE (<http://www.pearsonvue.com>) testing center. According to Cisco, the exam should last about 120 minutes. Be aware that when you register for the exam, you might be told to allow a certain amount of time to take the exam that is longer than the testing time indicated by the testing software when you begin. This discrepancy is because the testing center will want you to allow for some time to get settled and take the tutorial about the test engine.

Book Content Updates

Because Cisco occasionally updates exam topics without notice, Cisco Press might post additional preparatory content on the web page associated with this book at <http://www.ciscopress.com/title/9780136449621>. It is a good idea to check the website a couple of weeks before taking your exam to review any updated content that might be posted online. We also recommend that you periodically check back to this page on the Cisco Press website to view any errata or supporting book files that may be available.

Implementing Data Center Overlay Protocols

The adoption of server virtualization has been increasing rapidly. Server virtualization provides flexibility and agility in provisioning and placement of computing workloads. However, network connectivity has not kept pace with such innovations in the computing environment, although it still offers a rigid approach to provisioning transport services.

As a solution, network overlays abstract the details of the physical network, making it much faster to connect virtual machines (VMs) and other devices. Rather than provision paths on physical devices, overlays encapsulate traffic using protocols such as Overlay Transport Virtualization (OTV) or Virtual Extensible LAN (VXLAN) across the physical network. These newer protocols allow operators to move beyond the limitations of VLANs, which support only 4096 virtual networks, so that they can better support multitenant services.

This chapter covers the following key topics:

Overlay Transport Virtualization (OTV): This section provides an overview of overlay transportation, including Layer 2 MAC address routing along with a configuration example.

Virtual Extensible LAN (VXLAN) Overview: This section discusses the Layer 2 VLAN extension to provide multitenant flexibility, high segment scalability, and Layer 2 spanning tree improvement, along with a configuration example.

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 3-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes.”

Table 3-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
Overlay Transport Virtualization (OTV)	1–3
Virtual Extensible LAN (VXLAN) Overview	4–6

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark that question as wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. Which statement about Cisco Overlay Transport Virtualization is true?
 - a. OTV is a “MAC in IP” technique for supporting Layer 2 VLANs over any transport.
 - b. OTV is core and site transparent. Changes to the Layer 2 design of the sites are needed.
 - c. OTV cannot support multiple concurrent overlays.
 - d. OTV transport must support multicast.
2. What is an OTV joined interface?
 - a. A Layer 2 access or trunk interface
 - b. A logical multiaccess and multicast-capable interface
 - c. An interface to form OTV adjacencies with the other OTV edge devices belonging to the same VPN
 - d. A Layer 3 interface that uses multicast only to forward OTV traffic
3. How does OTV exchange MAC reachability information?
 - a. Uses OSPF as an internal control plane
 - b. Uses BGP as an internal control plane
 - c. Uses EIGRP as an internal control plane
 - d. Use IS-IS as an internal control plane
4. In current data center networking architecture, which network layer is used to transmit VXLAN packets or other overlay packets?
 - a. Overlay network
 - b. SD-WAN
 - c. Underlay network
 - d. MPLS
5. How many available IDs can be assigned to a VXLAN at any given time?
 - a. 4096
 - b. 160,000
 - c. 1 million
 - d. 16 million
6. Which statement about VXLAN high availability is correct?
 - a. For an anycast IP address, vPC VTEP switches can use the same VTEP IP address.
 - b. For an anycast IP address, vPC VTEP switches must use a secondary IP address on the loopback interface.
 - c. Distributed anycast gateways must be connected with vPC.
 - d. VTEP high availability will use unicast instead of multicast communications.

Foundation Topics

Overlay Transport Virtualization (OTV)

Overlay transportation introduces the concept of “MAC routing,” which means a control plane protocol is used to exchange MAC reachability information between network devices providing LAN extension functionality. This is a significant shift from Layer 2 switching that traditionally leverages data plane learning, and it is justified by the need to limit flooding of Layer 2 traffic across the transport infrastructure. As outlined in this chapter, Layer 2 communication between sites resembles routing more than switching. If the destination MAC address information is unknown, traffic is dropped (not flooded), preventing the waste of precious bandwidth across the WAN.

OTV also introduces the concept of the dynamic encapsulation for Layer 2 flows that need to be sent to remote locations. Each Ethernet frame is individually encapsulated into an IP packet and delivered across the transport network. This eliminates the need to establish virtual circuits, called *pseudowires*, between the data center locations. Immediate advantages include improved flexibility when adding or removing sites to the overlay, more optimal bandwidth utilization across the WAN (specifically when the transport infrastructure is multicast enabled), and independence from the transport characteristics (Layer 1, Layer 2, or Layer 3).

OTV provides a native built-in multihoming capability with automatic detection. Two or more devices can be leveraged in each data center to provide LAN extension functionality without running the risk of creating an end-to-end loop that would jeopardize the overall stability of the design. This is achieved by leveraging the same control plane protocol used for the exchange of MAC address information, without the need of extending the Spanning Tree Protocol (STP) across the overlay.

OTV Terminology

To understand how OTV works in an existing IP transport environment, let’s discuss the OTV interfaces and terms shown in Figure 3-1.

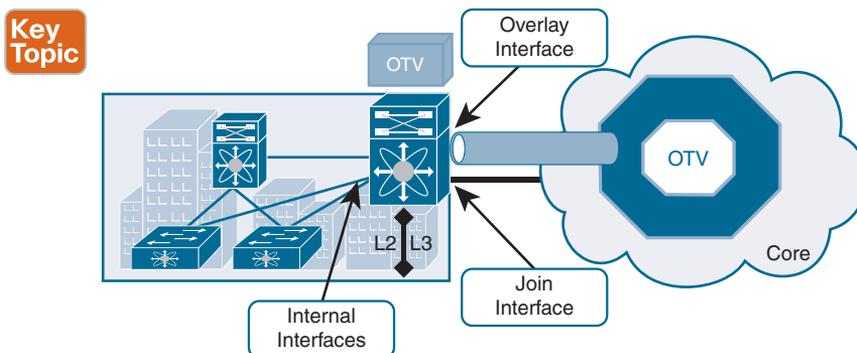


Figure 3-1 OTV Interfaces and Terms

- **Edge device (ED):** This device connects the site to the (WAN/MAN) core and is responsible for performing all the OTV functions.
 - An edge device receives Layer 2 traffic for all VLANs that need to be extended to remote locations and dynamically encapsulates the Ethernet frames into IP packets that are then sent across the OTV transport infrastructure.

- For resiliency, two OTV edge devices can be deployed on each site to provide redundancy.
- **Internal interfaces:** These are the L2 interfaces (usually 802.1q trunks) of the ED that face the site.
 - Internal interfaces are regular access or trunk ports.
 - Trunk configuration will extend more than one VLAN across the overlay. There is no need to apply OTV-specific configuration to these interfaces.
 - Typical Layer 2 functions (like local switching, spanning tree operation, data plane learning, and flooding) are performed on the internal interfaces.
- **Join interface:** This is the L3 interface of the ED that faces the core. The join interface is used by the edge device for different purposes:
 - “Join” the overlay network and discover the other remote OTV edge devices.
 - Form OTV adjacencies with the other OTV edge devices belonging to the same VPN.
 - Send/receive MAC reachability information.
 - Send/receive unicast and multicast traffic.
- **Overlay interface:** This is a logical multiaccess multicast-capable interface. It encapsulates Layer 2 frames in IP unicast or multicast headers.

Every time the OTV edge device receives a Layer 2 frame destined for a remote data center site, the frame is logically forwarded to the overlay interface. This instructs the edge device to perform the dynamic OTV encapsulation on the Layer 2 packet and send it to the join interface toward the routed domain.



OTV Control Plane Function

The principle of OTV is to build a control plane between the OTV edge devices to advertise MAC address reachability information instead of using data plane learning. However, before MAC reachability information can be exchanged, all OTV edge devices must become “adjacent” to each other from an OTV perspective.

Edge devices can be made adjacent in two ways, depending on the nature of the transport network interconnecting the various sites:

- If the transport is multicast enabled, a specific multicast group can be used to exchange the control protocol messages between the OTV edge devices.
- If the transport is not multicast enabled, an alternative deployment model is where one (or more) OTV edge device can be configured as an adjacency server to which all other edge devices register; this server communicates to them the list of devices belonging to a given overlay.

Multicast-Enabled Transport Infrastructure

If transport supports multicast, all OTV edge devices can be configured to join a specific any-source multicast (ASM) group where they simultaneously play the role of receiver and source. If the transport is owned by a service provider, the user will have to negotiate the use of this ASM group with the service provider.

Two important considerations for the OTV control plane protocol are as follows:

1. This protocol runs as an “overlay” control plane between OTV edge devices, which means there is no dependency with the routing protocol (IGP or BGP) used in the Layer 3 domain or in the transport infrastructure.
2. The OTV control plane is transparently enabled in the background after creating the OTV overlay interface and does not require explicit configuration. Tuning parameters, like timers, for the OTV protocol is allowed, but this is expected to be more of a corner case than a common requirement.

NOTE The routing protocol used to implement the OTV control plane is IS-IS. It was selected because it is a standards-based protocol, originally designed with the capability of carrying MAC address information in the TLV.

From a security perspective, it is possible to leverage the IS-IS HMAC-MD5 authentication feature to add an HMAC-MD5 digest to each OTV control protocol message. The digest allows authentication at the IS-IS protocol level, which prevents unauthorized routing messages from being injected into the network routing domain. At the same time, only authenticated devices will be allowed to successfully exchange OTV control protocol messages between them and hence to become part of the same overlay network.

The same control plane communication is also used to withdraw MAC reachability information. For example, if a specific network entity is disconnected from the network, or stops communicating, the corresponding MAC entry would eventually be removed from the CAM table of the OTV edge device. This occurs by default after 30 minutes on the OTV edge device. The removal of the MAC entry triggers an OTV protocol update so that all remote edge devices delete the same MAC entry from their respective tables.

Unicast-Only Transport Infrastructure (Adjacency-Server Mode)

OTV can be deployed with unicast-only transport. As previously described, a multicast-enabled transport infrastructure lets a single OTV update or Hello packet reach all other OTV devices by virtue of leveraging a specific multicast control group address.

The OTV control plane over a unicast-only transport works exactly the same way as OTV with multicast mode. The only difference is that each OTV device would need to create multiple copies of each control plane packet and unicast them to each remote OTV device part of the same logical overlay. Because of this head-end replication behavior, leveraging a multicast-enabled transport remains the recommended way of deploying OTV in cases where several DC sites are involved. At the same time, the operational simplification brought about by the unicast-only model (removing the need for multicast deployment) can make this deployment option very appealing in scenarios where LAN extension connectivity is required only between a few (two to three) DC sites.

To be able to communicate with all the remote OTV devices, each OTV node needs to know a list of neighbors to replicate the control packets to. Rather than statically configuring the list of all neighbors in each OTV node, a simple dynamic means is used to provide this

information. This is achieved by designating one (or more) OTV edge device to perform a specific role, named the *adjacency server*. Every OTV device wishing to join a specific OTV logical overlay needs to first “register” with the adjacency server (by sending OTV Hello messages to it). All other OTV neighbor addresses are discovered dynamically through the adjacency server. Consequently, when the OTV service needs to be extended to a new DC site, only the OTV edge devices for the new site need to be configured with the adjacency server addresses. No other sites need additional configuration. Figure 3-2 shows the differences between multicast-enabled transport and unicast-only transport.

Key Topic

3

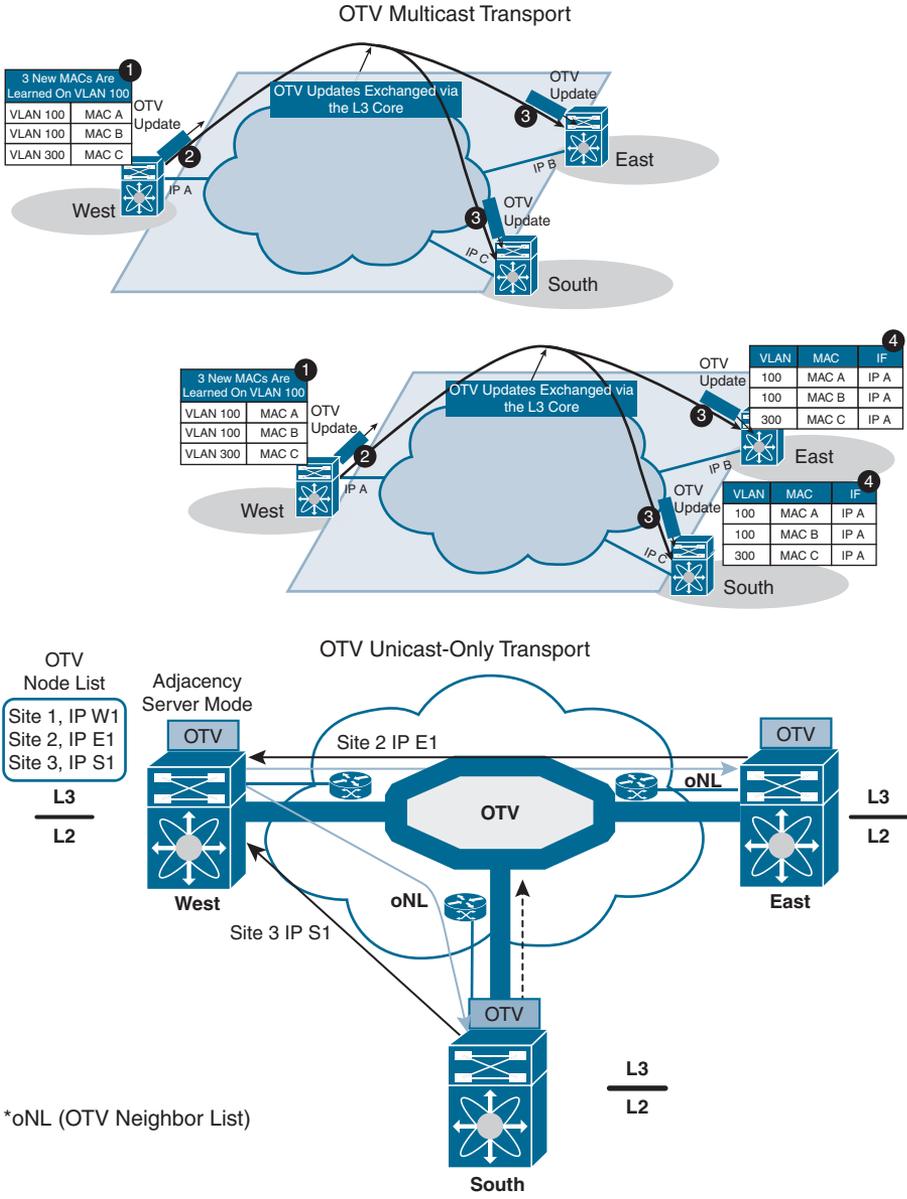


Figure 3-2 OTV Control Plane



OTV Data Plane Function

After the control plane adjacencies established between the OTV edge devices and MAC address reachability information are exchanged, traffic can start flowing across the overlay. Similar to any L2 Switch, data plane traffic can be

- Unicast traffic
- Multicast traffic
- Broadcast traffic

Unicast Traffic over OTV

In a Layer 2 switch, if PC1 is trying to communicate with PC2 and both belong to the same VLAN, the switch will perform Layer 2 MAC address lookup and forward the traffic to the destination local port. In OTV, an edge device will perform Layer 2 lookup, and the destination will be the remote edge IP address (as shown in Figure 3-3). The OTV edge will encapsulate the Layer 2 frames over Layer 3 and transport it to the OTV remote edge, as shown using the following steps:

- Step 1.** PC1 starts sending traffic to PC2.
- Step 2.** When traffic reaches the aggregation layer device (an OTV edge device), a usual Layer 2 lookup is performed to determine how to reach PC2.
- Step 3.** The MAC table points to a remote OTV edge IP address (in a local switch, MAC points to a local interface).
- Step 4.** The OTV edge device *encapsulates* the original Layer 2 frame. The source IP of the outer header is the IP address of its join interface, whereas the destination IP is the IP address of the join interface of the remote edge device. The OTV-encapsulated frame (a regular unicast IP packet) is carried across the transport infrastructure and delivered to the remote OTV edge device.
- Step 5.** The remote OTV edge device *decapsulates* the packet, exposing the original Layer 2 frame.
- Step 6.** The edge device performs another Layer 2 lookup on the original Ethernet frame and discovers that it is reachable through a physical interface, which means it is a MAC address local to the site.
- Step 7.** The frame is delivered to the MAC destination.

Given that Ethernet frames are carried across the transport infrastructure after being OTV encapsulated, some considerations around MTU are necessary. In the first implementation, the OTV encapsulation increases the overall MTU size of 42 bytes. This is a result of the operation of the edge device that removes the CRC and the 802.1Q fields from the original Layer 2 frame and adds an OTV Shim (containing the VLAN and Overlay ID information also) and an external IP header.

All OTV control and data plane packets originate from an OTV edge device with the Don't Fragment (DF) bit set. In a Layer 2 domain, the assumption is that all intermediate LAN segments support at least the configured interface MTU size of the host. This means that mechanisms like Path MTU Discovery (PMTUD) are not an option in this case.

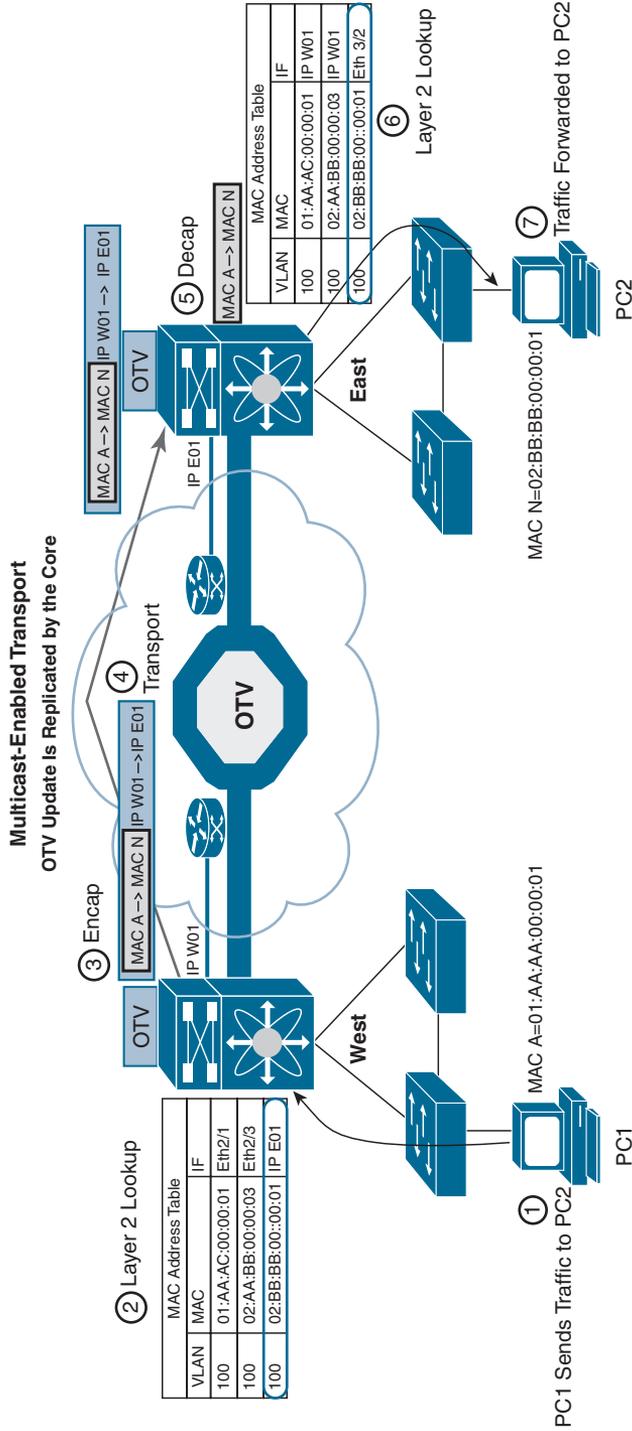


Figure 3-3 OTV Unicast Data Plane

Also, fragmentation and reassembly capabilities are not available on Nexus platforms. Consequently, Cisco recommends increasing the MTU size of all the physical interfaces along the path between the source and destination endpoints to account for introducing the extra 42 bytes by OTV.

NOTE This consideration is not OTV specific because the same challenge applies to other Layer 2 VPN technologies, such as EoMPLS or VPLS.

Multicast Traffic over OTV

In certain scenarios, there may be a requirement to establish Layer 2 multicast communication between remote sites. This is the case when a multicast source sending traffic to a specific group is deployed in a given VLAN 10 in site A, whereas multicast receivers belonging to the same VLAN 10 are placed in remote sites B to site N and need to receive traffic for that same group.

Similarly to what is done for the OTV control plane, you need to distinguish the two scenarios where the transport infrastructure is multicast enabled, or not, for the data plane.

For multicast-enabled transport, the Layer 2 multicast traffic must flow across the OTV overlay, and to avoid suboptimal head-end replication, a specific mechanism is required to ensure that multicast capabilities of the transport infrastructure can be leveraged.

The idea is to use a set of source-specific multicast (SSM) groups in the transport to carry these Layer 2 multicast streams. These groups are independent from the ASM group previously introduced to transport the OTV control protocol between sites.

For a unicast-only transporter, when multicast capabilities are not available in the transport infrastructure, Layer 2 multicast traffic can be sent across the OTV overlay by leveraging head-end replication from the OTV device deployed in the DC site where the multicast source is located. However, a specific mechanism based on IGMP snooping is still available to ensure Layer 2 multicast packets are sent only to remote DC sites where active receivers interested in that flow are connected.

Broadcast Traffic over OTV

It is important to highlight that a mechanism is required so that Layer 2 broadcast traffic can be delivered between sites across the OTV overlay. *Failure isolation* will detail how to limit the amount of broadcast traffic across the transport infrastructure, but some protocols, like Address Resolution Protocol (ARP), would always mandate the delivery of broadcast packets.

In the current OTV software release, when a multicast-enabled transport infrastructure is available, the current NX-OS software release broadcast frames are sent to all remote OTV edge devices by leveraging the same ASM multicast group in the transport already used for the OTV control protocol. Layer 2 broadcast traffic will then be handled exactly the same way as the OTV Hello messages shown in Figure 3-2.

For unicast-only transport infrastructure deployments, head-end replication performed on the OTV device in the site originating the broadcast would ensure traffic delivery to all the remote OTV edge devices that are part of the unicast-only list.

Failure Isolation

One of the main requirements of every LAN extension solution is to provide Layer 2 connectivity between remote sites without giving up the advantages of resiliency, stability, scalability, and so on, obtained by interconnecting sites through a routed transport infrastructure.

OTV achieves this goal by providing four main functions: Spanning Tree Protocol (STP) isolation, unknown unicast traffic suppression, ARP optimization, and broadcast policy control.

STP Isolation

OTV, by default, does not transmit STP bridge protocol data units (BPDUs) across the overlay, as shown in Figure 3-4. This native function does not require the use of an explicit configuration, such as BPDU filtering, and so on. Every site can then become an independent STP domain; STP root configuration, parameters, and the STP protocol flavor can be decided on a per-site basis.

STP isolation fundamentally limits the fate of sharing between data center sites: an STP problem in the control plane of a given site would not produce any effect on the remote data centers.

Limiting the extension of STP across the transport infrastructure potentially creates undetected end-to-end loops that would occur when at least two OTV edge devices are deployed in each site, inviting a common best practice to increase resiliency of the overall solution. *Multihoming* details how OTV prevents the creation of end-to-end loops without sending STP frames across the OTV overlay.

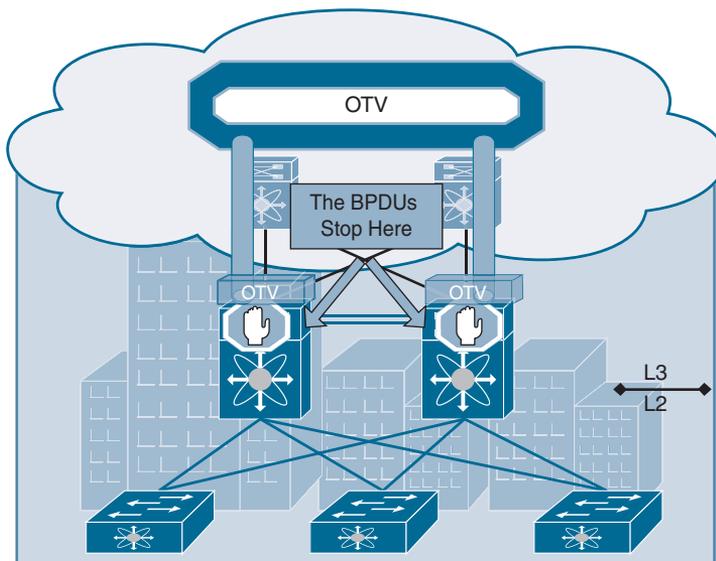


Figure 3-4 STP Isolation

Unknown Unicast Handling

The OTV control protocol will advertise MAC address reachability information between the OTV edge devices and mapping MAC address destinations to IP next hops. The consequence is that the OTV edge device starts behaving like a router instead of a Layer 2 bridge,

because it forwards Layer 2 traffic across the overlay if it has previously received information on how to reach that remote MAC destination. Figure 3-5 shows this behavior.

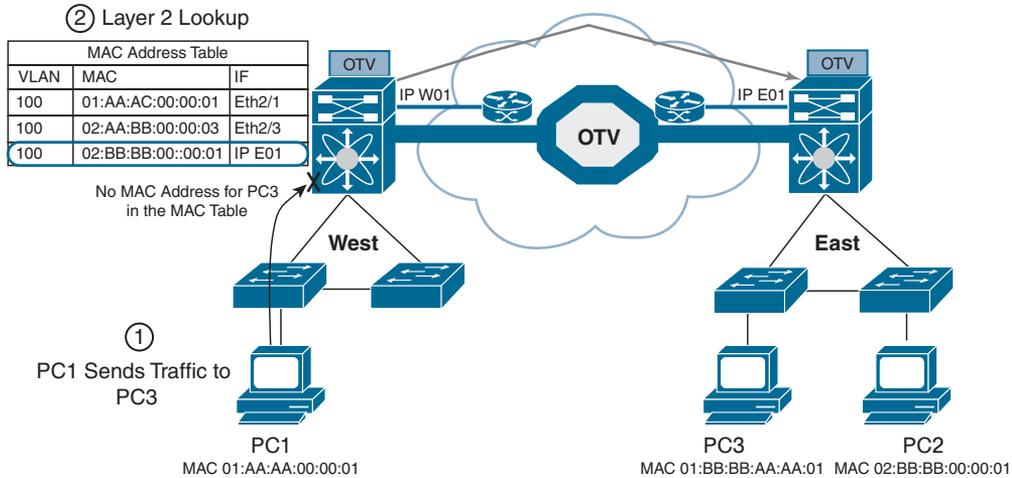


Figure 3-5 *Unknown Unicast Traffic*

When the OTV edge device receives a frame destined to PC3, it performs the usual Layer 2 lookup in the MAC table. Because it does not have information for the MAC of PC3, Layer 2 traffic is flooded out the internal interfaces, since they behave as regular Ethernet interfaces, but not via the overlay.

NOTE This OTV behavior is important to minimize the effects of a server misbehaving and generating streams directed to random MAC addresses. This could occur as a result of a DoS attack as well.

The assumption is that there are no silent or unidirectional devices in the network, so sooner or later the local OTV edge device will learn an address and communicate it to the remaining edge devices through the OTV protocol. To support specific applications, like Microsoft Network Load Balancing Services (NLBS), which require the flooding of Layer 2 traffic to functions, a configuration knob is provided to enable selective flooding. Individual MAC addresses can be statically defined so that Layer 2 traffic destined to them can be flooded across the overlay, or it can be broadcast to all remote OTV edge devices instead of being dropped. The expectation is that this configuration would be required in very specific corner cases so that the default behavior of dropping unknown unicast would be the usual operation model.

ARP Optimization

Another function that reduces the amount of traffic sent across the transport infrastructure is ARP optimization. ARP optimization will reduce the amount of broadcast traffic between sites.

IP ARP is a Layer 2 broadcast frame used to determine the MAC address of the host with a particular IP address. ARP requests are sent across the OTV overlay to all remote sites, with

the hope that they will reach the host with that particular IP. The intended host will respond to the originating host's ARP request using an ARP reply, which will pass via the original OTV edge device that forwarded the ARP request. That OTV edge device will record the ARP reply. OTV edge devices are capable of snooping ARP reply traffic and caching the contained mapping information in a local data table called ARP ND (Neighbor-Discovery). Any subsequent ARP broadcast requests that have a match in the ARP ND will be served from there and will not be sent across the overlay. Figure 3-6 shows an ARP optimization example.

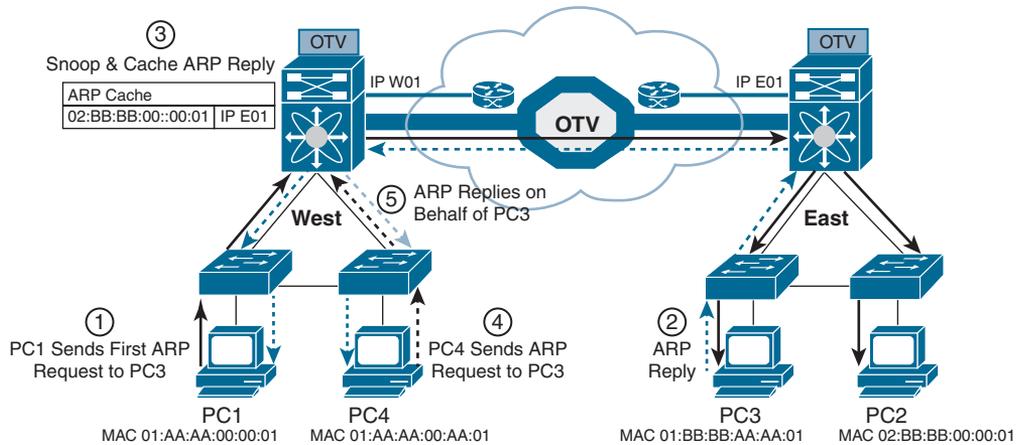
Key Topic


Figure 3-6 OTV ARP Optimization

One caveat to be aware of is the relationship between the MAC aging timer and the ARP cache timer. The ARP cache timer should always be lower than the MAC aging timer; otherwise, traffic might be *black-holed*, which means incoming or outgoing traffic is silently discarded or dropped. Using the default NX-OS values, and provided the default gateway resides on a Nexus, this should never be an issue with the default set values.

The Nexus defaults for these timers are

- OTV ARP aging timer: 480 seconds / 8 minutes
- MAC aging timer: 1800 seconds / 30 minutes

Broadcast Policy Control

In addition to the previously described ARP optimization, OTV will provide additional functionality, such as broadcast suppression, broadcast white listing, and so on, to reduce the amount of overall Layer 2 broadcast traffic sent across the overlay. Details will be provided upon future functional availability.

Key Topic

Multihoming OTV

One important function is multihoming where two (or more) OTV edge devices provide LAN extension services to a given site. As mentioned, this redundant node deployment, combined with the fact that STP BPDUs are not sent across the OTV overlay, may lead to the creation of an end-to-end loop, as shown in Figure 3-7.

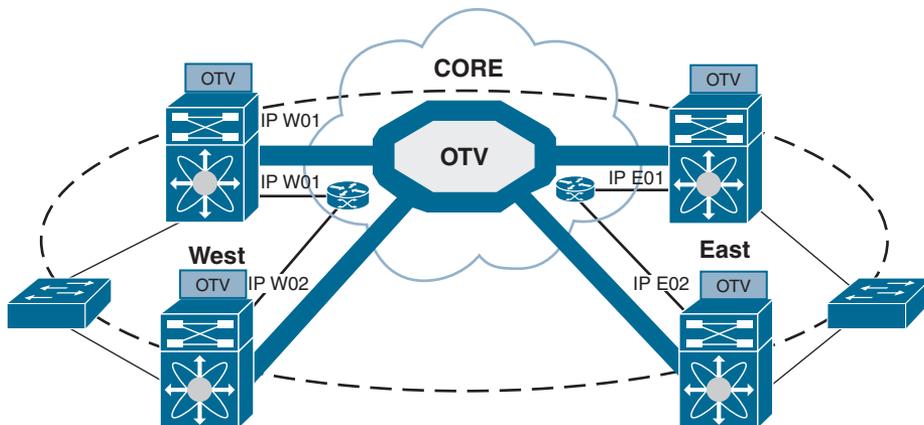


Figure 3-7 *Creation of an End-to-End STP Loop*

The concept of an Authoritative edge device (AED) is introduced to avoid the situation depicted in Figure 3-8. The AED has two main tasks:

1. Forwarding Layer 2 traffic (unicast, multicast, and broadcast) between the site and the overlay (and vice versa)
2. Advertising MAC reachability information to the remote edge devices

The AED role is negotiated, on a per-VLAN basis, between all the OTV edge devices belonging to the same site (that is, characterized by the same site identifier, or site ID). OTV uses the site adjacencies as input to determine Authoritative edge devices for the VLANs being extended from the site.

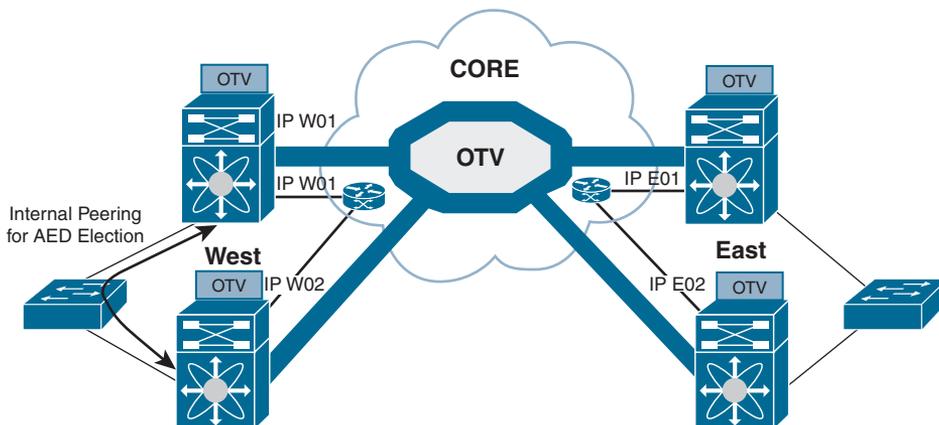


Figure 3-8 *Establishment of Internal Peering*

The site VLAN should be carried on multiple Layer 2 paths internal to a given site, to increase the resiliency of this internal adjacency (including vPC connections eventually established with other edge switches). However, the mechanism of electing an Authoritative edge device solely based on the communication established on the site VLAN may create

situations (resulting from connectivity issues or misconfiguration) where OTV edge devices belonging to the same site can fail to detect one another and thereby end up in an “active/active” mode (for the same data VLAN). This could ultimately result in the creation of a loop scenario.

To address this concern, each OTV device maintains dual adjacencies with other OTV edge devices belonging to the same DC site. OTV edge devices continue to use the site VLAN for discovering and establishing adjacency with other OTV edge devices in a site. This adjacency is called *site adjacency*.

In addition to the site adjacency, OTV devices also maintain a second adjacency, called *overlay adjacency*, established via the join interfaces across the Layer 3 network domain. To enable this new functionality, it is now mandatory to configure each OTV device with a site-identifier value. All edge devices that are in the same site must be configured with the same site identifier. This site identifier is advertised in IS-IS hello packets sent over both the overlay as well as on the site VLAN. The combination of the site identifier and the IS-IS system ID is used to identify a neighbor edge device in the same site.

The dual-site adjacency state (and not simply the site adjacency established on the site VLAN) is now used to determine the Authoritative edge device role for each extended data VLAN. All the OTV edge devices can now proactively inform their neighbors in a local site about their capability to become Authoritative edge devices and their forwarding readiness. In other words, if something happens on an OTV device that prevents it from performing its LAN extension functionalities, it can now inform its neighbor about this and let itself be excluded from the AED election process.

An explicit AED capability notification allows the neighbor edge devices to get a fast and reliable indication of failures and to determine AED status accordingly in the consequent AED election, rather than solely depending on the adjacency creation and teardown. The forwarding readiness may change due to local failures, such as the site VLAN or the extended VLANs going down or the join interface going down, or it may be intentional, such as when the edge device is starting up and/or initializing. Hence, the OTV adjacencies may be up, but the OTV device may not be ready to forward traffic. The edge device also triggers a local AED election when its forwarding readiness changes. As a result of its AED capability going down, it will no longer be AED for its VLANs.

The AED capability change received from a neighboring edge device in the same site influences the AED assignment and hence will trigger an AED election. If a neighbor indicates that it is not AED capable, it will not be considered as active in the site. An explicit AED capability down notification received over either the site or the overlay adjacency will bring the neighbor’s dual-site adjacency state down into an inactive state, and the resulting AED election will not assign any VLANs to that neighbor.

The dual-site adjacencies are used to negotiate the Authoritative edge device role. A deterministic algorithm is implemented to split the AED role for odd and even VLANs between two OTV edge devices, as shown in Figure 3-9. More specifically, the edge device identified by a lower system ID will become authoritative for all the even extended VLANs, whereas the device with a higher system ID will “own” the odd extended VLANs.

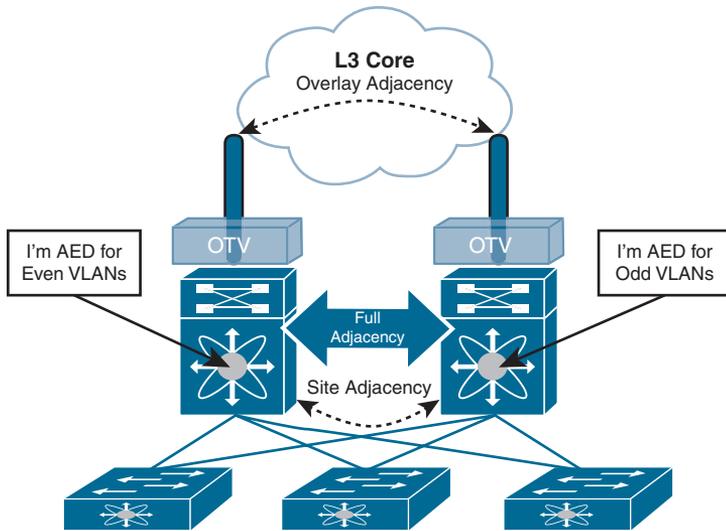


Figure 3-9 OTV AED VLAN Split

FHRP Isolation

One of the important capabilities introduced by OTV is to filter First Hop Redundancy Protocol (FHRP—HSRP, VRRP, and so on) messages *across* the logical overlay. This is required to allow for the existence of the same default gateway in different locations and optimize the outbound traffic flows (server-to-client direction). Figure 3-10 highlights the root of the problem.

Given that the *same* VLAN/IP subnet is available in different sites, the free exchange of FHRP messages across the OTV connection would lead to the election of a single default gateway. This would force traffic to follow a suboptimal path to reach the default gateway (in the site where it is deployed) each time it is required to be routed outside the subnet and the server is located in a different site.

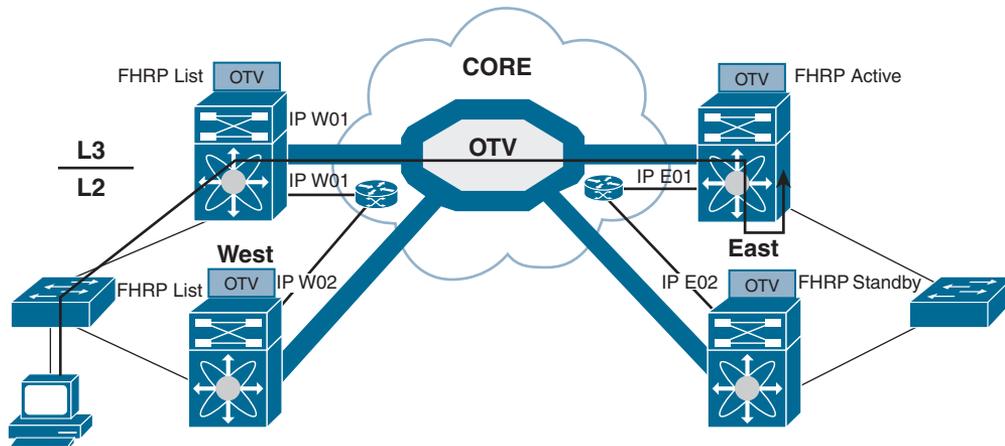


Figure 3-10 Suboptimal Outbound Routing

Figure 3-11 shows the deployment of independent default gateways in each data center site, to optimize and localize routing of outbound traffic flows.

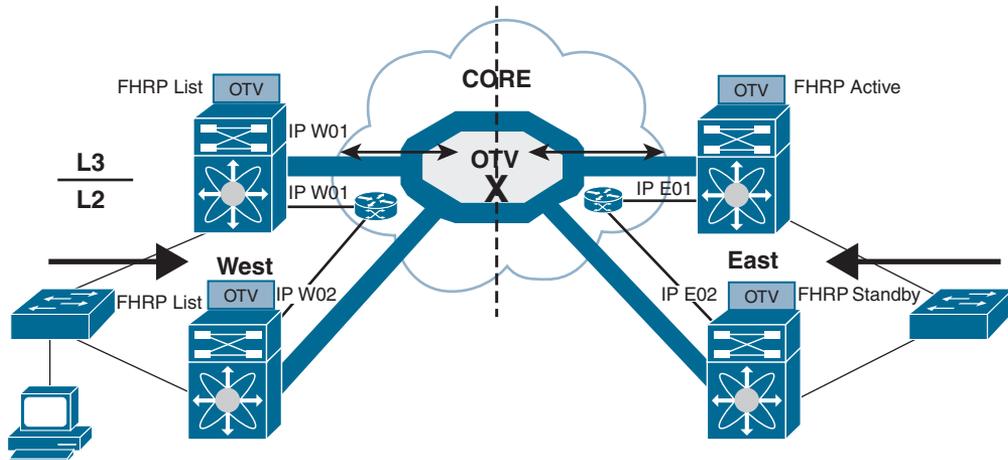


Figure 3-11 FHRP Isolation with OTV

It is critical that you enable the filtering of FHRP messages across the overlay because it allows the use of the same FHRP configuration in different sites. The end result is that the same default gateway is available and characterized by the same virtual IP and virtual MAC addresses in each data center. This means that the outbound traffic will be able to follow the optimal and shortest path, always leveraging the local default gateway.

It is important to stress how this outbound path optimization functionality should be deployed in conjunction with an equivalent one optimizing inbound traffic flows to avoid asymmetric traffic behavior (this would be highly undesirable, especially in deployments leveraging stateful services across data centers).

OTV Configurations and Verifications

Table 3-2 lists OTV default parameters. You can alter these parameters as necessary to optimize protocol functionality.

Table 3-2 OTV Default Settings

Parameters	Default
OTV feature	Disabled
Advertised VLANs	None
ARP and ND suppression	Enabled
Graceful restart	Enabled
Site VLAN	1
Site identifier	0x0
IS-IS overlay hello interval	20 seconds (Cisco NX-OS Release 6.2 or later) 4 seconds (Cisco NX-OS Release 5.2 through Cisco NX-OS Release 6.1)
IS-IS overlay hello multiplier	3

Parameters	Default
IS-IS site hello interval	3 seconds (Cisco NX-OS Release 6.2 or later) 1 second (Cisco NX-OS releases prior to 6.2)
IS-IS site hello multiplier	20 (Cisco NX-OS Release 6.2 or later) 10 (Cisco NX-OS releases prior to 6.2)
IS-IS CSNP interval	10 seconds
IS-IS LSP interval	33 milliseconds
Overlay route tracking	Disabled
Site BFD	Disabled

Table 3-3 covers the NX-OS feature license required for OTV. For more information, visit the Cisco NX-OS Licensing Guide.

Table 3-3 Feature-Based Licenses for Cisco NX-OS OTV

Platform	Feature License	Feature Name
Cisco Nexus 7000 Series	Transport Services Package LAN_TRANSPORT_ SERVICES_PKG	Overlay Transport Virtualization (OTV)

OTV has the following configuration recommendations and limitations:

- If the same device serves as the default gateway in a VLAN interface and the OTV edge device for the VLANs being extended, configure OTV on a device (Nexus 7000 VDC or switch) that is separate from the VLAN interfaces (SVIs).
- The site VLAN must not be extended into the OTV. This configuration is not supported, and this helps to avoid unexpected results.
- When possible, we recommend that you use a separate nondefault Nexus 7000 VDC for OTV to allow for better manageability and maintenance.
- An overlay interface will only be in an up state if the overlay interface configuration is complete and enabled (no shutdown). The join interface has to be in an up state.
- Configure the join interface and all Layer 3 interfaces that face the IP core between the OTV edge devices with the highest maximum transmission unit (MTU) size supported by the IP core. OTV sets the Don't Fragment (DF) bit in the IP header for all OTV control and data packets so the core cannot fragment these packets.
- Only one join interface can be specified per overlay. You can decide to use one of the following methods:
 - Configure a single join interface, which is shared across multiple overlays.
 - Configure a different join interface for each overlay, which increases the OTV reliability.

For a higher resiliency, you can use a port channel, but it is not mandatory. There are no requirements for 1-Gigabit Ethernet versus 10-Gigabit Ethernet or dedicated versus shared mode.

- The transport network must support PIM sparse mode (ASM) or PIM-Bidir multicast traffic.
- OTV is compatible with a transport network configured only for IPv4. IPv6 is not supported.
- Do not enable PIM on the join interface.
- ERSPAN ACLs are not supported for use with OTV.
- Ensure the site identifier is configured and is the same for all edge devices on a site. OTV brings down all overlays when a mismatched site identifier is detected from a neighbor edge device and generates a system message.
- You must upgrade all edge devices in the site and configure the site identifier on all edge devices in the site before traffic is restored. An edge device with an older Cisco NX-OS release in the same site can cause traffic loops. You should upgrade all edge devices in the site during the same upgrade window. You do not need to upgrade edge devices in other sites because OTV interoperates between sites with different Cisco NX-OS versions.
- For OTV fast convergence, remote unicast MAC addresses are installed in the OTV Routing Information Base (ORIB), even on non-AED VLANs.
- For OTV fast convergence, even non-AED OTV devices create a delivery source, delivery group (DS,DG) mapping for local multicast sources and send a join request to remote sources if local receivers are available. As a result, there are two remote data groups instead of one for a particular VLAN, source, group (V,S,G) entry.
- One primary IP address and no more than three secondary IP addresses are supported for OTV tunnel depolarization.

Tables 3-4 through 3-7 show some of the most-used OTV configuration commands with their purpose. For full commands, refer to the Nexus Interface Configuration Guide links provided in the “References” section at the end of the chapter.

Table 3-4 OTV Global-Level Commands

Command	Purpose
<code>feature otv</code>	Enables OTV for the device.
<code>interface overlay <i>interface</i></code>	Creates an OTV overlay interface and enters interface configuration mode.
<code>otv-isis <i>default</i></code>	Enters OTV router configuration mode.

Table 3-5 OTV Interface-Level Commands

Command	Purpose
<code>description [dstring]</code>	(Optional) Configures a description for the overlay network. The <i>dstring</i> is any case-sensitive, alphanumeric string up to 80 characters.
<code>otv join-interface interface</code>	<p>Joins the OTV overlay interface with a physical Layer 3 interface. You must configure an IP address on the physical interface.</p> <p>You can specify only one join interface per overlay. You can decide to use one of the following methods:</p> <ul style="list-style-type: none"> ■ A single join interface, which is shared across multiple overlays ■ A different join interface for each overlay, which increases the OTV reliability
<code>otv extend-vlan vlan-range</code>	Extends a range of VLANs over this overlay interface and enables OTV advertisements for these VLANs. The <i>vlan-range</i> is from 1 to 3967 and from 4048 to 4093.
<code>otv extend-vlan {add remove } vlan-range</code>	Displays the VLAN information for the overlay network.
<code>otv site-vlan vlan-id</code>	Configures a VLAN that all local edge devices communicate on. You must configure this VLAN ID to match on all local edge devices. We recommend that you use the same VLAN ID across all sites. The range is from 1 to 3967, and from 4048 to 4093. The default is 1.
<code>otv site-identifier id</code>	Configures the site identifier. You should configure this same site identifier on all local OTV edge devices. The site identifier should be unique across different sites. The range is from 0x1 to 0xffffffff. The default is 0x0. The format is either hexadecimal or MAC address format.
<code>otv adjacency-server unicast-only</code>	<p>(Optional) Configures the local edge device to act as an adjacency server.</p> <p>NOTE: If the two overlay interface numbers do not match between the two OTV sites configured to use unicast adjacency servers, the OTV adjacencies will not form, and OTV will not come up until the overlay interface numbers are changed to match.</p>
<code>otv use-adjacency-server primary-ip-address [secondary-ip-address] unicast-only</code>	(Optional) Configures the local edge device to use a remote adjacency server. The IP address format is in dotted-decimal notation. The <i>secondary-ip-address</i> argument is the IP address of the backup adjacency server if you have configured a backup adjacency server.
<code>otv control-group mcast-address</code>	Configures the multicast group address used by the OTV control plane for this OTV overlay network. The multicast group address is an IPv4 address in dotted decimal notation.

Command	Purpose
<code>otv data-group mcast-range1 [mcast-range2...]</code>	Configures one or more ranges of local IPv4 multicast group prefixes used for multicast data traffic. Use SSM multicast groups 232.0.0.0/8. The multicast group address is an IPv4 address in dotted-decimal notation. A subnet mask is used to indicate ranges of addresses. You can define up to eight data-group ranges.
<code>otv site-vlan vlan-id</code>	Configures a VLAN that all local edge devices communicate on. You must configure this VLAN ID on all local edge devices. The range is from 1 to 3967 and from 4048 to 4093. The default is 1.

Table 3-6 OTV Router-Level Commands

Command	Purpose
<code>otv isis csnp-interval seconds</code>	(Optional) Specifies the interval between CSNP PDUs on an interface. The <i>seconds</i> range is from 1 to 65,535. The default is 10 seconds.
<code>otv isis hello-interval seconds</code>	(Optional) Specifies the interval between hello PDUs on an interface. The <i>seconds</i> range is from 1 to 65,535. The default is 10 seconds.
<code>otv isis hello-multiplier multiplier</code>	(Optional) Specifies the multiplier that is used to calculate the interval within which hello PDUs must be received to keep the OTV adjacency up. The <i>multiplier</i> range is from 3 to 1000. The default is 3.
<code>otv isis metric metric</code>	(Optional) Configures the OTV metric on an interface. The <i>metric</i> range is from 1 to 16,777,215.
<code>otv isis priority dis-priority</code>	(Optional) Configures the OTV priority for DIS election on the interface. The <i>priority</i> range is from 1 to 127. The default is 64.

Table 3-7 OTV Global-Level Verification Commands

Command	Purpose
<code>show running-configuration otv [all]</code>	Displays the running configuration for OTV.
<code>show otv overlay [interface]</code>	Displays information about overlay interfaces.
<code>show otv adjacency [detail]</code>	Displays information about the adjacencies on the overlay network.
<code>show otv [overlay interface] [vlan [vlan-range] [authoritative detail]]</code>	Displays information about VLANs that are associated with an overlay interface.
<code>show otv isis site [database statistics]</code>	Displays the BFD configuration state on both local and neighboring edge devices.
<code>show otv site [all]</code>	Displays information about the local site.

Command	Purpose
<code>show otv [route [interface [neighbor-address ip-address]] [vlan vlan-range] [mac-address]]</code>	Displays information about the OTV routes.
<code>show otv mroute vlan vlan-id startup</code>	Displays the OTV multicast route information for a specific VLAN from the OTV Routing Information Base (ORIB).
<code>show forwarding distribution otv multicast route vlan vlan-id</code>	Displays Forwarding Information Base (FIB) OTV multicast route information for a specific VLAN.
<code>show otv vlan-mapping [overlay interface-number]</code>	Displays VLAN translation mappings from a local site to a remote site.
<code>show mac address-table</code>	Displays information about MAC addresses.

Figure 3-12 shows the OTV network topology with configurations.

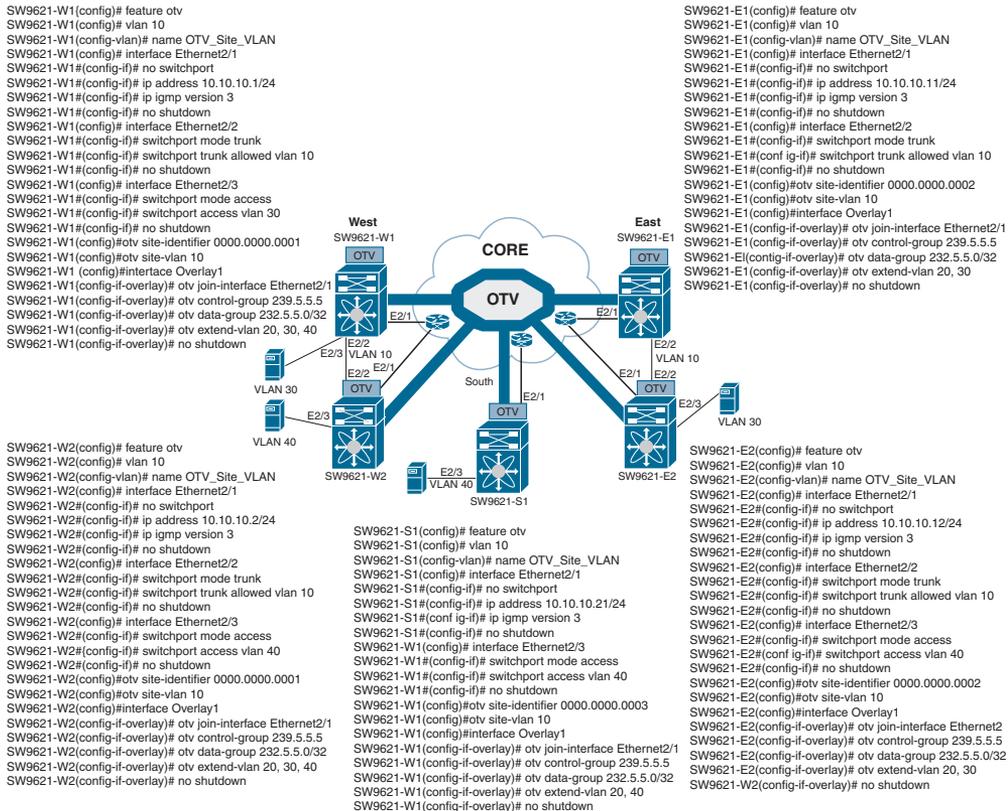


Figure 3-12 OTV Network Topology and Configurations

Example 3-1 shows switch one at the West site (SW9621-W1) with OTV information, neighbor adjacency, IS-IS underlay protocol, and multicast routing.

Example 3-1 SW9621-W1 OTV Results

```

SW9621-W1# show otv
OTV Overlay Information
Site Identifier 0000.0000.0001
Encapsulation-Format ip - gre
Overlay interface Overlay1
  VPN name          : Overlay1
  VPN state         : UP
  Extended vlans    : 20 30 40 (Total:3)
  Control group     : 239.5.5.5
  Data group range(s) : 232.5.5.0/32
  Broadcast group   : 239.5.5.5
  Join interface(s) : Eth2/1 (10.10.10.1)
  Site vlan         : 10 (up)
  AED-Capable      : yes
  Capability        : Multicast-Reachable
SW9621-W1# show otv adjacency
Overlay Adjacency database
Overlay-Interface Overlay1 :
Hostname          System-ID          Dest Addr      Up Time      State
SW9621-E2        fa16.3e31.f889      10.10.10.12   00:41:40    UP
SW9621-E1        fa16.3ec3.1f96      10.10.10.11   00:45:20    UP
SW9621-W2        fa16.3ed2.730a      10.10.10.2    00:55:39    UP
SW9621-S1        fa16.3ef4.5189      10.10.10.21   01:00:58    UP

SW9621-W1# show otv isis
ISIS process : default
Encap-Type for Multicast mode: GRE
VPN: Overlay1
fwd ready state FALSE
  System ID : fa16.3e56.e725  IS-Type : L1
  Version Number: 3
  Interoperability Flag: 0
  SAP : 439  Queue Handle : 15
  Maximum LSP MTU: 1392
  Graceful Restart enabled. State: Inactive
  Last graceful restart status : none
  Metric-style : advertise(wide), accept(narrow, wide)

```

```

Area address(es) :
  00
Process is up and running
VPN ID: 132
Incremental update routes during SPF run
Stale routes during non-graceful controlled restart
Interfaces supported by OTV-IS-IS :
  Overlay1
Level 1
Authentication type and keychain haven't been configured
Authentication check is specified
Address family IPv4 unicast :
  Number of interface : 1
  Adjacency check disabled
  Distance : 115
  tib-id : 0
Address family IPv6 unicast :
  Number of interface : 1
  Adjacency check disabled
  Distance : 115
  tib-id : 0
Address family MAC unicast :
  Number of interface : 1
  Adjacency check disabled
  Distance : 115
  tib-id : 0
L1 Next SPF: Inactive
AED Server Info:
Capability: 1
Priority: 0
AED Server Elected Value: fa16.3e56.e725
AED State: 0
AED Elected Operational: 4
Backup AED Server Info:
Backup AED Server Elected Value: fa16.3ed2.730a
Backup AED State: 0
Backup AED Elected Operational: 3
SW9621-W1# show ip mroute
IP Multicast Routing Table for VRF "default"

(*, 239.5.5.5/32), uptime: 01:28:26, otn ip
  Incoming interface: Ethernet2/1, RPF nbr: 10.10.10.1
  Outgoing interface list: (count: 1)
    Overlay1, uptime: 01:28:26, otn

```

Example 3-2 shows switch two at the West site (SW9621-W2) with OTV information.

Example 3-2 *SW9621-W2 OTV Results*

```

SW9621-W2# show otv
OTV Overlay Information
Site Identifier 0000.0000.0001
Encapsulation-Format ip - gre
Overlay interface Overlay1
  VPN name          : Overlay1
  VPN state         : UP
  Extended vlans    : 20 30 40 (Total:3)
  Control group     : 239.5.5.5
  Data group range(s) : 232.5.5.0/32
  Broadcast group   : 239.5.5.5
  Join interface(s) : Eth2/1 (10.10.10.2)
  Site vlan         : 10 (up)
  AED-Capable      : yes
  Capability        : Multicast-Reachable

```

Example 3-3 shows switch one at the East site (SW9621-E1) with OTV information. Be sure to pay attention to the AED status, which is shown as down, because there is no VLAN active at SW9621-E1.

Example 3-3 *SW9621-E1 OTV Results*

```

SW9621-E1# show otv
OTV Overlay Information
Site Identifier 0000.0000.0002
Encapsulation-Format ip - gre

Overlay interface Overlay1

  VPN name          : Overlay1
  VPN state         : UP
  Extended vlans    : 20 30 (Total:2)
  Control group     : 239.5.5.5
  Data group range(s) : 232.5.5.0/32
  Broadcast group   : 239.5.5.5
  Join interface(s) : Eth2/1 (10.10.10.11)
  Site vlan         : 10 (up)
  AED-Capable      : No (No extended vlan operationally up) <- no active VLAN
  Capability        : Multicast-Reachable

```

Example 3-4 shows switch two at the East site (SW9621-E2) with OTV information.

Example 3-4 SW9621-E2 OTV Results

```
SW9621-E2# show otv

OTV Overlay Information
Site Identifier 0000.0000.0002
Encapsulation-Format ip - gre

Overlay interface Overlay1
  VPN name          : Overlay1
  VPN state         : UP
  Extended vlans    : 20 30 (Total:2)
  Control group     : 239.5.5.5
  Data group range(s) : 232.5.5.0/32
  Broadcast group   : 239.5.5.5
  Join interface(s) : Eth2/3 (10.10.10.12)
  Site vlan         : 10 (up)
  AED-Capable      : Yes
  Capability        : Multicast-Reachable
```

Example 3-5 shows switch one at the South site (SW9621-S1) with OTV information. Again, be sure to pay attention to the VLAN site and the AED; both are shown as down, because there is a single OTV node at the South site.

Example 3-5 SW9621-S1 OTV Results

```
SW9621-S1# show otv

OTV Overlay Information
Site Identifier 0000.0000.0003
Encapsulation-Format ip - gre

Overlay interface Overlay1
  VPN name          : Overlay1
  VPN state         : UP
  Extended vlans    : 20 40 (Total:2)
  Control group     : 239.5.5.5
  Data group range(s) : 232.5.5.0/32
  Broadcast group   : 239.5.5.5
  Join interface(s) : Eth2/1 (10.10.10.21)
  Site vlan         : 10 (down)
  AED-Capable      : No (Site-VLAN is Down)
  Capability        : Multicast-Reachable
```

Example 3-6 shows connectivity results across the OTV network.

Example 3-6 *Connectivity Verifications from Devices Behind the Switches*

```

VLAN30-W1# ping 30.0.0.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.0.0.2, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/20/40 ms

VLAN30-W2# Router# ping 30.0.0.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 30.0.0.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

VLAN40-W1# Router# ping 40.0.0.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 40.0.0.2, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 3/20/40 ms

VLAN40-W2# Router# ping 40.0.0.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 40.0.0.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/3 ms

```

3

Virtual Extensible LAN (VXLAN) Overview

VXLAN is an extension to the Layer 2 VLAN. It was designed to provide the same VLAN functionality with greater extensibility and flexibility. VXLAN offers the following benefits:

- **VLAN flexibility in multitenant segments:** It provides a solution to extend Layer 2 segments over the underlying network infrastructure so that tenant workload can be placed across physical pods in the data center.
- **Higher scalability:** VXLAN uses a 24-bit segment ID known as the VXLAN network identifier (VNID), which enables up to 16 million VXLAN segments to coexist in the same administrative domain.
- **Improved network utilization:** VXLAN solved Layer 2 STP limitations. VXLAN packets are transferred through the underlying network based on its Layer 3 header and can take complete advantage of Layer 3 routing, equal-cost multipath (ECMP) routing, and link aggregation protocols to use all available paths.



VXLAN Encapsulation and Packet Format

VXLAN is a solution to support a flexible, large-scale multitenant environment over a shared common physical infrastructure. The transport protocol over the physical data center network is IP plus UDP.

VXLAN defines a MAC-in-UDP encapsulation scheme where the original Layer 2 frame has a VXLAN header added and is then placed in a UDP-IP packet. With this MAC-in-UDP encapsulation, VXLAN tunnels the Layer 2 network over the Layer 3 network. The VXLAN packet format is shown in Figure 3-13.

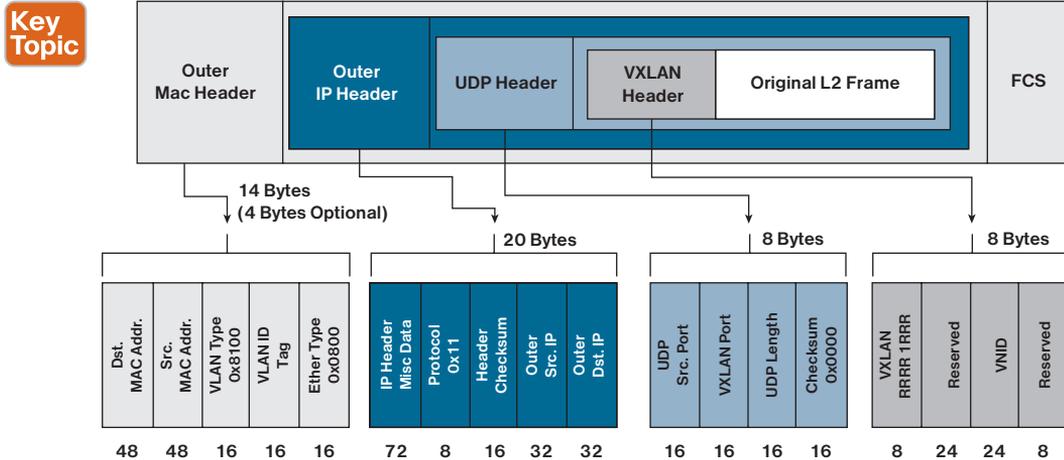


Figure 3-13 VXLAN Packet Format

As shown in Figure 3-13, VXLAN introduces an 8-byte VXLAN header that consists of a 24-bit VNID and a few reserved bits. The VXLAN header together with the original Ethernet frame goes in the UDP payload. The 24-bit VNID is used to identify Layer 2 segments and to maintain Layer 2 isolation between the segments. With all 24 bits in VNID, VXLAN can support 16 million LAN segments.

VXLAN Tunnel Endpoint

VXLAN uses the VXLAN tunnel endpoint (VTEP) to map tenants' end devices to VXLAN segments and to perform VXLAN encapsulation and decapsulation. Each VTEP function has two interfaces: one is a switch interface on the local LAN segment to support local end-point communication, and the other is an IP interface to the transport IP network.

Infrastructure VLAN is a unique IP address that identifies the VTEP device on the transport IP network. The VTEP device uses this IP address to encapsulate Ethernet frames and transmits the encapsulated packets to the transport network through the IP interface.

A VTEP device also discovers the remote VTEPs for its VXLAN segments and learns remote MAC Address-to-VTEP mappings through its IP interface. The functional components of VTEPs and the logical topology that is created for Layer 2 connectivity across the transport IP network are shown in Figure 3-14.

The VXLAN segments are independent of the underlying network topology; conversely, the underlying IP network between VTEPs is independent of the VXLAN overlay. It routes the encapsulated packets based on the outer IP address header, which has the initiating VTEP as the source IP address and the terminating VTEP as the destination IP address.

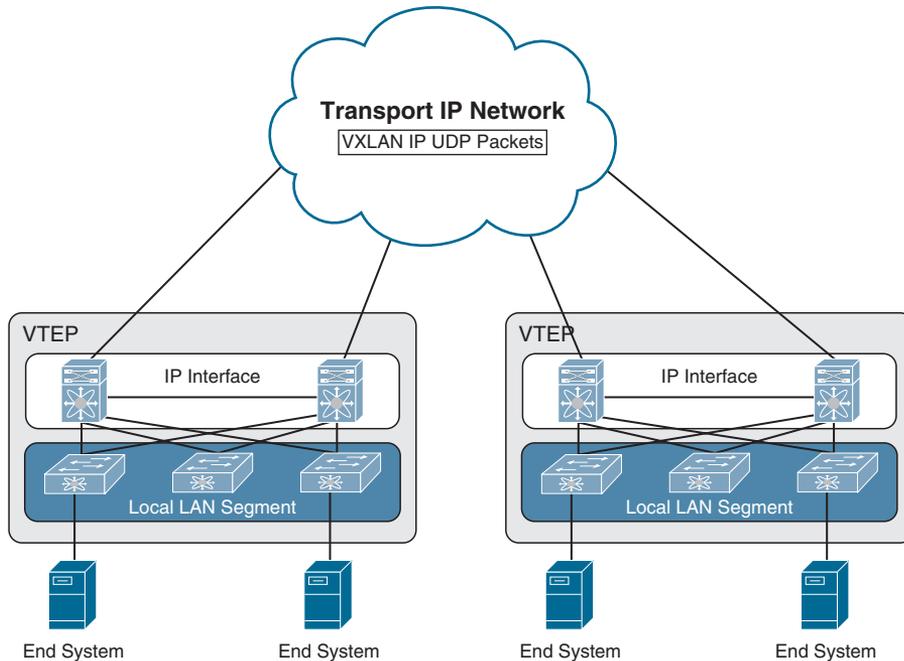


Figure 3-14 VXLAN Tunnel Endpoint (VTEP)

Virtual Network Identifier

A virtual network identifier (VNI) is a value that identifies a specific virtual network in the data plane. It is typically a 24-bit value part of the VXLAN header, which can support up to 16 million individual network segments. (Valid VNI values are from 4096 to 16,777,215.) There are two main VNI scopes:

- Network-wide scoped VNIs:** The same value is used to identify the specific Layer 3 virtual network across all network edge devices. This network scope is useful in environments such as within the data center where networks can be automatically provisioned by central orchestration systems.

Having a uniform VNI per VPN is a simple approach, while also easing network operations (such as troubleshooting). It also means simplified requirements on network edge devices, both physical and virtual devices. A critical requirement for this type of approach is to have a very large number of network identifier values given the network-wide scope.

- Locally assigned VNIs:** In an alternative approach supported as per RFC 4364, the identifier has local significance to the network edge device that advertises the route. In this case, the virtual network scale impact is determined on a per-node basis versus a network basis.

When it is locally scoped and uses the same existing semantics as an MPLS VPN label, the same forwarding behaviors as specified in RFC 4364 can be employed. This scope thus allows a seamless stitching together of a VPN that spans both an IP-based network overlay and an MPLS VPN.

This situation can occur, for instance, at the data center edge where the overlay network feeds into an MPLS VPN. In this case, the identifier may be dynamically allocated by the advertising device.

It is important to support both cases and, in doing so, ensure that the scope of the identifier be clear and the values not conflict with each other.

Key Topic

VXLAN Control Plane

Two widely adopted control planes are used with VXLAN: the VXLAN Flood and Learn Multicast-Based Control Plane and the VXLAN MPBGP EVPN Control Plane.

VXLAN Flood and Learn Multicast-Based Control Plane

Cisco Nexus switches utilize existing Layer 2 flooding mechanisms and dynamic MAC address learning to

- Transport broadcast, unknown unicast, and multicast (BUM) traffic
- Discover remote VTEPs
- Learn remote-host MAC addresses and MAC-to-VTEP mappings for each VXLAN segment

IP multicast is used to reduce the flooding scope of the set of hosts that are participating in the VXLAN segment. Each VXLAN segment, or VNID, is mapped to an IP multicast group in the transport IP network. Each VTEP device is independently configured and joins this multicast group as an IP host through the Internet Group Management Protocol (IGMP). The IGMP joins trigger Protocol Independent Multicast (PIM) joins and signaling through the transport network for the particular multicast group. The multicast distribution tree for this group is built through the transport network based on the locations of participating VTEPs. The multicast tunnel of a VXLAN segment through the underlying IP network is shown in Figure 3-15.

Key Topic

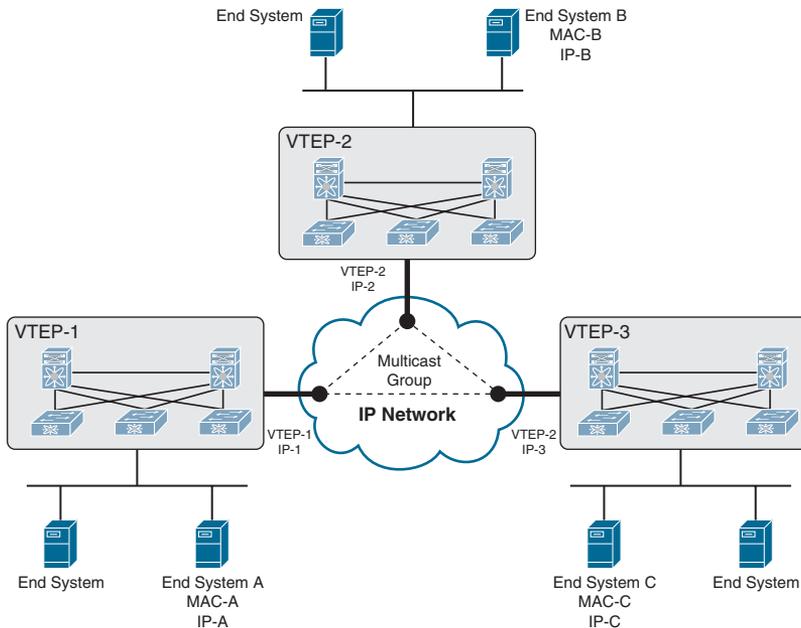


Figure 3-15 VXLAN Multicast Group in Transport Network

The multicast group shown in Figure 3-16 is used to transmit VXLAN broadcast, unknown unicast, and multicast traffic through the IP network, limiting Layer 2 flooding to those devices that have end systems participating in the same VXLAN segment. VTEPs communicate with one another through the flooded or multicast traffic in this multicast group.

Key Topic

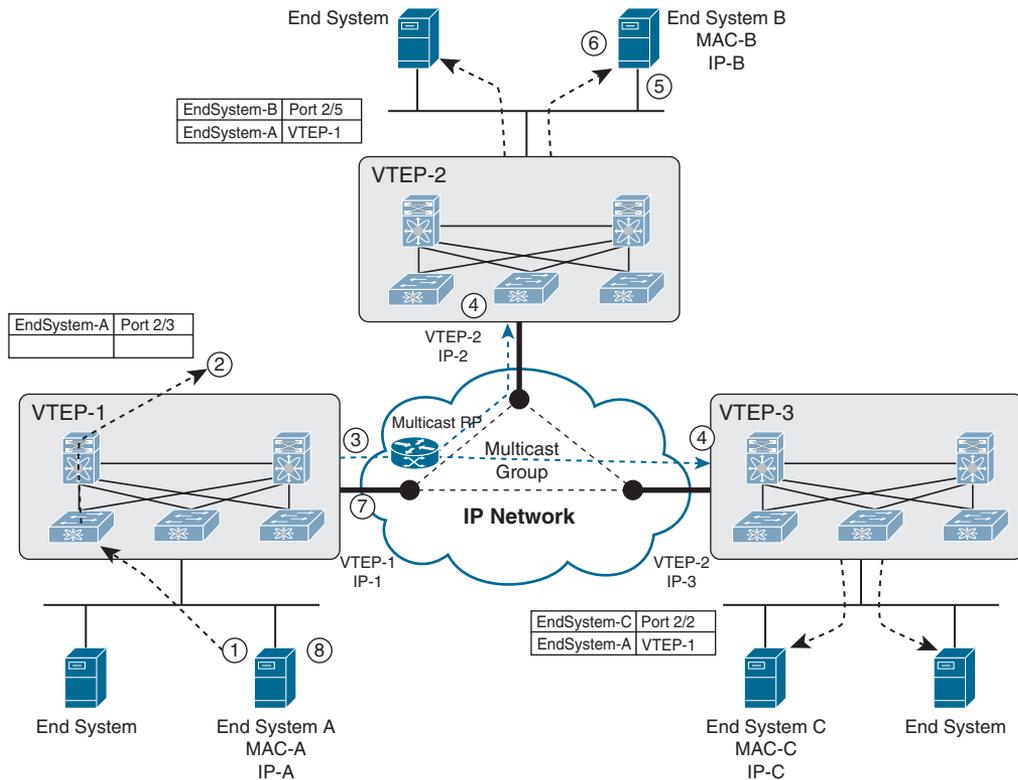


Figure 3-16 VXLAN Multicast Control Plane

As an example, if End System A wants to talk to End System B, it does the following:

1. End System A generates an ARP request trying to discover the End System B MAC address.
2. When the ARP request arrives at SW1, it will look up its local table, and if an entry is not found, it will encapsulate the ARP request over VXLAN and send it over the multicast group configured for the specific VNI.
3. The multicast RP receives the packet, and it forwards a copy to every VTEP that has joined the multicast group.
4. Each VTEP receives and deencapsulates the packet and learns the System A MAC address pointing to the remote VTEP address.
5. Each VTEP forwards the ARP request to its local destinations.
6. End System B generates the ARP reply. When SW2 VTEP2 receives it, it looks up its local table and finds an entry with the information that traffic destined to End System A must be sent to VTEP1 address. VTEP2 encapsulates the ARP reply with a VXLAN header and unicasts it to VTEP1.

7. VTEP1 receives and deencapsulates the packet and delivers it to End System A.
8. When the MAC address information is learned, additional packets are fed to the corresponding VTEP address.

VXLAN MPBGP EVPN Control Plane

The EVPN overlay specifies adaptations to the BGP MPLS-based EVPN solution so that it is applied as a network virtualization overlay with VXLAN encapsulation where

- The PE node role described in BGP MPLS EVPN is equivalent to the VTEP/network virtualization edge (NVE) device.
- VTEP information is distributed via BGP.
- VTEPs use control plane learning/distribution via BGP for remote MAC addresses instead of data plane learning.
- Broadcast, unknown unicast, and multicast (BUM) data traffic is sent using a shared multicast tree.
- A BGP route reflector (RR) is used to reduce the full mesh of BGP sessions among VTEPs to a single BGP session between a VTEP and the RR.
- Route filtering and constrained route distribution are used to ensure that the control plane traffic for a given overlay is distributed only to the VTEPs that are in that overlay instance.
- The host (MAC) mobility mechanism ensures that all the VTEPs in the overlay instance know the specific VTEP associated with the MAC.
- Virtual network identifiers (VNIs) are globally unique within the overlay.

The EVPN overlay solution for VXLAN can also be adapted to enable it to be applied as a network virtualization overlay with VXLAN for Layer 3 traffic segmentation. The adaptations for Layer 3 VXLAN are similar to L2 VXLAN, except the following:

- VTEPs use control plane learning/distribution via BGP of IP addresses (instead of MAC addresses).
- The virtual routing and forwarding instances are mapped to the VNI.
- The inner destination MAC address in the VXLAN header does not belong to the host but to the receiving VTEP that does the routing of the VXLAN payload. This MAC address is distributed via the BGP attribute along with EVPN routes.

VXLAN Gateways

VXLAN gateways are used to connect VXLAN and classic VLAN segments to create a common forwarding domain so that tenant devices can reside in both environments. The types of VXLAN gateways are

- **Layer 2 Gateway:** A Layer 2 VXLAN gateway is a device that encapsulates a classical Ethernet (CE) frame into a VXLAN frame and decapsulates a VXLAN frame into a CE frame. A gateway device transparently provides VXLAN benefits to a device that

does not support VXLAN; that device could be a physical host or a virtual machine. The physical hosts or VMs are completely unaware of the VXLAN encapsulation.

- **VXLAN Layer 3 Gateway:** Similar to traditional routing between different VLANs, a VXLAN router is required for communication between devices that are in different VXLAN segments. The VXLAN router translates frames from one VNI to another. Depending on the source and destination, this process might require decapsulation and reencapsulation of a frame. The Cisco Nexus device supports all combinations of decapsulation, route, and encapsulation. The routing can also be done across native Layer 3 interfaces and VXLAN segments.

You can enable VXLAN routing at the aggregation layer or on Cisco Nexus device aggregation nodes. The spine forwards only IP-based traffic and ignores the encapsulated packets. To help scaling, a few leaf nodes (a pair of border leaves) perform routing between VNIs. A set of VNIs can be grouped into a virtual routing and forwarding (VRF) instance (tenant VRF) to enable routing among those VNIs. If routing must be enabled among a large number of VNIs, you might need to split the VNIs between several VXLAN routers. Each router is responsible for a set of VNIs and a respective subnet. Redundancy is achieved with FHRP.

VXLAN High Availability

For high availability, a pair of virtual port channel (vPC) switches can be used as a logical VTEP device sharing an anycast VTEP address (shown in Figure 3-17).

The vPC switches provide vPCs for redundant host connectivity while individually running Layer 3 protocols with the upstream devices in the underlay network. Both will join the multicast group for the same VXLAN VNI and use the same anycast VTEP address as the source to send VXLAN-encapsulated packets to the devices in the underlay network, including the multicast rendezvous point and the remote VTEP devices. The two vPC VTEP switches appear to be one logical VTEP entity.

vPC peers must have the following identical configurations:

- Consistent mapping of the VLAN to the virtual network segment (VN-segment)
- Consistent NVE binding to the same loopback secondary IP address (anycast VTEP address)
- Consistent VNI-to-group mapping.

For the anycast IP address, vPC VTEP switches must use a secondary IP address on the loopback interface bound to the VXLAN NVE tunnel. The two vPC switches need to have the exact same secondary loopback IP address.

Both devices will advertise this anycast VTEP address on the underlay network so that the upstream devices learn the /32 route from both vPC VTEPs and can load-share VXLAN unicast-encapsulated traffic between them.

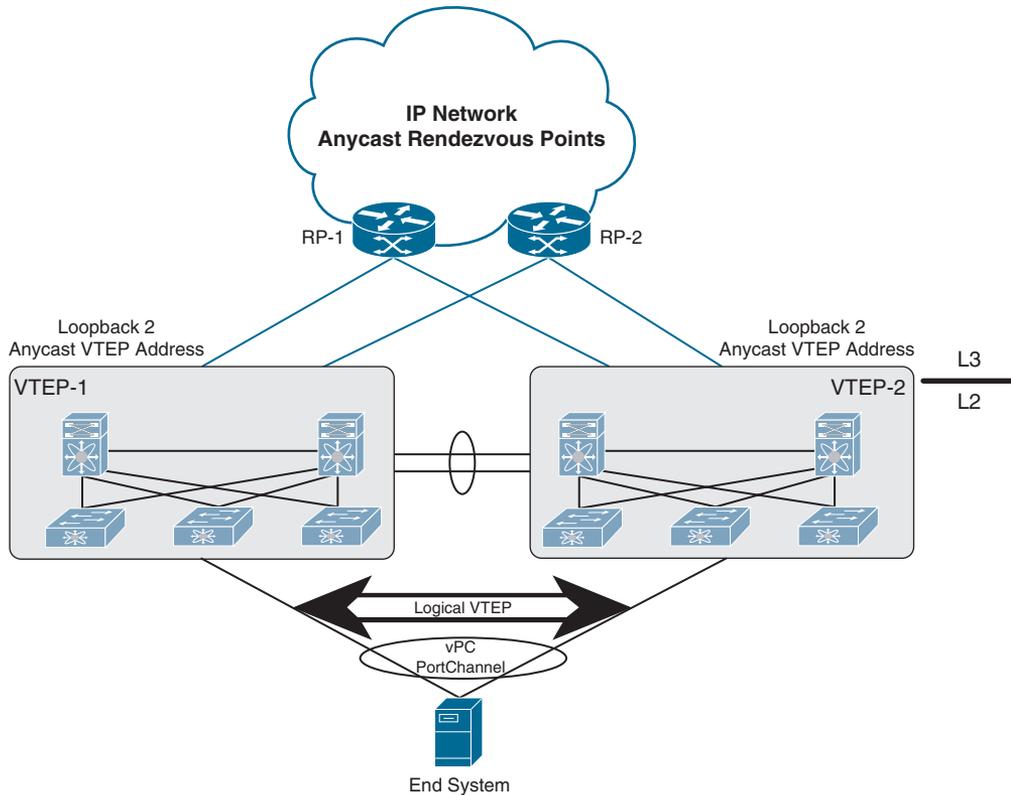


Figure 3-17 *VXLAN High Availability*

In the event of vPC peer-link failure, the vPC operational secondary switch will shut down its loopback interface bound to VXLAN NVE. This shutdown will cause the secondary vPC switch to withdraw the anycast VTEP address from its IGP advertisement so that the upstream devices in the underlay network start to send all traffic just to the primary vPC switch. The purpose of this process is to avoid a vPC active-active situation when the peer link is down. With this mechanism, the orphan devices connected to the secondary vPC switch will not be able to receive VXLAN traffic when the vPC peer link is down.

VXLAN Tenant Routed Multicast

Tenant Routed Multicast (TRM) brings the efficiency of multicast delivery to VXLAN overlays. It is based on standards-based next-gen control plane (ngMVPN) described in IETF RFCs 6513 and 6514. TRM enables the delivery of customer Layer 3 multicast traffic in a multitenant fabric, and this in an efficient and resilient manner.

While BGP EVPN provides a control plane for unicast routing, as shown in Figure 3-18, ngMVPN provides scalable multicast routing functionality. It follows an “always route” approach where every edge device (VTEP) with distributed IP Anycast Gateway for unicast becomes a designated router (DR) for multicast. Bridged multicast forwarding is present only on the edge devices (VTEP) where IGMP snooping optimizes the multicast forwarding to interested receivers. All other multicast traffic beyond local delivery is efficiently routed.

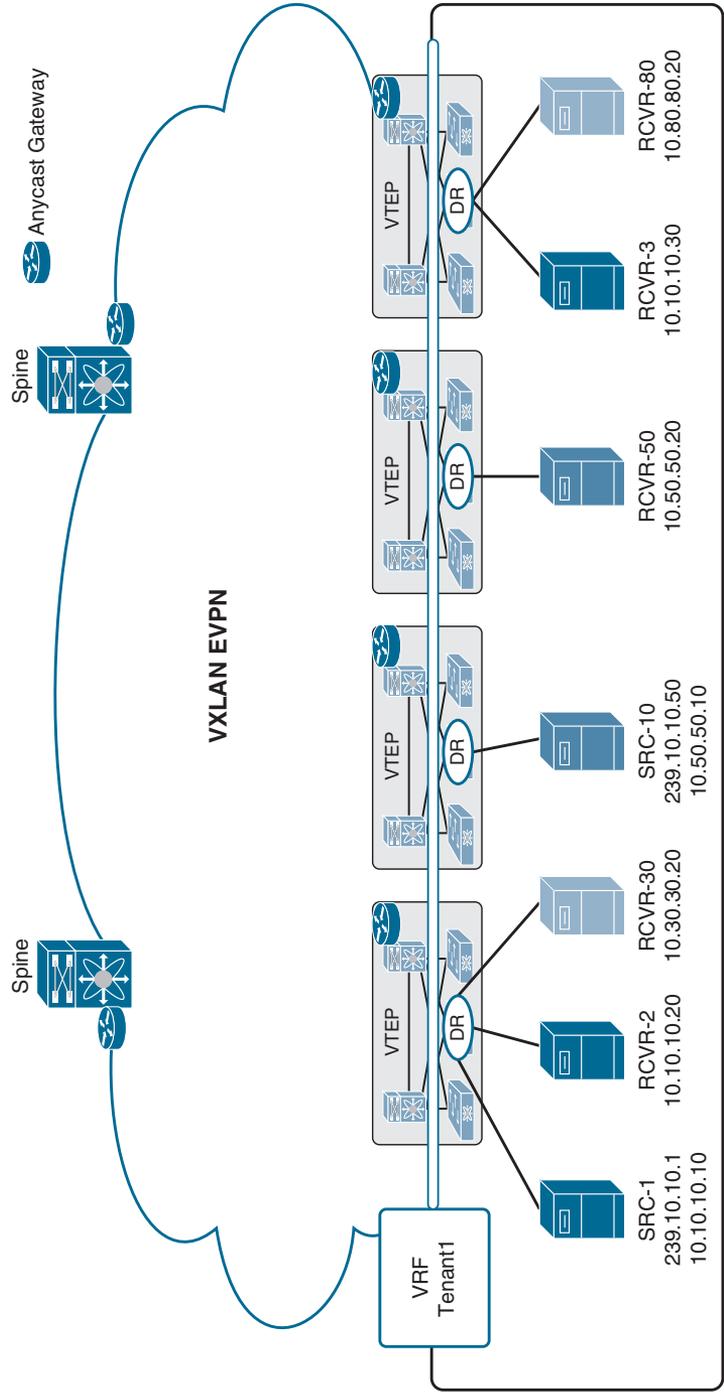


Figure 3-18 Tenant Routed Multicast (TRM)

With TRM enabled, multicast forwarding in the underlay is leveraged to replicate VXLAN-encapsulated routed multicast traffic. A Default Multicast Distribution Tree (Default-MDT) is built per VRF. This is an addition to the existing multicast groups for Layer 2 VNI broadcast, unknown unicast, and Layer 2 multicast replication group. The individual multicast group addresses in the overlay are mapped to the respective underlay multicast address for replication and transport. The advantage of using a BGP-based approach is that TRM can operate as a fully distributed overlay rendezvous point (RP), with the RP presence on every edge device (VTEP).

A multicast-enabled data center fabric is typically part of an overall multicast network. Multicast sources, receivers, and even the multicast rendezvous point might reside inside the data center but might also be inside the campus or externally reachable via the WAN. TRM allows seamless integration with existing multicast networks. It can leverage multicast rendezvous points external to the fabric. Furthermore, TRM allows for tenant-aware external connectivity using Layer 3 physical interfaces or subinterfaces.

VXLAN Configurations and Verifications

VXLAN requires a license. Table 3-8 shows the NX-OS feature license required for VXLAN. For more information, visit the Cisco NX-OS Licensing Guide.

Table 3-8 VXLAN Feature-Based Licenses for Cisco NX-OS

Platform	Feature License	Feature Name
Cisco Nexus 9000 Series switches	LAN_ENTERPRISE_SERVICES_PK	Cisco programmable fabric spine, leaf, or border leaf
Cisco Nexus 7000 Series switches and Cisco Nexus 7700 switches	LAN_ENTERPRISE_SERVICES_PKG Multiprotocol Label Switching (MPLS) Service Package (MPLS_PKG)	Cisco programmable fabric spine, leaf, border leaf, or border PE switch
Cisco Nexus 5600 switches	Layer 3 Base Services Package (LAN_BASE_SERVICES_PKG) Enterprise Services Package (LAN_ENTERPRISE_SERVICES_PKG)	Cisco programmable fabric spine, leaf, border leaf, or border PE switch

Tables 3-9 through 3-12 show the most-used VXLAN configuration commands along with their purpose. For full commands, refer to the Nexus VXLAN Configuration Guide.

Table 3-9 VXLAN Global-Level Commands

Command	Purpose
<code>feature nv overlay</code>	Enables the VXLAN feature.
<code>feature vn-segment-vlan-based</code>	Configures the global mode for all VXLAN bridge domains.
<code>vlan <i>vlan-id</i></code>	Specifies VLAN.
<code>vn-segment <i>vnid</i></code>	Specifies VXLAN virtual network identifier (VNID).
<code>bridge-domain <i>domain</i></code>	Enters the bridge domain configuration mode. It will create a bridge domain if it does not yet exist. Use from the global configuration mode.
<code>dot1q vlan <i>vni</i> <i>vni</i></code>	Creates mapping between VLAN and VNI. Use from the encapsulation profile configuration mode.

Command	Purpose
<code>encapsulation profile name_of_profile default</code>	Applies an encapsulation profile to a service profile. Use from the service instance configuration mode.
<code>encapsulation profile vni name_of_profile</code>	Creates an encapsulation profile. Use from the global configuration mode.
<code>service instance instance vni</code>	Creates a service instance. Use from the interface configuration mode.
<code>interface nve x</code>	Creates a VXLAN overlay interface that terminates VXLAN tunnels.
<code>mac address-table static mac-address vni vni-id interface nve x peer-ip ip-address</code>	Specifies the MAC address pointing to the remote VTEP. NOTE: Only 1 NVE interface is allowed on the switch.
<code>ip igmp snooping vxlan</code>	Enables IGMP snooping for VXLAN VLANs. You have to explicitly configure this command to enable snooping for VXLAN VLANs.
<code>ip igmp snooping disable-nve-static-router-port</code>	Configures IGMP snooping over VXLAN so that it does not include NVE as a static multicast router (mrouter) port using this global CLI command. The NVE interface for IGMP snooping over VXLAN is the mrouter port by default.

Table 3-10 Interface-Level Commands

Command	Purpose
<code>switchport vlan mapping enable</code>	Enables VLAN translation on the switch port. VLAN translation is disabled by default. NOTE: Use the no form of this command to disable VLAN translation.
<code>switchport vlan mapping vlan-id translated-vlan-id</code>	Translates a VLAN to another VLAN. The range for both the <i>vlan-id</i> and <i>translated-vlan-id</i> arguments is from 1 to 4094. <ul style="list-style-type: none"> You can configure VLAN translation between the ingress (incoming) VLAN and a local (translated) VLAN on a port. For the traffic arriving on the interface where VLAN translation is enabled, the incoming VLAN is mapped to a translated VLAN that is VXLAN enabled. On the underlay, this is mapped to a VNI; the inner dot1q is deleted and switched over to the VXLAN network. On the egress switch, the VNI is mapped to a translated VLAN. On the outgoing interface, where VLAN translation is configured, the traffic is converted to the original VLAN and egress out. NOTE: Use the no form of this command to clear the mappings between a pair of VLANs.
<code>switchport vlan mapping all</code>	Removes all VLAN mappings configured on the interface.

Table 3-11 Network Virtual Interface (NVE) Config Commands

Command	Purpose
<code>source-interface <i>src-if</i></code>	The source interface must be a loopback interface that is configured on the switch with a valid /32 IP address. The transient devices in the transport network and the remote VTEPs must know this /32 IP address. This is accomplished by advertising it through a dynamic routing protocol in the transport network.
<code>member vni <i>vni</i></code>	Associates VXLAN virtual network identifiers (VNIs) with the NVE interface.
<code>mcast-group <i>start-address</i> [<i>end-address</i>]</code>	Assigns a multicast group to the VNIs. NOTE: Used only for BUM traffic.
<code>ingress-replication protocol bgp</code>	Enables BGP EVPN with ingress replication for the VNI.
<code>ingress-replication protocol static</code>	Enables static ingress replication for the VNI.
<code>peer-ip <i>n.n.n.n</i></code>	Enables peer IP.

Table 3-12 VXLAN Global-Level Verification Commands

Command	Purpose
<code>show tech-support vxlan [<i>platform</i>]</code>	Displays related VXLAN tech-support information.
<code>show bridge-domain</code>	Shows the bridge domain.
<code>show logging level nve</code>	Displays the logging level.
<code>show tech-support nve</code>	Displays related NVE tech-support information.
<code>show run interface nve <i>x</i></code>	Displays NVE overlay interface configuration.
<code>show nve interface</code>	Displays NVE overlay interface status.
<code>show nve peers</code>	Displays NVE peer status.
<code>show nve peers <i>peer_IP_</i> <i>address</i> interface <i>interface_</i> <i>ID</i> counters</code>	Displays per-NVE peer statistics.
<code>clear nve peer-ip <i>peer-ip-address</i></code>	Clears stale NVE peers. Stale NVE peers are those that do not have MAC addresses learned behind them.
<code>show nve vni</code>	Displays VXLAN VNI status.
<code>show nve vni ingress-replication</code>	Displays the mapping of VNI to an ingress-replication peer list and uptime for each peer.
<code>show nve vni <i>vni_number</i> counters</code>	Displays per-VNI statistics.
<code>show nve vxlan-params</code>	Displays VXLAN parameters, such as VXLAN destination or UDP port.

Figure 3-19 shows the VXLAN network topology with configurations.

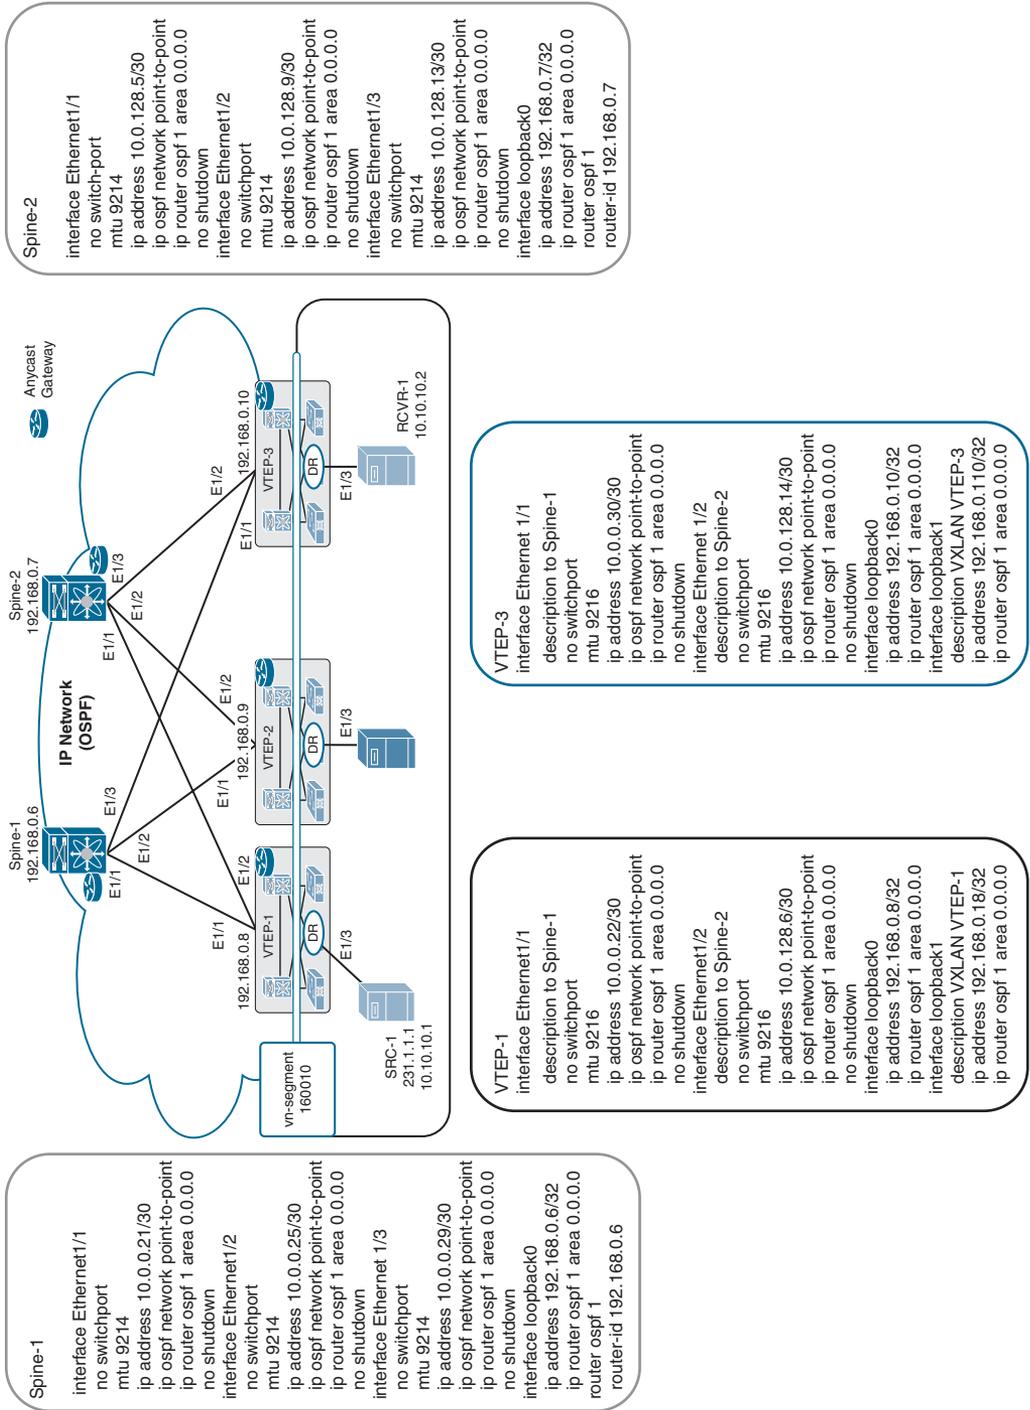


Figure 3-19 VXLAN Control Plane Topology

Example 3-7 shows the spine router (Spine-1 and Spine-2) OSPF and multicast routing configuration, VTEP (VTEP-1 and VTEP-3) multicast routing configuration, and multicast routing verification.

Example 3-7 *PIM Multicast Configurations and Verifications*

```

Spine-1 Config
Spine-1(config)# feature pim
Spine-1(config)# interface loopback1
Spine-1(config-if)# ip address 192.168.0.100/32
Spine-1(config-if)# ip pim sparse-mode
Spine-1(config-if)# ip router ospf 1 area 0.0.0.0
Spine-1(config)# ip pim rp-address 192.168.0.100
Spine-1(config)# ip pim anycast-rp 192.168.0.100 192.168.0.6
Spine-1(config)# ip pim anycast-rp 192.168.0.100 192.168.0.7
Spine-1(config)# interface E1/1
Spine-1(config-if)# ip pim sparse-mode
Spine-1(config)# interface E1/2
Spine-1(config-if)# ip pim sparse-mode
Spine-1(config)# interface E1/3
Spine-1(config-if)# ip pim sparse-mode
Spine-1(config)# interface loopback0
Spine-1(config-if)# ip pim sparse-mode
Spine-2 Config (PIM Redundancy)
Spine-2(config)# feature pim
Spine-2(config)# interface loopback1
Spine-2(config-if)# ip address 192.168.0.100/32
Spine-2(config-if)# ip pim sparse-mode
Spine-2(config-if)# ip router ospf 1 area 0.0.0.0
Spine-2(config)# ip pim rp-address 192.168.0.100
Spine-2(config)# ip pim anycast-rp 192.168.0.100 192.168.0.6
Spine-2(config)# ip pim anycast-rp 192.168.0.100 192.168.0.7
Spine-2(config)# interface E1/1
Spine-2(config-if)# ip pim sparse-mode
Spine-2(config)# interface E1/2
Spine-2(config-if)# ip pim sparse-mode
Spine-2(config)# interface E1/3
Spine-2(config-if)# ip pim sparse-mode
Spine-2(config)# interface loopback0
Spine-2(config-if)# ip pim sparse-mode
VTEP-1 PIM Config
VTEP-1(config)# feature pim
VTEP-1(config)# ip pim rp-address 192.168.0.100
VTEP-1 (config)# interface E1/1
VTEP-1 (config-if)# ip pim sparse-mode
VTEP-1 (config)# interface E1/2
VTEP-1 (config-if)# ip pim sparse-mode

```

```

VTEP-1 (config)# interface loopback0
VTEP-1 (config-if)# ip pim sparse-mode
VTEP-1 (config)# interface loopback1
VTEP-1 (config-if)# ip pim sparse-mode
VTEP-3 PIM Config
VTEP-3 (config)# feature pim
VTEP-3 (config)# ip pim rp-address 192.168.0.100
VTEP-3 (config)# interface E1/1
VTEP-3 (config-if)# ip pim sparse-mode
VTEP-3 (config)# interface E1/2
VTEP-3 (config-if)# ip pim sparse-mode
VTEP-3 (config)# interface loopback0
VTEP-3 (config-if)# ip pim sparse-mode
VTEP-3 (config)# interface loopback1
VTEP-3 (config-if)# ip pim sparse-mode
Spine 1 Verifications
Spine-1# show ip pim neighbor
PIM Neighbor Status for VRF "default"

```

Neighbor	Interface	Uptime	Expires	DR Priority	Bidir- Capable	BFD State
10.0.0.22	Ethernet1/1	00:02:21	00:01:23	1	yes	n/a
10.0.0.26	Ethernet1/2	00:01:50	00:01:20	1	yes	n/a
10.0.0.30	Ethernet1/3	00:00:37	00:01:38	1	yes	n/a

```

Spine-1# show ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP disabled
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
Anycast-RP 192.168.0.100 members:
  192.168.0.6* 192.168.0.7
RP: 192.168.0.100*, (0),
  uptime: 00:04:29  priority: 255,
  RP-source: (local),
  group ranges:
  224.0.0.0/4
Spine 2 Verifications
Spine-2# show ip pim neighbor
PIM Neighbor Status for VRF "default"

```

Neighbor	Interface	Uptime	Expires	DR Priority	Bidir- Capable	BFD State
10.0.128.6	Ethernet1/1	00:02:21	00:01:23	1	yes	n/a
10.0.128.10	Ethernet1/2	00:01:50	00:01:20	1	yes	n/a
10.0.128.14	Ethernet1/3	00:00:37	00:01:38	1	yes	n/a

```

Spine-2# show ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP disabled
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
Anycast-RP 192.168.0.100 members:
  192.168.0.6  192.168.0.7*
RP: 192.168.0.100*, (0),
  uptime: 00:04:16  priority: 255,
  RP-source: (local),
  group ranges:
  224.0.0.0/4

VTEP-1 Verifications
VTEP-1# show ip pim neighbor
PIM Neighbor Status for VRF "default"

```

Neighbor	Interface	Uptime	Expires	DR Priority	Bidir- Capable	BFD State
10.0.0.21	Ethernet1/1	00:03:47	00:01:32	1	yes	n/a
10.0.128.5	Ethernet1/2	00:03:46	00:01:37	1	yes	n/a

```

VTEP-1# show ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP disabled
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
RP: 192.168.0.100, (0),
  uptime: 00:03:53  priority: 255,
  RP-source: (local),
  group ranges:
  224.0.0.0/4

VTEP-3 Verifications
VTEP-3# show ip pim neighbor
PIM Neighbor Status for VRF "default"

```

Neighbor	Interface	Uptime	Expires	DR Priority	Bidir- Capable	BFD State
10.0.0.21	Ethernet1/1	00:03:47	00:21:32	1	yes	n/a
10.0.128.5	Ethernet1/2	00:03:46	00:03:37	1	yes	n/a

```

VTEP-3# show ip pim rp
PIM Neighbor Status for VRF "default"
Neighbor          Interface          Uptime    Expires    DR        Bidir-  BFD
                  Priority Capable State
10.0.0.29         Ethernet1/1        00:03:06  00:01:21  1         yes     n/a
10.0.128.13      Ethernet1/2        00:02:48  00:01:35  1         yes     n/a
Leaf-3(config)# show ip pim rp
PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP disabled
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
RP: 192.168.0.100, (0),
    uptime: 00:03:11  priority: 255,
    RP-source: (local),
    group ranges:
    224.0.0.0/4

```

Example 3-8 shows the VTEP (VETP-1 and VTEP-3) VXLAN and VXLAN Network Virtual Interface (NVE) configuration and status verification.

Example 3-8 VXLAN Configurations and Verifications

```

VTEP-1 Config
VTEP-1(config)# feature vn-segment-vlan-based
VTEP-1(config)# feature vn overlay
VTEP-1(config)# vlan 10
VTEP-1(config-vlan)# vn-segment 160010
VTEP-1(config)# vlan 20
VTEP-1(config-vlan)# vn-segment 160020
VTEP-1(config)# interface nve1
VTEP-1 (config-if)# source-interface loopback1
VTEP-1 (config-if)# member vni 160010 mcast-group 231.1.1.1
VTEP-1 (config-if)# member vni 160020 mcast-group 231.1.1.1
VTEP-1 (config-if)# no shutdown
VTEP-3 Config
VTEP-3(config)# feature vn-segment-vlan-based
VTEP-3(config)# feature vn overlay
VTEP-3(config)# vlan 10
VTEP-3(config-vlan)# vn-segment 160010
VTEP-3(config)# vlan 20
VTEP-3(config-vlan)# vn-segment 160020
VTEP-3(config)# interface nve1
VTEP-3(config-if)# source-interface loopback1

```

```

VTEP-3(config-if)# member vni 160010 mcast-group 231.1.1.1
VTEP-3(config-if)# member vni 160020 mcast-group 231.1.1.1
VTEP-3(config-if)# no shutdown

VTEP-1 Verifications
VTEP-1# show nve vni
Codes: CP - Control Plane      DP - Data Plane
      UC - Unconfigured       SA - Suppress ARP
      SU - Suppress Unknown Unicast

Interface VNI      Multicast-group  State Mode Type [BD/VRF]      Flags
-----
nve1      160010          231.1.1.1        Up   DP   L2   [10]
nve1      160020          231.1.1.1        Up   DP   L2   [20]

VTEP-1# show vxlan
Vlan      VN-Segment
====      =====
10         160010
20         160020

VTEP-1# ping 10.10.10.3
PING 10.10.10.3 (10.10.10.3) : 56 data bytes
64 bytes from 10.10.10.3: icmp_seq=0 ttl=254 time=8.114 ms
64 bytes from 10.10.10.3: icmp_seq=1 ttl=254 time=5.641 ms
64 bytes from 10.10.10.3: icmp_seq=2 ttl=254 time=6.213 ms
64 bytes from 10.10.10.3: icmp_seq=3 ttl=254 time=6.119 ms

VTEP-1# show nve peers
Interface Peer-IP      State LearnType Uptime      Router-Mac
-----
nve1      192.168.0.110      Up   DP           00:09:08    n/a

VTEP-1# show ip mroute
IP Multicast Routing Table for VRF "default"
(*, 231.1.1.1/32), uptime: 00:10:38, nve ip pim
  Incoming interface: Ethernet1/1, RPF nbr: 10.0.0.29
  Outgoing interface list: (count: 1)
    nve1, uptime: 00:10:38, nve
(192.168.0.18/32, 231.1.1.1/32), uptime: 00:02:34, ip mrrib pim
  Incoming interface: Ethernet1/2, RPF nbr: 10.0.128.13
  Outgoing interface list: (count: 1)
    nve1, uptime: 00:02:34, mrrib
(*, 232.0.0.0/8), uptime: 00:17:03, pim ip
  Incoming interface: Null, RPF nbr: 0.0.0.0
  Outgoing interface list: (count: 0)

VTEP-3 Verifications
VTEP-3# show nve vni
Codes: CP - Control Plane      DP - Data Plane
      UC - Unconfigured       SA - Suppress ARP
      SU - Suppress Unknown Unicast

```

```

Interface VNI      Multicast-group  State Mode Type [BD/VRF]      Flags
-----
nve1      160010          231.1.1.1       Up   DP  L2   [10]
nve1      160020          231.1.1.1       Up   DP  L2   [20]
VTEP-3# show vxlan
Vlan          VN-Segment
====          =====
10            160010
20            160020
VTEP-3# ping 10.10.10.1
PING 10.10.10.1 (10.10.10.1) : 56 data bytes
64 bytes from 10.10.10.1: icmp_seq=0 ttl=254 time=7.212 ms
64 bytes from 10.10.10.1: icmp_seq=1 ttl=254 time=6.243 ms
64 bytes from 10.10.10.1: icmp_seq=2 ttl=254 time=5.268 ms
64 bytes from 10.10.10.1: icmp_seq=3 ttl=254 time=6.397 ms
VTEP-1# show nve peers
Interface Peer-IP          State LearnType  Uptime  Router-Mac
-----
nve1      192.168.0.18      Up     DP          00:09:08 n/a
VTEP-3# show ip mroute
IP Multicast Routing Table for VRF "default"
(*, 231.1.1.1/32), uptime: 00:10:38, nve ip pim
  Incoming interface: Ethernet1/1, RPF nbr: 10.0.0.29
  Outgoing interface list: (count: 1)
    nve1, uptime: 00:10:38, nve
(192.168.0.18/32, 231.1.1.1/32), uptime: 00:02:34, ip mrib pim
  Incoming interface: Ethernet1/2, RPF nbr: 10.0.128.13
  Outgoing interface list: (count: 1)
    nve1, uptime: 00:02:34, mrib
(192.168.0.110/32, 231.1.1.1/32), uptime: 00:10:38, nve mrib ip pim
  Incoming interface: loopback1, RPF nbr: 192.168.0.110
  Outgoing interface list: (count: 1)
    Ethernet1/2, uptime: 00:09:39, pim
(*, 232.0.0.0/8), uptime: 00:17:03, pim ip
  Incoming interface: Null, RPF nbr: 0.0.0.0
  Outgoing interface list: (count: 0)

```

Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have a couple of choices for exam preparation: the exercises here, Chapter 20, “Final Preparation,” and the exam simulation questions in the Pearson Test Prep software online.

Review All Key Topics

Review the most important topics in the chapter, noted with the key topic icon in the outer margin of the page. Table 3-13 lists a reference to these key topics and the page numbers on which each is found.



Table 3-13 Key Topics for Chapter 3

Key Topic Element	Description	Page
Figure 3-1	OTV Interfaces and Terms	150
Section	OTV Control Plane Function	151
Figure 3-2	OTV Control Plane	153
Section	OTV Data Plane Function	154
Figure 3-3	OTV Unicast Data Plane	155
Figure 3-6	OTV ARP Optimization	159
Section	Multihoming OTV	159
Figure 3-7	Creation of an End-to-End STP loop	160
Section	VXLAN Encapsulation and Packet Format	173
Figure 3-13	VXLAN Packet Format	174
Figure 3-14	VXLAN Tunnel Endpoint (VTEP)	175
Section	VXLAN Control Plane	176
Figure 3-15	VXLAN Multicast Group in Transport Network	176
Figure 3-16	VXLAN Multicast Control Plane	177

Define Key Terms

Define the following key terms from this chapter, and check your answers in the Glossary.

virtual private network (VPN), any-source multicast (ASM), source-specific multicast (SSM), cyclic redundancy check (CRC), Path MTU Discovery (PMTUD), Network Load Balancing Services (NLBS), broadcast, unknown unicast, and multicast (BUM), Internet Group Management Protocol (IGMP), Protocol Independent Multicast (PIM), Ethernet VPN (EVPN), Spanning Tree Protocol (STP), bridge protocol data units (BPDUs), Media Access Control (MAC), local-area network (LAN), wide-area network (WAN), virtual LAN (VLAN), User Datagram Protocol (UDP), Internet Protocol (IP), virtual port channels (vPCs), Hot Standby Router Protocol (HSRP), virtual routing and forwarding (VRF), virtual device contexts (VDC), equal-cost multipath (ECMP), maximum transmission unit (MTU), Address Resolution Protocol (ARP), Cisco NX-OS, Cisco Nexus

References

Overlay Transport Virtualization (OTV): <https://www.cisco.com/c/en/us/solutions/data-center-virtualization/overlay-transport-virtualization-otv/index.html>

Cisco Overlay Transport Virtualization Technology Introduction and Deployment Considerations: https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI3_OTV_Intro.pdf

Cisco Nexus 7000 Series NX-OS OTV Configuration Guide, Release 8.0(x):
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/otv/config/cisco_nexus7000_otv_config_guide_8x.html

Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide, Release 9.3(x):
<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/vxlan/configuration/guide/b-cisco-nexus-9000-series-nx-os-vxlan-configuration-guide-93x.html>

Cisco Nexus 7000 Series NX-OS VXLAN Configuration Guide, Release 8.0(x):
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/vxlan/config/cisco_nexus7000_vxlan_config_guide_8x.html

Configure VXLAN: <https://www.cisco.com/c/en/us/support/docs/switches/nexus-9000-series-switches/118978-config-vxlan-00.html>

Cisco Live Design and Implementation of DCI Network BRKDCN-2657:
<https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2018/pdf/BRKDCN-2657.pdf>

A Summary of Cisco VXLAN Control Planes: Multicast, Unicast, MP-BGP EVPN: <https://blogs.cisco.com/perspectives/a-summary-of-cisco-vxlan-control-planes-multicast-unicast-mp-bgp-evpn-2>



Index

A

- AAA (authentication, authorization, and accounting), 794–796, 894
 - accounting, 896
 - authentication, 895, 908
 - authentication and authorization user login process, 797
 - authorization, 895–896, 908
 - global commands, 798–800
 - LDAP, 903–907
 - local services, 907–908
 - merging RADIUS and TACACS+ configurations, 910
 - NX-OS configurations, 798–801
 - passphrase and locking user account commands, 800
 - RADIUS, 898–900
 - remote services, 897–898
 - securing Cisco UCS, 865–866
 - server distribution, 909–910
 - server groups, 896
 - server monitoring, 896–897
 - service authentication options, 796–797
 - service configuration options, 896
 - TACACS+, 900–903
 - verification and monitoring commands, 801
- access-group command, 59
- accessing, Guest Shell for Cisco NX-OS, 725–726
- accounting, 896
- ACI (Cisco Application Centric Infrastructure), 196–197. *See also* contracts
 - APIC setup parameters, 206–209
 - APICs, 198–199, 201
 - benefits, 196–198
 - Cisco Nexus 9000 Series spine and leaf switches, 201–204
 - Cold Standby, 206
 - contracts, 231–233, 845–847
 - applying to an EPG, 849–850*
 - configuration parameters, 847–848*
 - creating, 848–849*
 - modifying, 849*
 - removing, 849*
 - design best practices, 245
 - fabric, 201–202
 - fabric discovery, 209–211
 - fabric DLB (dynamic load balancing), 243–244
 - fabric loop detection, 244
 - Fabric Policies Configuration page, 217–220
 - fabric upgrades, 212
 - initial setup, 204–209
 - intersubnet tenant traffic, 241–242
 - labs, 245–261

- management tenant, 237
 - in-band management access*, 237–238
 - out-of-band management access*, 238–239
- microsegmentation, 854–862
- Microsoft SCVMM integration, 223–224
- MIT (Management Information Tree), 213–214
 - AEP*, 216
 - application profiles*, 215, 230–231
 - bridge and domain subnets*, 214–215, 229
 - bridge domains*, 229–230
 - EPGs*, 215, 231
 - microsegmentation*, 215–216, 231
 - tenants*, 214
 - VRF objects*, 214, 228
- MOs (managed objects), 213
- multitier architecture, 202
- Policies Configuration page, 220–221
- policy identification and enforcement, 242
- policy model, 212–213
- tenants, 225–226, 227–228
- traffic storm control, 243
- Virtual Edge, 225
- VM Network page, 222–223
- VMM (Virtual Machine Manager), 222
- VMware vCenter integration, 224–225
- VRF contracts, 850–851
- VXLAN, 239–241
- vzAny, 233
- action statements, 716–717
- active zone set, 392–393
- address-family ipv6 command, 30
- address-family ipv6 unicast command, 14
- addressing
 - FC (Fibre Channel), 354–355
 - reserved FCIDs*, 356f
 - FCoE (Fibre Channel over Ethernet), 431–432
- adjacencies, OSPF, 6
- advertisements, LSAs, 6–8
- AED (authoritative edge device), 160
- AEP (attachable entity profile), 216
- AGs (application gateways), 613–614
- all-flash storage configurations, 696–697
- All-NVMe nodes, 697–698
- allocate fcoe-vlan-range command, 451
- allocate interface ethernet, 451
- Ansible, 747–748
 - authentication, 749
 - and Cisco NX-OS, 750–751
 - CLI prompt, 749
 - CLI tools, 750
 - configuration files, 749
 - inventory files, 748
 - jinja templates, 749
 - modules, 748
 - play, 749
 - playbooks, 749, 750–751
 - roles, 749
 - variable files, 749
 - variables, 749
 - workflows, 748
- Anycast RP, 52
- any-source multicast (ASM), 50
- APICs (Application Policy Infrastructure Controllers), 199–201
 - setup parameters, 206–209

APIs (application programming interface)

convert_to_ucs_python, 766–767

NX-API, 737–739

request and response elements,
739–741

Python CLI, 760–761

Python SDK for Cisco UCS, 765

Rest, 734–735

application profiles, 215, 230–231

architecture, UCS (Cisco Unified Computing System), 516–518

area authentication command, 13

area nssa command, 14

area stub command, 13

areas, OSPF, 9

NSSAs, 10

stub, 9–10

ARP optimization, 158–159

ASBR (autonomous system boundary router), 9

asynchronous mode, BFD, 36

attributes

BGP, 24

VSANs, 372

audit logs, 616

authentication, 749, 895, 908. *See also*

AAA (authentication, authorization, and accounting); keychain authentication; RBAC (role-based access control)

keychain, 884–885

LDAP, 903–907

MD5, 12

multiple services configuration for Cisco UCS, 884

OSPF, 11–12

RADIUS, 898–900

configuring in Cisco UCS,
878–882

RBAC (role-based access control),
801–803

Rest API, 735–736

TACACS+, 900–903

configuring in Cisco UCS,
878–882

two-factor, 869

VRRP, 75

authorization, 895–896, 908

LDAP, 903–907

user roles, 911

policies, 913

rules, 911–913

Auto Install, 662–665

upgrading server firmware, 673–675

**auto mode, FC (Fibre Channel),
354–355**

automation. *See also* Ansible; Cisco UCS Director; DCNM (Data Center Network Manager); POAP (PowerOn Auto Provisioning); PowerShell; Puppet; Python

Cisco Hyperflex, 690–691

AS (autonomous system), 24

Auto-RP, 52

Autozone, 395

B

backup policies, UCS (Cisco Unified Computing System), 648

configuring, 648–650

export fields, 649–650

reminder options, 650

backups, Cisco UCS, 642–643

creating, 643–646

importing, 650–652

running, 646–647

running fields, 647

Bash shell

- for Cisco NX-OS, 722–724
- feature RPMs, managing, 724
- patch RPMs, managing, 724–725

BB credits (buffer-to-buffer credits), 356–357**BDRs (backup designated routers), 11****best practices**

- for firmware updates, 659–661
- for image management, 656–659

BFD (Bidirectional Forwarding Detection), 36

- asynchronous mode, 36
- configuration limitations, 38
- default parameters, 37–38
- establishing neighbor relationships, 36
- feature enabling, 40
- global-level configurations, 38
- interface-level commands, 39
- neighbors, 40
- rapid detection of failures, 37
- routing level commands, 39
- verification commands, 39

bfd echo command, 39**bfd optimize subinterface command, 39****bfd slow-timer command, 38****BGP (Border Gateway Protocol), 23–24. *See also* MBGP (Multiprotocol BGP)**

- attributes, 24
- AS (autonomous system), 24
- configuration limitations, 29
- global-level configurations, 30
- hold time, 24
- keepalives, 24
- load balancing, 24
- path selection, 25

comparing pairs of paths, 25–27

determining the best-path change suppression, 27–28

determining the order of comparisons, 27

- peering, 24–25
- router configurations, 32, 35–36
- routing-level configurations, 30
- verification, 33–35
- verification and clear commands, 31–32

Bidir (bidirectional shared trees), 50**bidirectional shared trees, 46–48****bind interface ethernet command, 452****blade servers, UCS (Cisco Unified Computing System), 520–521****Blocking state, 162****boot sequence**

- Cisco MDS NX-OS, 481–482
- Cisco NX-OS, 280

BPDU Filter, 156**BPDU Guard, 156****Bridge Assurance, 154–156****bridge domains, 229–230****bridge-domain command, 182****broadcast networks, 11****broadcast policy control, 159****broadcast traffic over OTV, 156****BSRs (bootstrap routers), 52****C****Call Home, 504, 619–620. *See also* Smart Call Home****CFS (Cisco Fabric Services), 364–365**

- commands, 369–370
- fabric lock, 366–367
- features, 365–366

- merge, 368

- regions, 369

CFSofC, 367–368

CFSofIP, 367–368

chapter review, 932

**characteristics, of cloud computing,
267–268**

checkpoints, creating, 304–305

Cisco AVPair attribute, 868

Cisco HX Data Platform, 698–699

- enabling HX logical availability zones,
708–709

- expanding system clusters, 707–708

- high availability, 700–701

- HX cluster interfaces, 702

- HX server disk types, 705–706

- managing HX disks in the cluster,
703–706, 707

- Native Snapshots, 701–702

- ready clones, 701

- SEDs (self-encrypting drives), 702–703

- tolerated failures, 701

Cisco Hyperflex, 686–687, 688, 694

- for all-flash storage configurations,
696–697

- All-NVMe nodes, 697–698

- benefits, 689–690

- Cisco HX Data Platform, 698–699

- enabling HX logical availability
zones, 708–709*

- expanding system clusters,
707–708*

- high availability, 700–701*

- HX server disk types, 705–706*

- managing HX datastores,
706–707*

- managing HX disks in the cluster,
703–706*

- ready clones, 701*

- tolerated failures, 701*

- compute nodes, 699

- controller, 688

- converged nodes, 699

- datastores, 699

- drives, 699

- as and edge device, 694–696

- HX cluster interfaces, 702

- for hybrid storage configurations,
696–697

- for hyperconverged systems, 693–694,
696

- independent resource scaling, 692–693

- intelligent end-to-end automation,
690–691

- replication factor, 687

- SEDs (self-encrypting drives), 702–703

- storage clusters, 699

- unified management for all workloads,
691–692

- VICs (virtual interface cards), 690–692

- and VMware vSphere, 688–689

Cisco Hyperflex Anywhere, 695

Cisco Intersight, 629–630

- adding UCS, 633–636

- benefits, 630–632

- licensing, 632–633

- supported software, 632

- as telemetry data collection, 632

Cisco MDS 9000 Series switches.

See also Cisco MDS NX-OS

- Call Home, 504

- Cisco NX-OS

- in-band management*

- logical interface configuration,
479–480*

- setup utility, 472–478*

- displaying EPLD versions, 499–501
- EEM (Embedded Event Manager), 505
- fabric binding, 922–924
- LDAP configuration, 905–907
- port security, 915–917
 - configuring*, 917–920
- RBAC, 910
- RBAC configuration, 914–915
- RMON, 505
- RSPAN, 509–511
- SPAN, 505–509
- updating module EPLDs, 502–503
- upgrading EPLDs, 498–503
- verification of port security, 920–922
- verifying Cisco NX-OS version, 478–479
- Cisco MDS NX-OS**
 - boot sequence, 481–482
 - disruptive downgrades on Cisco MDS switch, 495–497
 - disruptive upgrades on Cisco MDS fabric switch, 487–489
 - nondisruptive downgrades on Cisco MDS fabric switch, 490–494
 - nondisruptive upgrades on Cisco MDS fabric switch, 482–487
 - software upgrade and downgrade, 480–481
 - system messages, 503–504
- Cisco Network Assurance Engine, 341–344**
- Cisco Nexus 5000 Series switches, FCoE configuration, 442–443**
- Cisco Nexus 7000 Series switches, FCoE configuration, 440–441**
- Cisco Nexus 9000 Series spine and leaf switches, 201–204**
- Cisco NSS3000 Series Network Storage System, 465–467**
 - hard drive and RAID configurations, 467
- Cisco NX-OS. *See also* Cisco MDS NX-OS**
 - AAA configurations, 798–801
 - and Ansible, 750–751
 - in-band management logical interface configuration, 479–480
 - Bash shell, 722–724
 - BGP support, 24
 - boot sequence, 280
 - channel modes, 118
 - Cisco Network Assurance Engine, 341–344
 - configuration save and backup, 303
 - console management connection, 279–280
 - control plane packets, 833–834
 - CoPP (Control Plane Policing), 831–833
 - classification*, 834–835
 - configuring*, 838–845
 - modular QoS command-line interface*, 836–838
 - DAI configurations, 813–814
 - data center infrastructure software lifecycle management, 287
 - DHCP snooping, 821
 - configuring*, 823–825
 - Option 82 data insertion*, 823
 - packet validation*, 822
 - trusted and untrusted sources*, 821–822
 - EPLDs (electrical programmable logical devices), 289–291
 - feature licenses, 13, 28–29, 56, 78, 164
 - GIR (Graceful Insertion and Removal), 291–295

- Guest Shell, 725
 - accessing*, 725–726
 - capabilities*, 726–728
 - managing*, 728–730
 - resources used*, 726
- images, 278
- keychain authentication, 884–885
- keychain configurations, 885–887
- MBGP support, 28
- NetFlow, 325–330
- network monitoring, 306
- Nexus config rollback and checkpoint, 303–305
- nondisruptive ISSU, 295–299
- NTP (Network Time Protocol),
 - displaying status, 312–313
- POAP (PowerOn Auto Provisioning), 283–287
- port security, 826–828
 - configuring*, 829–831
 - and port types*, 829
 - violations and actions*, 828
- and Puppet, 753–755
- Python, 758–759
 - CLI command APIs*, 760–761
 - compound commands*, 764
 - displaying information about*, 759
 - in interactive mode*, 761–762
 - noninteractive mode*, 762–764
 - package functions*, 760
- setup utility, 280–283
 - for Cisco MDS 9000 Series switches*, 472–478
- Smart Call Home, 324–325
- SMU (Software Maintenance Upgrade), 287–289
- SNMP (Simple Network Management Protocol), 317–319
 - configuration examples*, 323–324
 - global commands*, 320
 - MIBs*, 320–322
 - security models and levels*, 319
 - specific notification commands*, 322
 - verification commands*, 323
- SPAN (Switched Port Analyzer), 330–337
- streaming telemetry, 337–341
- supported authentication methods, 11–12
- supported LSAs, 7–8
- system messages, 306–307
- time management
 - NTP*, 307–313
 - PTP (Precision Time Protocol)*, 313–317
- user roles, 803–804
- verifying version on Cisco MDS 9000 Series switches, 478–479
- Cisco UCS 2304 IOM, 527–528
- Cisco UCS 5108 Blade Server Chassis, 520
- Cisco UCS 6300 Series Fabric Interconnects, 526–527
- Cisco UCS 6454 Fabric Interconnect, 524–526
- Cisco UCS Director, 782–783
 - automation and orchestration, 783–784
 - features and benefits, 784–785
 - PowerShell agent commands, 788–789
 - PowerShell agent installation, 787–788
 - system setup, 785–786
- Cisco UCS Mini solution, 523

- clear bgp all command, 31
- clear bgp all dampening command, 31
- clear bgp all flap-statistics command, 31
- clear ip ospf command, 15
- clear ospfv3 command, 15
- CLI prompt, Ansible, 749
- CLI tools, Ansible, 750
- clocks, PTP (Precision Time Protocol), 314
- cloud computing, 266–267
 - benefits, 267
 - characteristics, 267–268
 - deployment models
 - community cloud*, 274
 - hybrid cloud*, 273–274
 - private cloud*, 272
 - public cloud*, 272–273
 - IaaS (Infrastructure as a Service), 270–272
 - PaaS (Platform as a Service), 270
 - SaaS (Software as a Service), 269
- CM (Configuration Management), 303
- Cold Standby, ACI, 206
- collapsed-core topology, 350
- commands. *See also* show commands
 - AAA (authentication, authorization, and accounting), 800–801
 - access-group, 59
 - address-family ipv6, 30
 - address-family ipv6 unicast, 14
 - allocate fcoe-vlan-range, 451
 - area authentication, 13
 - area nssa, 14
 - area stub, 13
 - Bash, 722, 724–725
 - bfd echo, 39
 - bfd optimize subinterface, 39
 - bfd slow-timer, 38
 - bind interface ethernet, 452
 - bridge-domain, 182
 - CFS (Cisco Fabric Services), 369–370
 - Cisco UCS Director PowerShell agent, 788–789
 - clear bgp all, 31
 - clear bgp all dampening, 31
 - clear bgp all flap-statistics, 31
 - clear ip ospf, 15
 - clear ospfv3, 15
 - compound, 764
 - config terminal, 451
 - CoPP (Control Plane Policing), 840–844
 - DAI (Dynamic ARP Inspection), 814–815
 - description, 30, 166
 - device alias, 407–408
 - DHCP snooping, 824–825
 - dohost, 727
 - dot1q vni, 182
 - dr-priority, 58
 - encapsulation profile default, 183
 - fcoe fcf-priority, 451
 - fcoe fka-adv-period, 451
 - fcoe vloopback, 451
 - fcoe vsan, 452
 - feature bfd, 38
 - feature bgp, 30
 - feature fcoe, 451
 - feature hsrp, 79
 - feature lacp, 451
 - feature lldp, 451
 - feature nv-segment-vlan-based, 182
 - feature ospf, 13
 - feature ospfv3, 13
 - feature otv, 165

- feature pim, 57
- feature pim6, 57
- feature udld, 168
- feature vrrp, 79
- fex, 452
- fex associate, 453
- guestshell enable, 729
- hello-authentication ah-md5, 58
- hello-interval, 58
- hsrp, 79
- immediate-leave, 59
- ingress-replication protocol bgp, 184
- ingress-replication protocol static, 184
- install feature set fcoe, 451
- interface nve, 183
- interface overlay, 165
- interface vfc, 452
- ip igmp snooping vxlan, 183
- ip igmp ssm-translate, 58
- ip ospf, 14
- ip ospf bfd, 39
- ip ospf dead-interval, 14
- ip ospf hello-interval, 14
- ip ospf mtu-ignore, 14
- ip ospf passive-interface, 14
- ip ospf priority, 14
- ip ospf shutdown, 14
- ip pim auto-rp, 57
- ip pim bidir-rp-limit, 57
- ip pim bsr, 57
- ip pim spt-threshold infinity group-list, 57
- join-group, 58
- keychain, 885–887
- license fcoe, 451
- mcast-group, 184
- member nvi, 184
- name, 79
- neighbor remote-as, 30
- neighbor remote-as route-map, 30
- neighbor-policy prefix-list, 58
- network, 30
- NPV (N port virtualization), 411–412
- NTP (Network Time Protocol), 309–311
- ospf network, 14
- otv adjacency-server unicast-only, 166
- otv control-group, 166
- otv data-group, 167
- otv extend-vlan, 166
- otv isis csnp-interval, 167
- otv isis hello-interval, 167
- otv isis hello-multiplier, 167
- otv isis metric, 167
- otv isis priority, 167
- otv join-interface, 166
- otv site-identifier, 166
- otv site-vlan, 166, 167
- otv-isis, 165
- peer-ip, 184
- pim jp-policy, 59
- port security, 830–831
- preempt, 79
- PTP (Precision Time Protocol), 315–316
- query-interval, 59
- query-timeout, 59
- RBAC (role-based access control), 806–807
- register-rate-limit, 57
- router bgp, 30
- router ospf, 13
- router ospf bfd, 39
- router ospfv3, 13
- router-id, 13, 30

- routing multicast holddown, 58
- rp-address, 57
- run bash sudo su, 724
- SAN port channel, 385–386
- service instance, 183
- service-policy type network-qos
 - default-nq-7e-policy, 451
- show bfd neighbors, 39
- show bgp, 31
- show bgp all, 31
- show bgp community, 31
- show bgp convergence, 31
- show bgp sessions, 31
- show bgp statistics, 31
- show diff rollback-patch, 304
- show event manager system-policy, 716
- show feature, 453
- show feature-set, 453
- show file, 763
- show forwarding distribution otv multicast route vlan, 168
- show groups, 60
- show guestshell detail, 729–730
- show interface, 60
- show ip ospf, 14
- show ip ospf interface, 14
- show ip ospf route, 14
- show ip ospf statistics, 15
- show ip ospf traffic, 15
- show ip ospf virtual-links, 15
- show local-groups, 60
- show mac address-table, 168
- show mroute ip-address, 59
- show ospfv3, 15
- show ospfv3 interface, 15
- show otv, 167, 169–172
- show otv adjacency, 167
- show otv isis site, 167
- show otv mroute vlan startup, 167
- show otv overlay, 167
- show otv site, 167
- show otv vlan-mapping, 168
- show pim group-range, 59
- show pim rp, 59
- show port channel compatibility-parameters, 119
- show port-channel summary, 144–145
- show route, 60
- show running configuration, 15, 60
- show running configuration bgp, 31
- show running configuration otv, 167
- show running configuration pim, 60
- show spanning-tree summary, 113–115
- show vpc, 140–143
- shutdown, 30, 452
- shutdown lan, 452
- SNMP (Simple Network Management Protocol), 320, 322–323
- source-interface, 184
- spanning-tree bpdupfilter, 171
- spanning-tree bpduguard, 170
- spanning-tree cost, 171
- spanning-tree guard loop, 171
- spanning-tree guard root, 171
- spanning-tree loopguard default, 169
- spanning-tree mode rapid-pvst, 169
- spanning-tree pathcost method, 169
- spanning-tree port type edge, 170
- spanning-tree port type edge bpdu-guard default, 168
- spanning-tree port type edge default, 168
- spanning-tree port type edge trunk, 452
- spanning-tree port type network, 170

- spanning-tree port type network
 - default, 168
- spanning-tree port type normal, 170
- spanning-tree port-priority, 171
- spanning-tree vlan, 169
- spanning-tree vlan priority, 169
- spanning-tree vlan root primary, 169
- spanning-tree vlan root secondary, 169
- sparse-mode, 58
- startup-query-interval, 59
- sudo yum installed, 724
- switchport mode fex-fabric, 453
- switchport mode trunk, 452
- switchto vdc, 451
- system qos, 451
- timers, 30, 80
- track interface, 79
- udld, 171
- udld aggressive, 170
- udld message-time, 169
- vdc type storage, 451
- version value, 58
- vlan, 182
- vn-segment, 182
- VSAN (virtual storage-area network), 376–378
- vsan database, 452
- yum list available, 724
- zone mode enhanced vsan, 397
- zoning, 399–400
- community cloud deployments, 274**
- compatibility checks**
 - SAN port channels, 382–383
 - vPC, 127–128
- compound commands, 764**
- compute nodes, 699**
- config terminal command, 451**
- configuration files, Ansible, 749**
- configuring**
 - CoPP (Control Plane Policing), 838–845
 - DHCP snooping, 823–825
 - EEM (Embedded Event Manager), 717–718
 - fabric binding on MDS switches, 922–924
 - LDAP group mapping, 875–878
 - local encryption key, 703
 - NTP (Network Time Protocol), 312
 - port security, 829–831
 - port security on Cisco MDS 9000 Series switches, 917–920
 - RADIUS, 878–882
 - RBAC on MDS switches, 914–915
 - SANs (storage area networks), 596–597
 - Scheduler, 719–721
 - SPAN (Switched Port Analyzer), 505–509
 - switches with POAP, 772
 - TACACS+, 878–882
 - UCS backup policies, 648–650
- consistency checks, vPC, 127–128**
- console management connection, Cisco NX-OS, 279–280**
- contracts**
 - ACI, 231–233, 845–847
 - applying to an EPG, 849–850*
 - configuration parameters, 847–848*
 - creating, 848–849*
 - modifying, 849*
 - removing, 849*
 - exporting between private networks, 853
 - inter-tenant, 851–852

- multiple contracts unidirectional single filter, 854
 - single contract bidirectional reverse filter, 853
 - single contract unidirectional with multiple filters, 854
 - taboo, 233
 - VRF, 850–851
 - vzAny, 233
 - control plane**
 - OTV (overlay transport virtualization), 151
 - multicast-enabled transport infrastructure, 151–152*
 - unicast-only transport infrastructure, 152–153*
 - packets, 833–834
 - VXLAN (Virtual eXtensible LAN), 176–178
 - converged networks, 6**
 - converged nodes, 699**
 - convert_to_ucs_python API, 766–767**
 - CoPP (Control Plane Policing), 831–833**
 - classification, 834–835
 - commands, 840–844
 - configuring, 838–845
 - modular QoS command-line interface, 836–838
 - core-edge topology, 350–351**
 - core-edge-core topology, 351–353**
 - creating**
 - ACI contracts, 848–849
 - device aliases, 408–409
 - NPV (N port virtualization), 412–415
 - SAN port channels, 387–388
 - user roles, 808–809
 - VSANs (virtual storage-area networks), 379–380
 - zone profiles, 600–602
 - zones, 401–403
 - customizing your exams, 930–931**
-
- ## D
- DAI (Dynamic ARP Inspection), 810–813**
 - commands, 814–815
 - enabling, 816–821
 - network topology, 815–816
 - NX-OS configurations, 813–814
 - data fields, XML, 731**
 - data plane, OTV (overlay transport virtualization), 154**
 - broadcast traffic over OTV, 156
 - multicast traffic over OTV, 156
 - unicast traffic over OTV, 154–156
 - datastores**
 - adding in Cisco HX Data Platform, 706–707
 - mounting, 707
 - DCBX (Data Center Bridging Exchange), 424–426**
 - DCNM (Data Center Network Manager), 772–773**
 - benefits, 774–776
 - features for automation and REST APIs, 776–777
 - features for operations, 779
 - monitoring, visibility, and troubleshooting features, 778–779
 - web user interface, 779–782
 - dead-interval, OSPF, 5**
 - default parameters**
 - BFD, 37–38
 - HSRP, 77
 - IGMP, 54–55
 - MLD, 54–55

- OSPF, 12
- PIM, 55
- Rapid PVST+, 165–167
- VRRP, 78
- description command, 30, 166
- designated routers/forwarders, PIM (Protocol Independent Multicast), 53
- Developer Sandbox, NX-API, 741–742
- device aliases, 403
 - commands, 407–408
 - creation and verification, 408–409
 - distribution, 405–406
 - features, 403–404
 - modes, 404–405
 - versus zone aliases, 406–407
- device element, XML, 731
- DHCP discovery phase, POAP (PowerOn Auto Provisioning), 770–771
- DHCP snooping, 821
 - commands, 824–825
 - configuring, 823–825
 - Option 82 data insertion, 823
 - packet validation, 822
 - trusted and untrusted sources, 821–822
- direct firmware upgrades, 666–671
- direct updates, 659
- direct-attached topology (FCoE), 436–437
- Disabled state, 162
- disaster recovery, system restore on Cisco UCS, 652
- displaying
 - NTP status, 312–313
 - Python package information for Cisco NX-OS, 759
 - Scheduler jobs, 721–722
- disruptive upgrade/downgrade
 - on Cisco MDS fabric switch, 487–489
 - on Cisco MDS switch, 297–299, 495–497
 - fabric interconnect configuration, restoring on Cisco UCS, 653–654
 - on Nexus switches, 299–300
- distribution trees, 48
- DLB (dynamic load balancing), 243–244
- DME (data management engine), 612–613
- documents, XML (Extensible Markup Language), 732–733
- dohost command, 727
- domain ID distribution, FC (Fibre Channel), 360–361
- domain-level commands, vPC, 135
- dot1q vni command, 182
- downgrading, Cisco MDS NX-OS, 480–481
 - disruptive downgrades on Cisco MDS switch, 495–497
 - nondisruptive downgrades on Cisco MDS fabric switch, 490–494
- downlink connectivity, UCS (Cisco Unified Computing System), 546
- downloading SMU for Cisco NX-OS, 288–289
- DPVM (Dynamic Port VSAN Membership), 373
- drag-and-drop exam questions, 933
- dr-priority command, 58
- DRs (designated routers), 11
- dual-control plane, vPC (virtual port channel), 125
- dual-wire manager, UCS (Cisco Unified Computing System), 532–533
- duplicate frames prevention, vPC, 128–129
- dynamic router discovery, 72

E

E ports, 353–354

ECMP (equal-cost multipath), 24

edge devices, Hyperflex Edge, 694–696

edge ports, STP, 153

EEM (Embedded Event Manager), 505, 715

- action statements, 716–717
- configuring, 717–718
- event statements, 716
- policies, 715–716
- verifying the configuration, 718

encapsulation profile default command, 183

endpoints, 196–197

- UCS (Cisco Unified Computing System), 546

end-to-end automation, Cisco Hyperflex, 690–691

enhanced zoning, 397–398

ENodes (Ethernet Nodes), 429–431

EPGs (endpoint groups), 215, 231

- filters and subjects, 236–237

EPLDs (electrical programmable logical devices)

- for Cisco NX-OS, 289–291
- displaying current version, 498–499
- upgrading on Cisco MDS 9000 Series switches, 498–503

error codes, NX-API, 741

ERSPAN (Encapsulated Remote Switched Port Analyzer), 333–337

Ethernet

- traffic monitoring, 623–624, 625–626
- vNICs, 549–550

ETS (Enhanced Transmission Selection), 423–424

event statements, 716

expanding system clusters in Cisco HX Data Platform, 707–708

exporting contracts between private networks, 853

extensions, STP, 154

F

fabric binding, 922

- configuring on MDS switches, 922–924
- versus port security, 924–925

fabric discovery, ACI, 209–211

fabric infrastructure, UCS (Cisco Unified Computing System), 524

fabric interconnect configuration, restoring on Cisco UCS, 653–654

Fabric Policies Configuration page, ACI, 217–220

fabric reconfiguration, FC (Fibre Channel), 362

failure isolation, OTV (overlay transport virtualization), 157

FC (Fibre Channel), 349

- addressing, 355–356
 - reserved FCIDs*, 356
- CFS (Cisco Fabric Services), 364–365
 - fabric lock*, 366–367
 - features*, 365–366
 - merge*, 368
 - regions*, 369
- CFSofC, 367–368
- CFSofIP, 367–368
- device aliases, 403
 - commands*, 407–408
 - creation and verification*, 408–409
 - distribution*, 405–406

- features*, 403–404
- modes*, 404–405
- versus zone aliases*, 406–407
- domain ID distribution, 360–361
- fabric reconfiguration, 362
- FCID allocation, 362
- FLOGI (fabric login), 362
 - and FCNS databases*, 363–364
- flow control, 356–357
- NPIV (N port identifier virtualization), 410–411
- NPV (N port virtualization), 410–411
 - commands*, 411–412
 - creation and verification*, 412–415
- PLOGI (port login), 362
- port types, 353
 - auto mode*, 354–355
 - E*, 353–354
 - Fx*, 354
 - NP*, 354
 - TE*, 354
 - TF*, 354
 - TNP*, 354
- principal switch selection, 358–360
- PRLI (process login), 362
- SAN port channels, 381, 382
 - channel group configuration differences*, 385
 - commands*, 385–386
 - compatibility checks*, 382–383
 - creating*, 387–388
 - load balancing*, 383–384
 - modes*, 384
 - trunking*, 381–382
 - types of*, 381
- switched fabric initialization, 358
- topologies, 350
 - collapsed-core*, 350
 - core-edge*, 350–351
 - core-edge-core*, 351–353
- traffic monitoring, 624–625, 626–627
- VSANs (virtual storage-area networks), 370
 - advantages*, 373
 - attributes*, 372
 - commands*, 376–378
 - creating*, 379–380
 - DPVM (Dynamic Port VSAN Membership)*, 373
 - features*, 370–372
 - show commands*, 377–378
 - trunking*, 374–376
- zoning, 389
 - active zone set*, 392–393
 - Autozone*, 395
 - commands*, 399–400
 - creation and verification*, 401–403
 - enforcement*, 391–392
 - enhanced*, 397–398
 - fabric with two zones*, 390–391
 - features*, 389–390
 - full zone set*, 393–394
 - smart zoning*, 396–397
 - and VSANs*, 391
 - zone merge*, 395–396
- FCFs (Fibre Channel Forwarders)**, 429–430
- FCID allocation**, 362
- FCoE (Fibre Channel over Ethernet)**, 420–422
 - addressing, 431–432
 - benefits, 435
 - on Cisco Nexus 5000 Series switches, 442–443

- on Cisco Nexus 7000 Series switches, 440–441
- configuration and verification, 454–457
- configuration and verification commands, 451–453
- DCBX (Data Center Bridging Exchange), 424–426
- ENodes (Ethernet Nodes), 429–431
- Ethernet enhancements, 422
- ETS (Enhanced Transmission Selection), 423–424
- FCFs (Fibre Channel Forwarders), 429–430
- FIP (FCoE Initialization Protocol), 432–435
- forwarding, 431–432
- frame format, 426–428
- LEPs (link endpoints), 429–430
- miscellaneous configuration, 441–442
- multi-hop topology, 438–439
- NPV, 445–447
- over FEX, 444
- PFC (priority-based flow control), 422–423
- show commands, 453
- single-hop topology, 435–436
 - direct-attached*, 436–437
 - FEX*, 437–438
 - remote-attached*, 438
- verification, 448–451
- vFC (Virtual Fibre Channel), 428–429
- fcoe fcf-priority command**, 451
- fcoe fka-adv-period command**, 451
- fcoe veloopback command**, 451
- fcoe vsan command**, 452
- feature bfd command**, 38
- feature bgp command**, 30
- feature fcoe command**, 451
- feature hsrp command**, 79
- feature lacp command**, 451
- feature licenses**, Cisco NX-OS, 13, 28–29, 56, 78
- feature lldp command**, 451
- feature nv overlay command**, 182
- feature nv-segment-vlan-based command**, 182
- feature ospf command**, 13
- feature ospfv3 command**, 13
- feature otv command**, 165
- feature pim command**, 57
- feature pim6 command**, 57
- feature RPMs**, managing with Bash, 724
- feature udd command**, 168
- feature vrrp command**, 79
- FEs (fabric extenders)**, 527
- FEX**
 - FCoE configuration, 444
 - topology, 437–438
- fex associate command**, 453
- fex command**, 452
- FHRP (First Hop Redundancy Protocol)**, 68
- FHRP isolation**, 162–163
- FHS (First-Hop Security)**, 809–810
- fill-in-the-blank exam questions**, 933
- filters**, 236–237
- final preparation**
 - chapter review, 932
 - customizing your exams, 930–931
 - getting ready, 928–929
 - learning the question types, 932–935
 - suggested plan, 936
 - tools, 929–930
 - updating your exams, 931

FIP (FCoE Initialization Protocol), 432–435

firmware

Auto Install, 662–665, 673–675

best practices, 659–661

direct updates, 659

direct upgrades, 666–671

HUU (Cisco Host Upgrade Utility), 675–681

image management best practices, 656–659

infrastructure firmware fields, 671–672

T bundle, 655

updating on Cisco UCS, 654–661

version terminology, 661–662

first hop redundancy, IPv6, 76

FIs (fabric interconnects), 527

FLOGI (fabric login), 362

and FCNS databases, 363–364

flow control, FC (Fibre Channel), 356–357

flow monitors, 622

flow record definitions, 621

flows, 621

forwarding, FCoE (Fibre Channel over Ethernet), 431–432

Forwarding state, 162

frame format, FCoE (Fibre Channel over Ethernet), 426–428

full zone set, 393–394

Fx ports, 354

G

gateways, VXLAN, 178–179

GIR (Graceful Insertion and Removal), 291–295

global-level commands

OSPF, 13

OTV, 165, 167–168

STP, 167–170

vPC, 133–134

VXLAN, 182–183

global-level verification and process clear commands, OSPF, 14–15

global-level verification commands

port channel, 136

STP, 171–172

VXLAN, 184

graceful restart, OSPF, 6

group mapping, LDAP, 875–878

group pacing, 8

group-and-source specific query, IGMP (Internet Group Management Protocol), 43–44

groups, 41

VRRP (Virtual Router Redundancy Protocol), 74

Guest Shell for Cisco NX-OS, 725

accessing, 725–726

capabilities, 726–728

managing, 728–730

resources used, 726

guestshell enable command, 729

H

hard zoning, 392

HCI (Hyperconverged Infrastructure), 698

hello messages, PIM, 50

hello packets

HSRP, 70

OSPF, 5

hello-authentication ah-md5 command, 58

hello-interval command, 58

high availability

Cisco HX Data Platform, 700–701
VXLAN (Virtual eXtensible LAN),
179–180

hold time

**HSRP (Hot Standby Router Protocol),
68–69, 70–71**

configuration limitations, 78–79
default parameters, 77
full configurations, 82–83
global-level verification commands, 81
hello packets, 70
load sharing, 71–72
verification, 84–86
versions, 69–70
and vPC, 130

hsrp command, 79

**HUU (Cisco Host Upgrade Utility),
675–681**

HX cluster interfaces, 702

hybrid cloud deployments, 273–274

**hybrid storage configurations,
696–697**

**IaaS (Infrastructure as a Service),
270–272**

**IANA (Internet Assigned Numbers
Authority), 41**

identity pools, 571–572

IP pools, 574–576
MAC pools, 573–574
server pools, 576–577
universally unique identifier suffix
pools, 572–573

**IGMP (Internet Group Management
Protocol), 41**

configuration and verifications, 61–63
default parameters, 54–55
group-and-source specific query,
43–44
last member query response interval,
44
query-response process, 42–43
snooping, 44
TTL (time-to-live), 44

**image management best practices,
656–659**

immediate-leave command, 59

**importing backup data on Cisco UCS,
650–652**

**infrastructure firmware fields,
671–672**

**ingress-replication protocol bgp com-
mand, 184**

**ingress-replication protocol static com-
mand, 184**

initial setup

ACI, 204–209
UCS (Cisco Unified Computing
System), 536–540

install feature set fcoe command, 451

installing

Cisco UCS Director PowerShell agent,
787–788
Puppet, 754–755

intent-based networking, 341

interactive mode, Python, 761–762

interface nve command, 183

interface overlay command, 165

interface policies, 220

interface vfc command, 452

interface-level commands

- OSPF, 14
- port channel, 134
- PTP (Precision Time Protocol), 315–316
- STP, 170–171
- VRRP, 79–80
- VXLAN, 183
- internal interfaces, OTV, 151
- internal peering, OTV, 160–162
- Intersight. *See* Cisco Intersight
- intersubnet tenant traffic, 241–242
- inter-tenant contracts, 851–852
- inventory files, Ansible, 748
- I/O consolidation, 418, 420
- ip igmp snooping vxlan command, 183
- ip igmp ssm-translate command, 58
- ip ospf bfd command, 39
- ip ospf command, 14
- ip ospf dead-interval command, 14
- ip ospf hello-interval command, 14
- ip ospf mtu-ignore command, 14
- ip ospf passive-interface command, 14
- ip ospf priority command, 14
- ip ospf shutdown command, 14
- ip pim auto-rp command, 57
- ip pim bidir-rp-limit command, 57
- ip pim bsr command, 57
- ip pim spt-threshold infinity group-list command, 57
- IP pools, 574–576
- IPv6, first hop redundancy, 76
- ISSU (in-service software upgrade)
 - for Cisco NX-OS, 295–298–299
 - for Nexus 5000 switches, 295–296
 - for Nexus 7000 switches, 296–297
 - for Nexus 9000 switches, 297–299

J

- jinja templates, 749
- jobs
 - displaying, 721–722
 - scheduling, 718–719, 720–721
- join interface, OTV, 151
- join-group command, 58
- join-prune messages, PIM, 50–51
- JSON (JavaScript Object Notation), 733–734

K

- keepalives, 24
- keychain authentication, 884–885
 - Cisco NX-OS, 885–887
 - commands, 885–887
 - key selection, 888–889

L

- labs, ACI, 245–261
- last member query response interval, 44
- LAZ (logical availability zones), enabling, 708–709
- LDAP, 869–878, 903–907
 - creating providers, 870–874
 - deleting existing groups in Cisco UCS, 875
 - group mapping, 875–878
 - limitations, 904
 - on MDS switches, 905–907
 - modifying existing groups in Cisco UCS, 874–875
 - provider configurations, 870
 - provider fields, 871–872, 874–875

leaf switches, 201–204
 Learning state, 162
 LEPs (link endpoints), 429–430
 license fcoe command, 451
 licensing, Cisco Intersight, 632–633
 limitations
 of LDAP, 904
 of POAP, 767
 load balancing
 BGP, 24
 dynamic, 243–244
 port channels, 120–121
 SAN port channels, 383–384
 local AAA services, 907–908
 local encryption key, configuring, 703
 locally assigned VNIs, 175
 logging messages, Cisco NX-OS, 306–307
 Loop Guard, 156–157
 loops, 164
 LSA group pacing, 8
 LSAs (link-state advertisements), 6–8
 Cisco NX-OS support, 7–8
 opaque, 6

M

MAC pools, 573–574
 MAC routing, 150
 management tenant, 237
 in-band management access, 237–238
 out-of-band management access, 238–239
 managing Guest Shell for Cisco NX-OS, 728–730
 manifest, Puppet, 752, 756–758
 MBGP (Multiprotocol BGP), 28
 default BGP parameters

mcast-group command, 184
 MCEC (multichannel EtherChannel), 121
 MD5 authentication, 12
 MDT (multicast distribution tree), 45–46
 bidirectional shared trees, 46–48
 shared trees, 46
 member nvi command, 184
 merging RADIUS and TACACS+ configurations, 910
 messages
 PIM (Protocol Independent Multicast), 50–51
 syslog, 614–615
 MIBs (Management Information Bases), 320–322
 microsegmentation, 215–216, 231, 854–862
 Microsoft SCVMM, Cisco ACI integration, 223–224
 MIT (Management Information Tree), 213–214
 AEP, 216
 application profiles, 215, 230–231
 bridge and domain subnets, 214–215, 229
 bridge domains, 229–230
 EPGs, 215, 231
 microsegmentation, 215–216, 231
 tenants, 214, 225–226, 227–228
 VRF objects, 214, 228
 MLD (Multicast Listener Discovery), 44–45, 49
 default parameters, 54–55
 model-driven framework, 212
 modifying ACI contracts, 849
 modules, Ansible, 748
 MOs (managed objects), 213

mounting a datastore, 707

MSDP (Multicast Source Discovery Protocol), 49

Multicast IP, 41

groups, 41

IGMP (Internet Group Management Protocol), 41

configuration and verifications,
61–63

default parameters, 54–55

group-and-source specific query,
43–44

query-response process, 42–43

snooping, 44

TTL (time-to-live), 44

MDT (multicast distribution tree),
45–46

bidirectional shared trees, 46–48

shared trees, 46

MLD (Multicast Listener Discovery),
44–45

default parameters, 54–55

PIM (Protocol Independent Multicast),
48–49, 63–64

any-source multicast (ASM), 50

Bidir (bidirectional shared trees),
50

configuration limitations, 56

default parameters, 55

designated routers/forwarders,
53

distribution trees, 48

messages, 50–51

*multicast anycast RP configura-
tions and verification,* 67

*multicast auto RP configurations
and verification,* 66

*multicast BSRs RP configura-
tions and verification,* 65–66

multicast distribution modes, 57

*multicast global-level BGP verifi-
cation commands,* 59–60

multicast global-level commands,
57–58

*multicast interface-level
commands,* 58–59

*multicast static RP configura-
tions and verification,* 64–65

RPs, 52–53

source distribution trees, 48

SSM (source-specific multicast),
50

state refresh, 51–52

RPF (Reverse Path Forwarding), 53–54

multicast traffic over OTV, 156

multicast-enabled transport infrastruc-
ture, 151–152

multihoming OTV, 159–160

multi-hop topology, FCoE (Fibre Chan-
nel over Ethernet), 438–439

multiple authentication services
configuration, UCS (Cisco Unified
Computing System), 884

multiple contracts unidirectional single
filter, 854

multiple-choice exam questions, 932

multitier architecture, 202

N

name command, 79

named VSANs, 597–599

NAS (network-attached storage),
463–465

benefits, 465

Cisco NSS3000 Series Network
Storage System, 465–467

Native Snapshots, 701–702

neighbor relationships, BFD, 36

- neighbor remote-as command, 30
 - neighbor remote-as route-map command, 30
 - neighbor-policy prefix-list command, 58
 - NetFlow monitoring
 - for Cisco NX-OS, 325–330
 - for Cisco UCS, 621–622
 - network command, 30
 - network monitoring, Cisco NX-OS, 306
 - network ports, STP, 154
 - network requirements, POAP (PowerOn Auto Provisioning), 767–768
 - networks
 - broadcast, 11
 - point-to-point, 11
 - network-wide scoped VNIs, 175
 - Nexus 5000 switches, ISSU (in-service software upgrade), 295–296
 - Nexus 7000 switches, ISSU (in-service software upgrade), 296–297
 - Nexus 9000 switches, ISSU (in-service software upgrade), 297–299
 - NFS (Network File System), 461–463
 - NIST (National Institute of Standards and Technology), cloud computing definition, 267
 - nondisruptive ISSU (in-service software upgrade), for Cisco NX-OS, 295–298–299
 - nondisruptive upgrade/downgrade
 - on Cisco MDS fabric switch, 482–487, 490–494
 - on Nexus switches, 299–300
 - noninteractive mode, Python, 762–764
 - northbound interfaces, 614
 - NP ports, 354
 - NPIV (N port identifier virtualization), 409, 410–411
 - commands, 411–412
 - creation and verification, 412–415
 - NPV (N port virtualization), 409, 410–411
 - commands, 411–412
 - creation and verification, 412–415
 - NSSAs (not-so-stubby areas), 10
 - NTP (Network Time Protocol), 307–313
 - configuring, 312
 - default settings, 309
 - displaying status, 312–313
 - global commands, 309–311
 - NX-API, 737–739
 - Developer Sandbox, 741–742
 - error codes, 741
 - request and response elements, 739–741
-
- O
- opaque LSAs (link-state advertisements), 6
 - OSPF (Open Shortest Path First)
 - ABR
 - v2 verification*, 17–20
 - v3 verification*, 20–21
 - adjacencies, 6
 - areas, 9
 - NSSAs, 10
 - stub*, 9–10
 - authentication, 11–12
 - BDRs, 11
 - broadcast networks, 11
 - dead-interval, 5
 - default parameters, 12
 - differences between v2 and v3, 6
 - DRs, 11

- global-level commands, 13
- global-level verification and process
 - clear commands, 14–15
- graceful restart, 6
- hello packets, 5
- interface-level commands, 14
- LSAs, 6–8
- point-to-point networks, 11
- Puppet manifest, 756–757
- router configuration and verification, 21–23
- routing level commands, 13–14
- virtual links, 11
- ospf network command, 14**
- OTV (overlay transport virtualization), 150**
 - AED, 160
 - ARP optimization, 158–159
 - broadcast policy control, 159
 - configuration limitations, 164–165
 - configurations and verifications, 163–164
 - connectivity verifications, 173
 - control plane, 151
 - multicast-enabled transport infrastructure, 151–152*
 - unicast-only transport infrastructure, 152–153*
 - data plane, 154
 - broadcast traffic over OTV, 156*
 - multicast traffic over OTV, 156*
 - unicast traffic over OTV, 154–156*
 - edge device, 150
 - failure isolation, 157
 - FHRP isolation, 162–163
 - global-level commands, 165, 167–168
 - internal interfaces, 151
 - internal peering, 160–162
 - join interface, 151
 - multihoming, 159–160
 - overlay interface, 151
 - router-level commands, 167
 - STP isolation, 157
 - unknown unicast handling, 157–158
- otv adjacency-server unicast-only command, 166**
- otv control-group command, 166**
- otv data-group command, 167**
- otv extend-vlan command, 166**
- otv isis csnp-interval command, 167**
- otv isis hello-interval command, 167**
- otv isis hello-multiplier command, 167**
- otv isis metric command, 167**
- otv isis priority command, 167**
- otv join-interface command, 166**
- otv site-identifier command, 166**
- otv site-vlan command, 166, 167**
- otv-isis command, 165**
- overlay adjacency, 161**
- overlay interface, 151**

P

- PaaS (Platform as a Service), 270**
- patch RPMs, managing with Bash, 724–725**
- patching, 288**
- path selection, BGP (Border Gateway Protocol), 25**
 - comparing pairs of paths, 25–27
 - determining the best-path change suppression, 27–28
 - determining the order of comparisons, 27
- path-vector routing, 23**

- Pearson Test Prep software, accessing, 929
- peer gateway, vPC, 130–131
- peering, BGP, 24–25
- peer-ip command, 184
- PFC (priority-based flow control), 422–423
- PIM (Protocol Independent Multicast), 48–49, 63–64
 - any-source multicast (ASM), 50
 - Bidir (bidirectional shared trees), 50
 - configuration limitations, 56
 - default parameters, 55
 - designated routers/forwarders, 53
 - distribution trees, 48
 - messages, 50–51
 - multicast anycast RP configurations and verification, 67
 - multicast auto RP configurations and verification, 66
 - multicast BSRs RP configurations and verification, 65–66
 - multicast distribution modes, 57
 - multicast global-level BGP verification commands, 59–60
 - multicast global-level commands, 57–58
 - multicast interface-level commands, 58–59
 - multicast static RP configurations and verification, 64–65
 - RPs (Rendezvous Points), 52–53
 - source distribution trees, 48
 - state refresh, 51–52
- pim jp-policy command, 59
- playbooks, Ansible, 749, 750–751
- PLOGI (port login), 362
- POAP (PowerOn Auto Provisioning), 767
 - on Cisco NX-OS, 283–287
 - configuration script, 768
 - DHCP discovery phase, 770–771
 - limitations, 767
 - network requirements, 767–768
 - phases, 768–769
 - post-installation reload phase, 772
 - power-up phase, 770
 - script execution phase, 772
 - switch configuration, 772
 - USB discovery phase, 770
- point-to-point networks, 11
- policies, 196–197
 - EEM (Embedded Event Manager), 715–716
 - fabric, 216–217
 - interface, 220
 - SAN connectivity, 606–607
 - server, 580–583
 - for user roles, 913
- Policies Configuration page, ACI, 220–221
- policy model, ACI, 212–213
- port channel, 117–119. *See also* SAN port channels; vPC (virtual port channel)
 - global-level commands, 133–134
 - global-level verification commands, 136
 - interface-level commands, 134
 - load balancing, 120–121
 - recommendations, 132
 - switch configuration, 137–138
- port roles, Rapid PVST+, 162–164
- port security, 826–828
 - on Cisco MDS 9000 Series switches, 915–917
 - configuring, 917–920

- commands, 830–831
- configuring, 829–831
- versus fabric binding, 924–925
- and port types, 829
- verification, 920–922
- violations and actions, 828
- port types**
 - Rapid PVST+, 160–161
 - STP
 - network ports*, 154
 - spanning tree edge ports*, 153
- ports, FC (Fibre Channel), 353**
 - auto mode, 354–355
 - E ports, 353–354
 - Fx ports, 354
 - NP ports, 354
 - TE ports, 354
 - TF ports, 354
 - TNP ports, 354
- post-installation reload phase, POAP (PowerOn Auto Provisioning), 772**
- PowerShell, 787**
 - Cisco UCS Director PowerShell agent, 788–789
 - installing Cisco UCS Director PowerShell agent, 787–788
- power-up phase, POAP (PowerOn Auto Provisioning), 770**
- preempt command, 79**
- primary roles, vPC, 126–127**
- principal switch selection, FC (Fibre Channel), 358–360**
- private cloud deployments, 272**
- private VLANs, 570–571**
- PRLI (process login), 362**
- programming languages. *See* Python**
- prune messages, PIM, 50–51**
- pseudowires, 150**
- PTP (Precision Time Protocol), 313–317**
 - clocks, 314
 - default settings, 315
 - global commands, 315
 - interface-level commands, 315–316
 - verification commands, 315–316
- public cloud deployments, 272–273**
- Puppet, 751–752**
 - and Cisco NX-OS, 753–754
 - installing, 754–755
 - manifest, 752, 757–758
 - OSPF*, 756–757
 - resource types, 756
 - and UCS Manager integration, 757
 - workflow, 752–753
- Python, 758**
 - for Cisco NX-OS, 758–759
 - CLI command APIs*, 760–761
 - displaying information about*, 759
 - interactive mode*, 761–762
 - noninteractive mode*, 762–764
 - package functions*, 760
 - compound commands, 764
 - scripting, 758
- Python SDK for Cisco UCS, 764–766**
 - `convert_to_ucs_python` API, 766–767

Q

- QoS, 589**
 - policy configuration, 591
 - system classes, 589–591
- query-interval command, 59**
- query-response process, IGMP, 42–43**
- query-timeout command, 59**
- question types on exam, 932–935**

R

- RA (Router Advertisement) messages, 76–77
- rack servers, UCS (Cisco Unified Computing System), 521–522
- RADIUS, 898–900
 - configuring in Cisco UCS, 878–882
 - merging with TACACS+, 910
 - securing Cisco UCS, 866–868
- RAID configurations, Cisco NSS3000 Series Network Storage System, 467
- rapid detection of failures, BFD, 37
- Rapid PVST+, 158–160
 - configuration limitations, 167
 - default settings, 165–167
 - loops, 164
 - port roles, 162–164
 - port types, 160–161
- rate limiting, 834–835
- RBAC (role-based access control), 801–803, 910
 - configuring on MDS switches, 914–915
 - global commands, 806–807
 - NX-OS configurations, 805–806
 - user account and verification commands, 808
 - user roles and rules, 803–804
 - user roles, creating, 808–809
- Ready Clones, 701
- recommendations
 - for exam preparation, 936
 - port channeling, 132
 - Rapid PVST+, 167
 - for traffic monitoring, 625
- register messages, PIM, 51
- register-rate-limit command, 57
- remote authentication. *See also* authentication
 - RADIUS and TACACS+ configuration, 878–882
- remote services, AAA (authentication, authorization, and accounting), 897–898
- remote users role policy, UCS (Cisco Unified Computing System), 882–884
- remote-attached topology, 438
- removing ACI contracts, 849
- replication factor, 687
- requests
 - NX-API, 739–740
 - Rest API, 734–735
- reserved FCIDs, 356
- resource optimization, Cisco Hyperflex, 692–693
- resources, Puppet, 756
- responses
 - NX-API, 740
 - Rest API, 736–737
- Rest API
 - authentication, 735–736
 - DCNM features, 776–777
 - response, 736–737
- Rest APIs, 734–735
- restoring
 - backup data on Cisco UCS, 650–652
 - fabric interconnect configuration on Cisco UCS, 653–654
 - system configurations on Cisco UCS, 652
- RMON (Remote Network Monitoring), 505
- roles, Ansible, 749
- rollback, 303–304
 - triggers, 304

Root Guard, 157
 root switch, STP configuration,
 172–182
 router bgp command, 30
 router ospf bfd command, 39
 router ospf command, 13
 router ospfv3 command, 13
 router priority and preemption, VRRP,
 74–75
 router-id command, 13, 30
 routing level commands
 OSPF, 13–14
 OTV, 167
 routing multicast holddown command,
 58
 rp-address command, 57
 RPF (Reverse Path Forwarding),
 53–54
 RPs (Rendezvous Points), 52–53
 RSPAN (Remote SPAN), 509–511
 rules, 803–804, 911–913
 run bash sudo su command, 724
 running backup operations on Cisco
 UCS, 646–647

S

SaaS (Software as a Service), 269
 SAN port channels, 381, 382
 channel group configuration
 differences, 385
 commands, 385–386
 compatibility checks, 382
 creating, 387–388
 load balancing, 383–384
 modes, 384
 trunking, 381–382
 types of, 381
 SANs (storage area networks)
 configuring, 596–597
 connectivity, 593–596
 connectivity policies, 606–607
 Scheduler, 718–719
 clearing the log file, 720–721
 configuring, 719–721
 displaying jobs, 721–722
 jobs
 defining, 720
 starting, 720
 verifying the configuration, 721
 script execution phase, POAP
 (PowerOn Auto Provisioning), 772
 scripting, Python, 758
 interactive mode, 761–762
 noninteractive mode, 762–764
 secondary roles, vPC, 126–127
 security. *See also* authentication; port
 security; RBAC (role-based access
 control)
 AAA (authentication, authorization,
 and accounting), 865–866
 Cisco AVPair attribute, 868
 DAI (Dynamic ARP Inspection),
 810–813
 fabric binding, 922–924, 925
 FHS (First-Hop Security), 809–810
 microsegmentation, 854–862
 RADIUS, 866–868
 TACACS+, 866–868
 two-factor authentication, 869
 SEDs (self-encrypting drives), 702–703
 configuring a local encryption key,
 703
 SELs (system event logs), 615
 server distribution, AAA (authenticat-
 tion, authorization, and accounting),
 909–910

- server firmware, upgrading with Auto Install, 673–675
- server groups, AAA (authentication, authorization, and accounting), 896
- server policies, 580–583
- server pools, 576–577
- service authentication options, AAA (authentication, authorization, and accounting), 796–797
- service configuration options, AAA (authentication, authorization, and accounting), 896
- service instance command, 183
- service profile templates, 583–588
- service profiles, 577–580
- service-policy type network-qos default-nq-7e-policy command, 451
- Session Timeout Period, Cisco UCS, 869
- setup utility, Cisco NX-OS, 280–283
- severity levels for system messages, 307
- shared secrets, 884
- shared trees, 46
- show bfd neighbors command, 39
- show bgp all command, 31
- show bgp command, 31
- show bgp community command, 31
- show bgp convergence command, 31
- show bgp sessions command, 31
- show bgp statistics command, 31
- show commands
 - FCoE (Fibre Channel over Ethernet), 453
 - for NTP (Network Time Protocol), 311
 - SNMP (Simple Network Management Protocol), 323
 - VSAN (virtual storage-area network), 377–378
- show diff rollback-patch command, 304
- show event manager system-policy command, 716
- show feature command, 453
- show feature-set command, 453
- show file command, 763
- show forwarding distribution otv multicast route vlan command, 168
- show groups command, 60
- show guestshell detail command, 729–730
- show interface command, 60
- show ip mroute command, 67–68
- show ip ospf command, 14
- show ip ospf interface command, 14
- show ip ospf route command, 14
- show ip ospf statistics command, 15
- show ip ospf traffic command, 15
- show ip ospf virtual-links command, 15
- show local-groups command, 60
- show mac address-table command, 168
- show mroute ip-address command, 59
- show ospfv3 command, 15
- show ospfv3 interface command, 15
- show otv adjacency command, 167
- show otv command, 167, 169–172
- show otv isis site command, 167
- show otv mroute vlan startup command, 167
- show otv overlay command, 167
- show otv site command, 167
- show otv vlan-mapping command, 168
- show pim group-range command, 59
- show pim rp command, 59
- show port channel compatibility-parameters command, 119

- show port-channel summary command, 144–145
- show route command, 60
- show running configuration bgp command, 31
- show running configuration command, 15, 60
- show running configuration otv command, 167
- show running configuration pim command, 60
- show spanning-tree summary command, 113–115
- show vpc brief command, 143–144
- show vpc command, 140–143
- shutdown command, 30, 452
- shutdown lan command, 452
- simlet exam questions, 935
- simulation exam questions, 934
- single contract bidirectional reverse filter, 853
- single contract unidirectional with multiple filters, 854
- single-hop topology, FCoE (Fibre Channel over Ethernet), 435–436
 - direct-attached, 436–437
 - FEX, 437–438
 - remote-attached, 438
- site adjacency, 161
- SLAs (service-level agreements), 266
- Smart Call Home, 619–620
 - for Cisco NX-OS, 324–325
- smart zoning, 396–397
- SMU (Software Maintenance Upgrade), for Cisco NX-OS, 287–289
- SNMP (Simple Network Management Protocol), 317–319, 618–619
 - configuration examples, 323–324
 - global commands, 320
 - MIBs, 320–322
 - security models and levels, 319
 - specific notification commands, 322
 - verification commands, 323
- snooping, IGMP, 44
- soft zoning, 392
- software
 - updating on Cisco UCS, 654–661
 - upgrading/downgrading on Cisco MDS NX-OS, 480–481
- source distribution trees, 48
- source-interface command, 184
- SPAN (Switched Port Analyzer), 505–509, 622. *See also* ERSPAN (Encapsulated Remote Switched Port Analyzer)
 - for Cisco NX-OS, 330–337
- spanning-tree bpdudfilter command, 171
- spanning-tree bpduguard command, 170
- spanning-tree cost command, 171
- spanning-tree guard loop command, 171
- spanning-tree guard root command, 171
- spanning-tree loopguard default command, 169
- spanning-tree mode rapid-pvst command, 169
- spanning-tree pathcost method command, 169
- spanning-tree port type edge bpdudfilter default command, 168
- spanning-tree port type edge bpduguard default command, 168
- spanning-tree port type edge command, 170
- spanning-tree port type edge default command, 168

- spanning-tree port type edge trunk command, 452
- spanning-tree port type network command, 170
- spanning-tree port type network default command, 168
- spanning-tree port type normal command, 170
- spanning-tree port-priority command, 171
- spanning-tree vlan command, 169
- spanning-tree vlan priority command, 169
- spanning-tree vlan root primary command, 169
- spanning-tree vlan root secondary command, 169
- sparse-mode command, 58
- SPF (Shortest Path First) algorithm, 6
- spine switches, 201–204
- SSM (source-specific mulicast), 49
- startup-query-interval command, 59
- static RP (rendezvous point), 52
- statistics collection policy, UCS (Cisco Unified Computing System), 617–618
- storage, UCS (Cisco Unified Computing System), 592
- storage clusters, 699
- STP (Spanning Tree Protocol)
 - BPDU Filter, 156
 - BPDU Guard, 156
 - Bridge Assurance, 154–156
 - configurations and verifications, 164–165
 - default port cost, 163–164
 - extensions, 154
 - global-level commands, 167–170
 - global-level verification commands, 171–172
 - interface-level commands, 170–171
 - Loop Guard, 156–157
 - port channels, 117–119
 - port types, 153
 - edge ports*, 153
 - network ports*, 154
 - Rapid PVST+, 158–160
 - configuration limitations*, 167
 - default settings*, 165–167
 - loops*, 164
 - port roles*, 162–164
 - port types*, 160–161
 - Root Guard, 157
 - root switch configuration, 172–182
 - switch configuration, 113
 - switch global configuration, 115
 - switch interface configuration, 115
 - switch verification, 116–117
 - topology, 152
 - UDLD (Unidirectional Link Detection), 157–158
 - verification of configurations, 113–115
- STP isolation, 157
- streaming telemetry, 337–341
- structure, of XML, 731–732
- stub areas, 9–10
 - NSSAs, 10
- subjects, 236–237
- sudo yum installed command, 724
- switched fabric initialization, FC (Fibre Channel), 358
- switches, configuring with POAP, 772
- switchport mode fex-fabric command, 453
- switchport mode trunk command, 452
- switchto vdc command, 451

syntax, XML (Extensible Markup Language), 732–733

syslog messages, 614–615

system classes, 589–591

system messages

 Cisco MDS NX-OS, 503–504

 Cisco NX-OS, 306–307

 severity levels, 307

system qos command, 451

system restore, on Cisco UCS, 652

system setup, Cisco UCS Director, 785–786

T

T bundle, 655

taboo contracts, 233

TACACS+, 900–903

 configuring in Cisco UCS, 878–882

 merging with RADIUS, 910

 securing Cisco UCS, 866–868

TE ports, 354

telemetry data collection, 337–341, 632

tenants, 214, 225–226, 227–228

test preparation. *See* final preparation

testlet exam questions, 934–935

TF ports, 354

time management

 NTP (Network Time Protocol), 307–313

configuring, 312

default settings, 309

 PTP (Precision Time Protocol), 313–317

default settings, 315

global commands, 315

interface-level commands, 315–316

timers command, 30, 80

TNP ports, 354

tolerated failures, Cisco HX Data Platform, 701

topologies

 DAI (Dynamic ARP Inspection), 815–816

 FC (Fibre Channel)

collapsed-core, 350

core-edge, 350–351

core-edge-core, 351–353

 FCoE (Fibre Channel over Ethernet)

direct-attached, 436–437

FEX, 437–438

multi-hop, 438–439

remote-attached, 438

single-hop, 435–436

 STP, 152

 vPC, 122

 VRRP, 73

track interface command, 79

tracking, VRRP, 75–76

traffic flows, vPC, 124–125

traffic monitoring, 622–623, 628–629

 across Ethernet, 623–624, 625–626

 across Fibre Channel, 624–625, 626–627

 adding traffic sources, 628

traffic storms, 243

TRM (Tenant Routed Multicast), 180–182

trunking

 SAN port channels, 381–382

 VSANs (virtual storage-area networks), 374–376

TTL (time-to-live), 44

two-factor authentication, 869

U

UCS (Cisco Unified Computing System), 610, 612

AGs (application gateways), 613–614
 architecture, 516–518
 audit logs, 616
 backup operations, 642–643
 creating, 643–646
 imports, 650–652
 running, 646–647
 running fields, 637–647
 backup policies, 648
 configuring, 648–650
 export fields, 649–650
 reminder options, 650
 blade chassis slot to link mapping, 529
 blade servers, 520–521
 Call Home, 619–620
 chassis/FEX discovery, 556–557
 Cisco Intersight, 629–630, 633–636
 benefits, 630–632
 licensing, 632–633
 supported software, 632
 as telemetry data collection, 632
 Cisco UCS 2304 IOM, 527–528
 Cisco UCS 5108 Blade Server Chassis, 520
 Cisco UCS 6300 Series Fabric Interconnects, 526–527
 Cisco UCS 6454 Fabric Interconnect, 524–526
 Cisco UCS Mini solution, 523
 cluster verification CLI, 543–544
 components and connectivity, 518–519
 configuration management, 642–643

configuring subordinate fabric interconnect, 540–543
 connecting blade chassis fabric extenders to fabric interconnect, 528–529
 database health and hardware monitoring, 620–621
 determining the primary fabric interconnect, 544–545
 device discovery, 556
 direct updates, 659
 DME (data management engine), 612–613
 downlink connectivity, 546
 dual-wire manager, 532–533
 endpoints, 546
 Ethernet switching mode, 550–556
 fabric infrastructure, 524
 fabric interconnect configuration, restoring, 653–654
 fabric interconnect connectivity and configurations, 544
 fabric interconnect port modes, 547–548
 fault life cycle, 616–617
 FEX fabric, 529–531
 FEX virtual links, 531–532
 FI initialization, 539–540
 firmware/software upgrades, 654–661
 Auto Install, 662–665, 673–675
 best practices, 659–661
 direct upgrades, 666–671
 infrastructure firmware fields, 671–672
 version terminology, 661–662
 HUU (Cisco Host Upgrade Utility), 675–681
 identity pools, 571–572
 IP pools, 574–576
 MAC pools, 573–574

- server pools*, 576–577
- universally unique identifier suffix pools*, 572–573
- image management best practices, 656–659
- initial setup and management, 536–540
- initial setup for standalone UCS C-Series, 557–560
- LDAP providers and groups, 869–878
 - creating*, 870–874
 - deleting existing groups*, 875
 - group mapping*, 875–878
 - modifying existing groups*, 874–875
- monitoring policies, 616
- multiple authentication services configuration, 884
- NetFlow monitoring, 621–622
- network management, 563
- northbound interfaces, 614
- operating system installation on standalone UCS C-Series, 560–563
- Puppet integration, 757
- Python SDK, 764–766
- QoS, 589
 - policy configuration*, 591
 - system classes*, 589–591
- rack server discovery policy, 557
- rack servers, 521–522
- RADIUS configuration, 878–882
- remote users role policy, 882–884
- SANs (storage area networks)
 - configuration*, 596–597
 - connectivity*, 593–596
 - connectivity policies*, 606–607
- securing
 - with AAA*, 865–866
 - with RADIUS and TACACS+*, 866–868
 - with two-factor authentication*, 869
- SELs (system event logs), 615
- server policies, 580–583
- service profile templates, 583–588
- service profiles, 577–580
- Session Timeout Period, 869
- SNMP (Simple Network Management Protocol), 618–619
- statistics collection policy, 617–618
- storage, 592
- syslog messages, 614–615
- system restore, 652
- T bundle, 655
- TACACS+ configuration, 878–882
- traffic monitoring, 622–623, 628–629
 - across Ethernet*, 623–624, 625–626
 - across Fibre Channel*, 624–625, 626–627
 - adding traffic sources*, 628
- uplink connectivity, 546
- VICs (virtual interface cards), 535–536
- virtualization infrastructure, 533–535
- VLANs, 563–566
 - named*, 566–570
 - private*, 570–571
- vNICs (virtual network interface cards), 549–550
- VSANs (virtual storage-area networks), 597
 - named*, 597–599
- Web Session Refresh Period, 869
- World Wide Name pools, 603–605
- zone profiles, creating, 600–602
- zone sets, 600–601
- zones, 600

- UDLD (Unidirectional Link Detection), 157–158
 - udld aggressive command, 170
 - udld message-time command, 169
 - udldl command, 171
 - unicast traffic over OTV, 154–156
 - unicast-only transport infrastructure, 152–153
 - universally unique identifier suffix pools, 572–573
 - unknown unicast handling, 157–158
 - updating
 - firmware, 659
 - Auto Install*, 662–665
 - best practices*, 659–661
 - HUU (Cisco Host Upgrade Utility)*, 675–681
 - module EPLDs on Cisco MDS 9000 Series switches, 502–503
 - UCS firmware and software, 654–661
 - your exams, 931
 - upgrading. *See also* downgrading
 - ACI fabric, 212
 - Cisco MDS NX-OS, 480–481
 - Cisco NX-OS, 287–289, 291
 - nondisruptive ISSU*, 295–298–299
 - disruptive upgrade/downgrade
 - on Cisco MDS fabric switch*, 487–489
 - on Nexus switches*, 299–303
 - EPLD on Cisco MDS 9000 Series switches, 498–503
 - GIR (Graceful Insertion and Removal), 291–295
 - HUU (Cisco Host Upgrade Utility), 675–681
 - nondisruptive upgrade/downgrade
 - on Cisco MDS fabric switch*, 482–487
 - on Nexus switches*, 299–303
 - uplink connectivity, UCS (Cisco Unified Computing System), 546
 - USB discovery phase, POAP (PowerOn Auto Provisioning), 770
 - user roles, 911
 - for Cisco NX-OS, 803–804
 - creating, 808–809
 - policies, 913
 - rules, 911–913
- ## V
-
- variable files, Ansible, 749
 - variables, Ansible, 749
 - vdc type storage command, 451
 - verifying
 - BGP, 33–35
 - device aliases, 408–409
 - EEM configuration, 718
 - FCoE (Fibre Channel over Ethernet), 448–451
 - NPV (N port virtualization), 412–415
 - port security on Cisco MDS 9000 Series switches, 920–922
 - Scheduler configuration, 721
 - zones, 401–403
 - version value command, 58
 - vFC (Virtual Fibre Channel), 428–429
 - VIBs (VMware Installation Bundles), 688–689
 - VICs (virtual interface cards), 690–692
 - UCS (Cisco Unified Computing System), 535–536
 - Virtual Edge, 225
 - virtual links, OSPF, 11
 - vlan command, 182

- VLANs, 563–566**
 - named, 566–570
 - private, 570–571
- VMM (Virtual Machine Manager), 222**
- VMs (Virtual Machines), Ready Clones, 701**
- VMware vCenter, Cisco ACI integration, 224–225**
- VMware vSphere, 688**
- vNICs (virtual network interface cards), 549–550**
- VNIs (virtual network identifiers), 175–176**
- vn-segment command, 182**
- vPC (virtual port channel), 121–122, 138–139**
 - ARP synchronization, 130
 - components, 123–124
 - configuration consistency, 127–128
 - domain, 123
 - domain-level commands, 135
 - dual-control plane, 125
 - duplicate frames prevention, 128–129
 - fault-tolerant link, 123
 - global-level commands, 133–134
 - global-level verification commands, 136
 - HSRP gateway considerations, 130
 - interface-level commands, 134
 - member port, 123
 - non-vPC port, 123
 - peer gateway, 130–131
 - peer link, 123
 - peer switch, 123
 - peer-keepalive, 123
 - port channel configurations and verifications, 131–133
 - port channels, recommendations, 132
 - primary and secondary roles, 126–127
 - topologies, 122
 - traffic flows, 124–125
- VRF (Virtual Routing and Forwarding), 214, 228**
 - contracts, 850–851
- VRRP (Virtual Router Redundancy Protocol), 72**
 - authentication, 75
 - basic topology, 73
 - benefits, 73
 - default parameters, 78
 - dynamic router discovery, 72
 - full configurations, 82–83
 - global-level verification commands, 81
 - groups, 74
 - interface-level commands, 79–80
 - router priority and preemption, 74–75
 - tracking, 75–76
 - verification, 84–86
- vsan database command, 452**
- VSANs (virtual storage-area networks), 370, 428, 597**
 - advantages, 373
 - attributes, 372
 - commands, 376–378
 - creating, 379–380
 - DPVM (Dynamic Port VSAN Membership), 373
 - features, 370–372
 - named, 597–599
 - show commands, 377–378
 - trunking, 374–376
 - and zoning, 391
- VXLAN (Virtual eXtensible LAN), 173, 239–241**
 - configurations and verifications, 182, 189–191
 - control plane, 176–178
 - encapsulation, 173–174

- gateways, 178–179
- global-level commands, 182–183
- global-level verification commands, 184
- high availability, 179–180
- interface-level commands, 183
- packet format, 174
- PIM multicast configurations and verifications, 186–189
- TRM (Tenant Routed Multicast), 180–182
- tunnel endpoint, 174
- VNIs, 175–176
- vzAny, 233

W

- Web Session Refresh Period, Cisco UCS, 869
- web user interface, DCNM (Data Center Network Manager), 779–782
- World Wide Name pools, 603–605

X

- XML (Extensible Markup Language), 730–731
 - data fields, 731

- device element, 731
- documents, 732
- structure, 731–732
- syntax, 732–733

Y-Z

- yum list available command, 724
- zone mode enhanced vsan command, 397
- zone profiles, creating, 600–602
- zone sets, 600–601
- zones, 600
- zoning, 389
 - active zone set, 392–393
 - Autozone, 395
 - commands, 399–400
 - creation and verification, 401–403
 - enforcement, 391–392
 - enhanced, 397–398
 - fabric with two zones, 390–391
 - features, 389–390
 - full zone set, 393–394
 - smart, 396–397
 - and VSANs, 391
 - zone merge, 395–396