

Practice Tests



Flash Cards



Study Planner



Video Training

# Official Cert Guide

Advance your IT career with hands-on learning

# CCNP Data Center Application Centric Infrastructure

## DCACI 300-620

[ciscopress.com](http://ciscopress.com)

**Ammar Ahmadi, CCIE® No. 50928**

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



# **CCNP Data Center Application Centric Infrastructure DCACI 300-620**

**Official** Cert Guide

AMMAR AHMADI CCIE No. 50928

Cisco Press

# **CCNP Data Center Application Centric Infrastructure DCACI 300-620 Official Cert Guide**

Ammar Ahmadi

Copyright© 2021 Pearson Education, Inc.

Published by:  
Cisco Press  
Hoboken, NJ

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: 2020948500

ISBN-13: 978-0-13-660266-8

ISBN-10: 0-13-660266-5

## **Warning and Disclaimer**

This book is designed to provide information about the CCNP Implementing Cisco Application Centric Infrastructure DCACI 300-620 certification exam. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

## **Trademark Acknowledgments**

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

## **Special Sales**

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

For government sales inquiries, please contact [governmentsales@pearsoned.com](mailto:governmentsales@pearsoned.com).

For questions about sales outside the U.S., please contact [intlcs@pearson.com](mailto: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](mailto: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

**Copy Editor:** Kitty Wilson

**Alliances Manager, Cisco Press:** Arezou Gol

**Technical Editors:** Akhil Behl, Nikhil Behl

**Director, ITP Product Management:** Brett Bartow

**Editorial Assistant:** Cindy Teeters

**Executive Editor:** James Manly

**Cover Designer:** Chuti Prasertsith

**Managing Editor:** Sandra Schroeder

**Composition:** codeMantra

**Development Editor:** Ellie Bru

**Indexer:** Erika Millen

**Senior Project Editor:** Tonya Simpson

**Proofreader:** Donna Mulder



**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](http://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](http://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 Author

**Ammar Ahmadi, CCIE No. 50928**, has nearly a decade of experience in data center design, implementation, optimization, and troubleshooting. He currently consults for Cisco Gold partner AHEAD INC, where he has been designing and supporting large-scale ACI fabrics since the early days of ACI. Occasionally, he breaks from design work to produce network modernization roadmaps or demonstrate the possibilities of software-defined networking (SDN) to customers.

Ammar also owns and operates Networks Reimagined LLC, which focuses on SDN enablement and training. He can be reached at [ammar.ahmadi@networksreimagined.com](mailto:ammar.ahmadi@networksreimagined.com).

## About the Technical Reviewers

**Akhil Behl, CCIE No. 19564**, is a passionate IT executive with a key focus on cloud and security. He has more than 16 years of experience in the IT industry, working across several leadership, advisory, consultancy, and business development profiles with various organizations. His technology and business specialization includes cloud, security, infrastructure, data center, and business communication technologies. Currently he leads business development for cloud for a global systems integrator.

Akhil is a published author. Over the past few years, he has authored multiple titles on security and business communication technologies. He has contributed as technical editor to more than a dozen books on security, networking, and information technology. He has published several research papers in national and international journals, including *IEEE Xplore*, and presented at various IEEE conferences, as well as other prominent ICT, security, and telecom events. He is passionate about writing and mentoring.

He holds CCIE Emeritus (Collaboration and Security), Azure Solutions Architect Expert, Google Professional Cloud Architect, CCSK, CHFI, ITIL, VCP, TOGAF, CEH, ISM, CCDP, and many other industry certifications. He has a bachelor's degree in technology and an MBA.

**Nikhil Behl, CCIE No. 23335**, is a seasoned IT professional with exposure to a broad range of technologies. He has more than 15 years of experience working in the IT industry. He has worked in several ICT roles, including solutions architect, pre-sales lead, network architect, business consultant, and CISCO TAC engineer, and he has worked with system integration and managed network services.

Nikhil has expertise in various technologies, including cloud, core networking, data center networking, software-defined networking, Wi-Fi, SD-WAN, and Software-Defined Access. He actively participates in several industry conferences and IT forums as a speaker.

Nikhil holds CCIE (Enterprise Infrastructure), Azure Solutions Architect Expert, Cisco SD-WAN Blackbelt, CCNP (Enterprise), CCDP, CCNA, CCDA, JNCIA (Junos), JNCIS, and many other industry leading certifications. He has a bachelor's degree in computer applications.

## Dedication

I dedicate this book to my loving wife, Sophia, and my two children, Kiyana and Daniel. Sophia, your unrelenting support and patience made this book possible. I am forever grateful! Kiyana, you are full of life! You remind me every day that life is wonderful and that every moment should be cherished. Daniel, your nice big hugs have energized me at times when I really needed them. I look forward to spending more time with you and getting to know you more.

## Acknowledgments

First, I would like to thank all the people who helped get this book to print. My special thanks go to Pearson product manager James Manly for his patience and understanding during this long process. I am also grateful to development editor Eleanor Bru, whose keen eye for detail has contributed tremendously to the quality of this book. Thank you, Brett Bartow, for the opportunity. And thanks to all the other people who contributed behind the scenes.

Second, I would like to thank all the people who have helped me grow in my career. I thank Peter Thompson, whom I saw as a mentor early in my career at Cisco. You helped me make several tough decisions that each greatly influenced my career. Also, thank you to Ryan Alt from AHEAD and John Rider from Cisco for allowing me the freedom to pick my projects. And thanks to all the unnamed folks at AHEAD who have made the past few years enjoyable. Thank you, Anthony Wilde, for showing me that you actually can have your cake and eat it, too. I would also like to acknowledge my wife, Sophia, for never letting me fall into the trap of complacency.

Finally, I would like to thank my mom, whose greatest desire has always been for me and my siblings to succeed, be happy, and achieve our dreams.

## Contents at a Glance

Introduction xxv

### **Part I Introduction to Deployment**

- Chapter 1 The Big Picture: Why ACI? 2
- Chapter 2 Understanding ACI Hardware and Topologies 18
- Chapter 3 Initializing an ACI Fabric 42
- Chapter 4 Exploring ACI 96

### **Part II ACI Fundamentals**

- Chapter 5 Tenant Building Blocks 130
- Chapter 6 Access Policies 156
- Chapter 7 Implementing Access Policies 186
- Chapter 8 Implementing Tenant Policies 238

### **Part III External Connectivity**

- Chapter 9 L3Outs 288
- Chapter 10 Extending Layer 2 Outside ACI 358

### **Part IV Integrations**

- Chapter 11 Integrating ACI into vSphere Using VDS 392
- Chapter 12 Implementing Service Graphs 424

### **Part V Management and Monitoring**

- Chapter 13 Implementing Management 462
- Chapter 14 Monitoring ACI Using Syslog and SNMP 482
- Chapter 15 Implementing AAA and RBAC 514

### **Part VI Operations**

- Chapter 16 ACI Anywhere 552

## **Part VII Final Preparation**

Chapter 17 Final Preparation 566

Appendix A Answers to the “Do I Know This Already?” Questions 572

Appendix B CCNP Data Center Application Centric Infrastructure  
DCACI 300-620 Exam Updates 586

Glossary 589

Index 602

### **Online Elements**

Appendix C Memory Tables

Appendix D Memory Tables Answer Key

Appendix E Study Planner

Glossary

\*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	xxv
<b>Part I</b>	<b>Introduction to Deployment</b>	
<b>Chapter 1</b>	<b>The Big Picture: Why ACI?</b>	<b>2</b>
	“Do I Know This Already?” Quiz	2
	Foundation Topics	4
	Understanding the Shortcomings of Traditional Networks	4
	Network Management	4
	Scalability and Growth	5
	Network Agility	8
	Security	8
	Network Visibility	9
	Recognizing the Benefits of Cisco ACI	9
	Network Management Touchpoints	9
	Traffic Flow Optimizations	10
	Scalability Optimizations	10
	Programmability	11
	Stateless Network	11
	Multitenancy	11
	Zero-Trust Security	14
	Cross-Platform Integrations	15
	New Architectural Possibilities	15
	Integrated Health Monitoring and Enhanced Visibility	16
	Policy Reuse	16
	Exam Preparation Tasks	16
	Review All Key Topics	16
	Complete Tables and Lists from Memory	17
	Define Key Terms	17
<b>Chapter 2</b>	<b>Understanding ACI Hardware and Topologies</b>	<b>18</b>
	“Do I Know This Already?” Quiz	18
	Foundation Topics	21
	ACI Topologies and Components	21
	Clos Topology	21
	Standard ACI Topology	22
	ACI Stretched Fabric Topology	24
	ACI Multi-Pod Topology	25
	ACI Multi-Site Topology	26

	ACI Multi-Tier Architecture	28
	Remote Leaf Topology	30
	APIC Clusters	32
	APIC Cluster Scalability and Sizing	33
	Spine Hardware	36
	First-Generation Spine Switches	37
	Second-Generation Spine Switches	37
	Leaf Hardware	38
	First-Generation Leaf Switches	38
	Second-Generation Leaf Switches	39
	Exam Preparation Tasks	41
	Review All Key Topics	41
	Complete Tables and Lists from Memory	41
	Define Key Terms	41
<b>Chapter 3</b>	<b>Initializing an ACI Fabric</b>	<b>42</b>
	“Do I Know This Already?” Quiz	42
	Foundation Topics	44
	Understanding ACI Fabric Initialization	44
	Planning Fabric Initialization	45
	Understanding Cabling Requirements	45
	Connecting APICs to the Fabric	46
	Initial Configuration of APICs	47
	APIC OOB Configuration Requirements	47
	Out-of-Band Versus In-Band Management	48
	Configuration Information for Fabric Initialization	48
	Switch Discovery Process	49
	Fabric Discovery Stages	51
	Switch Discovery States	51
	Initializing an ACI Fabric	52
	Changing the APIC BIOS Password	52
	Configuring the APIC Cisco IMC	52
	Initializing the First APIC	53
	Discovering and Activating Switches	55
	Understanding Graceful Insertion and Removal (GIR)	58
	Initializing Subsequent APICs	59
	Understanding Connectivity Following Switch Initialization	59
	Basic Post-Initialization Tasks	63
	Assigning Static Out-of-Band Addresses to Switches and APICs	63

Applying a Default Contract to Out-of-Band Subnet	64
Upgrading an ACI Fabric	66
Understanding Schedulers	73
Enabling Automatic Upgrades of New Switches	74
Understanding Backups and Restores in ACI	75
Making On-Demand Backups in ACI	76
Making Scheduled Backups in ACI	79
Taking Configuration Snapshots in ACI	80
Importing Configuration Backups from Remote Servers	80
Executing Configuration Rollbacks	82
Pod Policy Basics	83
Configuring Network Time Protocol (NTP) Synchronization	84
Configuring DNS Servers for Lookups	90
Verifying COOP Group Configurations	92
Exam Preparation Tasks	93
Review All Key Topics	93
Complete Tables and Lists from Memory	94
Define Key Terms	94

## **Chapter 4 Exploring ACI 96**

“Do I Know This Already?” Quiz	96
Foundation Topics	98
ACI Access Methods	98
GUI	99
CLI	100
APIC CLI	100
Switch CLI	102
API	103
Management Access Modifications	103
Understanding the ACI Object Model	105
Learning ACI Through the Graphical User Interface	107
Exploring the Object Hierarchy by Using Visore	108
Why Understand Object Hierarchy Basics for DCACI?	110
Policy in Context	110
Integrated Health Monitoring and Enhanced Visibility	110
Understanding Faults	111
The Life of a Fault	113
Acknowledging Faults	115
Faults in the Object Model	116

Monitoring Policies in ACI	118
Customizing Fault Management Policies	120
Squelching Faults and Changing Fault Severity	121
Understanding Health Scores	124
Understanding Events	126
Squelching Events	127
Understanding Audit Logs	127
Exam Preparation Tasks	128
Review All Key Topics	128
Complete Tables and Lists from Memory	129
Define Key Terms	129

## **Part II ACI Fundamentals**

### **Chapter 5 Tenant Building Blocks 130**

“Do I Know This Already?” Quiz	130
Foundation Topics	132
Understanding the Basic Objects in Tenants	132
Tenants	133
Predefined Tenants in ACI	134
VRF Instances	135
Bridge Domains (BDs)	137
Endpoint Groups (EPGs)	137
Application Profiles	138
The Pain of Designing Around Subnet Boundaries	139
BDs and EPGs in Practice	141
Configuring Bridge Domains, Application Profiles, and EPGs	142
Classifying Endpoints into EPGs	146
APIC CLI Configuration of Tenant Objects	147
Contract Security Enforcement Basics	148
Contracts, Subjects, and Filters	148
Contract Direction	149
Contract Scope	150
Zero-Trust Using EPGs and Contracts	151
Objects Enabling Connectivity Outside the Fabric	151
External EPGs	151
Layer 3 Outside (L3Out)	153
Tenant Hierarchy Review	153
Exam Preparation Tasks	154
Review All Key Topics	154

Complete Tables and Lists from Memory 154

Define Key Terms 154

## **Chapter 6 Access Policies 156**

“Do I Know This Already?” Quiz 156

Foundation Topics 158

Pools, Domains, and AAEPs 158

VLAN Pools 159

Domains 160

Common Designs for VLAN Pools and Domains 161

Challenges with Overlap Between VLAN Pools 164

Attachable Access Entity Profiles (AAEPs) 165

Policies and Policy Groups 169

Interface Policies and Interface Policy Groups 169

Planning Deployment of Interface Policies 173

Switch Policies and Switch Policy Groups 174

Profiles and Selectors 176

Configuring Switch Profiles and Interface Profiles 179

Stateless Networking in ACI 182

Bringing It All Together 183

Access Policies Hierarchy in Review 183

Access Policies and Tenancy in Review 184

Exam Preparation Tasks 184

Review All Key Topics 184

Complete Tables and Lists from Memory 185

Define Key Terms 185

## **Chapter 7 Implementing Access Policies 186**

“Do I Know This Already?” Quiz 186

Foundation Topics 188

Configuring ACI Switch Ports 188

Configuring Individual Ports 188

Configuring Port Channels 196

Configuring Virtual Port Channel (vPC) Domains 201

Configuring Virtual Port Channels 204

Configuring Ports Using AAEP EPGs 208

Implications of Initial Access Policy Design on Capabilities 210

Configuring Access Policies Using Quick Start Wizards 211

The Configure Interface, PC, and VPC Wizard 211

The Configure Interface Wizard 211

Additional Access Policy Configurations	212
Configuring Fabric Extenders	212
Configuring Dynamic Breakout Ports	215
Configuring Global QoS Class Settings	217
Configuring DHCP Relay	219
Configuring MCP	221
Configuring Storm Control	223
Configuring CoPP	225
Modifying BPDU Guard and BPDU Filter Settings	230
Modifying the Error Disabled Recovery Policy	231
Configuring Leaf Interface Overrides	232
Configuring Port Channel Member Overrides	232
Exam Preparation Tasks	235
Review All Key Topics	235
Complete Tables and Lists from Memory	236
Define Key Terms	236
<b>Chapter 8</b>	<b>Implementing Tenant Policies 238</b>
“Do I Know This Already?” Quiz	238
Foundation Topics	241
ACI Endpoint Learning	241
Lookup Tables in ACI	241
Local Endpoints and Remote Endpoints	242
Understanding Local Endpoint Learning	243
Unicast Routing and Its Impact on Endpoint Learning	243
Understanding Remote Endpoint Learning	244
Understanding the Use of VLAN IDs and VNIDs in ACI	245
Endpoint Movements Within an ACI Fabric	247
Understanding Hardware Proxy and Spine Proxy	247
Endpoint Learning Considerations for Silent Hosts	248
Where Data Plane IP Learning Breaks Down	249
Endpoint Learning on L3Outs	249
Limiting IP Learning to a Subnet	249
Understanding Enforce Subnet Check	250
Disabling Data Plane Endpoint Learning on a Bridge Domain	250
Disabling IP Data Plane Learning at the VRF Level	251
Packet Forwarding in ACI	251
Forwarding Scenario 1: Both Endpoints Attach to the Same Leaf	251
Understanding Pervasive Gateways	252

Forwarding Scenario 2: Known Destination Behind Another Leaf	254
Verifying the Traffic Path Between Known Endpoints	254
Understanding Learning and Forwarding for vPCs	256
Forwarding Scenario 3: Spine Proxy to Unknown Destination	258
Forwarding Scenario 4: Flooding to Unknown Destination	261
Understanding ARP Flooding	262
Deploying a Multi-Tier Application	263
Configuring Application Profiles, BDs, and EPGs	264
Assigning Domains to EPGs	267
Policy Deployment Following BD and EPG Setup	267
<i>Mapping EPGs to Ports Using Static Bindings</i>	267
<i>Verifying EPG-to-Port Assignments</i>	269
<i>Policy Deployment Following EPG-to-Port Assignment</i>	270
Mapping an EPG to All Ports on a Leaf	270
Enabling DHCP Relay for a Bridge Domain	271
Whitelisting Intra-VRF Communications via Contracts	272
Planning Contract Enforcement	272
Configuring Filters for Bidirectional Application	273
<i>Configuring Subjects for Bidirectional Application of Filters</i>	275
<i>Understanding Apply Both Directions and Reverse Filter Ports</i>	277
<i>Verifying Subject Allocation to a Contract</i>	278
<i>Assigning Contracts to EPGs</i>	278
<i>Understanding the TCP Established Session Rule</i>	279
<i>Creating Filters for Unidirectional Application</i>	280
<i>Configuring Subjects for Unidirectional Application of Filters</i>	280
<i>Additional Whitelisting Examples</i>	282
<i>Verifying Contract Enforcement</i>	283
<i>Understanding the Stateful Checkbox in Filter Entries</i>	284
<i>Contract Scopes in Review</i>	284
Exam Preparation Tasks	285
Review All Key Topics	285
Complete Tables and Lists from Memory	287
Define Key Terms	287
<b>Part III External Connectivity</b>	
<b>Chapter 9 L3Outs</b>	<b>288</b>
“Do I Know This Already?” Quiz	288
Foundation Topics	291

L3Out Fundamentals	291
Stub Network and Transit Routing	291
Types of L3Outs	292
Key Functions of an L3Out	293
The Anatomy of an L3Out	293
Planning Deployment of L3Out Node and Interface Profiles	295
Understanding L3Out Interface Types	296
Understanding L3Out Bridge Domains	296
Understanding SVI Encap Scope	298
Understanding SVI Auto State	299
Understanding Prerequisites for Deployment of L3Outs	301
L3 Domain Implementation Examples	301
Understanding the Need for BGP Route Reflection	303
Implementing BGP Route Reflectors	304
Understanding Infra MP-BGP Route Distribution	305
Deploying L3Outs	307
Configuring an L3Out for EIGRP Peering	307
Deploying External EPGs	310
Verifying Forwarding Out an L3Out	312
Advertising Subnets Assigned to Bridge Domains via an L3Out	314
Enabling Communications over L3Outs Using Contracts	316
Deploying a Blacklist EPG with Logging	318
Advertising Host Routes Out an ACI Fabric	321
Implementing BFD on an EIGRP L3Out	321
Configuring Authentication for EIGRP	324
EIGRP Customizations Applied at the VRF Level	324
Configuring an L3Out for OSPF Peering	325
A Route Advertisement Problem for OSPF and EIGRP L3Outs	328
Implementing BFD on an OSPF L3Out	328
OSPF Customizations Applied at the VRF Level	329
Adding Static Routes on an L3Out	329
Implementing IP SLA Tracking for Static Routes	330
Configuring an L3Out for BGP Peering	334
Implementing BGP Customizations at the Node Level	337
Implementing Per-Neighbor BGP Customizations	339
Implementing BFD on a BGP L3Out	341
Implementing BGP Customizations at the VRF Level	342
Implementing OSPF for IP Reachability on a BGP L3Out	343

Implementing Hot Standby Router Protocol (HSRP)	344
IPv6 and OSPFv3 Support	344
Implementing Route Control	344
Route Profile Basics	344
Modifying Route Attributes to All Peers Behind an L3Out	346
Modifying Route Attributes to a Specific Peer Behind an L3Out	349
Assigning Different Policies to Routes at the L3Out Level	351
Configuring Inbound Route Filtering in ACI	352
Exam Preparation Tasks	353
Review All Key Topics	353
Complete Tables and Lists from Memory	356
Define Key Terms	356

## **Chapter 10 Extending Layer 2 Outside ACI 358**

“Do I Know This Already?” Quiz	358
Foundation Topics	361
Understanding Network Migrations into ACI	361
Understanding Network-Centric Deployments	361
Understanding Full-Mesh Network-Centric Contracts	362
Understanding Any EPG	364
Understanding Preferred Group Members	365
Disabling Contract Enforcement at the VRF Instance Level	367
Flooding Requirements for L2 Extension to Outside Switches	368
Understanding GARP-Based Detection	370
Understanding Legacy Mode	371
Endpoint Learning Considerations for Layer 2 Extension	371
Preparing for Network-Centric Migrations	372
Implementing Layer 2 Connectivity to Non-ACI Switches	372
Understanding EPG Extensions	372
Understanding Bridge Domain Extensions	374
Comparing EPG Extensions and BD Extensions	374
Implementing EPG Extensions	375
Implementing L2Outs	380
Migrating Overlapping VLANs into ACI	385
Understanding ACI Interaction with Spanning Tree Protocol	386
Remediating Against Excessive Spanning Tree Protocol TCNs	386
Configuring MST Instance Mappings in ACI	387
Understanding Spanning Tree Protocol Link Types	388
Using MCP to Detect Layer 2 Loops	388

Exam Preparation Tasks	389
Review All Key Topics	389
Complete Tables and Lists from Memory	390
Define Key Terms	390

## **Part IV Integrations**

### **Chapter 11 Integrating ACI into vSphere Using VDS 392**

“Do I Know This Already?” Quiz	392
Foundation Topics	394
Understanding Networking in VMware vSphere	394
Understanding vSphere Standard Switches	395
Understanding vSphere Distributed Switches	397
Understanding vSphere System Traffic	397
Impact of vCenter Failure on Production Traffic	399
Understanding Port Bindings in vSphere	400
Understanding Teaming and Failover Policies	400
Understanding VMM Integration	403
Planning vCenter VMM Integrations	403
What Happens After VDS Deployment?	405
Understanding Immediacy Settings	405
Connecting ESXi Servers to the Fabric	407
Configuring Connectivity to ESXi in UCS Domains	407
Integrating ACI into vSphere Using VDS	407
Prerequisites for VMM Integration with vSphere VDS	408
Configuring a VMM Domain Profile	408
Adding ESXi Hosts to a VDS	411
Pushing EPGs to vCenter as Distributed Port Groups	415
Assigning VMs to Distributed Port Groups	417
Less Common VMM Domain Association Settings	418
Enhanced LACP Policy Support	419
Exam Preparation Tasks	422
Review All Key Topics	422
Complete Tables and Lists from Memory	423
Define Key Terms	423

### **Chapter 12 Implementing Service Graphs 424**

“Do I Know This Already?” Quiz	424
Foundation Topics	426
Service Graph Fundamentals	426

Service Graphs as Concatenation of Functions	427
Service Graph Management Models	428
Understanding Network Policy Mode	428
Understanding Service Policy Mode	430
Understanding Service Manager Mode	432
When to Use Service Graphs	434
Choosing an L4–L7 Services Integration Method	435
Understanding Deployment Modes and the Number of BDs Required	435
Deploying Service Graphs for Devices in GoTo Mode	436
Deploying Service Graphs for Devices in GoThrough Mode	437
Deploying Service Graphs for One-Arm Load Balancers	437
Understanding Route Peering	438
Understanding Dynamic Endpoint Attach	439
Understanding Bridge Domain Settings for Service Graphs	439
Understanding Service Graph Rendering	440
Service Graph Implementation Workflow	441
Importing Device Packages	441
Identifying L4–L7 Devices to the Fabric	443
Creating Custom Function Profiles	444
Configuring a Service Graph Template	445
Configuring Device Selection Policies	446
Applying a Service Graph Template	446
Configuring Additional Service Graph Parameters	447
Monitoring Service Graphs and Devices	447
Service Graph Implementation Examples	447
Deploying an Unmanaged Firewall Pair in a Service Graph	447
Deploying Service Graphs for a Firewall in Managed Mode	453
Exam Preparation Tasks	460
Review All Key Topics	460
Complete Tables and Lists from Memory	461
Define Key Terms	461

## **Part V Management and Monitoring**

### **Chapter 13 Implementing Management 462**

“Do I Know This Already?” Quiz	462
Foundation Topics	464
Configuring Management in ACI	464
Understanding Out-of-Band Management Connectivity	464
Understanding In-Band Management Connectivity	465

Deploying In-Band and OOB Management Side by Side	467
Configuring In-Band Management	467
Configuring Access Policies for APIC In-Band Interfaces	468
Configuring the In-Band Management Bridge Domain	469
Configuring In-Band Management IP Addressing	470
Optionally Extending the In-Band Network Out of the Fabric	474
Optionally Setting Up Additional Connectivity	476
Whitelisting Desired Connectivity to and from an In-Band EPG	476
Evaluating APIC Connectivity Preferences	478
Out-of-Band Management Contracts in Review	479
Exam Preparation Tasks	481
Review All Key Topics	481
Memory Tables	481
Define Key Terms	481
<b>Chapter 14 Monitoring ACI Using Syslog and SNMP</b>	<b>482</b>
“Do I Know This Already?” Quiz	482
Foundation Topics	485
Understanding System Messages	485
Forwarding System Messages to Syslog Servers	487
Apply Necessary Contracts to Allow Syslog Forwarding	487
Configuring Syslog Monitoring Destination Groups	492
Configuring Syslog Sources for Desired Monitoring Policies	494
Verify Syslog Forwarding to Desired Syslog Servers	498
Using SNMP in ACI	500
ACI Support for SNMP	501
ACI SNMP Configuration Caveats	502
Configuring ACI for SNMP	502
Apply Necessary Contracts for SNMP	503
Associate an SNMP Policy with a Pod Policy	504
Associate SNMP Contexts with Desired VRF Instances	506
Configure SNMP Monitoring Destination Groups	507
Configure SNMP Sources for All Desired Monitoring Policies	508
Verify SNMP Forwarding to Desired SNMP Servers	509
Exam Preparation Tasks	511
Review All Key Topics	511
Complete Tables and Lists from Memory	512
Define Key Terms	512

**Chapter 15 Implementing AAA and RBAC 514**

- “Do I Know This Already?” Quiz 514
- Foundation Topics 516
- Implementing Role-Based Access Control (RBAC) 516
  - Understanding Security Domains 517
  - Understanding Privileges and Roles 519
  - Creating Local Users and Assigning Access 521
  - Tweaking Roles and User Access 525
  - Custom RBAC Rules 528
  - A Common RBAC Pitfall 531
- Integrating with External AAA Servers 532
  - Configuring ACI for TACACS+ 532
  - Configuring ISE to Authenticate and Authorize Users for ACI 536
  - Expected Cisco AV Pair Formatting for ACI 538
  - Configuring ACI for RADIUS 540
  - Configuring ACI for LDAP 541
  - AAA Authentication Policy Settings 547
  - Regaining Access to the Fabric via Fallback Domain 550
- Exam Preparation Tasks 550
- Review All Key Topics 550
- Complete Tables and Lists from Memory 551
- Define Key Terms 551

**Part VI Operations**

**Chapter 16 ACI Anywhere 552**

- “Do I Know This Already?” Quiz 552
- Foundation Topics 555
- ACI Multi-Site Fundamentals 555
  - Interconnecting ACI Fabrics with ACI Multi-Site 555
  - New ACI Multi-Site Constructs and Configuration Concepts 557
  - Locally Governed Versus MSO-Governed Configurations 557
  - Schemas and Templates in Practice 557
- Building Primary and Disaster Recovery
  - Data Centers with ACI 558
  - Centralized Orchestration and Management of Multiple Fabrics 559
  - Tweaking Broadcast and Stretch Settings on a Per-BD Basis 560
  - Cross-Data Center Ingress Routing Optimizations 561
  - Simultaneous or Independent Policy Deployment to Sites 561

Building Active/Active Data Centers with ACI	562
VMM Integrations Applicable to Multiple Data Centers	563
Stateful-Services Integration in ACI Multi-Pod and Multi-Site	563
Extending ACI to Remote Locations and Public Clouds	564
Extending ACI into Public Clouds with ACI Multi-Site	564
Extending ACI into Bare-Metal Clouds with vPod	564
Integrating Remote Sites into ACI Using Remote Leaf Switches	564
Exam Preparation Tasks	565
Review All Key Topics	565
Memory Tables	565
Define Key Terms	565

## **Part VII Final Preparation**

### **Chapter 17 Final Preparation 566**

Getting Ready	566
Tools for Final Preparation	567
Pearson Cert Practice Test Engine and Questions on the Website	567
Accessing the Pearson Test Prep Software Online	567
Accessing the Pearson Test Prep Software Offline	568
Customizing Your Exams	568
Updating Your Exams	569
Premium Edition	569
Suggested Plan for Final Review/Study	570
Summary	570

Appendix A	Answers to the “Do I Know This Already?” Questions	572
------------	--	-----

Appendix B	CCNP Data Center Application Centric Infrastructure DCACI 300-620 Exam Updates	586
------------	--	-----


Glossary	589
----------	-----

Index	602
-------	-----

## **Online Elements**

Appendix C	Memory Tables
Appendix D	Memory Tables Answer Key
Appendix E	Study Planner
	Glossary

## 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

Welcome to the brave new world of Cisco ACI! This book strives to help you to:

- Understand the benefits of Cisco ACI and unlock its often-untapped potential
- Gain the expertise necessary to design, deploy, and support single-pod ACI fabrics
- Pass the Implementing Cisco Application Centric Infrastructure DCACI 300-620 exam.

The order of these three objectives is very important. An exam candidate who has an in-depth understanding of the fundamentals of a solution not only has an easier time on exam day but is also, arguably, a more capable engineer. That is why this book places an extraordinary amount of emphasis on the fundamentals of ACI rather than tips and tricks, corner-case scenarios, and platform-specific caveats.

This does not mean that this book is lacking in coverage of the DCACI blueprint. On the contrary, this book covers all the exam topics and then some. It does so with plain language and example after example of how particular features can be deployed and how they fit into the bigger picture of enabling ACI to be the data center SDN platform of the future.

## Perspectives on the DCACI 300-620 Exam

In June 2019, Cisco announced that it was making substantial changes to certification products at all levels.

Cisco Application Centric Infrastructure (ACI) is a case in point for why these changes were necessary. Previous Cisco Certified Network Professional (CCNP) certifications followed a monolithic approach that necessitated major changes at both the CCNP and Cisco Certified Network Associate (CCNA) levels before a newer solution like ACI could be retrofitted into an overall curriculum. It commonly took several years for even immensely popular products (like ACI) to make it into the CCNP—and some never made it.

Newer Cisco certifications, on the other hand, take a more modular approach and encourage specialization in solutions most relevant to candidate job roles. If, for example, you are only interested in ACI, you can just take the DCACI 300-620 exam and obtain a specialist designation instead of a CCNA or CCNP. In the case of ACI, the Cisco certification evolution translates into greater depth of coverage without having content dispersed into a daunting number of exams alongside unrelated content.

One challenge that remains is that designing a certification covering all facets of a network product can require candidates to learn several thousand pages of content. This would unnecessarily discourage exam takers. Cisco has therefore divided coverage of ACI into two main exams:

- The DCACI 300-620 exam covers the fundamentals of ACI single-pod fabrics, such as endpoint learning, forwarding, management, monitoring, and basic integrations.

In addition to being a specialization exam, the DCACI 300-620 exam also counts as a concentration toward the CCNP Data Center certification.

- The Implementing Cisco Application Centric Infrastructure—Advanced (300-630 DCACIA) exam addresses the implementation of more advanced ACI architectures, such as ACI Multi-Pod and ACI Multi-Site. It also covers route leaking, advanced contract implementation, and service insertion via policy-based redirect (PBR).

The DCACI 300-620 exam addresses at least 70% of the concepts a typical ACI engineer deals with on a day-to-day basis and provides an excellent on ramp for engineers seeking to build the foundational knowledge necessary to implement the most complex of ACI designs.

As you might have noticed, one essential topic still missing from the blueprints of these two exams is network automation. Cisco has released a dedicated exam for data center automation that includes ACI, called the Automating and Programming Cisco Data Center Solutions (300-635 DCAUTO) exam. Therefore, this book does not cover network automation, opting instead to serve as a tool to help engineers build a solid foundation in ACI.

## Who Should Read This Book?

This book has been written with you in mind!

For engineers new to ACI, this book attempts to demystify the complex language of ACI by using unambiguous wording and a wide range of examples. It includes detailed configuration steps and can even be used as a lab guide. This book recognizes ACI newcomers as a significant part of its target audience and has been written to be the most comprehensive and up-to-date book on ACI while also being the easiest to read.

For more advanced engineers who have experience with ACI but need a guide to prepare for the DCACI 300-620 exam or to address knowledge gaps, this book is comprehensive enough to address the topics on the exam while also taking a look under the hood of ACI to enable these engineers to better appreciate how ACI works.

This book can also help network automation engineers build a solid foundation of ACI design and implementation concepts. Even though this book does not cover automation in ACI, it does address, in detail, how some of the most significant and often-used objects interact with one another.

This book is not an introduction to general networking and does expect readers to understand the basics of switching and routing. But this book does not assume that readers have any prior knowledge of ACI or even basic knowledge of data center overlay technologies. For this reason, this book can be used as a network engineer's first introduction to ACI.

## 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.

To access the companion website, start by establishing a login at [www.ciscopress.com](http://www.ciscopress.com) and registering your book. To do so, simply go to [www.ciscopress.com/register](http://www.ciscopress.com/register) and enter the ISBN of the print book: 9780136602668. 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 (PTP) App

You have two options for installing and using the Pearson Test Prep application: a web app and a desktop app. To use the Pearson Test Prep application, start by finding the access 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 access 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](http://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 by Amazon.
- **Other bookseller eBooks:** Note that if you purchase an eBook 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 access code because it is the only means with which you can access the QA content with the book.

Once you have the access code, to find instructions about both the Pearson Test Prep web app and the desktop app, follow these steps:

- Step 1.** Open this book's companion website.
- Step 2.** Click the **Practice Exams** button.
- Step 3.** Follow the instructions listed there for installing the desktop app and for using the web app.

If you want to use the web app only at this point, just navigate to [www.pearsonestprep.com](http://www.pearsonestprep.com), establish a free login if you do not already have one, and register this book's practice tests using the access code you just found. The process should take only a couple of minutes.

**NOTE** Amazon eBook (Kindle) customers: It is easy to miss Amazon's email that lists your Pearson Test Prep 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 eBook customers: As of the time of publication, only the publisher and Amazon supply Pearson Test Prep access codes when you purchase their eBook 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 are interested in learning. Chapters 1 through 16 cover topics that are relevant to the DCACI 300-620 exam:

- **Chapter 1, “The Big Picture: Why ACI?”:** This chapter describes some of the challenges inherent in traditional network switches and routers and how ACI is able to solve these challenges.
- **Chapter 2, “Understanding ACI Hardware and Topologies”:** This chapter addresses the prominent ACI topologies in use today as well as ACI hardware platforms.
- **Chapter 3, “Initializing an ACI Fabric”:** This chapter covers planning parameters that are important for fabric initialization, the fabric initialization process itself, and some common post-initialization tasks, such as assignment of static out-of-band IP addresses to ACI nodes as well as making fabric backups and restoring configurations.
- **Chapter 4, “Exploring ACI”:** This chapter explores ACI access methods, the ACI object model, and some basic fabric health monitoring and fault management concepts.
- **Chapter 5, “Tenant Building Blocks”:** This chapter examines from a conceptual viewpoint the various objects present under the tenant hierarchy and how they relate to one another.
- **Chapter 6, “Access Policies”:** This chapter examines the concepts behind configuration of switch downlinks to servers, external switches, and routers. It also addresses how switch port configurations tie in with the tenant hierarchy.
- **Chapter 7, “Implementing Access Policies”:** This chapter focuses on configuration of individual switch ports, port channels, vPCs, and fabric extenders (FEX) down to servers, external switches, and routers.

- **Chapter 8, “Implementing Tenant Policies”:** This chapter covers endpoint learning and forwarding in ACI as well as deployment of multitier applications and the enforcement of contracts to whitelist data center communications.
- **Chapter 9, “L3Outs”:** This chapter examines implementation of ACI route peering with outside Layer 3 devices as well as inbound and outbound route filtering.
- **Chapter 10, “Extending Layer 2 Outside ACI”:** This chapter addresses ACI Layer 2 connectivity with non-ACI switches and interaction with Spanning Tree Protocol. It also provides basic coverage of network migrations into and out of ACI.
- **Chapter 11, “Integrating ACI into vSphere Using VDS”:** This chapter addresses implementation of the most popular ACI integration and why it is important.
- **Chapter 12, “Implementing Service Graphs”:** This chapter tackles the introduction of firewalls and load balancers into ACI fabrics using service graphs.
- **Chapter 13, “Implementing Management”:** This chapter revisits the topic of in-band and out-of-band management in ACI and dives into the implementation of in-band management.
- **Chapter 14, “Monitoring ACI Using Syslog and SNMP”:** This chapter covers how ACI can forward faults and other monitoring information to syslog or SNMP servers.
- **Chapter 15, “Implementing AAA and RBAC”:** This chapter dives into role-based access control and how multitenancy can be enforced from a management perspective.
- **Chapter 16, “ACI Anywhere”:** This chapter provides a primer on additional ACI solutions within the ACI portfolio, including ACI Multi-Pod and ACI Multi-Site, which allow extension of ACI policies between data centers, between remote locations, and between public clouds.

## How to Use This Book

The questions for each certification exam are a closely guarded secret. However, Cisco has published exam blueprints that list the topics you must know to *successfully* complete the exams. Table I-1 lists the exam topics listed in the DCACI 300-620 exam blueprint along with a reference to the book chapter that covers each topic. These are the same topics you should be proficient in when designing and implementing ACI fabrics in the real world.

**Table I-1** CCNP DCACI 300-620 Exam Topics and Chapter References

Exam Topic	Chapter(s) in Which Topic Is Covered
1.0 ACI Fabric Infrastructure	
1.1 Describe ACI topology and hardware	2
1.2 Describe ACI Object Model	4
1.3 Utilize faults, event record, and audit log	4
1.4 Describe ACI fabric discovery	3

Exam Topic	Chapter(s) in Which Topic Is Covered
1.5 Implement ACI policies 1.5.a access 1.5.b fabric	5, 6, 7
1.6 Implement ACI logical constructs 1.6.a tenant 1.6.b application profile 1.6.c VRF 1.6.d bridge domain (unicast routing, Layer 2 unknown hardware proxy, ARP flooding) 1.6.e endpoint groups (EPG) 1.6.f contracts (filter, provider, consumer, reverse port filter, VRF enforced)	5, 8, 9, 10
2.0 ACI Packet Forwarding	
2.1 Describe endpoint learning	8
2.2 Implement bridge domain configuration knob (unicast routing, Layer 2 unknown hardware proxy, ARP flooding)	8
3.0 External Network Connectivity	
3.1 Implement Layer 2 out (STP/MCP basics)	10
3.2 Implement Layer 3 out (excludes transit routing and VRF route leaking)	9
4.0 Integrations	
4.1 Implement VMware vCenter DVS integration	11
4.2 Describe resolution immediacy in VMM	11
4.3 Implement service graph (managed and unmanaged)	12
5.0 ACI Management	
5.1 Implement out-of-band and in-band	3, 13
5.2 Utilize syslog and snmp services	14
5.3 Implement configuration backup (snapshot/config import export)	3
5.4 Implement AAA and RBAC	15
5.5 Configure an upgrade	3
6.0 ACI Anywhere	
6.1 Describe multipod	16
6.2 Describe multisite	16

Each version of the exam may emphasize different topics, and some topics are rather broad and generalized. The goal of this book is to provide 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 CCNP data center engineer.

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

Note that as ACI features and solutions continue to evolve, 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, you should check Cisco.com to verify the current list of topics to ensure that you are prepared to take the exam. You can view the current exam topics on any current Cisco certification exam by visiting the Cisco.com website and choosing Menu > Training & Events and 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/9780136602668>. It's a good idea to check the website a couple weeks before taking the exam to be sure you have up-to-date content.

## Figure Credits

Figure 11-03: Screenshot of a VMkernel adapter with management services enabled © 2020 VMware, Inc

Figure 11-4: Screenshot of selecting Ephemeral - No Binding as the port binding type © 2020 VMware, Inc

Figure 11-5: Screenshot of teaming and failover settings for port groups © 2020 VMware, Inc

Figure 11-6: Screenshot of data center, cluster, and ESXi host hierarchy in vCenter © 2020 VMware, Inc

Figure 11-11: Screenshot of validating VDS creation in vCenter © 2020 VMware, Inc

Figure 11-12: Screenshot of navigating to the Add and Manage Hosts Wizard in vCenter © 2020 VMware, Inc

Figure 11-13: Screenshot of selecting add hosts © 2020 VMware, Inc

Figure 11-14: Screenshot of clicking new hosts © 2020 VMware, Inc

Figure 11-15: Screenshot of choosing the hosts to add on the Select New Hosts page © 2020 VMware, Inc

Figure 11-16: Screenshot of assigning uplinks to a VDS © 2020 VMware, Inc

Figure 11-17: Screenshot of the Manage VMkernel Adapters page © 2020 VMware, Inc

Figure 11-18: Screenshot of the Manage VM Networking page © 2020 VMware, Inc

Figure 11-19: Screenshot of confirming the addition of ESXi hosts to the VDS © 2020 VMware, Inc

Figure 11-22: Screenshot of verifying distributed port group generation in vCenter © 2020 VMware, Inc

Figure 11-23: Screenshot of reassigning a VM vNIC to a distributed port group © 2020 VMware, Inc

Figure 11-25: Screenshot of verifying the result of custom EPG naming and delimiter modification © 2020 VMware, Inc

Figure 11-26: Screenshot of verifying the result of active uplinks and standby uplinks settings © 2020 VMware, Inc

Figure 11-28: Screenshot of assigning ESXi host uplinks to a link aggregation group © 2020 VMware, Inc

Figure 11-30: Screenshot of verifying distributed port group mapping to uplinks © 2020 VMware, Inc

*This page intentionally left blank*



## CHAPTER 2

# Understanding ACI Hardware and Topologies

### This chapter covers the following topics:

**ACI Topologies and Components:** This section describes the key hardware components and acceptable topologies for ACI fabrics.

**APIC Clusters:** This section covers available APIC hardware models and provides an understanding of APIC cluster sizes and failover implications.

**Spine Hardware:** This section addresses available spine hardware options.

**Leaf Hardware:** This section outlines the leaf platforms available for deployment in ACI fabrics.

This chapter covers the following exam topics:

- 1.1 Describe ACI topology and hardware
- 6.1 Describe Multi-Pod
- 6.2 Describe Multi-Site

ACI is designed to allow small and large enterprises and service providers to build massively scalable data centers using a relatively small number of very flexible topologies.

This chapter details the topologies with which an ACI fabric can be built or extended. Understanding supported ACI topologies helps guide decisions on target-state network architecture and hardware selection.

Each hardware component in an ACI fabric performs a specific set of functions. For example, leaf switches enforce security rules, and spine switches track all endpoints within a fabric in a local database.

But not all ACI switches are created equally. Nor are APICs created equally. This chapter therefore aims to provide a high-level understanding of some of the things to consider when selecting hardware.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows 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 2-1 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?’ Questions.”

**Table 2-1** “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
ACI Topologies and Components	1–5
APIC Clusters	6
Spine Hardware	7, 8
Leaf Hardware	9, 10

**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. An ACI fabric is being extended to a secondary location to replace two top-of-rack switches and integrate a handful of servers into a corporate ACI environment. Which solution should ideally be deployed at the remote location if the deployment of new spines is considered cost-prohibitive and direct fiber links from the main data center cannot be dedicated to this function?
  - a. ACI Multi-Site
  - b. ACI Remote Leaf
  - c. ACI Multi-Tier
  - d. ACI Multi-Pod
2. Which of the following is a requirement for a Multi-Pod IPN that is not needed in an ACI Multi-Site ISN?
  - a. Increased MTU support
  - b. OSPF support on last-hop routers connecting to ACI spines
  - c. End-to-end IP connectivity
  - d. Multicast PIM-Bidir
3. Which of the following connections would ACI definitely block?
  - a. APIC-to-leaf cabling
  - b. Leaf-to-leaf cabling
  - c. Spine-to-leaf cabling
  - d. Spine-to-spine cabling
4. Which of the following are valid reasons for ACI Multi-Site requiring more specialized spine hardware? (Choose all that apply.)
  - a. Ingress replication of BUM traffic
  - b. IP fragmentation
  - c. Namespace normalization
  - d. Support for PIM-Bidir for multicast forwarding

5. Which of the following options best describes border leaf switches?
  - a. Border leaf switches provide Layer 2 and 3 connectivity to outside networks.
  - b. Border leaf switches connect to Layer 4–7 service appliances, such as firewalls and load balancers.
  - c. Border leaf switches are ACI leaf switches that connect to servers.
  - d. Border leaf switches serve as the border between server network traffic and FCoE storage traffic.
6. Which of the following statements is accurate?
  - a. A three-node M3 cluster of APICs can scale up to 200 leaf switches.
  - b. Sharding is a result of the evolution of what is called horizontal partitioning of databases.
  - c. The number of shards distributed among APICs for a given attribute is directly correlated to the number of APICs deployed.
  - d. A standby APIC actively synchronizes with active APICs and has a copy of all attributes within the APIC database at all times.
7. Out of the following switches, which are spine platforms that support ACI Multi-Site? (Choose all that apply.)
  - a. Nexus 93180YC-EX
  - b. Nexus 9364C
  - c. Nexus 9736C-FX line card
  - d. Nexus 9396PX
8. Which of the following is a valid reason for upgrading a pair of Nexus 9336PQ ACI switches to second-generation Nexus 9332C spine hardware? (Choose all that apply.)
  - a. Namespace normalization for ACI Multi-Site support
  - b. Support for 40 Gbps leaf-to-spine connectivity
  - c. Support for CloudSec
  - d. Support for ACI Multi-Pod
9. True or false: The Nexus 93180YC-FX leaf switch supports MACsec.
  - a. True
  - b. False
10. Which of the following platforms is a low-cost option for server CIMC and other low-bandwidth functions that rely on RJ-45 connectivity?
  - a. Nexus 9336C-FX2
  - b. Nexus 93180YC-FX
  - c. Nexus 9332C
  - d. Nexus 9348GC-FXP

## Foundation Topics

### ACI Topologies and Components

Like many other current data center fabrics, ACI fabrics conform to a Clos-based leaf-and-spine topology.


In ACI, leaf and spine switches are each responsible for different functions. Together, they create an architecture that is highly standardized across deployments. Cisco has introduced several new connectivity models and extensions for ACI fabrics over the years, but none of these changes break the core ACI topology that has been the standard from day one. Any topology modifications introduced in this section should therefore be seen as slight enhancements that help address specific use cases and not as deviations from the standard ACI topology.

### Clos Topology

In his 1952 paper titled “A Study of Non-blocking Switching Networks,” Bell Laboratories researcher Charles Clos formalized how multistage telephone switching systems could be built to forward traffic, regardless of the number of calls served by the overall system.


The mathematical principles proposed by Clos also help address the challenge of needing to build highly scalable data centers using relatively low-cost switches.

Figure 2-1 illustrates a three-stage Clos fabric consisting of one layer for ingress traffic, one layer for egress traffic, and a central layer for forwarding traffic between the layers. Multi-stage designs such as this can result in networks that are not oversubscribed or that are very close to not being oversubscribed.



**Figure 2-1** *Conceptual View of a Three-Stage Clos Topology*

Modern data center switches forward traffic at full duplex. Therefore, there is little reason to depict separate layers for ingress and egress traffic. It is possible to fold the top layer from the three-tier Clos topology in Figure 2-1 into the bottom layer to achieve what the industry refers to as a “folded” Clos topology, illustrated in Figure 2-2.



**Figure 2-2** *Folded Clos Topology*

As indicated in Figure 2-2, a leaf switch is an ingress/egress switch. A spine switch is an intermediary switch whose most critical function is to perform rapid forwarding of traffic between leaf switches. Leaf switches connect to spine switches in a full-mesh topology.


**NOTE** At first glance, a three-tier Clos topology may appear to be similar to the traditional three-tier data center architecture. However, there are some subtle differences. First, there are no physical links between leaf switches in the Clos topology. Second, there are no physical links between spine switches. The elimination of cross-links within each layer simplifies network design and reduces control plane complexity.

### Standard ACI Topology

An ACI fabric forms a Clos-based spine-and-leaf topology and is usually depicted using two rows of switches. Depending on the oversubscription and overall network throughput requirements, the number of spines and leaf switches will be different in each ACI fabric.

**NOTE** In the context of the Implementing Cisco Application Centric Infrastructure DCACI 300-620 exam, it does not matter whether you look at a given ACI fabric as a two-tiered Clos topology or as a three-tiered folded Clos topology. It is common for the standard ACI topology to be referred to as a two-tier spine-and-leaf topology.

Figure 2-3 shows the required components and cabling for an ACI fabric. Inheriting from its Clos roots, no cables should be connected between ACI leaf switches. Likewise, ACI spines being cross-cabled results in ACI disabling the cross-connected ports. While the topology shows a full mesh of cabling between the spine-and-leaf layers, a fabric can operate without a full mesh. However, a full mesh of cables between layers is still recommended.



**Figure 2-3** Standard ACI Fabric Topology

In addition to optics and cabling, the primary hardware components required to build an ACI fabric are as follows:

**Key Topic**

- Application Policy Infrastructure Controllers (APICs):** The APICs are the brains of an ACI fabric and serve as the single source of truth for configuration within the fabric. A clustered set of (typically three) controllers attaches directly to leaf switches and provides management, policy programming, application deployment, and health monitoring for an ACI fabric. Note in Figure 2-3 that APICs are not in the data path or the forwarding topology. Therefore, the failure of one or more APICs does not halt packet forwarding. An ACI fabric requires a minimum of one APIC, but an ACI fabric with one APIC should be used only for lab purposes.
- Spine switches:** ACI spine switches are Clos intermediary switches that have a number of key functions. They exchange routing updates with leaf switches via Intermediate System-to-Intermediate System (IS-IS) and perform rapid forwarding of packets between leaf switches. They provide endpoint lookup services to leaf switches through the Council of Oracle Protocol (COOP). They also handle route reflection to leaf switches using Multiprotocol BGP (MP-BGP), allowing external routes to be distributed across the fabric regardless of the number of tenants. (All three of these are control plane protocols and are covered in more detail in future chapters.) Spine switches also serve as roots for multicast trees within a fabric. By default, all spine switch interfaces besides the mgmt0 port are configured as fabric ports. *Fabric ports* are the interfaces that are used to interconnect spine and leaf switches within a fabric.
- Leaf switches:** Leaf switches are the ingress/egress points for traffic into and out of an ACI fabric. As such, they are the connectivity points for endpoints, including servers and appliances, into the fabric. Layer 2 and 3 connectivity from the outside world into an ACI fabric is also typically established via leaf switches. ACI security policy enforcement occurs on leaf switches. Each leaf switch has a number of high-bandwidth uplink ports preconfigured as fabric ports.

In addition to the components mentioned previously, optional hardware components that can be deployed alongside an ACI fabric include fabric extenders (FEX). Use of FEX solutions in ACI is not ideal because leaf hardware models currently on the market are generally low cost and feature heavy compared to FEX technology.

FEX attachment to ACI is still supported to allow for migration of brownfield gear into ACI fabrics. The DCACI 300-620 exam does not cover specific FEX model support, so neither does this book.

**NOTE** There are ways to extend an ACI fabric into a virtualized environment by using ACI Virtual Edge (AVE) and Application Virtual Switch (AVS). These are software rather than hardware components and are beyond the scope of the DCACI 300-620 exam.

Engineers may sometimes dedicate two or more leaf switches to a particular function. Engineers typically evaluate the following categories of leaf switches as potential options for dedicating hardware:

### Key Topic

- **Border Leaf:** *Border leaf* switches provide Layer 2 and 3 connectivity between an ACI fabric and the outside world. Border leaf switches are sometimes points of policy enforcement between internal and external endpoints.
- **Service Leaf:** *Service leaf* switches are leaf switches that connect to Layer 4–7 service appliances, such as firewalls and load balancers.
- **Compute Leaf:** *Compute leaf* switches are ACI leaf switches that connect to servers. Compute leaf switches are points of policy enforcement when traffic is being sent between local endpoints.
- **IP Storage Leaf:** *IP storage leaf* switches are ACI leaf switches that connect to IP storage systems. IP storage leaf switches can also be points of policy enforcement for traffic to and from local endpoints.


There are scalability benefits associated with dedicating leaf switches to particular functions, but if the size of the network does not justify dedicating leaf switches to a function, consider at least dedicating a pair of leaf switches as border leaf switches. Service leaf functionality can optionally be combined with border leaf functionality, resulting in the deployment of a pair (or more) of collapsed border/service leaf switches in smaller environments.

Cisco publishes a Verified Scalability Guide for each ACI code release. At the time of this writing, 500 is considered the maximum number of leaf switches that can be safely deployed in a single fabric that runs on the latest code.

## ACI Stretched Fabric Topology

A *stretched ACI fabric* is a partially meshed design that connects ACI leaf and spine switches distributed in multiple locations. The stretched ACI fabric design helps lower deployment costs when full-mesh cable runs between all leaf and spine switches in a fabric tend to be cost-prohibitive.

Figure 2-4 shows a stretched ACI fabric across two sites.



**Figure 2-4** ACI Stretched Fabric Topology

A stretched fabric amounts to a single administrative domain and a single availability zone. Because APICs in a stretched fabric design tend to be spread across sites, cross-site latency is an important consideration. APIC clustering has been validated across distances of 800 kilometers between two sites.

A new term introduced in Figure 2-4 is *transit leaf*. A *transit leaf* is a leaf switch that provides connectivity between two sites in a stretched fabric design. Transit leaf switches connect to spine switches in both sites. No special configuration is required for transit leaf switches. At least one transit leaf switch must be provisioned in each site for redundancy reasons.

While stretched fabrics simplify extension of an ACI fabric, this design does not provide the benefits of newer topologies such as ACI Multi-Pod and ACI Multi-Site and stretched fabrics are therefore no longer commonly deployed or recommended.

## ACI Multi-Pod Topology

### Key Topic


The *ACI Multi-Pod* topology is a natural evolution of the ACI stretched fabric design in which spine and leaf switches are divided into pods, and different instances of IS-IS, COOP, and MP-BGP protocols run inside each pod to enable a level of control plane fault isolation.

Spine switches in each pod connect to an interpod network (IPN). Pods communicate with one another through the IPN. Figure 2-5 depicts an ACI Multi-Pod topology.

### Key Topic

An ACI Multi-Pod IPN has certain requirements that include support for OSPF, end-to-end IP reachability, DHCP relay capabilities on the last-hop routers that connect to spines in each pod, and an increased maximum transmission unit (MTU). In addition, a Multi-Pod IPN needs to support forwarding of multicast traffic (PIM-Bidir) to allow the replication of broadcast, unknown unicast, and multicast (BUM) traffic across pods.

One of the most significant use cases for ACI Multi-Pod is active/active data center design. Although ACI Multi-Pod supports a maximum round-trip time latency of 50 milliseconds between pods, most Multi-Pod deployments are often built to achieve active/active functionality and therefore tend to have latencies of less than 5 milliseconds.



**Figure 2-5** ACI Multi-Pod Topology

**NOTE** Another solution that falls under the umbrella of ACI Multi-Pod is Virtual Pod (vPod). ACI vPod is not a new topology per se. It is an extension of a Multi-Pod fabric in the form of a new pod at a remote location where at least two ESXi servers are available, and deployment of ACI hardware is not desirable. ACI vPod components needed at the remote site for this solution include virtual spine (vSpine) appliances, virtual leaf (vLeaf) appliances, and the Cisco ACI Virtual Edge. ACI vPod still requires a physical ACI footprint since vPod is managed by the overall Multi-Pod APIC cluster.

On the issue of scalability, it should be noted that as of the time of writing, 500 is the maximum number of leaf switches that can be safely deployed within a single ACI fabric. However, the Verified Scalability Guide for the latest code revisions specifies 400 as the absolute maximum number of leaf switches that can be safely deployed in each pod. Therefore, for a fabric to reach its maximum supported scale, leaf switches should be deployed across at least 2 pods within a Multi-Pod fabric. Each pod supports deployment of 6 spines, and each Multi-Pod fabric currently supports the deployment of up to 12 pods.

Chapter 16, “ACI Anywhere,” covers ACI Multi-Pod in more detail. For now, understand that Multi-Pod is functionally a single fabric and a single availability zone, even though it does not represent a single network failure domain.

## ACI Multi-Site Topology




*ACI Multi-Site* is a solution that interconnects multiple ACI fabrics for the purpose of homogenous policy deployment across ACI fabrics, homogenous security policy deployment across on-premises ACI fabrics and public clouds, and cross-site stretched subnet capabilities, among others.

**Key Topic**

In an ACI Multi-Site design, each ACI fabric has its own dedicated APIC cluster. A clustered set of three nodes called Multi-Site Orchestrator (MSO) establishes API calls to each fabric independently and can configure tenants within each fabric with desired policies.

**NOTE** Nodes forming an MSO cluster have traditionally been deployed as VMware ESXi virtual machines (VMs). Cisco has recently introduced the ability to deploy an MSO cluster as a distributed application (.aci format) on Cisco Application Services Engine (ASE). Cisco ASE is a container-based solution that provides a common platform for deploying and managing Cisco data center applications. ASE can be deployed in three form factors: a physical form factor consisting of bare-metal servers, a virtual machine form factor for on-premises deployments via ESXi or Linux KVM hypervisors, and a virtual machine form factor deployable within a specific Amazon Web Services (AWS) region.

Figure 2-6 shows an ACI Multi-Site topology that leverages a traditional VM-based MSO cluster.



**Figure 2-6** ACI Multi-Site Topology

**Key Topic**

As indicated in Figure 2-6, end-to-end communication between sites in an ACI Multi-Site design requires the use of an intersite network (ISN). An ACI Multi-Site ISN faces less stringent requirements compared to ACI Multi-Pod IPNs. In an ISN, end-to-end IP connectivity between spines across sites, OSPF on the last-hop routers connecting to the spines, and increased MTU support allowing VXLAN-in-IP encapsulation are all still required. However, ACI Multi-Site does not dictate any cross-site latency requirements, nor does it require support for multicast or DHCP relay within the ISN.

ACI Multi-Site does not impose multicast requirements on the ISN because ACI Multi-Site has been designed to accommodate larger-scale ACI deployments that may span the globe. It is not always feasible or expected for a company that has a global data center footprint to also have a multicast backbone spanning the globe and between all data centers.

**Key  
Topic**

Due to the introduction of new functionalities that were not required in earlier ACI fabrics, Cisco introduced a second generation of spine hardware. Each ACI fabric within an ACI Multi-Site design requires at least one second-generation or newer piece of spine hardware for the following reasons:

- **Ingress replication of BUM traffic:** To accommodate BUM traffic forwarding between ACI fabrics without the need to support multicast in the ISN, Multi-Site-enabled spines perform ingress replication of BUM traffic. This function is supported only on second-generation spine hardware.
- **Cross-fabric namespace normalization:** Each ACI fabric has an independent APIC cluster and therefore an independent brain. When policies and parameters are communicated between fabrics in VXLAN header information, spines receiving cross-site traffic need to have a way to swap remotely significant parameters, such as VXLAN network identifiers (VNIDs), with equivalent values for the local site. This function, which is handled in hardware and is called *namespace normalization*, requires second-generation or newer spines.

Note that in contrast to ACI Multi-Site, ACI Multi-Pod *can* be deployed using first-generation spine switches.


For ACI Multi-Site deployments, current verified scalability limits published by Cisco suggest that fabrics with stretched policy requirements that have up to 200 leaf switches can be safely incorporated into ACI Multi-Site. A single ACI Multi-Site deployment can incorporate up to 12 fabrics as long as the total number of leaf switches in the deployment does not surpass 1600.

Each fabric in an ACI Multi-Site design forms a separate network failure domain and a separate availability zone.

## ACI Multi-Tier Architecture

Introduced in Release 4.1, ACI Multi-Tier provides the capability for vertical expansion of an ACI fabric by adding an extra layer or tier of leaf switches below the standard ACI leaf layer.

With the Multi-Tier enhancement, the standard ACI leaf layer can also be termed the Tier 1 leaf layer. The new layer of leaf switches that are added to vertically expand the fabric is called the Tier 2 leaf layer. Figure 2-7 shows these tiers. APICs, as indicated, can attach to either Tier 1 or Tier 2 leaf switches.




**Figure 2-7** ACI Multi-Tier Topology

**NOTE** The topology shown in Figure 2-7 goes against the requirement outlined earlier in this chapter, in the section “Standard ACI Topology,” *not* to cross-connect leaf switches. The ACI Multi-Tier architecture is an exception to this rule. Leaf switches within each tier, however, still should never be cross-connected.

An example of a use case for ACI Multi-Tier is the extension of an ACI fabric across data center halls or across buildings that are in relatively close proximity while minimizing long-distance cabling and optics requirements. Examine the diagram in Figure 2-8. Suppose that an enterprise data center has workloads in an alternate building. In this case, the company can deploy a pair of Tier 1 leaf switches in the new building and expand the ACI fabric to the extent needed within the building by using a Tier 2 leaf layer. Assuming that 6 leaf switches would have been required to accommodate the port requirements in the building, as Figure 2-8 suggests, directly cabling these 6 leaf switches to the spines as Tier 1 leaf switches would have necessitated 12 cross-building cables. However, the use of an ACI Multi-Tier design enables the deployment of the same number of switches using 4 long-distance cable runs.

ACI Multi-Tier can also be an effective solution for use within data centers in which the cable management strategy is to minimize inter-row cabling and relatively low-bandwidth requirements exist for top-of-rack switches. In such a scenario, Tier 1 leaf switches can be deployed end-of-row, and Tier 2 leaf switches can be deployed top-of-rack.



**Figure 2-8** Extending an ACI Fabric by Using ACI Multi-Tier in an Alternative Location

**NOTE** ACI Multi-Tier *might not* be a suitable solution if the amount of bandwidth flowing upstream from Tier 2 leaf switches justifies the use of dedicated uplinks to spines.

Not all ACI switch platforms support Multi-Tier functionality.

## Remote Leaf Topology

### Key Topic

For remote sites in which data center endpoints may be deployed but their number and significance do not justify the deployment of an entirely new fabric or pod, the ACI *Remote Leaf* solution can be used to extend connectivity and ensure consistent policies between the main data center and the remote site. With such a solution, leaf switches housed at the remote site communicate with spines and APICs at the main data center over a generic IPN. Each Remote Leaf switch can be bound to a single pod.


There are three main use cases for Remote Leaf deployments:

- **Satellite/small colo data centers:** If a company has a small data center consisting of several top-of-rack switches and the data center may already have dependencies on a main data center, this satellite data center can be integrated into the main data center by using the Remote Leaf solution.
- **Data center extension and migrations:** Cross-data center migrations that have traditionally been done through Layer 2 extension can instead be performed by deploying a pair of Remote Leafs in the legacy data center. This approach often has cost benefits compared to alternative Layer 2 extension solutions if there is already an ACI fabric in the target state data center.
- **Telco 5G distributed data centers:** Telecom operators that are transitioning to more distributed mini data centers to bring services closer to customers but still desire centralized management and consistent policy deployment across sites can leverage Remote Leaf for these mini data centers.

In addition to these three main use cases, disaster recovery (DR) is sometimes considered a use case for Remote Leaf deployments, even though DR is a use case more closely aligned with ACI Multi-Site designs.

In a Remote Leaf solution, the APICs at the main data center deploy policy to the Remote Leaf switches as if they were locally connected.

Figure 2-9 illustrates a Remote Leaf solution.



**Figure 2-9** Remote Leaf Topology and IPN Requirements

IPN requirements for a Remote Leaf solution are as follows:

- **MTU:** The solution must support an end-to-end MTU that is at least 100 bytes higher than that of the endpoint source traffic. Assuming that 1500 bytes has been configured for data plane MTU, Remote Leaf can be deployed using a minimum MTU of 1600 bytes. An IPN MTU this low, however, necessitates that ACI administrators lower the ACI fabricwide control plane MTU, which is 9000 bytes by default.
- **Latency:** Up to 300 milliseconds latency between the main data center and remote location is acceptable.
- **Bandwidth:** Remote Leaf is supported with a minimum IPN bandwidth of 100 Mbps.
- **VTEP reachability:** A Remote Leaf switch logically associates with a single pod if integrated into a Multi-Pod solution. To make this association possible, the Remote Leaf should be able to route traffic over the IPN to the VTEP pool of the associated pod. Use of a dedicated VRF for IPN traffic is recommended where feasible.

- **APIC infra IP reachability:** A Remote Leaf switch needs IP connectivity with all APICs in a Multi-Pod cluster at the main data center. If an APIC has assigned itself IP addresses from a VTEP range different than the pod VTEP pool, the additional VTEP addresses need to also be advertised over the IPN.
- **OSPF support on upstream routers:** Routers northbound of both the Remote Leaf switches and the spine switches need to support OSPF and must be able to encapsulate traffic destined to directly attached ACI switches using VLAN 4. This requirement exists only for directly connected devices and does not extend end-to-end in the IPN.
- **DHCP relay:** The upstream router directly connected to Remote Leaf switches needs to enable DHCP relay to relay DHCP packets to the APIC IP addresses in the infra tenant. The DHCP relay configuration needs to be applied on the VLAN 4 subinterface or SVI.

Note that unlike a Multi-Pod IPN, a Remote Leaf IPN does not require Multicast PIM-Bidir support. This is because the Remote Leaf solution uses headend replication (HER) tunnels to forward BUM traffic between sites.

In a Remote Leaf design, traffic between known local endpoints at the remote site is switched directly, whether physically or virtually. Any traffic whose destination is in ACI but is unknown or not local to the remote site is forwarded to the main data center spines.

**NOTE** Chapter 16 details MTU requirements for IPN and ISN environments for ACI Multi-Pod and ACI Multi-Site. It also covers how to lower control plane and data plane MTU values within ACI if the IPN or ISN does not support high MTU values. Although it does not cover Remote Leaf, the same general IPN MTU concepts apply.

Not all ACI switches support Remote Leaf functionality. The current maximum verified scalability number for Remote Leaf switches is 100 per fabric.

## APIC Clusters

The ultimate size of an APIC cluster should be directly proportionate to the size of the Cisco ACI deployment. From a management perspective, any active APIC controller in a cluster can service any user for any operation. Controllers can be transparently added to or removed from a cluster.

### Key Topic

APICs can be purchased either as physical or virtual appliances. Physical APICs are 1 rack unit (RU) Cisco C-Series servers with ACI code installed and come in two different sizes: M for medium and L for large. In the context of APICs, “size” refers to the scale of the fabric and the number of endpoints. Virtual APICs are used in ACI mini deployments, which consist of fabrics with up to two spine switches and four leaf switches.

### Key Topic

As hardware improves, Cisco releases new generations of APICs with updated specifications. At the time of this writing, Cisco has released three generations of APICs. The first generation of APICs (M1/L1) shipped as Cisco UCS C220 M3 servers. Second-generation APICs (M2/L2) were Cisco UCS C220 M4 servers. Third-generation APICs (M3/L3) are shipping as UCS C220 M5 servers.

Table 2-2 details specifications for current M3 and L3 APICs.

**Table 2-2** M3 and L3 APIC Specifications

Component	M3	L3
Processor	2x 1.7 GHz Xeon scalable 3106/85W 8C/11MB cache/DDR4 2133MHz	2x 2.1 GHz Xeon scalable 4110/85W 8C/11MB cache/DDR4 2400MHz
Memory	6x 16 GB DDR4-2666-MHz RDIMM/PC4-21300/single rank/ x4/1.2v	12x 16 GB DDR4-2666-MHz RDIMM/PC4-21300/single rank/ x4/1.2v
Hard drive	2x 1 TB 12G SAS 7.2K RPM SFF HDD	2x 2.4 TB 12G SAS 10K RPM SFF HDD (4K)
Network cards	1x Cisco UCS VIC 1455 Quad Port 10/25G SFP28 CNA PCIE	1x Cisco UCS VIC 1455 Quad Port 10/25G SFP28 CNA PCIE

Note in Table 2-2 that the only differences between M3 and L3 APICs are the sizes of their CPUs, memory, and hard drives. This is because fabric growth necessitates that increased transaction rates be supported, which drives up compute requirements.

Table 2-3 shows the hardware requirements for virtual APICs.

**Table 2-3** Virtual APIC Specifications

Component	Virtual APIC
Processor	8 vCPUs
Memory	32 GB
Hard drive*	300 GB HDD 100 GB SSD
Supported ESXi hypervisor version	6.5 or above

\* A VM is deployed with two HDDs.

## APIC Cluster Scalability and Sizing

APIC cluster hardware is typically purchased from Cisco in the form of a bundle. An APIC bundle is a collection of one or more physical or virtual APICs, and the bundle that needs to be purchased depends on the desired target state scalability of the ACI fabric.

Table 2-4 shows currently available APIC cluster hardware options and the general scalability each bundle can individually achieve.

**Table 2-4** APIC Hardware Bundles

Part Number	Number of APICs	General Scalability
APIC-CLUSTER-XS (ACI mini bundle)	1 M3 APIC, 2 virtual APICs, and 2 Nexus 9332C spine switches	Up to 2 spines and 4 leaf switches
APIC-CLUSTER-M3	3 M3 APICs	Up to 1200 edge ports
APIC-CLUSTER-L3	3 L3 APICs	More than 1200 edge ports

APIC-CLUSTER-XS specifically addresses ACI mini fabrics. ACI mini is a fabric deployed using two Nexus 9332C spine switches and up to four leaf switches. ACI mini is suitable for lab deployments, small colocation deployments, and deployments that are not expected to span beyond four leaf switches.

APIC-CLUSTER-M3 is designed for medium-sized deployments where the number of server ports connecting to ACI is not expected to exceed 1200, which roughly translates to 24 leaf switches.

APIC-CLUSTER-L3 is a bundle designed for large-scale deployments where the number of server ports connecting to ACI exceeds or will eventually exceed 1200.

Beyond bundles, Cisco allows customers to purchase individual APICs for the purpose of expanding an APIC cluster to enable further scaling of a fabric. Once a fabric expands beyond 1200 edge ports, ACI Verified Scalability Guides should be referenced to determine the optimal number of APICs for the fabric.


According to Verified Scalability Guides for ACI Release 4.1(1), an APIC cluster of three L3 APICs should suffice in deployments with up to 80 leaf switches. However, the cluster size would need to be expanded to four or more APICs to allow a fabric to scale up to 200 leaf switches.

**NOTE** Cisco recommends against deployment of APIC cluster sizes of 4 and 6. Current recommended cluster sizes are 3, 5, or 7 APICs per fabric.

Each APIC cluster houses a distributed multi-active database in which processes are active on all nodes. Data, however, is distributed or sliced across APICs via a process called *database sharding*. *Sharding* is a result of the evolution of what is called horizontal partitioning of databases and involves distributing a database across multiple instances of the schema. Sharding increases both redundancy and performance because a large partitioned table can be split across multiple database servers. It also enables a scale-out model involving adding to the number of servers as opposed to having to constantly scale up servers through hardware upgrades.

ACI shards each attribute within the APIC database to three nodes. A single APIC out of the three is considered active (the leader) for a given attribute at all times. If the APIC that houses the active copy of a particular slice or partition of data fails, the APIC cluster is able to recover via the two backup copies of the data residing on the other APICs. This is why the deployment of a minimum of three APICs is advised. Any APIC cluster deployed with fewer than three APICs is deemed unsuitable for production uses. Note that only the APIC that has been elected leader for a given attribute can modify the attribute.

Figure 2-10 provides a conceptual view of data sharding across a three-APIC cluster. For each data set or attribute depicted, a single APIC is elected leader. Assume that the active copy indicates that the APIC holding the active copy is leader for the given attribute.



**Figure 2-10** *Data Sharding Across Three APICs*

For a portion of a database to allow writes (configuration changes), a quorum of APICs housing the pertinent database attributes undergoing a write operation must be healthy and online. Because each attribute in an APIC database is sharded into three copies, a quorum is defined as two copies. If two nodes in a three-node APIC cluster were to fail simultaneously, the remaining APIC would move the entire database into a read-only state, and no configuration changes would be allowed until the quorum was restored.


When an APIC cluster scales to five or seven APICs, the sharding process remains unchanged. In other words, the number of shards of a particular subset of data does not increase past three, but the cluster further distributes the shards. This means that cluster expansion past three APICs does not increase the redundancy of the overall APIC database.

Figure 2-11 illustrates how an outage of Data Center 2, which results in the failure of two APICs, could result in portions of the APIC database moving into a read-only state. In this case, the operational APICs have at least two shards for Data Sets 1 and 3, so administrators can continue to make configuration changes involving these database attributes. However, Data Set 2 is now in read-only mode because two replicas of the attribute in question have been lost.

As Figure 2-11 demonstrates, increasing APIC cluster size to five or seven does not necessarily increase the redundancy of the overall cluster.

A general recommendation in determining APIC cluster sizes is to deploy three APICs in fabrics scaling up to 80 leaf switches. If recoverability is a concern, a standby APIC can be added to the deployment. A total of five or seven APICs should be deployed for scalability purposes in fabrics expanding beyond 80 leaf switches.

If, for any reason, a fabric with more than three APICs is bifurcated, the APIC cluster attempts to recover this split-brain event. Once connectivity across all APICs is restored, automatic reconciliation takes place within the cluster, based on timestamps.



**Figure 2-11** *Impact of APIC Failures in a Five-Node Cluster*

What would happen if Data Center 1 in Figure 2-11 failed instead of Data Center 2, and all shards for a specific subset of data resided in Data Center 1 at the time of the outage? In such a scenario, the failure of three APICs could lead to the hypothetical loss of all three shards of a specific subset of data. To ensure that a total loss of a given pod does not result in the loss of all shards for a given attribute, Cisco recommends that no more than two APICs be placed in a single pod.

**NOTE** Standby APICs allow an administrator to commission an APIC to allow recoverability of a fabric during failure scenarios in which the APIC quorum has been lost. When a standby APIC is deployed in a fabric, it acts as a passive player. It does not actively service users or configure ACI switches. It also does not synchronize data with active APICs. When first deploying a controller as a standby APIC, at least three APICs in the cluster need to be active.

## Spine Hardware

Cisco ACI spine hardware options includes Nexus 9300 Series fixed form factor switches as well as Nexus 9500 modular switches. Not all switches in the noted switch families can be deployed in ACI mode.

The primary factors that guide spine purchasing decisions are desired port bandwidths, feature requirements, hardware generation, and the required number of target state ports.

Whereas a fixed spine switch has a limited number of ports, a port in a modular platform can scale with the addition of more line cards to a chassis. For this reason, modular chassis are more suitable for fabrics that require massive scale.

Fixed spine platforms satisfy the scalability requirements of small to medium fabrics without problem.

## First-Generation Spine Switches

As noted earlier in this chapter, first-generation spine switches are not supported as spines interconnecting ACI fabrics in ACI Multi-Site deployments. Other new solutions, such as Remote Leaf and ACI Multi-Tier also require second-generation spine switches. Understanding first-generation spine platforms is, however, beneficial for historical purposes because a large number of ACI deployments still contain first-generation hardware.

First-generation ACI spine switch models on the market at the time of this writing have model numbers that end in PQ. Table 2-5 lists first-generation Nexus spine switches.

### Key Topic

**Table 2-5** First-Generation Spine Switches

Characteristic	Nexus 9336PQ	Nexus 9736PQ
Form factor	2 RU fixed switch	Line card for modular chassis
Supported modular platforms	N/A	Nexus 9504 Nexus 9508 Nexus 9516
40 Gigabit Ethernet ports	36 ports	36 ports
100 Gigabit Ethernet ports	N/A	N/A
ACI Multi-Pod support	Yes	Yes
CloudSec support	No	No
Remote Leaf support	No	No
ACI Multi-Tier support	No	No
ACI Multi-Site support	No	No

Even though first-generation spine switches do not support namespace normalization or ingress replication of BUM traffic, they can coexist with second-generation spine switches within a fabric. This coexistence enables companies to integrate fabrics into ACI Multi-Site without having to decommission older spines before the regular hardware refresh cycle.

**NOTE** First-generation spine switches can no longer be ordered from Cisco.

## Second-Generation Spine Switches

In addition to providing support for ACI Multi-Site, Remote Leaf, and ACI Multi-Tier, second-generation spine switch ports operate at both 40 Gigabit Ethernet and 100 Gigabit Ethernet speeds and therefore enable dramatic fabric bandwidth upgrades.

Second-generation spine switches also support MACsec and CloudSec. MACsec enables port-to-port encryption of traffic in transit at line rate. CloudSec enables cross-site encryption at line rate, eliminating the need for intermediary devices to support or perform encryption. Cross-site encryption is also referred to as *VTEP-to-VTEP encryption*.

Second-generation ACI spine switch models on the market at the time of this writing have model numbers that end in C, EX, and FX. Table 2-6 provides additional details about second-generation spine platforms.

**Table 2-6** Second-Generation Spine Switches

Characteristic	Nexus 9364C	Nexus 9332C	Nexus 9732C-EX	Nexus 9736C-FX
Form factor	2 RU fixed	1 RU fixed	Line card for modular chassis	Line card for modular chassis
Supported modular platforms	N/A	N/A	Nexus 9504 Nexus 9508 Nexus 9516	Nexus 9504 Nexus 9508 Nexus 9516
40/100 Gigabit Ethernet ports	64	32	32	36
ACI Multi-Pod support	Yes	Yes	Yes	Yes
CloudSec support	Last 16 ports	Last 8 ports	N/A	All ports
Remote Leaf support	Yes	Yes	Yes	Yes
ACI Multi-Tier support	Yes	Yes	Yes	Yes
ACI Multi-Site support	Yes	Yes	Yes	Yes

In addition to the hardware listed in Table 2-6, Nexus 9732C-FX line cards will be supported as ACI spine line cards in the near future.

New spine switches with 100/400 Gigabit Ethernet ports are also on the horizon. The Nexus 9316D-GX is already available and is supported as an ACI spine. This platform is also in the roadmap for support as a leaf switch. The 100/400 Gigabit Ethernet Nexus 93600CD-GX switch, which is supported as an ACI leaf, is also in the roadmap for use as a spine.

Cisco uses the term *cloud scale* to refer to the newer Nexus switch models that contain the specialized ASICs needed for larger buffer sizes, larger endpoint tables, and visibility into packets and flows traversing the switch without impacting CPU utilization. Second-generation ACI spine switches fall into the category of cloud-scale switches.

## Leaf Hardware

Cisco ACI leaf hardware options include Nexus 9300 Series fixed form factor switches. Not all switches in the noted switch families can be deployed in ACI mode.

The primary factors that guide leaf purchasing decisions are the desired port bandwidths, feature requirements, hardware generation, and the required number of target state ports.

### First-Generation Leaf Switches

First-generation ACI leaf switches are Nexus 9300 Series platforms that are based on the Application Leaf Engine (ALE) ASICs.

The hardware resources that enable whitelisting of traffic are ternary content-addressable memory (TCAM) resources, referred to as the *policy CAM*.

Policy CAM sizes vary depending on the hardware. The policy CAM size and behavior limitations in first-generation switches tended to sometimes limit whitelisting projects.

There are also a number of other capability differences between first- and second-generation leaf hardware, such as handling of Layer 4 operations and multicast routing.

**NOTE** The majority of first-generation leaf switches can no longer be ordered from Cisco. All Nexus 9300 Series ACI leaf switches whose model numbers end in PX, TX, PQ, PX-E, and TX-E are considered first-generation leaf switches.

2

## Second-Generation Leaf Switches

Second-generation ACI leaf switches are Nexus 9300 Series platforms that are based on cloud-scale ASICs. Second-generation leaf switches support Remote Leaf and ACI Multi-Tier, have significantly larger policy CAM sizes, and offer enhanced hardware capabilities and port speeds.

**NOTE** MACsec is supported on all ports with speeds greater than or equal to 10 Gbps on Nexus 9300 ACI switches whose model numbers end in FX. Check specific support levels for other platforms.

ACI leaf switches whose model numbers end in EX, FX, FX2, and FXP are considered second-generation leaf switches. Table 2-7 provides details about second-generation switches that have 1/10 Gigabit Ethernet copper port connectivity for servers.

**Key  
Topic**

**Table 2-7** Second-Generation 1/10 Gigabit Ethernet Copper Leaf Switches

Characteristic	Nexus 93108TC-EX	Nexus 9348GC-FXP	Nexus 93108TC-FX	Nexus 93216TC-FX2
Form factor	1 RU fixed	1 RU fixed	1 RU fixed	2 RU fixed
100 Mbps and 1 Gigabit Ethernet copper ports	N/A	48	N/A	N/A
100 Mbps and 1/10 Gigabit Ethernet copper ports	48	N/A	48	96
10/25 Gigabit Ethernet ports	N/A	N/A	4	N/A
40/100 Gigabit Ethernet ports	6	2	6	12
ACI Multi-Pod support	Yes	Yes	Yes	Yes
Remote Leaf support	Yes	Yes	Yes	Yes
Can be used as a Tier 1 leaf	Yes	Yes	Yes	Yes
Can be used as a Tier 2 leaf	Yes	Yes	Yes	Yes

The Nexus 9348GC-FXP switch has 48 ports, offering 100 Mbps or 1 Gigabit Ethernet connectivity. These ports have RJ-45 connections, eliminating the need for transceivers. Due to its low cost and support for cloud-scale features, the Nexus 9348GC-FXP is an ideal replacement for Fabric Extenders.

**NOTE** Support for ACI Multi-Site is dependent on spine switches in the fabric and not leaf switches. Also, at the time of writing, CloudSec is most relevant to spine switches.

Table 2-8 details second-generation switches that provide 1/10/25 Gigabit Ethernet fiber port connectivity for servers.

**Key  
Topic**

**Table 2-8** Second-Generation 1/10/25 Gigabit Ethernet Fiber Leaf Switches

Characteristic	Nexus 93180YC-EX	Nexus 93180YC-FX	Nexus 93240YC-FX2	Nexus 93360YC-FX2
Form factor	1 RU fixed	1 RU fixed	1.2 RU fixed	2 RU fixed
1/10/25 Gigabit Ethernet ports	48	48	48	96
40/100 Gigabit Ethernet ports	6	6	12	12
ACI Multi-Pod support	Yes	Yes	Yes	Yes
Remote Leaf support	Yes	Yes	Yes	Yes
Can be used as a Tier 1 leaf	Yes	Yes	Yes	Yes
Can be used as a Tier 2 leaf	Yes	Yes	Yes	Yes

Table 2-9 lists details on the only second-generation switch available at the time of writing that provides 40/100 Gigabit Ethernet connectivity for servers.

**Key  
Topic**

**Table 2-9** Second-Generation 40/100 Gigabit Ethernet Leaf Switches

Characteristic	Nexus 9336C-FX2
Form factor	1 RU fixed
40/100 Gigabit Ethernet ports	36
ACI Multi-Pod support	Yes
Remote Leaf support	Yes
Can be used as a Tier 1 leaf	Yes
Can be used as a Tier 2 leaf	Yes

## 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: Chapter 17, “Final Preparation,” and the exam simulation questions in the Pearson Test Prep Software Online.

## Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topic icon in the outer margin of the page. Table 2-10 lists these key topics and the page number on which each is found.



**Table 2-10** Key Topics for Chapter 2

Key Topic Element	Description	Page Number
List	Describes APICs, spine switches, and leaf switches	23
List	Describes some functions engineers commonly evaluate when deciding whether to dedicate leaf switches to functions	24
Paragraph	Describes ACI Multi-Pod	25
Paragraph	Calls out requirements for an ACI Multi-Pod IPN	25
Paragraph	Describes ACI Multi-Site	26
Paragraph	Explains APIC cluster separation in ACI Multi-Site fabrics and MSO communication with each cluster	27
Paragraph	Calls out requirements for an ACI Multi-Site ISN	27
Paragraph	Explains why ACI Multi-Site requires the use of at least one Gen 2 spine in each site	28
Paragraph	Describes Remote Leaf	30
Paragraph	Explains the significance of sizes in APIC purchases and the relevance of M versus L models	32
Paragraph	Explains APIC hardware generations and correlation with UCS C-Series server generations	32
Table 2-5	Lists first-generation spine switches	37
Table 2-6	Lists second-generation spine switches	38
Table 2-7	Lists second-generation 1/10 Gigabit Ethernet copper leaf switches	39
Table 2-8	Lists second-generation 1/10/25 Gigabit Ethernet fiber leaf switches	40
Table 2-9	Lists second-generation 40/100 Gigabit Ethernet leaf switches	40

## Complete Tables and Lists from Memory

There are no memory tables or lists in this chapter.

## Define Key Terms

Define the following key terms from this chapter and check your answers in the glossary:

fabric port, border leaf, service leaf, compute leaf, IP storage leaf, stretched ACI fabric, transit leaf, ACI Multi-Pod, ACI Multi-Site, sharding



# Index

## Symbols

---

| (pipe) delimiter character, 416  
802.1Q standard (IEEE), 7  
802.1X authentication, 173, 175

## A

---

- AAA (access, authentication, and accounting), 106. *See also* access methods; access policies
- EIGRP (Enhanced Interior Gateway Routing Protocol) authentication, 324
- external AAA server integration, 532
- AAA authentication policy settings, 547–550*
  - ACI configuration for LDAP, 541–547*
  - ACI configuration for RADIUS, 540–541*
  - ACI configuration for TACACS+, 532–536*
  - Cisco AV pair formatting, 538–540*
  - ISE (Identity Service Engine) configuration, 536–538*
- RBAC (role-based access control) access, 133
- user access
- granting, 523–525*
  - modifying, 523–525–528*
- aaa role, 520
- aaaSessionLR class, 128
- AAEPs (attachable access entity profiles), 165–169, 382
- configuration, 167–169, 190–191
  - definition of, 589
  - EPGs (endpoint groups), 207–210
  - in-band management configuration, 468
  - in L3Out domains, 302
  - overview of, 165–166
- Access (802.1P) port-binding mode, 268
- Access (Untagged) port-binding mode, 268
- access, authentication, and accounting. *See* AAA (access, authentication, and accounting)
- access encapsulation VLANs, 246
- access methods
- API (application programming interface), 103
  - CLIs (command-line interfaces)
    - APIC CLI, 100–102*
    - overview of, 100*
    - switch CLI, 102–103*
  - GUI (graphical user interface), 99, 107–108
  - management access modification, 103–105
  - overview of, 98–99

access policies. *See also* AAA (access, authentication, and accounting); domains; profiles

AAEPs (attachable access entity profiles), 165–169, 382  
*configuration*, 167–169, 190–191  
*definition of*, 589  
*EPGs (endpoint groups)*, 207–210  
*in-band management*, 468  
*in L3Out domains*, 302  
*overview of*, 165–166

ACI switch port configurations  
*AAEP EPGs (endpoint groups)*, 207–210  
*implications of*, 210  
*individual port configuration*, 188–196  
*overview of*, 186–188  
*port channel configuration*, 196–201  
*vPC (virtual port channel) configuration*, 204–207  
*vPC (virtual port channel) domains*, 201–204

for APIC in-band interfaces, 468–469

BPDU Guard/BPDU Filter, 230–231

configuration with quick start wizards, 211  
*Configure Interface, PC, and VPC Wizard*, 211  
*Configure Interface wizard*, 211–212

CoPP (Control Plane Policing), 225–229

definition of, 158, 589

DHCP relay configuration, 219–221

dynamic breakout port configuration, 215–217

Error Disabled Recovery Policy, 231–232

FEX (fabric extender) configuration, 212–215

hierarchy of, 183–184

interface policies and interface policy groups, 169–174  
*CDP (Cisco Discovery Protocol)*, 190  
*definition of*, 593  
*deployment of*, 173–174  
*EIGRP (Enhanced Interior Gateway Routing Protocol)*, 321–323  
*link level*, 189  
*LLDP (Link Layer Discovery Protocol)*, 189  
*OSPF (Open Shortest Path First)*, 326–327  
*overview of*, 169–170  
*port channel*, 198  
*types of*, 170–173  
*vPC (virtual port channel)*, 205–206

leaf interface override policy, 232, 594

MCP (MisCabling Protocol), 221–222, 594

overview of, 158

port channel member overrides, 232–235, 596

QoS (quality of service) global class settings, 217–219

selectors, 176–179

stateless networking, 182–183

storm control, 223–224

switch policies and switch policy groups, 174–176

tenancy and, 184

VLAN pools, 159–160  
*creating*, 159–160  
*designs for*, 161–163  
*overlap between*, 164–165

- static versus dynamic VLAN allocation*, 159, 598
  - VLAN ranges in*, 160
- access-admin role, 520
- access-list access-list-inbound extended permit icmp any any command, 460
- Account Expires setting (User Identity), 523
- Account Status setting (User Identity), 523
- accounting. *See* AAA (access, authentication, and accounting)
- ACI (Application Centric Infrastructure), benefits of, 9–16
  - architectural possibilities, 15
  - cross-platform integrations, 15
  - health monitoring and enhanced visibility, 16
  - multitenancy, 11–14
  - network management touchpoints, 9–10
  - policy reuse, 16
  - programmability, 11
  - scalability optimizations, 10
  - stateless networks, 11
  - traffic flow optimizations, 10
  - zero-trust security, 14–15
- ACI Anywhere
  - ACI Multi-Pod
    - active/active data center deployments*, 562–564
    - definition of*, 589
    - overview of*, 15, 25–26
    - stateful-services integration*, 563–564
  - ACI Multi-Site, 555–558
    - cross-site connectivity*, 555–557
    - extending ACI to remote locations and public clouds*, 564
    - locally governed configurations*, 557
    - MSO-governed configurations*, 557
    - primary and disaster recovery data center deployments*, 558–562
    - schemas*, 557–558
    - stateful-services integration*, 563–564
    - terminology for*, 557
  - extending ACI to remote locations and public clouds, 564–565
    - with ACI Multi-Site*, 564
    - with remote leaf switches*, 564–565
    - with vPod*, 564
  - overview of, 552
- ACI fabric initialization
  - backups and restores
    - configuration backups*, importing, 80–82
    - configuration exports*, 75
    - configuration imports*, 75–76
    - configuration rollbacks*, 82–83
    - configuration snapshots*, 80
    - on-demand backups*, 76–79
    - scheduled backups*, 79–80
  - Cisco FabricPath, 7
  - connectivity outside, 151–153
    - external EPGs (endpoint groups)*, 151–153
    - Layer 3 Outside (L3Out)*, 153
  - endpoint movements within ACI fabric, 247
  - fabric policies, 106, 592
  - fabric ports, 23, 592
  - FEX (fabric extenders), 24, 212–215
  - FTEP (fabric tunnel endpoint), 60, 592

- importance of, 42
- initialization process
  - APIC BIOS password, changing*, 52
  - APIC Cisco IMC (Integrated Management Controller) configuration*, 52–53
  - connectivity following switch initialization*, 59–62
  - GIR (Graceful Insertion and Removal)*, 58
  - initialization of first APIC*, 53–55
  - initialization of subsequent APICs*, 58
  - switch activation*, 55–58
  - verification of*, 59–62
- interconnecting with ACI Multi-Site, 557
- overview of, 44–45
- planning, 45, 47–48
  - APIC connectivity*, 46–47
  - APIC OOB configuration requirements*, 47–48
  - basic configuration parameters*, 48–49
  - cabling requirements*, 45
  - fabric discovery states*, 51–52
  - out-of-band versus in-band management*, 48
  - switch discovery process*, 49–51
  - switch discovery states*, 51–52
- post-initialization tasks
  - automatic upgrades of new switches*, 74–75
  - COOP (Council of Oracle Protocol) group configuration*, 92
  - default contracts*, 64–66
  - DNS (Domain Name System) server configuration*, 90–92
  - fabric upgrades*, 66–71
  - NTP (Network Time Protocol) synchronization*, 84–89
  - pod policy*, 83
  - schedulers*, 73, 597
  - static out-of-band addresses, assignment of*, 63–64
- .aci format**, 27
- ACI MIB Support List**, 502
- ACI Multi-Pod**
  - active/active data center deployments, 562–564
    - overview of*, 562–563
    - stateful-services integration*, 563–564
    - VMM integrations for multiple data centers*, 563
  - definition of, 589
  - overview of, 15, 25–26
  - stateful-services integration, 563–564
- ACI Multi-Site, 555–558**
  - cross-site connectivity, 555–557
  - definition of, 589
  - extending ACI to remote locations and public clouds, 564
  - locally governed configurations, 557
  - MSO (Multi-Site Orchestrator), 15, 27, 555, 557
  - MSO-governed configurations, 557
  - overview of, 15, 26–28
  - primary and disaster recovery data center deployments, 558–562
    - centralized orchestration and management capabilities*, 559
    - cross-data center ingress routing optimization*, 561
    - overview of*, 558–559
    - per-bridge domain broadcast and stretch settings*, 560
    - policy deployment to sites*, 561

- schemas, 557–558
- stateful-services integration, 563–564
- terminology for, 557
- ACI Multi-Tier, 28–30**
- ACI Virtual Edge (AVE), 24, 564**
- acknowledging faults, 115–116
- acquisitions, multitenancy and, 13
- activation of switches, 55–58
- active controllers, number of, 48
- Active Interval setting (EIGRP Context per Address Family), 324**
- Active state (fabric node discovery), 52**
- active/active data center deployments, 562–564**
  - overview of, 562–563
  - stateful-services integration, 563–564
  - VMM integrations for multiple data centers, 563
- adapters, VMkernel, 395, 397, 405, 406, 413–414**
- Adaptive Security Appliance (ASA), 430**
- Add and Manage Hosts wizard, 411–414**
- Add vCenter Controller page (vCenter Domain Creation wizard), 409–410**
- Add VMM Domain Association page, 418**
- Address Resolution Protocol. *See* ARP (Address Resolution Protocol)**
- addresses**
  - IP (Internet Protocol), 325–326, 470–474, 597
    - dynamic*, 472–474
    - secondary*, 325–326, 597
    - static*, 470–472
  - VIP, 245
- Admin Distance Preference (OSPF Timer policy), 329**
- admin role, 520**
- Admin State parameter**
  - BGP peer connectivity profiles, 341
  - Global MCP Policy, 222
- administrative separation, 133**
- Advertise Subnet option (OSPF), 329**
- Advertised Externally domain scope, 474–476**
- advertisement**
  - of host routes, 321
  - of subnets assigned to bridge domains, 314–315
- agents, SNMP (Simple Network Management Protocol), 500**
- agility, network, 8**
- all security domain, 517**
- all/aaa/read-all value, 539**
- all//admin value, 539**
- all/admin/(16005) value, 539**
- all/fabric-admin[access-admin/ value, 540**
- Allow Self AS parameter (BGP peer connectivity profiles), 339**
- Allowed Self AS Count parameter (BGP peer connectivity profiles), 340**
- all//read\_all value, 539**
- Amazon Web Services (AWS), 27**
- Any EPG (vzAny), 364–365, 589**
- API (application programming interface), 103**
- APIC (Application Policy Infrastructure Controller) clusters, 9–10, 11**
  - access policies, 106
  - ACI fabric initialization and
    - APIC Cisco IMC (Integrated Management Controller) configuration*, 52–53
    - BIOS password, changing*, 52
    - connectivity following switch initialization*, 59–62

- default contracts*, 64–66
- initialization of first APIC*, 53–55
- initialization of subsequent APICs*, 58
- static out-of-band addresses, assignment of*, 63–64
- verification of*, 59–62
- APIC-CLUSTER-L3, 33–34
- APIC-CLUSTER-M3, 33–34
- APIC-CLUSTER-XS, 33–34
- CLI (command-line interface), 100–102
  - bash shell interactions*, 102
  - configuration modes*, 101
  - navigation in*, 101
  - running configuration, viewing*, 102
  - vPC (virtual port channel) configuration*, 207
- components, 46–47
- configuration, 47–48
- connecting to fabric, 46–48
- definition of, 23, 589
- fabric policies, 106
- for in-band management
  - access policies for*, 468–469
  - connectivity preferences*, 478–479
- scalability and sizing, 33–36
- specifications for, 26–28
- tenants and, 106, 147
- APIC OOB Addresses and Default Gateway parameter (ACI fabric initialization), 49
- `apic#fallback\username`, 549–550
- Application Policy Infrastructure Controller. *See* APIC (Application Policy Infrastructure Controller) clusters
- application profiles, 527
  - advantages of, 138–139
  - configuration, 142–145
  - definition of, 138, 589
  - multitier application deployment, 264–266
  - topology, 138–139
- application programming interface. *See* APIC (Application Policy Infrastructure Controller) clusters
- Application Services Engine (ASE), 27
- Application Virtual Switch (AVS), 24
- applications, multitier. *See* multitier application deployment
- Apply Both Directions feature, 277–278, 476
- architecture, ACI (Application Centric Infrastructure), 9–10, 15
- ARP (Address Resolution Protocol), 241
  - ARP table, 241, 242
  - flooding, 262–263, 368
  - gleaning, 248–249
- ARP Flooding field (bridge domain settings), 369
- AS override parameter (BGP peer connectivity profiles), 339
- ASA (Adaptive Security Appliance), 430
- ASE (Application Services Engine), 27
- ASNs (autonomous system numbers), 303, 555
- Atomic import mode, 76
- attachable access entity profiles. *See* AAEPs (attachable access entity profiles)
- Attribute parameter (LDAP providers), 542
- audit logs
  - definition of, 127, 590
  - reviewing, 127–128

authentication. *See* AAA (access, authentication, and accounting)

Authorization Port parameter (RADIUS providers), 540

Authorization Protocol parameter  
RADIUS providers, 540  
TACACS+ providers, 532

Auto State, L3Out, 299–300

autonomous system numbers (ASNs), 303, 555

AV pair formatting, 538–540

AVE (ACI Virtual Edge), 24, 564

AVS (Application Virtual Switch), 24

AWS (Amazon Web Services), 27

## B

---

### backups

- configuration exports, 75
- configuration imports, 75–76
- configuration rollbacks, 82–83
- configuration snapshots, 80
- importing, 80–82
- on-demand, 76–79
- scheduled, 79–80

bandwidth, Remote Leaf deployments, 31

Bandwidth Preference (OSPF Timer policy), 329

Bandwidth Reference setting (OSPF timer policy), 329

bare-metal clouds, extending ACI to, 564

Base DN parameter (LDAP providers), 542

bash, 102

BD Multicast Addresses parameter (ACI fabric initialization), 49

### BDs (bridge domains)

- advertisement of subnets assigned to, 314–315
- configuration, 142–145
- definition of, 137, 590
- DHCP (Dynamic Host Configuration Protocol), 271
- disabling endpoint learning on, 250–251
- extensions, 374–375
- external, 161, 380–385, 591
- in-band management, 469
- L3Out, 296–298, 594
- multitier application deployment, 264–266
- practical application of, 141–142
- service graph settings, 439–440

Beacon Probing, 402

Bell Laboratories, 21

Best Effort import mode, 76

BFD (Bidirectional Forwarding Detection) configuration

- for BGP peering, 339–341
- for EIGRP L3Out, 321–324
- for OSPF L3Outs, 328–329
- switch policies, 174

BFD Interface Profile Creation Wizard, 322–323

BGP (Border Gateway Protocol), 106

ASNs (autonomous system numbers), 303, 555

BGP EVPN router IDs, 556, 590

iBGP peerings, 303–304

L3Out for BGP peering

*BFD (Bidirectional Forwarding Detection) implementation, 339–341*

*BGP timer policies, 338–339, 342*

- L3Out Creation Wizard*, 334–337
  - node-level customizations*, 337–339
  - OSPF for IP reachability*, 343
  - peer connectivity profiles*, 339–341
  - per-neighbor customizations*, 339–341
  - VRF-level customizations*, 342–343
  - route reflection
    - definition of*, 597
    - implementation of*, 304–305
    - infra MP-BGP route distribution*, 305–306
    - need for*, 303–304
  - BGP Peer Prefix Policy parameter (BGP peer connectivity profiles), 341
  - Bidirectional Forwarding Detection. *See* BFD (Bidirectional Forwarding Detection) configuration
  - Bidirectional Forwarding Detection parameter (BGP peer connectivity profiles), 340
  - Bind DN parameter (LDAP providers), 542
  - binding
    - dynamic, 159
    - static, 159
  - bindings, port, 400, 596
  - BIOS password, changing, 52
  - Blacklist port usage, 195
  - blacklisting, 14, 318–320, 590
  - bocks, port, 198–199
  - Border Gateway Protocol. *See* BGP (Border Gateway Protocol)
  - border leaf switches, 24, 590
  - bounce entry, 247
  - BPDUs (bridge protocol data units), 386–388
    - BPDU Filter, 230–231
    - BPDU Guard, 230–231
    - PVST+ (Per-VLAN Spanning Tree), 387–388
    - Rapid PVST+ (Rapid Per-VLAN Spanning Tree), 387–388
    - TCNs (topology change notifications), 386–387
  - bridge domains. *See* BDs (bridge domains)
  - bridge protocol data units. *See* BPDUs (bridge protocol data units)
  - bridge-domain BD-CRITICAL-STUFF command, 147
  - broadcast, unknown unicast, and multicast forwarding. *See* BUM (broadcast, unknown unicast, and multicast) forwarding
  - broadcast settings, 560
  - Broadcom, 103
  - BUM (broadcast, unknown unicast, and multicast) forwarding, 25, 223
    - ingress replication of, 28
    - preventing, 560
    - storm control, 223–224
- ## C
- 
- cabling, ACI fabric initialization, 45
  - CAM (content-addressable memory), 362
  - cat /etc/ntp.conf command, 86–87
  - CDP (Cisco Discovery Protocol)
    - configuration, 190
    - interface policies, 171
  - centralized organization, 559
  - channels, port. *See* port channels

- Cisco Adaptive Security Appliance (ASA), 430
- Cisco Application Services Engine (ASE), 27
- Cisco AV pair formatting, 538–540
- Cisco Cloud APIC, 564
- Cisco Cloud Service Router (CSR) 1000V, 564
- Cisco FabricPath, 7
- Cisco Firepower Management Center (FMC), 432
- Cisco IMC (Integrated Management Controller), 52–53, 100, 589
- Cisco Nexus 1000v distributed virtual switch, 9
- Cisco Overlay Transport Virtualization (OTV), 560
- Cisco Tetration, 272
- Cisco Unified Computing System (UCS), 218
- Citrix NetScaler Management and Analytics System (MAS), 432
- Class of Service (CoS), 218
- classes, 590
  - aaaSessionLR, 128
  - definition of, 109
  - faultDelegate, 116–118
  - faultInst, 116–118
  - monCommonPol, 118
  - monEPGPol, 118
  - monFabricPol, 118
  - monInfraPol, 118
- Cleared faults, 112
- Clearing Interval timer, 115
- CLIs (command-line interfaces)
  - APIC CLI, 100–102
    - bash shell interactions*, 102
    - configuration modes*, 101
    - navigation in*, 101
    - running configuration, viewing*, 102
    - tenant object configuration*, 147
    - vPC (virtual port channel) configuration*, 207
  - overview of, 100
  - switch CLI, 102–103
- Clos, Charles, 21
- Clos topology, 21–22
- Cloud APIC, 564
- Cloud Service Router (CSR) 1000V, 564
- CloudSec encryption, 556, 590
- clusters, APIC. *See* APIC (Application Policy Infrastructure Controller) clusters
  - clusters, MSO (Multi-Site Orchestrator), 555
- CNA (converged network adapters), 218
- commands. *See individual commands*
- common security domain, 517
- common tenants, 135
- common//read-all value, 540
- Communication faults, 113
- components. *See hardware components*
- compute leaf switches, 24, 590
- concrete devices, 444, 449–450, 590
- configuration backups, importing, 80–82
- configuration exports, 75
- Configuration faults, 113
- configuration imports, 75–76
- configuration rollbacks, 82–83
- configuration snapshots, 80
- configure command, 101
- Configure Interface, PC, and VPC Wizard, 211–212
- configure terminal command, 101

- Confirm New Password dialog box, 52
- conformity, template, 557, 599
- Console Authentication Realm setting (AAA policy), 549
- constructs, ACI Multi-Site, 557
- consumers, 149–150, 590
- content-addressable memory (CAM), 362
- contexts, SNMP (Simple Network Management Protocol), 506
- contexts, VRF. *See* VRF (virtual routing and forwarding) instances
- contracts, 148–151
  - ACI fabric initialization, 64–66
  - Apply Both Directions feature, 277–278
  - assigning to EPGs, 278–279
  - definition of, 148–149, 590
  - direction of, 149–150
  - filters, 148–149, 592
    - for bidirectional application, 273–275*
    - for unidirectional application, 280*
  - L3Out for EIGRP peering, 316–318
  - for out-of-band (OOB) management, 479–480
  - overview of, 272
  - planning, 272–273
  - Reverse Filter Ports feature, 277–278
  - scope of, 150–151, 284–285, 590
  - service graphs, 427–428
  - SNMP (Network Management Protocol), 503
  - Stateful feature, 284
  - subjects, 278
    - for bidirectional application, 275–276*
    - definition of, 148–149*
    - for unidirectional application, 280–281*
    - verification of, 278*
  - syslog forwarding, 487–491
  - TCP established session rule, 279–280
  - verifying enforcement of, 283–284
  - VRF-level enforcement, 367
  - whitelisting, 282–283
  - zero-trust networks and, 151
- Control Plane Policing (CoPP)
  - configuration, 225–229
  - switch policies, 174
- controller ID parameter (ACI fabric initialization), 49
- controller name parameter (ACI fabric initialization), 49
- Controller port usage, 195
- controller role, 57
- converged network adapters (CNA), 218
- COOP (Council of Oracle Protocol), 23, 62, 92, 243
- Coordinated Universal Time (UTC), 89
- CoPP (Control Plane Policing)
  - configuration, 225–229
  - interface policies, 172
  - prefilter switch policies, 175
  - switch policies, 174
- CoS (Class of Service), 218
- Council of Oracle Protocol. *See* COOP (Council of Oracle Protocol)
- Create Access Port Selector, 191
- Create Application EPG wizard, 265–266
- Create Application Profile wizard, 144–146
- Create Bridge Domain wizard, 142–144, 264–265
- Create Bridged Outside wizard, 380–385

- Create Configuration Export Policy wizard, 76–78, 80–81
- Create Contract page, 275
- Create Contract Subject page, 275–276
- Create Date and Time Policy window, 85
- Create DHCP Relay Label wizard, 271
- Create DNS Profile option, 90
- Create External Management Network Instance Profiles command, 64–65, 491
- Create Filter wizard, 488
- Create LDAP Group Map Rules wizard, 543–547
- Create LDAP Provider wizard, 543
- Create Leaf Breakout Port Group page, 216
- Create Leaf Interface Profile wizard, 179–180
- Create Leaf Profile wizard, 180–181
- Create Local User wizard, 521–523
- Create Login Domain page, 535
- Create New Password dialog box, 52
- Create Out-of-Band Contract page, 489
- Create Physical Domain wizard, 163
- Create RBAC Rule wizard, 529–531
- Create Remote Location wizard, 77
- Create Role page, 520
- Create Security Domain window, 517
- Create SNMP Monitoring Destination Group command, 507
- Create Static Node Management Addresses page, 63, 470–471
- Create Syslog Monitoring Destination Group command, 492
- Create Syslog Source wizard, 494–495
- Create Tenant wizard, 133–134

- Create vCenter Credential page (vCenter Domain Creation wizard), 409
- Create vCenter Domain page (vCenter Domain Creation wizard), 409, 410
- Create VLAN Pool wizard, 159–160
- Create VPC Explicit Protection Group wizard, 202–203
- Create VRF wizard, 135–137
- Critical faults, 112
- cross-fabric namespace normalization, 28
- cross-platform integrations, 15
- cross-site connectivity, 555–557
- CSR (Cloud Service Router) 1000V, 564
- Custom EPG Name field, 416
- Custom Filter parameter (LDAP providers), 542
- custom function profiles, 444–445
- custom RBAC (role-based access control) rules, 528–531

## D

---

### data center deployments

- with ACI Multi-Site, 558–562
- active/active data centers, 562–564
  - overview of, 562–563*
  - stateful-services integration, 563–564*
  - VMM integrations for multiple data centers, 563*
- extension and migration, Remote Leaf deployments for, 30
- multitenancy, 12–13
- primary and disaster recovery data centers, 558–562
  - centralized orchestration and management capabilities, 559*

- cross-data center ingress routing optimization*, 561
- overview of*, 558–559
- per-bridge domain broadcast and stretch settings*, 560
- traditional networks
  - agility of*, 8
  - management of*, 4–5
  - scalability and growth of*, 5–7
  - security of*, 8–9
  - three-tier design of*, 5–7
  - visibility of*, 9
- data plane policing (DPP)**, 172
- database sharding**, 34–35, 597–598
- date/time, NTP (Network Time Protocol) synchronization**, 84–89, 106
- Datetime Format object**, 89
- DCACI 300-620 exam**
  - final review/study, 570
  - preparation for, 566–570
    - Pearson IT Certification test prep software*, 567–570
    - tips and guidelines*, 566–567
  - updates to, 586–587
- DCACIA 300–630 exam**, 361
- Default Authentication Realm setting (AAA policy)**, 548
- Default Firmware Version setting**, 74–75
- delimiter characters**, 416
- demilitarized zone (DMZ)**, 8
- deployment immediacy**, 268, 405–407, 590
- design, VLAN pools and domains**, 161–163
  - hybrid approach, 162–163
  - single VLAN pool for each type of domain, 161
  - single VLAN pool per function, 162
- destination groups**
  - SNMP (Simple Network Management Protocol), 507–508
  - syslog monitoring, 492–493
- Development/tenant-admin/ value**, 540
- device managers**, 433–434, 590
- device packages, importing**, 430, 441–442
- device selection policies**, 446
- devices, deploying service graphs for**
  - devices in GoThrough Mode, 436, 437
  - devices in GoTo Mode, 435, 436–437
  - one-arm load balancers, 437
- DHCP (Dynamic Host Configuration Protocol) relay**, 32, 219–221, 271
- DHT (distributed hash table)**, 92
- Differentiated Services Code Point (DSCP)**, 218
- direction, of contracts**, 149–150
- Disable Connected Check parameter (BGP peer connectivity profiles)**, 340
- Disable Peer AS Check parameter (BGP peer connectivity profiles)**, 340
- Disabled/Decommissioned state (fabric node discovery)**, 52
- disabling**
  - contract enforcement at VRF instance level, 367
  - endpoint learning
    - on bridge domain*, 250–251
    - at VRF level*, 251
- disaster recovery (DR)**
  - data center deployments with ACI Multi-Site, 558–562
    - centralized orchestration and management capabilities*, 559
    - cross-data center ingress routing optimization*, 561–562

- overview of, 558–559*
  - per-bridge domain broadcast and stretch settings, 560*
- Remote Leaf deployments, 31
- Discovering state (fabric node discovery), 52
- Discovery port usage, 195
- discovery process, ACI fabric initialization
  - fabric discovery stages, 51
  - switch discovery and activation, 55–58
  - switch discovery process, 49–51
  - switch discovery states, 51–52
- distinguished names (DNs), 109, 591
- distributed hash table (DHT), 92
- distributed port groups, 397
  - assigning VMs (virtual machines) to, 417–418
  - pushing EPGs (endpoint groups) to vCenter as, 415–416
- Distributed Resource Scheduler (DRS), 399
- DMZ (demilitarized zone), 8
- DNs (distinguished names), 109, 591
- DNS (Domain Name System), 90–92, 106
- dns command, 101
- Domain Name System. *See* DNS (Domain Name System)
- domain profiles (VMM), 408–410
- domains
  - assigning to EPGs, 267
  - associating VLAN pools with, 163–164
  - BDs (bridge domains)
    - advertisement of subnets assigned to, 314–315*
    - configuration, 142–145*
    - definition of, 137*
  - DHCP (Dynamic Host Configuration Protocol), 271
    - external, 161, 591*
    - L3Out, 296–298, 594*
    - multitier application deployment, 264–266*
    - practical application of, 141–142*
  - creating, 163
  - definition of, 591
  - designs for, 161–163
    - hybrid approach, 162–163*
    - single VLAN pool for each type of domain, 161*
    - single VLAN pool per function, 162*
  - external bridged, 380–385
  - external routed, 161, 591
  - fallback, 549–550
  - fault, 113
  - Fibre Channel, 161
  - L3 domain implementation examples, 301–302
  - mapping to security domain, 526–528
  - physical, 161, 595
  - security, 517–518
    - assigning tenants to, 517–518*
    - creating, 517*
    - definition of, 517, 597*
    - mapping domains to, 526–528*
    - predefined, 517*
  - types of, 160–161
  - UCS (Unified Computing System) domains, 218
  - VMM (Virtual Machine Manager), 161, 418, 599
  - vPC (virtual port channel) domain, 201–204
- DPP (data plane policing), 172
- DR. *See* disaster recovery (DR)

Drop option (Multi Destination Flooding), 370

DRS (Distributed Resource Scheduler), 399

DSCP (Differentiated Services Code Point), 218

DTEP (dynamic tunnel endpoint), 61, 591

DWDM, 173

dynamic breakout ports, 215–217, 591

dynamic endpoint attach, 439

Dynamic Host Configuration Protocol. *See* DHCP (Dynamic Host Configuration Protocol) relay

dynamic IP addressing, 472–474

dynamic port bindings, 400

dynamic tunnel endpoint (DTEP), 61, 591

dynamic VLAN allocation, 159, 591

## E

---

east-west traffic, 7

eBGP Distance parameter (BGP timer policy), 342

eBGP/iBGP Max ECMP parameter (BGP timer policy), 342

ECMP (equal-cost multipathing), 325

EDT (Error Detect Timeout), 175

EIGRP (Enhanced Interior Gateway Routing Protocol) peering, L3Out configuration for, 305–306, 324–325

advertisement of host routes, 321

advertisement of subnets assigned to bridge domains, 314–315

authentication, 324

BFD (Bidirectional Forwarding Detection) implementation, 321–324

blacklist EPG (endpoint group) with logging, 318–320

contracts, 316–318

EIGRP Context per Address Family, 324–325

external EPG (endpoint group) deployment, 310–312

forwarding verification, 312–314

L3Out Creation Wizard, 307–309

route advertisement, 328

VRF-level customizations, 324–325

electronic programmable logic device (EPLD) upgrades, 66

Enable Host Route Leak parameter (BGP timer policy), 342

Enable MCP PDU Per VLAN setting (Global MCP Policy), 222

Enable Name Lookup for Router IDs setting (OSPF timer policy), 329

Enable SSL parameter (LDAP providers), 542

Encap Scope, L3Out, 298–299

encapsulation, port, 246

encryption, CloudSec, 556, 590

end command, 101

endpoint groups. *See* EPGs (endpoint groups)

endpoint learning. *See also* endpoints; EPGs (endpoint groups)

definition of, 591

disabling

- on bridge domain, 250–251*
- at VRF level, 251*

endpoint movements within ACI fabric, 247

Enforce Subnet Check feature, 250

hardware proxy, 247–248

on L3Outs, 249

Limit IP Learning to Subnet feature, 249–250

- local versus remote endpoints, 242–243, 594, 596
- lookup tables, 241–242
- network migrations into ACI, 371–372
- overview of, 241
- remote, 244–245
- for silent hosts, 245–247
- spine proxy, 247–248
- unicast routing and, 243–244
- VLAN IDs, 245–247, 253–254
- VNIDs (VXLAN network identifiers), 244, 245–247
- endpoints. *See also* endpoint learning; EPGs (endpoint groups)**
  - definition of, 591
  - dynamic endpoint attach, 439
  - endpoint tables, 242
  - local versus remote, 242–243, 594, 596
  - local versus remote endpoints, 596
  - movements within ACI fabric, 247
  - number of, 4
  - TEPs (tunnel endpoints)
    - overlay multicast*, 555, 595
    - overlay unicast*, 555, 595
    - PTEP (physical tunnel endpoint)*, 60, 595
    - TEP pools*, 49, 50, 57, 59, 599
    - verifying addresses of*, 60–62
  - verification of traffic path between, 254–256
- Enforce Bootscript Version Validation setting, 74**
- Enforce Subnet Check feature, 250**
- Enhanced Interior Gateway Routing Protocol. *See* EIGRP (Enhanced Interior Gateway Routing Protocol) peering, L3Out configuration for enhanced LACP policy support, 419–421**
- Enter Password dialog box, 52**
- Environmental faults, 113**
- EPG-Login-Servers, 384**
- EPGs (endpoint groups), 151–153**
  - AAEP (attachable access entity profile) configuration, 207–210
  - Any EPG (vzAny), 364–365, 589
  - blacklist EPG with logging, 318–320
  - classifying endpoints into, 146–147
  - configuration, 142–145
  - contracts, assigning to, 278–279
  - definition of, 137–138
  - domains, assigning to, 267
  - extensions
    - bridge domain extensions compared to*, 374–375
    - definition of*, 372–373, 375–380
  - external, 151–153, 294, 310–312, 591
  - in-band management, 470–471, 476–477
  - for multitier application deployment
    - DHCP (Dynamic Host Configuration Protocol) relay*, 271
    - EPG configuration*, 264–266
    - EPG mapping to ports*, 267–269, 270
    - policy deployment*, 270
    - verification of EPG-to-port assignments*, 269–270
  - out-of-band (OOB) management, 480, 491
  - port usage, 195
  - practical application of, 141–142
  - pushing to vCenter as distributed port groups, 415–416
  - shadow, 428–429, 597
- ephemeral port bindings, 400**

- EPLD (electronic programmable logic device) upgrades, 66
- equal-cost multipathing (ECMP), 325
- equipment flash configuration policies, 176
- Error Detect Timeout (EDT), 175
- Error Disabled Recovery Policy, 231–232
- established session rule (TCP), 279–280
- ESXi
  - servers
    - adding to VDS (vSphere distributed switches), 411–414*
    - connecting in UCS domains, 407*
    - connecting to fabric, 407*
    - definition of, 395*
    - standard switch architecture, 395*
    - VMkernel interface, 397–399*
    - vSphere distributed switches (VDS), 397*
  - VMs (virtual machines), 27
- Ethernet ports, 46
- events
  - definition of, 126
  - event records, 591
  - snuffling, 127, 598
  - viewing, 126
- exam, DCACI 300-620
  - final review/study, 570
  - preparation for, 566–570
    - Pearson IT Certification test prep software, 567–570*
    - tips and guidelines, 566–567*
  - updates to, 586–587
- exports, configuring, 75
- Extensible Markup Language (XML), 103
- extension, Layer 2. *See also* BDs (bridge domains); EPGs (endpoint groups)
  - ACI interaction with STP (Spanning Tree Protocol), 386–388
    - link types, 388*
    - MCP (MisCabling Protocol), 388*
    - MST (Multiple Spanning Tree) instance mappings, 387–388*
    - TCNs (topology change notifications), 386–387*
  - Layer 2 connectivity to non-ACI switches
    - bridge domain extensions, 374–375*
    - EPG (endpoint group) extensions, 372–373, 375–380*
    - L2Outs, 380–385*
    - migration of overlapping VLANs into ACI, 385–386*
    - overview of, 372*
  - network migrations into ACI
    - Any EPG (vzAny), 364–365, 589*
    - contract enforcement at VRF instance level, 367*
    - endpoint learning considerations, 371–372*
    - flooding requirements, 368–370*
    - full-mesh network-centric contracts, 362–364*
    - GARP-based detection, 370–371*
    - Legacy mode, 371*
    - network-centric deployments, 361–362*
    - overview of, 361*
    - preferred group members, 365–367*
    - preparation for, 372*
    - security enforcement in, 362*
  - overview of, 358

- external (L3Out) bridge domains, 296–298
- external AAA server integration, 532
  - AAA authentication policy settings, 547–550
    - configuration options*, 548–549
    - fallback domain*, 550
    - login domains and global AAA settings*, 549–550
- ACI configuration for LDAP, 541–547
  - configuration parameters*, 541–543
  - Create LDAP Group Map Rules wizard*, 543–547
  - Create LDAP Provider wizard*, 543
- ACI configuration for RADIUS, 540–541
- ACI configuration for TACACS+, 532–536
- Cisco AV pair formatting, 538–540
- ISE (Identity Service Engine) configuration, 536–538
- external advertisement, inb bridge domain, 474–476
- external BDs (bridge domains), 591
- external bridge domains, 161, 380–385
- external bridged networks, 372, 375, 380, 591
- External Distance setting (EIGRP Context per Address Family), 325
- External EPG Creation Wizard, 310–312
- External EPG Networks page (Create Bridged Outside wizard), 382
- external EPGs (endpoint groups), 151–153, 294, 310–312, 591
- External Management Network Instance Profile folder, 64–65

- external management network instance profiles, 491
- external routed domains, 161, 591
- external VLAN ranges, 160

## F

---

- Fabric Connectivity Infra page, 556
- fabric extenders (FEX), 212–215
- Fabric External port usage, 195
- fabric ID parameter (ACI fabric initialization), 48
- fabric initialization. *See* ACI fabric initialization
- Fabric Membership view, 57–58
- fabric name parameter (ACI fabric initialization), 48
- fabric policies, 106, 592
- fabric ports, 23, 195, 592
- fabric tunnel endpoint (FTEP), 592
- fabric-admin role, 520
- fabric-equipment privilege, 519
- FabricPath, 7
- Failover Order setting, 403
- failover policies, 400–403
- Fallback Check setting (AAA policy), 549
- fallback domain, 549–550
- Fallback setting, 403
- Fast Select Hot Standby Ports setting, 196, 201
- fault domains, 13–14
- Fault Severity Assignment Policies folder, 122
- faultDelegate class, 116–118
- faultInst class, 116–118
- faults
  - acknowledging, 115–116
  - definition of, 592

- domains, 113
- fault management policies, 120–121
- fault MOs, 116–118
- isolation of, 133
- lifecycle, 113–115, 592
  - phases, 114–115*
  - timers, 115*
- overview of, 111
- policies, 120–121
- severity levels, 111–112, 121–124
- types of, 113
- viewing, 111
- FC port channel interface policy groups, 171
- FEX (fabric extenders), 24, 212–215
- FHRP (first-hop redundancy protocol), 6
- Fibre Channel (FC), 172
  - domains, 161
  - interface policy groups, 171
  - switch policies, 175
- Filter Type parameter (LDAP providers), 542
- filters
  - creating
    - for bidirectional application, 273–275*
    - for unidirectional application, 280*
  - definition of, 148–149, 592
  - filter entries, 148
  - in-band management, 476–477
  - Stateful feature, 284
  - for syslog forwarding, 487–491
- finite state machine (FSM), 111
- Firepower Management Center (FMC), 432
- firewall-domain, 162, 164
- firewalls
  - interface policies, 173
  - service graphs
    - for managed devices, 453–460*
    - unmanaged firewall pair deployment, 447–452*
- first-generation leaf switches, 38–39
- first-generation spine switches, 37
- first-hop redundancy protocol (FHRP), 6
- Flash Card mode, Pearson IT Certification test prep software, 568–569
- floating SVIs (switch virtual interfaces), 296, 592
- Flood in BD option (Multi Destination Flooding), 370
- Flood in Encapsulation option (Multi Destination Flooding), 370
- flooding, 368–370
  - ARP (Address Resolution Protocol), 262–263, 369
  - Forwarding setting, 368
  - L2 Unknown Unicast, 368
  - L3 Unknown Multicast, 369
  - Multi Destination, 370
- flow control, PFC (priority-based flow control), 172
- FMC (Firepower Management Center), 432
- folded Clos topology, 21–22
- forwarding, 62
  - L3Out for EIGRP peering, 312–314
  - packet, 251
    - both endpoints attach to same leaf, 251–254*
    - flooding to unknown destination, 261–263*
    - known destination behind another leaf, 254–258*

- spine proxy to unknown destination*, 258–261
- system messages to syslog servers
  - contracts for*, 487–491
  - steps for*, 487
  - syslog monitoring destination groups*, 492–493
  - syslog sources*, 494–498
  - verification of*, 498–500
- Forwarding field (bridge domain settings), 368
- forwarding scale profiles, 175
- forwarding tag trees. *See* FTag trees
- FQDNs (fully qualified domain names), 126
- FSM (finite state machine), 111
- FTag topology, 261–262
- FTag trees, 261, 262, 592
- FTEP (fabric tunnel endpoint), 60, 592
- full-mesh network-centric contracts, 362–364
- fully qualified domain names (FQDNs), 126
- function nodes, 445–446
- function profiles
  - creating custom, 444–445
  - definition of, 592
  - examples of, 430–433, 442
- functions, service graphs as
  - concatenation of, 427–428

## G

---

- GARP (Gratuitous ARP), 241, 370–371
- gateways
  - for in-band management, 469–470
  - pervasive, 252–254, 595

- Gigabit Ethernet dedicated management ports, 46
- GIR (Graceful Insertion and Removal), 58
- Global AES Encryption Settings for All Configuration Import and Export page, 78
- global MIB (Management Information Base) scope, 502, 592
- GoThrough mode, devices in, 436, 437
- GoTo mode, devices in, 435, 436–437
- Graceful Convergence setting, 196, 201
- Graceful Insertion and Removal (GIR), 58
- Graceful Restart Helper setting
  - BGP timer policy, 338
  - OSPF timer policy, 329
- granting access, 523–525
- graphical user interface. *See* GUI (graphical user interface)
- graphs, service. *See* service graphs
- Gratuitous ARP (GARP), 241, 370–371
- grep command, 102
- GUI (graphical user interface), 99, 107–108

## H

---

- hardware components. *See* APIC (Application Policy Infrastructure Controller) clusters; switches
- hardware proxy, 247–248
- hardware proxy forwarding, 62
- hashes, 402
- headend replication (HER) tunnels, 32
- health monitoring, 16
  - audit logs
    - definition of*, 127
    - reviewing*, 127–128

- events
    - definition of*, 126
    - event records*, 126
    - snelching*, 127, 598
    - viewing*, 126
  - faults
    - acknowledging*, 115–116
    - definition of*, 592
    - domains*, 113
    - fault management policies*, 120–121
    - fault MOs*, 116–118
    - lifecycle*, 113–115, 592
    - overview of*, 111
    - severity levels*, 111–112, 121–124
    - snelching*, 121–124
    - types of*, 113
    - viewing*, 111
  - health scores, 122–125, 592–593
    - overview of, 110–111
    - policies for, 118–120
  - Health Score Evaluation policy, 125
  - health scores, 122–125, 592–593
  - HER (headend replication) tunnels, 32
  - hierarchy
    - access policies, 183–184
    - tenant, 153–154
  - Hold Interval (BGP timer policy), 338
  - Host Name parameter
    - LDAP providers, 542
    - RADIUS providers, 540
    - TACACS+ providers, 532
  - host routes, advertisement of, 321
  - hosts
    - ESXi
      - adding to VDS (vSphere distributed switches)*, 411–414
      - connecting in UCS domains*, 407
      - connecting to fabric*, 407
      - definition of*, 395
      - standard switch architecture*, 395
      - VMkernel interface*, 397–399
      - vSphere distributed switches (VDS)*, 397
    - silent, 245–247, 248–249
  - HSRP (Hot Standby Router Protocol), 344
  - HTTP (Hypertext Transfer Protocol), 67
  - hybrid mode, service graph
    - deployment, 428, 432–434
- 
- iBGP Distance parameter (BGP timer policy), 342
  - iBGP peerings, 303–304
  - Identity Service Engine (ISE), 536–538
  - IDs
    - BGP EVPN router, 556, 590
    - node, 11, 595
    - VLAN, 245–247, 253–254
    - VNIDs, 244, 245–247
  - IEEE (Institute of Electrical and Electronics Engineers) 802.1Q standard, 7
  - IFM (intra-fabric messaging), 50–51, 593
  - IGMP (Internet Group Management Protocol) snooping, 368
  - Ignore Fault window, 122
  - IMC (Integrated Management Controller), 52–53, 100
  - immediacy settings, VMM (Virtual Machine Manager), 405–407
  - Import Mode parameter, 75
  - Import Type parameter, 75

**imports**

- ACI Multi-Site, 557
- configuration backups, 80–82
- configuration imports, 75–76
- device packages, 430, 441–444

**Inactive state (fabric node discovery), 52****in-band management, 462**

- configuration
  - access policies for APIC in-band interfaces, 468–469*
  - APIC connectivity preferences, 478–479*
  - bridge domains, 469*
  - in-band network, extending out of fabric, 474–476*
  - IP addressing, 470–474*
  - steps for, 467–468*
  - whitelisting desired connectivity to/from, 476–477*

- connectivity, 465–467
- deployment alongside out-of-band (OOB) management, 467
- out-of-band (OOB) management compared to, 48

**inbound route filtering, 352–353****Info faults, 112****informs, SNMP (Simple Network Management Protocol), 501****infra IP reachability, 32****Infra port usage, 195****infra tenants, 134–135, 292****infrastructure VLAN, 49, 593****ingress replication, 28****ingress routing optimization**

- cross-data center ingress routing optimization, 561–562
- primary and disaster recovery data center deployments, 561

**Initial Delay setting (Global MCP Policy), 222****initialization of ACI fabric. *See* ACI fabric initialization****Institute of Electrical and Electronics Engineers 802.1Q standard, 7****integrated health monitoring. *See* health monitoring****Integrated Management Controller. *See* IMC (Integrated Management Controller)****interface bridge-domain  
BD-CRITICAL-STUFF command, 147****interface policies, 169–174**

- CDP (Cisco Discovery Protocol), 190
- definition of, 593
- deployment of, 173–174
- EIGRP (Enhanced Interior Gateway Routing Protocol), 321–323
- link level, 189
- LLDP (Link Layer Discovery Protocol), 189
- OSPF (Open Shortest Path First), 326–327
- overview of, 169–170
- port channel, 198
- Storm Control, 223–224
- types of, 170–173
- vPC (virtual port channel), 205–206

**interface policy groups, 169–174**

- definition of, 593
- overview of, 169–170
- types of, 170–171

**interface profiles, 179, 294**

- configuration, 179–182
- definition of, 593
- logical, 295, 594

**interface selectors, 179, 593**

interleak, 305, 593

Intermediate System-to-Intermediate System (IS-IS), 23, 60, 106

Internal Distance setting (EIGRP Context per Address Family), 325

internal VLAN ranges, 160

Internet Group Management Protocol (IGMP) snooping, 368

interpod network (IPN), 25

intersite L3Out, 557, 593

intersite networks. *See* ISNs (intersite networks), ACI Multi-Site

intra-fabric messaging (IFM), 50–51, 593

IP (Internet Protocol)

- addresses, 470–474
  - dynamic*, 472–474
  - secondary*, 325–326, 597
  - static*, 470–472
- IP hash, routes based on, 402
- IP storage leaf switches, 24, 593
- IPv6, 241, 344
- service-level agreement (SLA) tracking, 330–334

ip address command, 147

IP Address parameter

- LDAP providers, 542
- RADIUS providers, 540
- TACACS+ providers, 532

IP storage leaf switches, 24, 593

IPN (interpod network), 25

ISE (Identity Service Engine), 536–538

IS-IS (Intermediate System-to-Intermediate System), 23, 60, 106

ISNs (intersite networks), ACI Multi-Site, 27, 555–558

- cross-site connectivity, 555–557
- extending ACI to remote locations and public clouds, 564

- locally governed configurations, 557
- MSO-governed configurations, 557
- primary and disaster recovery data center deployments, 558–562
- schemas, 557–558
- stateful-services integration, 563–564
- terminology for, 557

ITSM (IT service management), 110

## J-K

---

JavaScript Object Notation (JSON), 103

JSON (JavaScript Object Notation), 103

Keepalive Interval (BGP timer policy), 338

Key parameter

- Global MCP Policy, 222
- RADIUS providers, 540
- TACACS+ providers, 532

keychain policy, 324

KVM (Keyboard Video Mouse), 100

## L

---

L2 interface policy, 172

L2 Unknown Unicast field (bridge domain settings), 368, 593

L2Outs, 380–385

L3 APIC (Application Policy Infrastructure Controller) specifications, 33

L3 domain implementation, 301–302

L3 Unknown Multicast Flooding field (bridge domain settings), 369

L3Out Creation Wizard

- AAEPs (attachable access entity profiles), 302

configuration for BGP peering,  
334–337

configuration for EIGRP peering,  
307–309

configuration for OSPF peering,  
325–327

## L3Outs, 153, 284

anatomy of, 293–295

configuration for BGP peering

- BFD (Bidirectional Forwarding Detection) implementation*, 339–341
- BGP timer policies*, 338–339, 342
- L3Out Creation Wizard*, 334–337
- node-level customizations*, 337–339
- OSPF for IP reachability*, 343
- peer connectivity profiles*, 339–341
- per-neighbor customizations*, 339–341
- VRF-level customizations*, 342–343

configuration for EIGRP peering,  
305–306

- advertisement of host routes*, 321
- advertisement of subnets assigned to bridge domains*, 314–315
- authentication*, 324
- BFD (Bidirectional Forwarding Detection) implementation*, 321–324
- blacklist EPG (endpoint group) with logging*, 318–320
- contracts*, 316–318
- EIGRP Context per Address Family*, 324–325

- external EPG (endpoint group) deployment*, 310–312
- forwarding verification*, 312–314
- L3Out Creation Wizard*, 307–309
- route advertisement*, 328
- VRF-level customizations*, 324–325

configuration for OSPF peering, 329

- BFD (Bidirectional Forwarding Detection) implementation*, 328–329
- L3Out Creation Wizard*, 325–327
- OSPFv3 support*, 344
- route advertisement*, 328
- VRF-level customizations*, 329

definition of, 593–594

endpoint learning on, 249

HSRP (Hot Standby Router Protocol), 344

interface types, 296

intersite, 557, 593

IPv6 and, 344

key functions of, 293

L3Out BDs (bridge domains), 296–298, 594

logical node and interface profiles, 295, 594

node and interface profiles, 294

prerequisites for

- BGP (Border Gateway Protocol) route reflection*, 303–305, 597
- infra MP-BGP route distribution*, 305–306
- L3 domain implementation examples*, 301–302
- overview of*, 301

- route control
  - inbound route filtering*, 352–353
  - overview of*, 344–345
  - policies for*, 351–352
  - route attributes to all peers*, 346–349
  - route attributes to specific peer*, 349–351
  - route profiles*, 344–345, 597
- static routes on
  - IP service-level agreement (SLA) tracking*, 330–334
  - overview of*, 329–330
- stub network and transit routing, 291–292
- SVI Auto State, 299–300
- SVI Encap Scope, 298–299
- types of, 292–293
- VLAN pool assignment, 301–302
- L4-L7 devices, identifying to fabric**, 443–444
- L4-L7 services**, 106, 435–436
- LACP (Link Aggregation Control Protocol)**, 202, 419–421
- LAN on motherboard (LOM) ports**, 47
- latency, Remote Leaf deployments**, 31
- Layer 2 extension to non-ACI switches**
  - ACI interaction with STP (Spanning Tree Protocol), 386–388
  - link types*, 388
  - MCP (MisCabling Protocol)*, 388
  - MST (Multiple Spanning Tree) instance mappings*, 387–388
  - TCNs (topology change notifications)*, 386–387
- Layer 2 connectivity to non-ACI switches**
  - bridge domain extensions*, 374–375
  - EPG (endpoint group) extensions*, 372–373, 375–380
  - L2Outs*, 380–385
  - migration of overlapping VLANs into ACI*, 385–386
  - overview of*, 372
- network migrations into ACI**
  - Any EPG (vzAny)*, 364–365, 589
  - contract enforcement at VRF instance level*, 367
  - endpoint learning considerations*, 371–372
  - flooding requirements*, 368–370
  - full-mesh network-centric contracts*, 362–364
  - GARP-based detection*, 370–371
  - Legacy mode*, 371
  - network-centric deployments*, 361–362
  - overview of*, 361
  - preferred group members*, 365–367
  - preparation for*, 372
  - security enforcement in*, 362
- overview of**, 358
- Layer 3 Outside**. *See* L3Outs
- LDAP (Lightweight Directory Access Protocol)**
  - ACI configuration for, 541–547
  - Create LDAP Group Map Rules wizard, 543–547
  - Create LDAP Provider wizard, 543
- leaf access port policy groups**, 170, 191
- leaf breakout port groups**, 171
- leaf interface override policy**, 232, 594
- leaf selectors**, 181
- leaf switches**, 9–10
  - border, 24, 590
  - cabling requirements, 45

- compute, 24, 590
- definition of, 22, 594
- first-generation, 38–39
- initialization of, 51
- IP storage, 24, 593
- leaf access port policy groups, 191
- overview of, 23
- packet forwarding scenarios, 251
  - both endpoints attach to same leaf, 251–254*
  - flooding to unknown destination, 261–263*
  - known destination behind another leaf, 254–258*
  - spine proxy to unknown destination, 258–261*
- pervasive gateways, 252–254, 595
- pervasive routes, 258–260, 595
- profiles, 179–182
- purchasing decisions, 38
- second-generation, 39–40
- service, 24, 597
- transit, 25, 599
- leaf topologies. *See* topologies
- leaking, route, 466
- learning, endpoint. *See* endpoint learning
- Legacy mode, 371
- Lightweight Directory Access Protocol. *See* LDAP (Lightweight Directory Access Protocol)
- Limit IP Learning to Subnet feature, 249–250
- Link Aggregation Control Protocol (LACP), 202
- link debounce interval, 189, 594
- Link Layer Discovery Protocol. *See* LLDP (Link Layer Discovery Protocol)
- link level interface policies, 171, 189
- Link Status Only option, 402
- links, STP (Spanning Tree Protocol), 388
- LLDP (Link Layer Discovery Protocol)
  - interface policies, 171, 189
  - neighbor discovery, 50
- load balancing, 600
  - one-arm load balancers, 437
  - vSphere, 401, 407, 419
- Local Distance parameter (BGP timer policy), 342
- local endpoints, 242–243, 594
- local users
  - access
    - assigning, 523–525*
    - modifying, 523–525*
  - creating, 521–523
- Local-AS Number Config parameter (BGP peer connectivity profiles), 341
- Local-AS Number parameter (BGP peer connectivity profiles), 341
- locally governed ACI Multi-Site configurations, 557
- locked configuration items (function profiles), 445
- logical devices, 444, 446, 594. *See also* device selection policies
- logical node and interface profiles, 295, 594
- login domains, 549–550
- login IDs, 522
- logit severity command, 498–499
- logs
  - audit logs
    - definition of, 127*
    - reviewing, 127–128*
  - blacklist EPG (endpoint group) with logging, 318–320

longest prefix match (LPM), 175  
 lookups, 90–92, 241–242  
 Loop Detect Multiplication Factor  
 setting (Global MCP Policy), 222  
 Loop Protection Action setting (Global  
 MCP Policy), 222  
 loopback interfaces, 60  
 LPM (longest prefix match), 175

## M

---

M3 APIC (Application Policy  
 Infrastructure Controller)  
 specifications, 33  
 MAC address tables, 241  
 MAC hashes, routes based on, 402  
 MACsec, 173  
 Maintenance state (fabric node  
 discovery), 52  
 Major faults, 112  
 Manage VMkernel Adapters page (Add  
 and Manage Hosts wizard), 413  
 managed devices, deploying service  
 graphs for, 453–460  
 managed mode, 428, 430–432  
 managed objects (MOs), 109, 594  
 management  
 in-band, 462  
   *access policies for APIC in-band  
   interfaces, 468–469*  
   *APIC connectivity preferences,  
   478–479*  
   *bridge domains, 469*  
   *compared to out-of-band (OOB)  
   management, 48*  
   *configuration steps for, 467–468*  
   *connectivity, 465–467*  
   *deployment alongside out-of-  
   band (OOB) management, 467*

*in-band network, extending out  
   of fabric, 474–476*  
   *IP addressing, 470–474*  
   *whitelisting desired connectivity  
   to/from, 476–477*  
 management access modification,  
 103–105  
 MIBs (Management Information  
 Bases), 501–502, 594  
 MIM (Management Information  
 Model), 105–107  
 multitenancy, 13  
 out-of-band (OOB), 462  
   *compared to in-band  
   management, 48*  
   *connectivity, 464–465*  
   *contracts, 479–480*  
   *deployment alongside in-band  
   management, 467*  
 service graph management models  
   *definition of, 597*  
   *network policy mode, 428–430,  
   595*  
   *overview of, 428*  
   *service manager mode, 428,  
   430–432–434, 597*  
   *service policy mode, 428,  
   430–432, 597*  
   *templates, 597*  
 of traditional networks, 4–7  
**Management and Analytics System  
 (MAS), 432**  
**Management EPG parameter**  
 LDAP providers, 542  
 RADIUS providers, 541  
 TACACS+ providers, 533  
 managers, SNMP (Simple Network  
 Management Protocol), 500  
 mandatory configuration items  
 (function profiles), 445

- many-to-many relationships, 153
- mapping ports, 191–194, 267–269
- marking traffic, 218
- Maximum AS Limit setting (BGP timer policy), 338
- Maximum ECMP setting (OSPF timer policy), 329
- Maximum Path Limit setting (EIGRP Context per Address Family), 325
- maximum transmission units (MTUs), 25, 31
- MCP (MisCabling Protocol), 388
  - configuration, 221–222
  - definition of, 594
  - interface policies, 172
- MEC (Multichassis EtherChannel) technology, 6
- Merge import type, 75
- mergers, multitenancy and, 13
- messages, system. *See* system messages
- Metric Style setting (EIGRP Context per Address Family), 325
- mgmt security domain, 517
- mgmt tenants, 135
- MIBs (Management Information Bases), 501–502
  - definition of, 594
  - scope of, 502, 592
  - VRF-specific, 600
- microsegmentation, 146
- Migrate VM Networking page (Add and Manage Hosts wizard), 414
- migration, VLAN, 385–386
- MIM (Management Information Model), 105–107, 594
- Minor faults, 112
- MisCabling Protocol (MCP), 221–222, 388, 594

## modes

- Pearson IT Certification test prep software, 568–569
- service graph deployment
  - network policy mode*, 428–430, 595
  - service manager mode*, 428, 430–432–434, 597
  - service policy mode*, 428, 430–432, 597

monCommonPol class, 118

monEPGPol class, 118

monFabricPol class, 118

monInfraPol class, 118

## monitoring

- health. *See* health monitoring
- message forwarding to syslog servers, 485–492
  - ACI severity levels*, 485
  - contracts for*, 487–491
  - NX-OS severity levels*, 485
  - steps for*, 487
  - syslog monitoring destination groups*, 492–493
  - syslog sources*, 494–498
  - verification of*, 498–500
- service graphs, 447
- SNMP (Simple Network Management Protocol)
  - ACI support for*, 501–502
  - client group policies*, 504–506
  - configuration caveats*, 502
  - contexts*, 506
  - contracts*, 503
  - informs*, 501
  - SNMP agents*, 500
  - SNMP managers*, 500
  - SNMP monitoring destination groups*, 507–508

*SNMP sources for desired monitoring policies*, 508  
*steps for*, 502–503  
*traps*, 501  
*verification of*, 509–511  
*versions of*, 500–501  
**monitoring destination group**, 594  
**monitoring sources**, 594  
**monPolDn attribute**, 116–117, 594  
**MOQuery**, 110, 128, 499–500, 510–511, 595  
**MOs (managed objects)**, 109, 594  
**MP-BGP (Multiprotocol BGP)**, 23, 305–306  
**MP-BGP EVPN (Multiprotocol Border Gateway Protocol Ethernet Virtual Private Network)**, 12  
**MSO (Multi-Site Orchestrator)**, 15, 27, 555, 557  
**MST (Multiple Spanning Tree) instance mappings**, 387–388  
**MTUs (maximum transmission units)**, 25, 31  
**Multi Destination Flooding field (bridge domain settings)**, 370, 595  
**Multicast PIM-Bidir support**, 32  
**Multichassis EtherChannel (MEC) technology**, 6  
**Multiple Spanning Tree (MST) instance mappings**, 387–388  
**Multiprotocol BGP (MP-BGP)**, 23  
**Multiprotocol Border Gateway Protocol Ethernet Virtual Private Network (MP-BGP EVPN)**, 12  
**Multi-Site**. *See* ACI Multi-Site  
**Multi-Site Orchestrator (MSO)**, 15, 27, 555, 557  
**multitenancy**, 11–14, 549–550, 595.  
*See also* tenants

**multitier application deployment**  
 application profiles, 264–266  
 BD (bridge domain) configuration, 264–266  
 DHCP (Dynamic Host Configuration Protocol) relay, 270  
 domain assignment, 267  
 EPG (endpoint group) configuration, 264–266  
 EPG (endpoint group) mapping to ports, 267–269, 270  
 overview of, 263  
 policy deployment, 270  
 verification of EPG-to-port assignments, 269–270

## N

---

**names, distinguished**, 109, 591  
**namespace normalization**, 28  
**navigation, APIC CLI (command-line interface)**, 101  
**ND (Neighbor Discovery)**, 241  
**nesting tenants**, 134  
**NetFlow**, 171, 174  
**netstat command**, 86–87  
**Network Failure Detection setting**, 402  
**network management touchpoints**, 9–10  
**network migrations into ACI**  
 Any EPG (vzAny), 364–365, 589  
 contract enforcement at VRF instance level, 367  
 endpoint learning considerations, 371–372  
 flooding requirements, 368–370  
*ARP Flooding*, 369  
*Forwarding setting*, 368  
*L2 Unknown Unicast*, 368

- L3 Unknown Multicast Flooding*, 369
- Multi Destination Flooding*, 370
- full-mesh network-centric contracts, 362–364
- GARP-based detection, 370–371
- Legacy mode, 371
- network-centric deployments, 361–362
- overview of, 361
- preferred group members, 365–367
- preparation for, 372
- security enforcement in, 362
- network policy mode, 428–430, 595
- Network Time Protocol. *See* NTP (Network Time Protocol)
- network-centric deployments, 361–362
- networks. *See also* VLANs (virtual LANs); vSphere
  - software-defined, 4
  - traditional
    - agility of*, 8
    - management of*, 4–5
    - scalability and growth of*, 5–7
    - security of*, 8–9
    - three-tier design of*, 5–7
    - visibility of*, 9
- New Distributed Port Group wizard, 400
- Next-hop Self parameter (BGP peer connectivity profiles), 340
- Nexus 1000v distributed virtual switch, 9
- Nexus 9332C switches, 38
- Nexus 9336C-FX2 switches, 40, 215
- Nexus 9336PQ switches, 37
- Nexus 9364C switches, 38
- Nexus 9732C-EX switches, 38
- Nexus 9736C-FX switches, 38
- Nexus 9736PQ switches, 37
- Nexus 93180YC-EX switches, 40, 215
- Nexus 93180YC-FX switches, 40, 45
- Nexus 93240YC-FX2 switches, 40
- Nexus 93360YC-FX2 switches, 40
- NIC loads, routes based on, 402
- no password `pwd-strength-check` command, 522–523
- node ID, 11, 595
- node management, 472–473
- node profiles, 294, 295
- Node Registration wizard, 57
- node-level BGP (Border Gateway Protocol) customizations, 337–339
- non-ACI switches, Layer 2 extension to
  - ACI interaction with STP (Spanning Tree Protocol), 386–388
    - link types*, 388
    - MCP (MisCabling Protocol)*, 388
    - MST (Multiple Spanning Tree) instance mappings*, 387–388
    - TCNs (topology change notifications)*, 386–387
- Layer 2 connectivity to non-ACI switches
  - bridge domain extensions*, 374–375
  - EPG (endpoint group) extensions*, 372–373, 375–380
  - L2Outs*, 380–385
  - migration of overlapping VLANs into ACI*, 385–386
  - overview of*, 372
- network migrations into ACI
  - Any EPG (vzAny)*, 364–365, 589
  - contract enforcement at VRF instance level*, 367
  - endpoint learning considerations*, 371–372

- flooding requirements*, 368–370
  - full-mesh network-centric contracts*, 362–364
  - GARP-based detection*, 370–371
  - Legacy mode*, 371
  - network-centric deployments*, 361–362
  - overview of*, 361
  - preferred group members*, 365–367
  - preparation for*, 372
  - security enforcement in*, 362
  - overview of, 358
  - north-south traffic, 7
  - Notify Switches setting, 403
  - NSX-T VDS, 395
  - NTP (Network Time Protocol), 84–89, 106
  - ntpstat command, 87
  - N-VDS, 395
  - nw-svc-admin role, 520
  - nw-svc-params role, 520
  - NX-OS, system message severity levels, 485
- O**
- 
- object model, 96
    - correlation of menus to, 107–108
    - exploring with Visore, 108–110
    - faults in, 116–118
    - importance of, 110
    - overview of, 105–107
    - topRoot object, 106
  - On Demand immediacy, 407
  - on-demand backups, 76–79
  - one-arm load balancers, 437
  - one-to-many relationships, 153
  - online access, Pearson IT Certification test prep software, 567
  - Operational faults, 113
  - operations, ACI Anywhere. *See* ACI Anywhere
  - ops role, 520
  - OSPF (Open Shortest Path First)
    - configuration for BGP L3Out, 344
    - L3Out for OSPF peering
      - BFD (Bidirectional Forwarding Detection) implementation*, 328–329
      - L3Out Creation Wizard*, 325–327
      - OSPF timer policy*, 329
      - OSPFv3 support*, 344
      - route advertisement*, 328
      - VRF-level customizations*, 329
    - Remote Leaf deployments, 32
  - OTV (Overlay Transport Virtualization), 560
  - out-of-band (OOB) management, 462
    - ACI fabric initialization and
      - automatic upgrades of new switches*, 74–75
      - default contracts*, 64–66
      - fabric upgrades*, 66–71
      - schedulers*, 73, 597
      - static out-of-band addresses, assignment of*, 63–64
    - APICs (Application Policy Infrastructure Controllers), 47–48
    - connectivity, 464–465
    - contracts, 479–480
    - deployment alongside in-band management, 467
    - in-band management compared to, 48
  - outside EPGs (endpoint groups), 151–153

overlapping VLANs, migration into  
ACI, 385–386

overlay, 135

overlay multicast TEP, 555, 595

Overlay Transport Virtualization  
(OTV), 560

overlay unicast TEP, 555, 595

overrides

leaf interface, 232, 594

port channel member, 232–235, 596

## P

---

packages, device. *See* device packages,  
importing

packet forwarding scenarios, 251

both endpoints attach to same leaf,  
251–254

flooding to unknown destination,  
261–263

known destination behind another  
leaf, 254–258

*learning and forwarding for  
vPCs, 256–258*

*verification of traffic path  
between known endpoints,  
254–256*

spine proxy to unknown destination,  
258–261

partnerships, multitenancy and, 13

Passive Participation option (OSPF),  
329

Password parameter (LDAP providers),  
542

Password Strength parameter (ACI  
fabric initialization), 49

Password Strength Policy, 522–523

Password/Confirm Password  
parameter (BGP peer connectivity  
profiles), 340

passwords

BIOS, 52

strength of, 522–523

PBR (policy-based redirect), 141, 427

PC/VPC override policy groups, 171

Pearson IT Certification test prep  
software, 567–568

modes, 568–569

offline access, 568

online access, 567

overview of, 567

Premium Edition, 569–570

updating, 569

peer connectivity profiles (BGP),  
339–341

per-bridge domain settings, 560, 561

per-neighbor BGP (Border Gateway  
Protocol) customizations, 339–341

pervasive gateways, 252–254, 595

pervasive routes, 258–260, 595

Per-VLAN Spanning Tree (PVST+),  
387–388

PFC (priority-based flow control), 172

physical domains, 161, 595

physical tunnel endpoint (PTEP), 60,  
595

PI VLANs (platform-independent  
VLANs), 246, 253–254, 595

Ping Check setting (AAA policy), 548

ping command, 364, 474

planning

ACI fabric initialization, 42, 44–45,  
47–48, 51

*APIC connectivity, 46–47*

*APIC OOB configuration  
requirements, 47–48*

*basic configuration parameters,  
48–49*

*cabling requirements, 45*

- fabric discovery stages*, 51–52
- out-of-band versus in-band management*, 48
- switch discovery process*, 49–51
- switch discovery states*, 51–52
- VMM (Virtual Machine Manager) integrations, 403–405
- platform-independent VLANs (PI VLANs)**, 246, 253–254, 595
- platforms, cross-platform integrations**, 15
- Pod 1 TEP Pool parameter (ACI fabric initialization)**, 49
- pod ID parameter (ACI fabric initialization)**, 48
- pod policy**, 83
- pod policy groups**, 83–86, 595
- pod profiles**, 83–84, 596
- pod selectors**, 83, 596
- PoE (Power over Ethernet)**, 172, 175
- policies. *See also* access policies; tenant policies**
  - AAA authentication, 547–550
  - date/time, 84–89
  - definition of, 110
  - device selection, 446
  - fabric, 106, 592
  - fault management, 120–121
  - health monitoring, 118–120
  - keychain, 324
  - LACP (Link Aggregation Control Protocol), 419–421
  - password strength, 522–523
  - pod, 83
  - Policy Universe, 105, 106, 109, 596
  - primary and disaster recovery data center deployments, 561–562
  - reuse of, 16
  - route control, 351–352
  - SNMP (Simple Network Management Protocol), 504–506
  - timer
    - BGP (Border Gateway Protocol)*, 337–339, 342
    - OSPF (Open Shortest Path First)*, 329
  - VMware vSphere, 400–403
- policing**, 218
- Policy Universe**, 105, 106, 109, 596
- policy-based redirect (PBR)**, 141, 427
- pools**
  - TEP (tunnel endpoint), 49, 50, 57, 59, 599
  - VLAN (virtual LAN), 159–160, 599
    - creating*, 159–160
    - designs for*, 161–163
    - overlap between*, 164–165
    - static versus dynamic VLAN allocation*, 159, 598
    - VLAN ranges in*, 160
- port bindings**, 400, 596
- port blocks**, 198–199
- port channel member overrides**, 232–235, 596
- port channels**
  - configuration, 196–201
  - interface policies and interface policy groups, 170, 171
  - member overrides, 232–235, 596
  - members, 172
  - vPCs (virtual port channels), 6
    - configuration*, 204–207
    - domains*, 201–204
    - learning and forwarding for*, 256–258
    - VIP addresses*, 245

**port encapsulation VLANs, 246, 596**

**port groups**

definition of, 395–397, 596

distributed, 397

*assigning VMs (virtual machines) to, 417–418*

*pushing EPGs (endpoint groups) to vCenter as, 415–416*

**Port Local Scope, 385–386**

**Port parameter**

LDAP providers, 542

TACACS+ providers, 532

**port security interface policies, 173**

**PortFast, 387**

**ports**

ACI switch port configurations

*AAEP EPGs (endpoint groups), 207–210*

*implications of, 210*

*individual port configuration, 188–196*

*overview of, 186–188*

*port channel configuration, 196–201*

*vPC (virtual port channel) configuration, 204–207*

*vPC (virtual port channel) domain configuration, 201–204*

APIC (Application Policy Infrastructure Controller), 46–47, 589

bindings, 400, 596

blocks, 198–199

dynamic breakout ports, 215–217, 591

encapsulation, 246

fabric, 23, 592

mappings, 191–194, 267–269

port groups, 395–397

*definition of, 596*

*distributed, 397, 415–418*

Port Local Scope, 385–386

usage types, 195

VMware vSphere, 400

**Power over Ethernet (PoE), 172, 175**

**power supplies, 46**

**Practice Exam mode, Pearson IT**

**Certification test prep software, 568–569**

**preferred group members, 365–367, 596**

**Prefix Suppression setting (OSPF timer policy), 329**

**Premium Edition, Pearson IT**

**Certification test prep software, 569–570**

**Pre-Provision immediacy, 405–406**

**primary data center deployments, 558–562**

centralized orchestration and management capabilities, 559

cross-data center ingress routing optimization, 561–562

overview of, 558–559

per-bridge domain broadcast and stretch settings, 560

**priority flow control interface policies, 172**

**priority levels, 218**

**priority-based flow control (PFC), 172**

**private networks. See VRF (virtual routing and forwarding) instances**

**privileges**

definition of, 596

expanding, 523–525

for predefined roles, 519–521

**Production/EPG-Admin/ value, 540**

**profiles, 176–179**

- AAEPs (attachable access entity profiles), 165–169, 382
  - configuration, 167–169, 190–191*
  - definition of, 589*
  - EPGs (endpoint groups), 207–210*
  - in-band management, 468*
  - in L3Out domains, 302*
  - overview of, 165–166*
- application, 527
  - advantages of, 138–139*
  - configuration, 142–145*
  - definition of, 138, 589*
  - multitier application deployment, 264–266*
  - topology, 138–139*
- custom function, 444–445
- DNS (Domain Name System) server, 86–92
- external management network instance, 491
- FEX (fabric extenders), 213–214
- forwarding scale, 175
- function, 592
  - custom, 444–445*
  - examples of, 430–433, 442*
- interface, 179–182
- logical node and interface, 295, 594
- pod, 83–84, 596
- route, 344–345, 597
- switch, 179–182, 598
- TACACS+, 537–538
- VMM (Virtual Machine Manager) domain, 408–410
- programmability, ACI (Application Centric Infrastructure), 11**
- providers, 149–150, 596**

**proxy, hardware, 247–248**

- proxy TEP addresses, 61**
- PTEP (physical tunnel endpoint), 60, 595**
- public clouds, extending ACI to, 564–565**
  - with ACI Multi-Site, 564
  - with remote leaf switches, 564–565
  - with vPod, 564
- PVST+ (Per-VLAN Spanning Tree), 387–388**
- Python SDK (software development kit), 103**

## Q-R

---

- QoS (quality of service), 217–219**
- queueing, 218**
- RADIUS, ACI configuration for, 540–541**
- Raised state (faults), 114**
- Raised-Clearing state (faults), 115**
- ranges, VLAN, 160**
- Rapid PVST+ (Rapid Per-VLAN Spanning Tree), 387–388**
- RARP (Reverse ARP), 403**
- RAT (Resource Allocation Timeout), 175**
- RBAC (role-based access control), 133, 514**
  - access, 523–525–528
  - common pitfalls in, 531
  - custom rules, 528–531
  - local users
    - assigning access to, 523–525*
    - creating, 521–523*
    - modifying access for, 523–525*
  - multitenancy and, 13
  - overview of, 516–517

- privileges
  - expanding, 523–525*
  - for predefined roles, 519–521*
- roles
  - modifying, 523–525–528*
  - predefined, 519–521*
- rules, 596
- security domains, 517–518
  - assigning tenants to, 517–518*
  - creating, 517*
  - definition of, 517*
  - predefined, 517*
- read queries, 501
- read-all role, 520
- recovery, Error Disabled Recovery Policy, 231–232
- relay, DHCP (Dynamic Host Configuration Protocol), 219–221, 271
- reload command, 100
- remote endpoint learning, 244–245
- remote endpoints, 242–243, 596
- remote leaf switches, 564–565
- Remote Leaf topology, 30–32
- remote locations, extending ACI to,
  - 564–565
  - with ACI Multi-Site, 564
  - with remote leaf switches, 564–565
  - with vPod, 564
- Remote User Login Policy setting (AAA policy), 548
- remoteleaf role, 57
- Remove All Private AS parameter (BGP peer connectivity profiles), 340
- Remove Private AS parameter (BGP peer connectivity profiles), 340
- Replace import type, 75
- Replace Private AS with Local AS parameter (BGP peer connectivity profiles), 341
- representational state transfer. *See* REST (representational state transfer) API
- resolution immediacy, 405–407, 596
- Resource Allocation Timeout (RAT), 175
- REST (representational state transfer) API, 11, 98
- restores
  - configuration backups, importing, 80–82
  - configuration exports, 75
  - configuration imports, 75–76
  - configuration rollbacks, 82–83
  - configuration snapshots, 80
  - on-demand backups, 76–79
  - scheduled backups, 79–80
- Retaining state (faults), 115
- Retention Interval timer, 115
- Retries parameter
  - LDAP providers, 542
  - RADIUS providers, 540
  - TACACS+ providers, 532
- Reverse ARP (RARP), 403
- Reverse Filter Ports feature, 277–278, 476
- RIB (Routing Information Base), 241, 242
- RJ-45 connectors, 46
- role-based access control. *See* RBAC (role-based access control)
- roles
  - definition of, 596
  - modifying, 523–525–528
  - predefined, 519–521
- Rollback to This Configuration option, 82
- rollbacks, configuration, 82–83

**roots, L3Out, 294**  
**route advertisement**  
   for EIGRP (Enhanced Interior Gateway Routing Protocol) L3Outs, 328  
   for OSPF (Open Shortest Path First) L3Outs, 328  
**route control**  
   inbound route filtering, 352–353  
   overview of, 344–345  
   policies for, 351–352  
   route attributes to all peers, 346–349  
   route attributes to specific peer, 349–351  
   route profiles, 344–345, 597  
**Route Control Profile parameter (BGP peer connectivity profiles), 341**  
**route leaking, 466**  
**route peering, 438–439**  
**route profiles, 597**  
**route reflection, BGP (Border Gateway Protocol)**  
   definition of, 597  
   implementation of, 304–305  
   infra MP-BGP route distribution, 305–306  
   need for, 303–304  
**routed mode, devices in, 435**  
**Routing Information Base (RIB), 241, 242**  
**rules, RBAC (role-based access control), 528–531, 596**  
**running configuration, viewing, 102**

## S

---

**satellite/small colo data centers, Remote Leaf deployments for, 30**  
**Save & Exit Setup dialog box, 52**

**scalability**  
   ACI (Application Centric Infrastructure), 10  
   APIC (Application Policy Infrastructure Controller) clusters, 33–36  
   traditional networks, 5–7  
**Schedule Controller Upgrade page, 67–69**  
**Schedule Node Upgrade window, 70**  
**scheduled backups, 79–80**  
**schedulers, 73, 597**  
**schemas, 557–558, 597**  
**scope**  
   of contracts, 150–151, 284–285, 590  
   of MIBs (Management Information Bases), 502, 592  
   Port Local, 385–386  
   SVI Encap, 298–299  
**SCP (Secure Copy Protocol), 67**  
**scripts, setup-clean-config.sh, 100**  
**SDN (software-defined networking), 4**  
**secondary IP addresses, 325–326, 597**  
**secondary keyword, 147**  
**second-generation leaf switches, 39–40**  
**second-generation spine switches, 37–38**  
**Secure Copy Protocol (SCP), 67**  
**security domains, 517–518**  
   assigning tenants to, 517–518  
   creating, 517  
   definition of, 517, 597  
   mapping domains to, 526–528  
   predefined, 517  
**seed leaf initialization, 51**  
**Select New Hosts page (Add and Manage Hosts wizard), 412**

- selectors, 83, 176–179
- Self Nexthop, 324
- Send Community parameter (BGP peer connectivity profiles), 340
- Send Extended Community parameter (BGP peer connectivity profiles), 340
- serial ports (RJ-45 connector), 46
- Server Monitoring parameter
  - LDAP providers, 542
  - RADIUS providers, 541
  - TACACS+ providers, 533
- servers
  - DNS (Domain Name System), 90–92
  - ESXi
    - adding to VDS (vSphere distributed switches)*, 411–414
    - connecting in UCS domains*, 407
    - connecting to fabric*, 407
    - definition of*, 395
    - standard switch architecture*, 395
    - VMkernel interface*, 397–399
    - vSphere distributed switches (VDS)*, 397
  - external AAA server integration, 532
    - AAA authentication policy settings*, 547–550
    - ACI configuration for LDAP*, 541–547
    - ACI configuration for RADIUS*, 540–541
    - ACI configuration for TACACS+*, 532–536
    - Cisco AV pair formatting*, 538–540
    - ISE (Identity Service Engine) configuration*, 536–538
  - syslog, forwarding system messages to
    - contracts for*, 487–491
    - steps for*, 487
    - syslog monitoring destination groups*, 492–493
    - syslog sources*, 492–493
    - verification of*, 498–500
  - service chains, 440, 450, 452, 597
  - service graphs, 424
    - bridge domain settings for, 439–440
    - as concatenation of functions, 427–428
    - contracts for, 427–428
    - definition of, 597
    - deployment
      - for devices in GoThrough Mode*, 436, 437
      - for devices in GoTo Mode*, 435, 436–437
      - for one-arm load balancers*, 437
    - dynamic endpoint attach, 439
    - implementation examples
      - service graphs for managed devices*, 453–460
      - unmanaged firewall pair deployment*, 447–452
    - implementation workflow
      - custom function profiles*, 444–445
      - device package imports*, 430, 441–442
      - device selection policy configuration*, 446
      - L4-L7 device identification*, 443–444
      - overview of*, 441
      - service graph monitoring*, 447
      - service graph parameters*, 447
      - service graph template configuration*, 445–446

- service graph template instantiation*, 446–447
- L4-L7 services integration, 435–436
- management models
  - network policy mode*, 428–430, 595
  - overview of*, 428
  - service manager mode*, 428, 430–434, 597
  - service policy mode*, 428, 430–432, 597
- monitoring, 447
- overview of, 426–427
- parameters, 447
- with policy-based redirect (PBR), 427
- rendering, 440–441
- route peering, 438–439
- templates
  - configuration*, 445–446
  - definition of*, 597
  - instantiation of*, 446–447
- when to use, 434–435
- without policy-based redirect (PBR), 427
- service leaf switches, 24, 597
- service manager mode, 428, 430–432, 597
- service policy mode, 428, 430–432, 597
- service VM orchestration, 428
- service-level agreement (SLA) tracking, 330–334
- setup-clean-config.sh, 100
- severity levels
  - ACI, 485
  - faults, 111–112, 121–124
  - NX-OS, 485
- shadow EPGs (endpoint groups), 428–429, 597
- sharding, database, 34–35, 597–598
- shared configuration items (function profiles), 445
- shared service, 293
- shell:domains value assignments, 539–540
- shells
  - bash, 102
  - Broadcom, 103
  - VSH (Virtual Shell), 103
  - vsh\_lc, 103
- show bgp ipv4 unicast vrf all command, 304
- show bgp process detail vrf all command, 304
- show bgp sessions vrf overlay-1 command, 306
- show bgp vpnv4 unicast vrf overlay-1 command, 304
- show copp policy command, 226–229
- show endpoint ip command, 252, 258
- show fex command, 214
- show interface ethernet 1/45–46 status command, 194
- show interface tunnel 1–20 command, 62
- show ip arp vrf command, 314
- show ip bgp command, 348
- show ip dhcp relay command, 271
- show ip eigrp neighbors command, 313
- show ip int brief command, 60
- show ip int brief vrf command, 61, 270
- show ip route command, 136–137, 258
- show ip route vrf command, 136–137, 270, 313
- show isis adjacency detail vrf overlay-1 command, 60
- show isis dteps vrf overlay-1 command, 61, 259

- show lacp interface command, 233–235
- show ntp peers command, 87–88
- show ntp peer-status command, 87–88
- show ntp statistics peer ipaddr command, 88
- show port-channel summary command, 199–200, 206–207, 214, 233
- show run command, 536, 541
- show run tenant DCACI command, 147
- show running-config all command, 195, 207
- show running-config command, 102
- show running-config dns command, 102
- show running-config leaf-interface-profile command, 182, 215
- show running-config vlan-domain DCACI-Domain command, 169
- show snmp command, 509–510
- show snmp summary command, 510
- show system internal epm endpoint ip command, 256
- show system internal epm vpc command, 257
- show track command, 333
- show vlan extended command, 246–247, 270, 386
- show vpc command, 204, 206
- show vrf command, 136–137
- show vrf Production:MP detail extended command, 247
- show zoning-rule command, 284
- silent hosts, 245–247, 248–249
- Simple Network Management Protocol. *See* SNMP (Simple Network Management Protocol)
- site IDs, 555, 598
- sites, 555, 598
- sizing APIC (Application Policy Infrastructure Controller) clusters, 33–36
- SLA (service-level agreement) tracking, 330–334
- slow drain interface policies, 173
- snapshots, configuration, 80
- SNAT (source NAT), 436
- SNMP (Simple Network Management Protocol)
  - ACI support for, 501–502
  - agents, 500
  - client group policies, 504–506
  - configuration caveats, 502
  - contexts, 506
  - contracts, 503
  - informs, 501
  - managers, 500
  - monitoring destination groups, 507–508
  - notifications, 598
  - sources for desired monitoring policies, 508
  - steps for, 502–503
  - traps, 501, 598
  - verification of, 509–511
  - versions of, 500–501
- snooping, IGMP (Internet Group Management Protocol), 368
- Soaking Interval timer, 115
- Soaking state (faults), 114
- Soaking-Clearing state (faults), 114
- software, Pearson IT Certification test prep software, 567–568
  - modes, 568–569
  - offline access, 568
  - online access, 567
  - overview of, 567
  - Premium Edition, 569–570
  - updating, 569
- software development, multitenancy and, 13
- software-defined networking (SDN), 4

- source NAT (SNAT), 436
- sources, SNMP (Simple Network Management Protocol), 508
- SPAN (Switched Port Analyzer), 219
- spanning tree interface policies, 172
- spanning-tree link type point-to-point command, 388
- spanning-tree link type shared command, 388
- spanning-tree port type edge command, 387
- spanning-tree port type edge trunk command, 387
- spine hardware, 23, 36
- spine proxy, 61, 247–248
- spine switches
  - definition of, 22, 598
  - first-generation, 37
  - second-generation, 37–38
  - spine initialization, 51
- spine-and-leaf architecture, 9–10
- split horizon, 324
- snelching, 127, 598
- SSL Certificate Validation Level parameter (LDAP providers), 542
- Stale Interval (BGP timer policy), 338
- standard ACI topology, 22–24
- standby controller parameter (ACI fabric initialization), 48
- stateful-services integration, 563–564
- stateless networking, 11, 182–183
- states
  - faults, 114–115
  - switch discovery, 51–52
- static bindings, mapping EPGs (endpoint groups) to ports with, 267–269
- static IP addressing, 470–472
- Static Node Management Addresses folder, 64
- static out-of-band addresses, 63–64
- static port bindings, 400
- static port channeling, 204–205
- static routes, adding on L3Outs, 329–330–334
- static VLAN allocation, 159, 598
- storm control, 172, 223–224
- STP (Spanning Tree Protocol), ACI
  - interaction with, 386–388
  - link types, 388
  - MCP (MisCabling Protocol), 388
  - MST (Multiple Spanning Tree) instance mappings, 387–388
  - TCNs (topology change notifications), 386–387
- stretch settings, 560
- stretched fabric topology, 24–25, 598
- stretched objects, 557, 598
- stub network and transit routing, 291–292
- Study mode, Pearson IT Certification test prep software, 568–569
- subjects
  - for bidirectional application, 275–276
  - definition of, 148–149, 598
  - for unidirectional application, 280–281
  - verification of, 278
- subnets, 469–470
  - advertising, 314–315
  - limiting endpoint learning to, 249–250
  - subnet boundaries, designing around, 139–141
- Suspend Individual Port setting, 197, 201
- SVIs (switch virtual interfaces), 147
  - Auto State, 299–300
  - Encap Scope, 298–299
  - floating, 296, 592
- switch CLI, 102–103
- switch policies, 174–176

**switch policy groups, 174–176****switch port configurations**

AAEP EPGs (endpoint groups),  
207–210

implications of, 210

individual port configuration, 188–196

*AAEPs (attachable access entity  
profiles), 190–191, 589*

*CDP interface policy, 190*

*common control settings,  
196–198*

*leaf access port policy groups,  
191*

*link level interface policy, 189*

*LLDP interface policy, 189*

*port blocks, 198–199*

*port mappings, 191–194*

*port usage types, 195*

*verification of, 194*

overview of, 186–188

port channel configuration, 196–201

vPC (virtual port channel)  
configuration, 204–207

*with APIC CLI, 207*

*static port channeling, 204–205*

*verification of, 206–207*

*vPC interface policy groups,  
205–206*

vPC (virtual port channel) domains,  
201–204

**switch profiles, 179–182, 598**

**switch selectors, 179, 598**

**switch virtual interfaces. *See* SVIs  
(switch virtual interfaces)**

**Switched Port Analyzer (SPAN), 219**

**switches. *See also* non-ACI switches,  
Layer 2 extension to**

ACI fabric initialization and  
*automatic upgrades of new  
switches, 74–75*

*connectivity following switch  
initialization, 59–62*

*default contracts, 64–66*

*discovery and activation, 55–58*

*discovery process, 49–51*

*discovery states, 51–52*

*static out-of-band addresses,  
assignment of, 63–64*

*verification of, 59–62*

**ACI switch port configurations**

AAEP EPGs (endpoint groups),  
207–210

implications of, 210

individual port configuration,  
188–196

overview of, 186–188

port channel configuration,  
196–201

vPC (virtual port channel)  
configuration, 204–207

vPC (virtual port channel)  
domains, 201–204

**CLI (command-line interface), 100–102**  
leaf, 9–11

*border, 24, 590*

*cabling requirements, 45*

*compute, 24, 590*

*definition of, 22, 594*

*first-generation, 38–39*

*initialization of, 51*

*IP storage, 24, 593*

*leaf access port policy groups, 191*

*overview of, 23*

*pervasive gateways, 252–254,  
595*

*pervasive routes, 258–260, 595*

*profiles, 179–182*

*purchasing decisions, 38*

*remote, 564–565*

- second-generation*, 39–40
- service*, 24, 597
- transit*, 25, 599
- Nexus 1000v distributed virtual switch, 9
- packet forwarding scenarios, 251
  - both endpoints attach to same leaf*, 251–254
  - flooding to unknown destination*, 261–263
  - known destination behind another leaf*, 254–258
  - spine proxy to unknown destination*, 258–261
- spine
  - definition of*, 22, 598
  - first-generation*, 37
  - overview of*, 23
  - second-generation*, 37–38
  - spine initialization*, 51
- SVIs (switch virtual interfaces), 147
  - Auto State*, 299–300
  - Encap Scope*, 298–299
  - floating*, 296
- VDS (vSphere distributed switch), VMM integration with, 392, 405
  - advantages of*, 403
  - enhanced LACP policy support*, 419–421
  - EPGs, pushing to vCenter as distributed port groups*, 415–416
  - ESXi connectivity*, 407
  - ESXi hosts, adding to VDS*, 411–414
  - immediacy settings*, 405–407
  - NSX-T VDS*, 395
  - N-VDS*, 395
  - planning*, 403–405
  - prerequisites for*, 403, 407
  - VDS deployment*, 405
  - VM assignment to distributed port groups*, 417–418
  - VMM domain association settings*, 418
  - VMM domain profiles*, 408–410
- vSwitches (vSphere standard switches), 395, 600
- Switch-Facing-Interface policy**, 375
- Symmetric Hashing setting**, 197
- synchronization, NTP (Network Time Protocol)**, 84–89, 106
- syntax, system message**, 485–492
- syslog monitoring destination groups**, 492–493
- syslog servers, forwarding system messages to**
  - contracts for, 487–491
  - steps for, 487
- syslog monitoring destination groups, 492–493
- syslog sources, 494–498
- verification of, 498–500
- System Health panel**, 124
- system messages**
  - ACI severity levels, 487
  - definition of, 598
  - forwarding to syslog servers
    - contracts for*, 487–491
    - steps for*, 487
  - syslog monitoring destination groups*, 492–493
  - syslog sources*, 494–498
  - verification of*, 498–500
- NX-OS severity levels, 485
- syntax, 485–492

## T

---

### tables

DHT (distributed hash table), 92  
lookup, 241–242

TACACS+, ACI configuration for, 532–536

TCNs (topology change notifications), 386–387

TCP (Transmission Control Protocol), 279–280

teaming, VMware vSphere, 400–403

Telco 5G distributed data centers,  
Remote Leaf deployments for, 30

template conformity, 557, 599

### templates

ACI Multi-Site, 557–558

definition of, 598

service graphs

*configuration*, 445–446

*definition of*, 597

*instantiation of*, 446–447

tenant DCACI command, 147

tenant policies. *See also* access policies;  
endpoint learning

multitier application deployment, 263

*application profiles*, 264–266

*BD (bridge domain)*

*configuration*, 264–266

*DHCP (Dynamic Host*

*Configuration Protocol) relay*,  
271

*domain assignment*, 267

*EPG (endpoint group)*

*configuration*, 264–266

*EPG (endpoint group) mapping  
to ports*, 267–269, 270

*overview of*, 263

*policy deployment*, 270

*verification of EPG-to-port  
assignments*, 269–270

packet forwarding scenarios, 251

*both endpoints attach to same  
leaf*, 251–254

*flooding to unknown  
destination*, 261–263

*known destination behind  
another leaf*, 254–258

*spine proxy to unknown  
destination*, 258–261

tenant-admin role, 520

tenant-a-l3domain, 162, 164

tenant-a-pdomain, 162, 164

tenant-b-l3domain, 162, 164

tenant-b-pdomain, 162, 164

tenant-c-l3domain, 162, 164

tenant-c-pdomain, 162, 164

tenant-ext-admin role, 520

tenants. *See also* contracts; EPGs  
(endpoint groups); L3Outs; tenant  
policies; VRF (virtual routing and  
forwarding) instances

access policies and, 184

APIC CLI configuration, 147

application profiles, 527

*advantages of*, 138–139

*configuration*, 142–145

*definition of*, 138

*topology*, 138–139

assigning to security domains,  
517–518

BDs (bridge domains)

*advertisement of subnets*

*assigned to*, 314–315

*configuration*, 142–145

*definition of*, 137, 590

*DHCP (Dynamic Host*

*Configuration Protocol)*, 271

- disabling endpoint learning on,*  
250–251
- extensions,* 374–375
- external,* 161, 380–385, 591
- in-band management,* 469
- L3Out,* 296–298, 594
- multitier application*  
*deployment,* 264–266
- practical application of,* 141–142
- service graph settings,* 439–440
- connectivity outside the fabric,  
151–153
  - external EPGs (endpoint*  
*groups),* 151–153
  - Layer 3 Outside (L3Out),* 153
- creating, 133–134
- definition of, 133, 599
- hierarchy, 153–154
- multitenancy, 11–14, 549–550, 595
- nesting, 134
- predefined, 134–135
- use cases for, 133
- tenant-security privilege,** 519
- TEPs (tunnel endpoints)**
  - overlay multicast, 555, 595
  - overlay unicast, 555, 595
  - PTEP (physical tunnel endpoint), 60,  
595
  - TEP pools, 49, 50, 57, 59, 599
  - verifying addresses of, 60–62
  - VTEP, 31–32, 245, 256–258
- Tetration,** 272
- three-tier Clos topology,** 22
- three-tier network design,** 5–7
- tier-2-leaf role,** 57
- time/date, NTP (Network Time**  
**Protocol) synchronization,** 84–89,  
106
- Timeout parameter**
  - LDAP providers, 542
  - RADIUS providers, 540
  - TACACS+ providers, 532
- timer policy**
  - BGP (Border Gateway Protocol), 337–  
339, 342
  - OSPF (Open Shortest Path First), 329
- timers, lifecycle,** 115
- timestamps,** 485
- topologies,** 18, 21
  - ACI Multi-Pod
    - active/active data center*  
*deployments,* 562–564
    - overview of,* 25–26
    - stateful-services integration,*  
563–564
  - ACI Multi-Site, 555–558
    - cross-site connectivity,* 555–557
    - extending ACI to remote*  
*locations and public clouds,*  
564
    - locally governed configurations,*  
557
    - MSO-governed configurations,*  
557
    - overview of,* 26–28
    - primary and disaster recovery*  
*data center deployments,*  
558–562
    - schemas,* 557–558
    - stateful-services integration,*  
563–564
    - terminology for,* 557
  - ACI Multi-Tier, 28–30
  - ACI stretched fabric, 24–25, 598
  - Clos, 21–22
  - Remote Leaf, 30–32
  - standard ACI, 22–24

- topology change notifications (TCNs), 386–387
- topRoot object, 106
- tracking, service-level agreement, 330–334
- traditional networks
  - agility of, 8
  - management of, 4–5
  - scalability and growth of, 5–7
  - security of, 8–9
  - three-tier design of, 5–7
  - visibility of, 9
- transit leaf switches, 25, 599
- Transmission Frequency setting (Global MCP Policy), 222
- transparent mode, devices in, 436
- traps, SNMP (Simple Network Management Protocol), 501, 598
- Trunk port-binding mode, 268
- tunnel endpoints. *See* TEPs (tunnel endpoints)
- tunnel interfaces, 62

## U

---

- UCS (Unified Computing System) domains, 218, 407
- underlay, 135
- Undiscovered state (fabric node discovery), 52
- unicast routing, 243–244, 371
- Unified Computing System (UCS) domains, 218, 407
- Unknown state (fabric node discovery), 52
- unmanaged firewall pairs, 447–452
- unmanaged mode, 428–430
- Unsupported state (fabric node discovery), 52

- updates
  - DCACI 300-620 exam, 586–587
  - Pearson IT Certification test prep software, 569
- upgrades
  - ACI fabric, 66–71
  - EPLD (electronic programmable logic device), 66
  - switches, 74–75
- USB ports, 46
- uSeg EPGs (endpoint groups), 146
- User Certificate Attribute setting (User Identity), 522–523
- user identity groups, 536
- User Identity page (Create Local User wizard), 522–523
- User VRF L3Outs (VRF-lite), 292
- users. *See also* tenants
  - local
    - assigning access to*, 523–525
    - creating*, 521–523
    - modifying access for*, 523–525
  - user identity groups, 536
- UTC (Coordinated Universal Time), 89

## V

---

- vCenter
  - failure, 399
  - pushing EPGs (endpoint groups) to, 415–416
- vCenter Domain Creation wizard, 408–410
- VDS (vSphere distributed switch), VMM integration with, 392
  - advantages of, 403
  - architecture, 397
  - definition of, 600
  - enhanced LACP policy support, 419–421

- EPGs, pushing to vCenter as
  - distributed port groups, 415–416
- ESXi hosts, adding to VDS, 411–414
- ESXi servers
  - adding to VDS, 411–414*
  - connecting in UCS domains, 407*
  - connecting to fabric, 407*
- immediacy settings, 405–407
- NSX-T VDS, 395
- N-VDS, 395
- planning, 403–405
- prerequisites for, 403, 407
- VDS deployment, 405
- VM assignment to distributed port groups, 417–418
- VMM domain association settings, 418
- VMM domain profiles, 408–410
- Verified Scalability Guide, 24, 26, 34
- VGA video ports, 46
- VIC 1455 ports, 46–47
- VIP addresses, 245
- virtual APIC (Application Policy Infrastructure Controller)
  - specifications, 33
- virtual LANs. *See* VLANs (virtual LANs)
- virtual leaf (vLeaf), 564
- Virtual Machine Manager integrations.
  - See* VMM (Virtual Machine Manager) integration
- virtual machines. *See* VMs (virtual machines)
- virtual network interface cards (vNICs), 395, 400, 406, 414, 417, 465, 599
- Virtual Pod (vPod), 26
- virtual port channels. *See* vPCs (virtual port channels)
- virtual routing and forwarding.
  - See* VRF (virtual routing and forwarding) instances
- Virtual Shell (VSH), 103
- virtual spine (vSpine), 564
- virtualeaf roles, 57
- virtualspine roles, 57
- Visore, 108–110, 599
- VLAN (virtual LAN) pools, 159–160, 599
  - creating, 159–160
  - designs for, 161–163
    - hybrid approach, 162–163*
    - single VLAN pool for each type of domain, 161*
    - single VLAN pool per function, 162*
  - in L3Out domains, 301–302
  - overlap between, 164–165
  - static versus dynamic VLAN allocation, 159, 598
  - VLAN ranges in, 160
- VLANs (virtual LANs)
  - access encapsulation, 246
  - dynamic VLAN allocation, 591
  - infrastructure, 593
  - overlapping, 385–386
  - PI VLANs (platform-independent VLANs), 246, 253–254, 595
  - pools, 159–160, 599
    - creating, 159–160*
    - designs for, 161–163*
    - in L3Out domains, 301–302*
    - overlap between, 164–165*
    - static versus dynamic VLAN allocation, 159, 598*
    - VLAN ranges in, 160*
  - port encapsulation, 246, 596
  - types of, 245–247
  - VLAN IDs, 245–247, 253–254
- vLeaf (virtual leaf), 564

- vmk0 interface, 397–399
- VMkernel, 395, 397–399, 405, 406, 413–414, 599
- VMM (Virtual Machine Manager)
  - integration, 161, 599
  - advantages of, 403
  - definition of, 392
  - domains, 599
  - ESXi servers
    - connecting in UCS domains, 407*
    - connecting to fabric, 407*
  - immediacy settings, 405–407
  - for multiple data centers, 563
  - planning, 403–405
  - prerequisites for, 403, 407
  - VDS deployment, 405
  - VMM (Virtual Machine Manager) integration, 161
- vmm-admin role, 520
- VMM-NONPROD, 163, 164
- VMM-PROD, 162, 164
- VMM-VOICE, 163, 164
- vmnic, 395, 413, 599
- vMotion, 399
- VMs (virtual machines)
  - assigning to distributed port groups, 417–418
  - VMware ESXi, 27
- VMware ESXi. *See* ESXi
- VMware vSphere
  - definition of, 392
  - distributed port groups, 397
    - assigning VMs (virtual machines) to, 417–418*
    - pushing EPGs (endpoint groups) to vCenter as, 415–416*
  - load balancing, 401, 407, 419, 600
  - overview of, 394–395
  - port bindings, 400
  - port groups, 395–397
  - system traffic
    - impact of vCenter failure on, 399*
    - overview of, 397–399*
  - teaming and failover policies, 400–403
  - VMM integration with vSphere VDS, 392
    - advantages of, 403*
    - architecture, 397*
    - definition of, 600*
    - enhanced LACP policy support, 419–421*
    - EPGs, pushing to vCenter as distributed port groups, 415–416*
    - ESXi hosts, adding to VDS, 411–414*
    - ESXi servers, 407, 411–414*
    - immediacy settings, 405–407*
    - NSX-T VDS, 395*
    - N-VDS, 395*
    - planning, 403–405*
    - prerequisites for, 403, 407*
    - VDS deployment, 405*
    - VM assignment to distributed port groups, 417–418*
    - VMM domain association settings, 418*
    - VMM domain profiles, 408–410*
  - vSwitches (vSphere standard switches), 395, 494–498, 600
- vNICs (virtual network interface cards), 395, 400, 406, 414, 417, 465, 599
- VNIDs (VXLAN network identifiers), 10, 28, 244, 245–247
- vpc domain explicit 21 leaf 101 102 command, 203

VPC Explicit Protection Group wizard, 202

VPC interface policy groups, 170

vPC peer dead interval, 203, 599

vPCs (virtual port channels), 6
 

- configuration, 204–207
  - with APIC CLI*, 207
  - static port channeling*, 204–205
  - verification of*, 206–207
  - vPC interface policy groups*, 205–206
- domains, 201–204
- learning and forwarding for, 256–258
- VIP addresses, 245
- vPC peer dead interval, 203, 599

vPod, 26, 564

VRF (virtual routing and forwarding) instances, 135–137
 

- BGP customizations, 342–343
- creating, 135–137
- data plane learning, disabling, 251
- definition of, 135, 599
- EIGRP customizations, 324–325
- for in-band management, 465–466, 474
- multitenancy and, 12–13
- OSPF customizations, 328–329
- for out-of-band (OOB) management, 465
- preferred group members, 365–367
- SNMP contexts for, 506
- User VRF L3Outs (VRF-lite), 292
- VRF-specific MIB, 600

vrf context command, 147

vrf member command, 147

VSH (Virtual Shell), 103

vsh\_lc, 103

vSphere. *See* VMware vSphere

vSpine (virtual spine), 564

vSwitches (vSphere standard switches), 395, 494–498, 600

VTEP, 31–32, 245, 256–258

VXLAN network identifiers (VNIDs), 10, 28, 165, 244, 245–247

## W

---

Warning faults, 112

Weight for routes from this neighbor parameter (BGP peer connectivity profiles), 340

whitelisting, 14–15, 148–151, 427
 

- contracts
  - additional whitelisting examples*, 282–283
  - Apply Both Directions feature*, 277–278
  - assigning to EPGs*, 278–279
  - definition of*, 148–149
  - direction of*, 149–150
  - filters*, 148–149, 273–275, 280
  - overview of*, 272
  - planning enforcement of*, 272–273
  - Reverse Filter Ports feature*, 277–278
  - scope of*, 150–151, 284–285
  - Stateful feature*, 284
  - subjects*, 148–149, 278
  - subjects for bidirectional application*, 275–276
  - subjects for unidirectional application*, 280–281
  - TCP established session rule*, 279–280
  - verifying enforcement of*, 283–284
  - zero-trust networks and*, 151

definition of, 600  
for in-band management, 476–477

## wizards

Add and Manage Hosts, 411–414  
 BFD Interface Profile Creation, 322–323  
 Configure Interface, PC, and VPC, 211–212  
 Create Application EPG, 265–266  
 Create Application Profile, 144–146  
 Create Bridge Domain, 142–144, 264–265  
 Create Bridged Outside, 380–385  
 Create Configuration Export Policy, 76–78, 80–81  
 Create DHCP Relay Label, 271  
 Create Filter, 488  
 Create LDAP Group Map Rules, 543–547  
 Create LDAP Provider, 543  
 Create Leaf Interface Profile, 179–180  
 Create Leaf Profile, 180–181  
 Create Local User, 521–523  
 Create Physical Domain, 163  
 Create RBAC Rule, 529–531  
 Create Remote Location, 77  
 Create Syslog Source, 494–495

Create Tenant, 133–134  
 Create VLAN Pool, 159–160  
 Create VPC Explicit Protection Group, 202–203  
 Create VRF, 135–137  
 External EPG Creation, 310–312  
 L3Out Creation  
   *AAEPs (attachable access entity profiles), 302*  
   *configuration for BGP peering, 334–337*  
   *configuration for EIGRP peering, 307–309*  
   *configuration for OSPF peering, 325–327*  
 New Distributed Port Group, 400  
 Node Registration, 57  
 vCenter Domain Creation, 408–410  
 VPC Explicit Protection Group, 202

## X-Y-Z

---

XML (Extensible Markup Language), 103  
 zAny, 364–365, 589  
 ZeroMQ, 92, 243  
 zero-penalty forwarding decision, 248  
 zero-trust security, 14–15, 151, 600