



Official Cert Guide

Learn, prepare, and practice for exam success



CCNA Data Center DCICT 200-155

NAVAID SHAMSEE, CCIE® NO. 12625

DAVID KLEBANOV, CCIE® NO. 13791

HESHAM FAYED, CCIE® NO. 9303

AHMED AFROSE

OZDEN KARAKOK, CCIE® NO. 6331

ciscopress.com

Exclusive Offer – 40% OFF

Cisco Press Video Training

livelessons®

ciscopress.com/video

Use coupon code CPVIDEO40 during checkout.



Video Instruction from Technology Experts



Advance Your Skills

Get started with fundamentals, become an expert, or get certified.



Train Anywhere

Train anywhere, at your own pace, on any device.



Learn

Learn from trusted author trainers published by Cisco Press.

Try Our Popular Video Training for FREE!

ciscopress.com/video

Explore hundreds of **FREE** video lessons from our growing library of Complete Video Courses, LiveLessons, networking talks, and workshops.

Cisco Press

ciscopress.com/video

CCNA Data Center DCICT 200-155 Official Cert Guide

NAVAID SHAMSEE, CCIE No. 12625

DAVID KLEBANOV, CCIE No. 13791

HESHAM FAYED, CCIE No. 9303

AHMED AFROSE

OZDEN KARAKOK, CCIE No. 6331

Cisco Press

800 East 96th Street
Indianapolis, IN 46240

CCNA Data Center

DCICT 200-155 Official Cert Guide

Navaid Shamsee, CCIE No. 12625

David Klebanov, CCIE No. 13791

Hesham Fayed, CCIE No. 9303

Ahmed Afrose

Ozden Karakok, CCIE No. 6331

Copyright© 2017 Pearson Education, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing December 2016

Library of Congress Control Number: 2016949728

ISBN-13: 978-1-58720-591-0

ISBN-10: 1-58720-591-2

Warning and Disclaimer

This book is designed to provide information about the 200-155 DCICT exam for CCNA Data Center certification. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

Trademark Acknowledgments

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

Special Sales

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

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

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

Feedback Information

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

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

We greatly appreciate your assistance.

Editor-in-Chief: Mark Taub

Copy Editor: Gill Editorial Services

Product Line Manager: Brett Bartow

Technical Editor: David Burns

Business Operation Manager, Cisco Press: Ronald Fligge

Editorial Assistant: Vanessa Evans

Executive Editor: Mary Beth Ray

Designer: Chuti Prasertsith

Managing Editor: Sandra Schroeder

Composition: Tricia Bronkella

Development Editor: Ellie Bru

Indexer: Ken Johnson

Project Editor: Mandie Frank

Proofreader: Debbie Williams



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

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

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Authors

Navaid Shamsee, CCIE No.12625, is a senior solutions architect in the Cisco Services organization. He holds a master's degree in telecommunication and a bachelor's degree in electrical engineering. He is also a triple CCIE in routing and switching, service provider, and data center technologies. Navaid has extensive experience in designing, implementing and securing many large-scale enterprise and service provider data centers. In Cisco, Navaid is focused on the security of data center, cloud, and software-defined networking technologies. You can reach Navaid on Twitter: @NavaidShamsee.

David Klebanov, CCIE No.13791 (Routing and Switching), is leading Technical Marketing organization at Viptela, the Software Defined Wide Area Network (SD-WAN) company. David has more than 15 years of diverse industry experience architecting and deploying complex network environments. In his work, David sets strategic direction for industry-leading network platforms, which transform the world of wide area communications for enterprises and service providers alike. David also takes great pride in speaking at industry events, releasing publications, and working on patents. You can reach David on Twitter: @DavidKlebanov.

Hesham Fayed, CCIE No.9303 (Routing and Switching/Data Center), is a consulting systems engineer for data center and virtualization based in California. Hesham has been with Cisco for more than 11 years and has 19 years of experience in the computer industry, working with service providers and large enterprises. His main focus is working with customers in the western region of the United States to address their challenges by doing end-to-end data center architectures.

Ahmed Afrose is a solutions architect at Cisco Cloud and Networking Services (C&NS) Innovation and Delivery team. He is responsible for providing architectural design guidance and leading complex multitech service deliveries. Furthermore, he is involved in demonstrating the Cisco value propositions in data center analytics, application automation, software-defined data centers, and Cisco Unified Computing System (UCS). Ahmed has a bachelor's degree in information systems. He started his career with Sun Microsystem-based technologies and has 17 years of diverse experience in the industry. He's been working for Cisco Systems for 7+ years and was directly responsible for establishing Cisco UCS Advanced Services delivery capabilities while evangelizing the product in the EMEA region. You can reach Ahmed on Twitter: @ahmedafrose.

Ozden Karakok, CCIE No. 6331, is a technical leader from the data center products and technologies team in the Technical Assistant Center (TAC). She has been with Cisco Systems for 17 years and specializes in storage area and data center networks. Prior to joining Cisco, Ozden spent five years working for a number of Cisco's large customers in various telecommunication roles. She is a Cisco Certified Internetwork Expert in routing and switching, SNA/IP, and storage. A frequent speaker at Cisco and data center events, she serves as a member of the patent committee at Cisco Services. Ozden holds a degree in computer engineering from Istanbul Bogazici University. Currently, she is focused on Application Centric Infrastructure (ACI) and software-defined storage (SDS). You can reach Ozden on Twitter: @okarakok.

About the Technical Reviewers

David Burns has in-depth knowledge of routing and switching technologies, network security, data center, and mobility. Dave is currently the Vice President of Engineering & Architecture for QTS DataCenters where he is focused on driving strategy and execution for connectivity, infrastructure, security, and data center and cloud services being delivered to QTS customers and partners. Prior to joining QTS, Dave led various Engineering & Architecture teams for both in the U.S. and internationally-focused on Service Providers and Enterprise IT, covering the spectrum of technologies ranging from security, network, data center, cloud, and connectivity. Dave holds various sales, industry, and Cisco technical certifications, including CISSP, CCNP Security (formerly CCSP), CCDP, and multiple associate-level certifications. Dave is also a published author with Cisco Press and active technical reviewer for multiple titles. Dave has a Masters in Business Administration and a Bachelor of Science in Telecommunications Engineering Technology.

Dedications

Navaid Shamsee: To my parents for their guidance and prayers. To my wife, Hareem, for her love and support, and to my children, Ahsan, Rida, and Maria.

David Klebanov: This book is dedicated to my gorgeous wife, Tanya, and to our two wonderful daughters, Lia and Maya. I wanted you to know how grateful I am to have you in my life, and how your love and support inspire me each and every day. Tanya, your journey to become OTR/L is an inspiration to anyone pursuing personal dreams and a true testament that if you put your mind to it, you can accomplish anything! You are my pride, and I applaud you for your success. Lia, as you become a young adult, keep dreaming big, and remember that life is like a box of chocolates—you never know what you are going to get. Let your talent and perseverance pave the way. Maya, you may still be young, but you have a big heart. Thank you for filling our house with joy and laughter.

This book is also dedicated to all of you who have chosen the path of technology. Together we can make this world a better place.

Hesham Fayed: This book is dedicated to my lovely wife, Eman, and my beautiful children, Ali, Laila, Hamza, Fayrouz, and Farouk. Without your support and encouragement, I would never have been able to finish this book. I can't forget to acknowledge my parents; your guidance, education, and encouragement to strive to be better is what helped in my journey.

Ahmed Afrose: I dedicate this book to my parents, Hilal and Faiziya. Without them, none of this would have been possible. I also dedicate this to those people who are deeply curious and advocate self-advancement.

Ozden Karakok: To my loving husband, Tom, for his endless support, encouragement, and love. Merci beaucoup, Askim.

To Remi and Mira, for being the most incredible miracles of my life and being my number-one source of happiness.

To my wonderful parents, who supported me in every chapter of my life and are an inspiration for life.

To my awesome sisters, Gulden and Cigdem, for their unconditional support and loving me just the way I am.

To the memory of Steve Ritchie, for being the best mentor ever and passionate technology lover.

Acknowledgments

First and foremost, I would like to say thank you to my wife Hareem for her unconditional support throughout my career and writing this book.

To my co-authors, Afrose, David, Hesham, and Ozden: It is my pleasure and honor working with you again on the second edition of this book. Without your support, this book would not have been possible.

I would also like to say thank you to David Burns for reviewing this book and providing us feedback. His input was very valuable and it helped us improve the quality of content.

A big thank-you to the entire Cisco Press team for all their support in getting this book published. Special thanks to Mary Beth and Ellie for keeping us on track and guiding us in the journey of writing this book.

—Navaid Shamsee

My deepest gratitude goes to all the people who have shared my journey over the years and who have inspired me to always strive for more. I would also like to thank my fellow authors for the wonderful team spirit and the entire Cisco Press team for ensuring our success.

—David Klebanov

I'd like to thank my co-authors, Navaid, Afrose, Ozden, and David, for working as a team to complete the second edition of this book. Mary Beth and Ellie, thank you both for your patience and support through my second publication. It has been an honor working with you both, and I have learned a lot during this process.

I want to thank my family for their support and patience while I was working on this book.

—Hesham Fayed

Navaid Shamsee has been a good friend and colleague. I am thankful to him for this awesome opportunity. He has helped me achieve a milestone in my career by offering me the opportunity to be associated with this publication and the world-renowned Cisco Press team.

I'm humbled by this experienced and talented team of co-authors: Navaid, David, Hesham, and Ozden. It was a pleasure working with a team of like-minded professionals.

I am also thankful to all our professional editors, especially Mary Beth Ray and Ellie Bru, for their patience and guidance every step of the way. A big thank-you to all the folks involved in production, publishing, and bringing this book to the shelves.

To our technical reviewer, David Burns, for his keen attention to detail; it helped gauge depth and consistency and improved the overall quality of this certification guide.

—Ahmed Afrose

This book would never have become a reality without the help, support, and advice of a great number of people.

I would like to thank my great co-authors, Navaid, David, Afrose, and Hesham. Thank you for your excellent collaboration, hard work, and priceless time. I truly enjoyed working with each one of you on our second book. I really appreciated Navaid taking the lead and being the glue of this diverse team. It was a great pleasure and honor working with such talented engineers. It would have been impossible to finish this book without your support.

To our technical reviewer, David Burns: Thank you for providing us with your valuable feedback, suggestions, hard work, and quick turnaround. Your excellent input helped us improve the quality and accuracy of the content. It was a great pleasure for all of us to work with you.

To Mary Beth Ray, Ellie Bru, Mandie Frank, and the Cisco Press team: A big thank-you to the entire Cisco Press team for all their support in getting this book published. Special thanks to Mary Beth and Ellie for their endless encouragement and support.

To the extended teams at Cisco: Thank you for being patient while our minds were in the book. Thank you for believing in and supporting us on this journey. Thank you for the innovative organization and development team at Cisco.

To all extended family and friends: Thank you for your patience and endless support.

—Ozden Karakok

Contents at a Glance

Introduction xxxii

Getting Started 3

Part I Data Center Networking Technologies 9

Chapter 1 Data Center Networking 11

Chapter 2 Management and Monitoring of Cisco Nexus Devices 47

Chapter 3 Unified Fabric Overview 77

Part I Review

Part II Network Virtualization 113

Chapter 4 Cisco Nexus 1000V and Virtual Switching 115

Chapter 5 Data Center Overlay Networks 159

Chapter 6 Virtualizing Cisco Network Devices 189

Part II Review

Part III Cisco Unified Computing 229

Chapter 7 Cisco UCS Architecture 231

Chapter 8 Cisco UCS Manager 295

Chapter 9 Cisco Unified Computing System Pools, Policies, Templates, and Service Profiles 327

Chapter 10 Administration, Management, and Monitoring of Cisco UCS 363

Chapter 11 Server Virtualization Solutions 403

Part III Review

Part IV Application Centric Infrastructure 429

Chapter 12 ACI Architecture 431

Chapter 13 ACI Logical Model and Policy Framework 475

Chapter 14 Operating ACI 515

Part IV Review

Part V Automation and Orchestration 563

Chapter 15 Cloud Computing 565

Chapter 16 UCS Director 583

Chapter 17 Understanding and Troubleshooting UCSD Workflows 631

Part V Review

Part VI Final Preparation 657

Chapter 18 Final Preparation 659

Appendix A Answers to the “Do I Know This Already?” Quizzes 679

Appendix B Memory Tables 697

Appendix C Memory Tables Answer Key 709

Glossary 721

Index 752

On the Companion Website:

Appendix D Study Planner

Contents

Introduction xxxii

Getting Started 3

- A Brief Perspective on the CCNA Data Center Certification Exam 3
- Suggestions for How to Approach Your Study with This Book 3
 - This Book Is Compiled from 17 Short Read-and-Review Sessions 4
 - Practice, Practice, Practice—Did I Mention: Practice? 4
 - In Each Part of the Book You Will Hit a New Milestone 5
 - Use the Final Preparation Chapter to Refine Skills 6
 - Set Goals and Track Your Progress 6
- Other Small Tasks Before Getting Started 7

Part I Data Center Networking Technologies 9

Chapter 1 Data Center Networking 11

- “Do I Know This Already?” Quiz 11
- Foundation Topics 15
- Port Channel 15
 - What Is Port Channel? 15
 - Benefits of Using Port Channels 15
 - Port Channel Compatibility Requirements 16
 - Link Aggregation Control Protocol 16
 - Port Channel Modes 17
 - Configuring Port Channel 18
 - Port Channel Load Balancing 19
 - Verifying Port Channel Configuration 19
- Virtual Port Channel 20
 - What Is Virtual Port Channel? 20
 - Benefits of Using vPC 22
 - Components of vPC 22
 - vPC Data Plane Operation 24
 - vPC Control Plane Operation 26
 - vPC Limitations 27
 - Configuration Steps of vPC 27
 - Verification of vPC 29
- FabricPath 29
 - Spanning Tree Protocol 29
 - What Is FabricPath? 31

Benefits of FabricPath	31
Components of FabricPath	32
FabricPath Frame Format	34
FabricPath Control Plane	35
FabricPath Data Plane	37
Conversational Learning	38
FabricPath Packet Flow Example	38
<i>Unknown Unicast Packet Flow</i>	39
<i>Known Unicast Packet Flow</i>	39
Virtual Port Channel Plus	40
FabricPath Interaction with Spanning Tree	41
Configuring FabricPath	42
Verifying FabricPath	43
Reference List	44
Exam Preparation Tasks	45
Review All Key Topics	45
Complete Tables and Lists from Memory	45
Define Key Terms	45

Chapter 2 Management and Monitoring of Cisco Nexus Devices 47

“Do I Know This Already?” Quiz	47
Foundation Topics	50
Operational Planes of a Nexus Switch	50
Data Plane	50
<i>Store-and-Forward Switching</i>	51
<i>Cut-Through Switching</i>	51
<i>Nexus 5500 Data Plane Architecture</i>	51
Control Plane	53
<i>Nexus 5500 Control Plane Architecture</i>	54
<i>Control Plane Policing</i>	55
<i>Control Plane Analyzer</i>	57
Management Plane	58
Nexus Management and Monitoring Features	59
Out-of-Band Management	59
<i>Console Port</i>	60
<i>Connectivity Management Processor</i>	61
<i>Management Port (mgmt0)</i>	61
In-Band Management	62

	Simple Network Management Protocol	62
	<i>SNMP Notifications</i>	63
	<i>SNMPv3</i>	63
	Remote Monitoring	64
	<i>RMON Alarms</i>	64
	<i>RMON Events</i>	65
	Syslog	65
	Embedded Event Manager	66
	<i>Event Statements</i>	66
	<i>Action Statements</i>	66
	<i>Policies</i>	66
	Generic Online Diagnostics	67
	Smart Call Home	68
	NX-API	68
	Role-Based Access Control	70
	User Roles	71
	Rules	71
	User Role Policies	72
	RBAC Characteristics and Guidelines	72
	Privilege Levels	72
	Reference List	73
	Exam Preparation Tasks	74
	Review All Key Topics	74
	Complete Tables and Lists from Memory	74
	Define Key Terms	74
Chapter 3	Unified Fabric Overview	77
	“Do I Know This Already?” Quiz	77
	Foundation Topics	80
	Challenges of Today’s Data Center Networks	80
	Cisco Unified Fabric Principles	82
	Convergence of Network and Storage	82
	Scalability and Growth	88
	Security and Intelligence	93
	Inter-Data Center Unified Fabric	97
	Fibre Channel over IP	98
	Overlay Transport Virtualization	100
	Locator ID Separation Protocol	104

Reference List	106
Exam Preparation Tasks	107
Review All Key Topics	107
Complete Tables and Lists from Memory	107
Define Key Terms	107

Part I Review 108

Part II Network Virtualization 113

Chapter 4 Cisco Nexus 1000V and Virtual Switching 115

“Do I Know This Already?” Quiz	116
Foundation Topics	119
Evolution of Virtual Switching	119
Before Server Virtualization	119
Server Virtualization with Static VMware vSwitch	121
<i>Virtual Network Components</i>	122
<i>Virtual Access Layer</i>	123
Standard VMware vSwitch Overview	124
<i>Standard VMware vSwitch Operations</i>	125
<i>Standard VMware vSwitch Configuration</i>	127
VMware vDS Overview	128
<i>VMware vDS Configuration</i>	129
<i>VMware vDS Enhancements</i>	131
<i>VMware vSwitch and vDS</i>	131
Cisco Nexus 1000V Virtual Networking Solution	132
Cisco Nexus 1000V System Overview	134
Cisco Nexus 1000V Salient Features and Benefits	134
Cisco Nexus 1000V Series Virtual Switch Architecture	137
<i>Cisco Nexus 1000V Virtual Supervisory Module</i>	137
<i>Cisco Nexus 1000V Virtual Ethernet Module</i>	138
<i>Cisco Nexus 1000V Component Communication</i>	140
<i>Cisco Nexus 1000V Management Communication</i>	140
<i>Cisco Nexus 1000V Port Profiles</i>	141
<i>Types of Port Profiles in Cisco Nexus 1000V</i>	142
Cisco Nexus 1000V Administrator View and Roles	142
Cisco Nexus 1000V Verifying Initial Configuration	144
Cisco Nexus 1000V VSM Installation Methods	144

	Initial VSM Configuration Verification	145
	<i>Verifying VMware vCenter Connectivity</i>	145
	<i>Verifying Nexus 1000V Module Status</i>	147
	Cisco Nexus 1000V VEM Installation Methods	148
	<i>Initial VEM Status on ESX or ESXi Host</i>	148
	<i>Verifying VSM Connectivity from vCenter Server</i>	148
	<i>Verifying VEM Agent Running</i>	148
	<i>Verifying VEM Uplink Interface Status</i>	149
	<i>Verifying VEM Parameters with VSM Configuration</i>	150
	<i>Validating VM Port Groups and Port Profiles</i>	150
	<i>Verifying Port Profile and Groups</i>	152
	Key New Technologies Integrated with Cisco Nexus 1000V	153
	<i>What Is the Difference Between VN-Link and VN-Tag?</i>	153
	<i>What Is VXLAN?</i>	153
	<i>What Is vPath Technology?</i>	154
	<i>What Is Cisco Application Virtual Switch (AVS)?</i>	154
	Reference List	155
	Exam Preparation Tasks	156
	Review All Key Topics	156
	Complete Tables and Lists from Memory	156
	Define Key Terms	157
Chapter 5	Data Center Overlay Networks	159
	“Do I Know This Already?” Quiz	159
	Foundation Topics	162
	Overlay Transport Virtualization	162
	Introduction to Overlay Transport Virtualization	162
	OTV Terminology	164
	OTV Control Plane	166
	<i>Multicast-Enabled Transport Infrastructure</i>	166
	<i>Unicast-Enabled Transport Infrastructure</i>	167
	Data Plane for Unicast Traffic	168
	Data Plane for Multicast Traffic	170
	Failure Isolation	172
	<i>STP Isolation</i>	172
	<i>Unknown Unicast Handling</i>	172
	ARP Optimization	173

OTV Multihoming	173
<i>First Hop Redundancy Protocol Isolation</i>	174
OTV Configuration Example with Multicast Transport	174
OTV Configuration Example with Unicast Transport	177
Virtual Extensible LAN	179
VXLAN Terminology	180
VTEP	180
VNID	181
VXLAN Gateway Types	181
<i>VXLAN Bridging</i>	181
<i>VXLAN Routing</i>	181
VXLAN Packet Forwarding	183
VXLAN VTEP Discovery and Address Learning	183
<i>VXLAN Data Plane Learning</i>	184
<i>VXLAN Unicast-Only Mode</i>	184
<i>VXLAN Using Control Plane Protocol</i>	184
Network Virtualization Using Generic Routing Encapsulation	184
Reference List	185
Exam Preparation Tasks	186
Review All Key Topics	186
Define Key Terms	186
Chapter 6 Virtualizing Cisco Network Devices	189
“Do I Know This Already?” Quiz	189
Foundation Topics	192
Describing VDCs on the Cisco Nexus 7000 Series Switch	192
VDC Deployment Scenarios	194
<i>Horizontal Consolidation Scenarios</i>	194
<i>Vertical Consolidation Scenarios</i>	195
<i>VDCs for Service Insertion</i>	196
Understanding Different Types of VDCs	196
Interface Allocation	198
VDC CPU Shares and Hardware Resources	199
VDC Administration	199
VDC Requirements	200
Verifying VDCs on the Cisco Nexus 7000 Series Switch	201
Describing Layer 3 Virtualization Within VDCs	204

	Describe the Benefits and Perform Simple Troubleshooting of VDC STP	207
	Verifying STP Operation	209
	Describing Network Interface Virtualization	210
	Cisco Nexus 2000 FEX Terminology	210
	Nexus 2000 Series Fabric Extender Connectivity	211
	VN-Tag Overview	212
	Cisco Nexus 2000 FEX Packet Flow	213
	Cisco Nexus 2000 FEX Port Connectivity	214
	Cisco Nexus 2000 FEX Configuration on the Nexus 7000 Series	216
	Cisco Adapter FEX Technology	219
	Reference List	221
	Exam Preparation Tasks	222
	Review All Key Topics	222
	Define Key Terms	222
	Part II Review	224
Part III	Cisco Unified Computing	229
Chapter 7	Cisco UCS Architecture	231
	“Do I Know This Already?” Quiz	232
	Foundation Topics	235
	Evolution of Server Computing in a Nutshell	235
	What Is a Socket?	236
	What Is a Core?	236
	What Is Hyperthreading?	237
	Understanding Server Processor Numbering	237
	Value of Cisco UCS in the Data Center	238
	One Unified System	239
	Unified Fabric	240
	Unified Management and Operations	240
	Cisco UCS Stateless Computing	241
	Intelligent Automation, Cloud-Ready Infrastructure	242
	Describing the Cisco UCS Product Family	243
	Cisco UCS Computer Hardware Naming	243
	Cisco UCS Fabric Interconnects	244
	<i>What Is a Unified Port?</i>	247
	<i>What Are QSA Modules and Breakout Cables?</i>	247

Cisco UCS Fabric Interconnect—Expansion Modules	247
Cisco UCS 5108 Blade Server Chassis	248
<i>Cisco UCS 5108 Blade Server Chassis—Power Supply and Redundancy Modes</i>	248
Cisco UCS M4308 Chassis	249
Cisco UCS I/O Modules (FEX)	250
Cisco UCS B-Series Blade Servers	251
Cisco UCS B-Series Best Practices for Populating DRAM	253
Cisco Adapters (Mezzanine Cards) for UCS B-Series Blade Servers	254
<i>Cisco Converged Network Adapters (CNA)</i>	254
<i>Cisco Virtual Interface Cards</i>	255
<i>Cisco UCS VIC 1340/1380</i>	256
<i>Cisco UCS Port Expander</i>	257
<i>What Is Cisco Adapter FEX?</i>	257
<i>What Is Cisco Virtual Machine-Fabric Extender?</i>	257
<i>Cisco UCS Storage Accelerator Adapters</i>	258
Cisco UCS C-Series Rackmount Servers	258
Cisco UCS C-Series Dense Storage Rack Server	261
Cisco UCS C-Series Best Practices for Populating DRAM	261
Cisco Adapters for UCS C-Series Rackmount Servers	262
<i>Cisco UCS C-Series RAID Adapter Options</i>	262
<i>Cisco UCS C-Series Virtual Interface Cards and OEM CNA Adapters</i>	265
<i>Cisco UCS VIC 1387 mLOM</i>	266
<i>What Is SR-IOV (Single Root-IO Virtualization)?</i>	266
<i>Cisco UCS C-Series Network Interface Card</i>	267
<i>Cisco UCS C-Series Host Bus Adapter</i>	268
<i>What Is N_Port ID Virtualization (NPIV)?</i>	269
<i>Cisco UCS Storage Accelerator Adapters</i>	270
Cisco UCS E-Series Servers	271
Cisco UCS Mini	272
Cisco UCS M-Series Servers	272
Cisco UCS Software	274
Cisco UCS Manager	274
Cisco UCS Central	274
Cisco UCS StorMagic SvSAN	274
Cisco UCS Director	274
Cisco IMC Supervisor	275

	<i>System Menu</i>	276
	<i>Policies Menu</i>	276
	<i>Administration Menu</i>	276
	Cisco UCS Platform Emulator	277
	Cisco goUCS Automation Tool	277
	Cisco UCS Connectivity Architecture	277
	Cisco UCS 5108 Chassis to Fabric Interconnect Physical Connectivity	277
	Cisco UCS C-Series Rackmount Server Physical Connectivity	278
	Cisco UCS 6200 Fabric Interconnect to LAN, SAN Connectivity	280
	<i>Cisco UCS Fabric Interconnect Switching Modes</i>	281
	<i>Ethernet Switching Mode</i>	281
	<i>Fibre Channel Switching Mode</i>	282
	Cisco UCS I/O Module and Fabric Interconnect Connectivity	283
	Cisco UCS I/O Module Architecture	285
	Fabric and Adapter Port Channel Architecture	287
	Cisco Integrated Management Controller Architecture	287
	Reference List	289
	Exam Preparation Tasks	290
	Review All Key Topics	290
	Complete Tables and Lists from Memory	292
	Define Key Terms	292
Chapter 8	Cisco UCS Manager	295
	“Do I Know This Already?” Quiz	296
	Foundation Topics	299
	Cabling a Cisco UCS Fabric Interconnect HA Cluster	299
	Cisco UCS Fabric Interconnect HA Architecture	299
	Cisco UCS Fabric Interconnect Cluster Connectivity	300
	Initial Setup Script for Primary Fabric Interconnect	302
	Initial Setup Script for Subordinate Secondary Fabric Interconnect	303
	Verify Cisco UCS Fabric Interconnect Cluster Setup	303
	Changing Cluster Addressing via Command-Line Interface	305
	<i>Command Modes</i>	305
	<i>Changing Cluster IP Addresses via Command-Line Interface</i>	306
	<i>Changing Static IP Addresses via Command-Line Interface</i>	306
	Cisco UCS Manager Operations	307
	Cisco UCS Manager Functions	307
	Cisco UCS Manager GUI Layout	310

Cisco UCS Manager: Navigation Pane—Tabs	313
<i>Equipment Tab</i>	313
<i>Servers Tab</i>	314
<i>LAN Tab</i>	314
<i>SAN Tab</i>	315
<i>VM Tab</i>	315
<i>Admin Tab</i>	317
<i>Storage Tab</i>	317
Basic Port Roles in the Cisco UCS Fabric Interconnects	318
Cisco UCS Manager Hardware Discovery	319
Verifying Device Discovery in Cisco UCS Manager	321
Reference List	323
Exam Preparation Tasks	324
Review All Key Topics	324
Define Key Terms	325

Chapter 9 Cisco Unified Computing System Pools, Policies, Templates, and Service Profiles 327

“Do I Know This Already?” Quiz	327
Foundation Topics	331
Cisco UCS Service Profiles	331
Cisco UCS Hardware Abstraction and Stateless Computing	331
Contents of a Cisco UCS Service Profile	333
Other Benefits of Cisco UCS Service Profiles	335
<i>Ease of Incremental Deployments</i>	335
<i>Plan and Preconfigure Cisco UCS Solutions</i>	335
<i>Ease of Server Replacement</i>	336
<i>Right-Sized Server Deployments</i>	336
Cisco UCS Templates	336
Organizational Awareness with Service Profile Templates	337
Cisco UCS Service Profile Templates	338
Cisco UCS vNIC Templates	339
Cisco UCS vHBA Templates	339
Cisco UCS Logical Resource Pools	340
UUID Identity Pools	341
MAC Address Identity Pools	342
WWN Address Identity Pools	343
<i>WWNN Pools</i>	343
<i>WWPN Pools</i>	344

Cisco UCS Physical Resource Pools	345
Server Pools	345
Cisco UCS Server Qualification and Pool Policies	345
Server Auto-Configuration Policies	346
Creation of Policies for Cisco UCS Service Profiles and Service Profile Templates	347
Frequently Used Policies Explained	347
<i>Equipment Tab—Global Policies</i>	347
<i>BIOS Policy</i>	348
<i>Boot Policy</i>	348
<i>Local Disk Configuration Policy</i>	349
<i>Maintenance Policy</i>	350
<i>Scrub Policy</i>	351
<i>Host Firmware Packages</i>	351
<i>Adapter Policy</i>	352
Cisco UCS Chassis and Blade Power Capping	353
What Is Power Capping?	353
<i>Cisco UCS Power Capping Strategies</i>	353
Utilizing Cisco UCS Service Profile Templates	355
Creating Cisco UCS Service Profile Templates	355
Modifying a Cisco UCS Service Profile Template	356
Utilizing Cisco UCS Service Profile Templates for Different Workloads	357
Creating Service Profiles from Templates	357
Reference List	359
Exam Preparation Tasks	360
Review All Key Topics	360
Complete Tables and Lists from Memory	361
Define Key Terms	361
Chapter 10 Administration, Management, and Monitoring of Cisco UCS	363
“Do I Know This Already?” Quiz	363
Foundation Topics	367
Cisco UCS Administration	367
Organization and Locales	367
Role-Based Access Control	369

Authentication Methods	370
LDAP Providers	370
LDAP Provider Groups	370
Domain, Realm, and Default Authentication Settings	371
Communication Management	372
UCS Firmware, Backup, and License Management	373
Firmware Terminology	373
Firmware Definitions	374
Firmware Update Guidelines	375
Host Firmware Packages	375
Cross-Version Firmware Support	376
Firmware Auto Install	376
UCS Backups	377
Backup Automation	378
License Management	379
Cisco UCS Management	380
Operational Planes	380
In-Band Versus Out-of-Band Management	381
Remote Accessibility	382
Direct KVM Access	382
Advanced UCS Management	383
Cisco UCS Manager XML API	383
goUCS Automation Toolkit	384
Using the Python SDK	386
Multi-UCS Management	387
Cisco UCS Monitoring	387
Cisco UCS System Logs	388
Fault, Event, and Audit Logs (Syslogs)	388
System Event Logs	391
UCS SNMP	391
Fault Suppression	393
Collection and Threshold Policies	394
Call Home and Smart Call Home	396
Reference List	400
Exam Preparation Tasks	401
Review All Key Topics	401
Define Key Terms	401

Chapter 11	Server Virtualization Solutions	403
	“Do I Know This Already?” Quiz	403
	Foundation Topics	406
	Brief History of Server Virtualization	406
	Server Virtualization Components	407
	Hypervisor	407
	Virtual Machines	408
	Virtual Switching	409
	Shared Storage	410
	Management Tools	411
	Types of Server Virtualization	411
	Full Virtualization	412
	Paravirtualization	413
	Operating System Virtualization	414
	Server Virtualization Benefits and Challenges	415
	Server Virtualization Benefits	415
	Server Virtualization Challenges	416
	Reference List	421
	Exam Preparation Tasks	422
	Review All Key Topics	422
	Define Key Terms	422
	Part III Review	424
Part IV	Application Centric Infrastructure	429
Chapter 12	ACI Architecture	431
	“Do I Know This Already?” Quiz	432
	Foundation Topics	436
	What’s an Application Programming Interface (API)?	436
	Network Management Options	436
	Cisco Software-Defined Networking (SDN)	438
	The Policy-Based Model	445
	OpFlex	447
	Spine-Leaf Data Center Design	447
	ACI Physical Topology and Software	449
	ACI Fabric Hardware—Cisco Nexus 9000 Family	452
	Application Policy Infrastructure Controller (APIC)	456
	VXLAN Forwarding	464
	Mapping Database	467

- ACI Fabric Provisioning and Startup Discovery 468
- Reference List 470
- Exam Preparation Tasks 472
- Review All Key Topics 472
- Complete Tables and Lists from Memory 473
- Define Key Terms 473

Chapter 13 ACI Logical Model and Policy Framework 475

- “Do I Know This Already?” Quiz 475
- Foundation Topics 479
- Abstraction 479
- ACI Policy Object Model 480
- Different Models 487
- ACI Logical Constructs 488
 - Bridge Domains and Subnets 492
 - In-Band and Out-of-Band Management Networks 493
 - Endpoint Groups 497
 - Policy Enforcement 502
 - Multicast Policy Enforcement 504
 - Microsegmentation 505
 - Intra-EPG Endpoint Isolation 505
 - Application Profiles 506
 - Contracts 506
 - Labels, Filters, and Subjects 508
- Reference List 509
- Exam Preparation Tasks 511
- Review All Key Topics 511
- Define Key Terms 512

Chapter 14 Operating ACI 515

- “Do I Know This Already?” Quiz 515
- Foundation Topics 518
- ACI and Three-Tier Application Model 518
- Telemetry 520
 - Atomic Counters 520
 - Latency Metrics 522
 - Health Scores and Health Monitoring 522
 - Faults 523
 - Summary of Cisco ACI Telemetry 526

Open Policy Framework	528
Basic or Advanced GUI?	530
Menu Bar and Submenu Bar	531
Navigation and Work Pane	532
System Tab	533
Tenants Tab	533
Fabric Tab	534
VM Networking Tab	534
L4-L7 Services Tab	534
Admin Tab	534
Operations Tab	534
Hypervisor Integration	535
Virtual Machine Manager Domain Main Components	536
Attachable Entity Profile (AEP)	537
VMware vSphere Distributed Switch and VMware vShield Integration	538
<i>VMware vCenter and vShield Domain Operational Workflow</i>	540
VMware AVS Integration—Cisco Application Virtual Switch (AVS)	542
Local Switching Mode	543
No Local Switching Mode	543
Microsoft SCVMM and Azure Pack Integration	544
OpenStack	546
Layer 4 to Layer 7 Services Integration	549
Service Graphs	550
Device Packages	552
Microsegmentation	553
Why Microsegmentation Matters	554
Main Features	555
Reference List	555
Exam Preparation Tasks	556
Review All Key Topics	556
Complete Tables and Lists from Memory	557
Define Key Terms	557
Part IV Review	558

Part V Automation and Orchestration 563

Chapter 15 Cloud Computing 565

“Do I Know This Already?” Quiz	565
Foundation Topics	568
What Is Cloud Computing?	568
Brief History of Cloud Computing	569
Cloud Computing and Virtualization	570
Cloud Computing Deployment Models	571
Private Cloud	571
Public Cloud	572
Hybrid Cloud	573
Community Cloud	574
Cloud Computing Services Models	574
Infrastructure as a Service	574
Platform as a Service	576
Software as a Service	576
Application Programming Interfaces	577
Reference List	579
Exam Preparation Tasks	580
Review All Key Topics	580
Define Key Terms	580

Chapter 16 UCS Director 583

“Do I Know This Already?” Quiz	583
Foundation Topics	587
What Is UCS Director?	587
Benefits of UCS Director	588
Features and Functions of UCS Director	589
UCS Director Solution Overview	591
UCS Director Components	591
Cisco ONE Enterprise Cloud Suite	592
UCS Director Deployment	594
Single-Node Deployment	594
Multinode Deployment	595
Deploying Baremetal Agent	596
UCS Director Personalities	600
UCSD Administration	600
Shell Administration	600
UCS Director REST API	601

Installing and Managing a License	601
Mail Setup	602
Configuring System Parameters	602
Multitenancy	603
User Roles	605
Users and Group Management	606
Authentication Preference	607
Lightweight Directory Access Protocol (LDAP) Integration	607
Support Information	607
Unified Infrastructure Management	608
Site	608
Pods	609
Virtual Account	610
Physical Account	611
Multidomain Managers	612
Managed Network Elements	613
Rack Accounts	614
Device Discovery	614
Policies and Policy-Based Provisioning	615
Credential Policies	616
Computing Policies	616
Network Policies	617
Storage Policies	617
System Policies	617
Virtual Data Center (VDC)	618
Catalogs	619
Application Container	620
Self-Service Provisioning	620
ACI Integration	621
Resource Groups	621
Resource Tags	622
Service Offerings	622
Tenant Profile	623
Chargeback	623
Budget Policies	623
Cost Model	623

Reporting and Monitoring	624
Dashboard	624
Summary	624
Reports	624
CloudSense Reports	625
Report Builder	625
Reference List	626
Exam Preparation Tasks	627
Review All Key Topics	627
Complete Tables and Lists from Memory	627
Define Key Terms	628

Chapter 17 Understanding and Troubleshooting UCSD Workflows 631

“Do I Know This Already?” Quiz	631
Foundation Topics	634
Orchestration	634
UCS Director Orchestrator	635
Understanding UCS Director Workflows	635
Task	636
Input and Output	636
Workflow	638
Activities	638
Service Request	639
Task Libraries	639
Sample Workflow	640
Using Workflow Designer	640
Creating Workflows	641
Validating Workflow	645
Executing Workflows	645
Managing Workflows	646
Workflow Templates	646
Workflow Troubleshooting	647
Workflow Resources	649
Open Automation	650
Reference List	650
Exam Preparation Tasks	651
Review All Key Topics	651
Define Key Terms	651

Part V Review 652

Part VI Final Preparation 657**Chapter 18 Final Preparation 659**

Advice About the Exam Event 659

Learn the Question Types Using the Cisco Certification Exam
Tutorial 659

Think About Your Time Budget Versus the Number of Questions 663

A Suggested Time-Check Method 664

Miscellaneous Pre-Exam Suggestions 664

Exam-Day Advice 665

Pearson IT Certification Practice Test Engine and Questions 665

Companion Website 666

Install the Software 666

Activate and Download the Practice Exam 667

Activate Other Exams 667

Premium Edition 668

Exam Review 669

Take Practice Exams 669

Practice Taking the DCICT Exam 670

Advice on How to Answer Exam Questions 671

Take Other Practice Exams 672

Find Knowledge Gaps Through Question Review 673

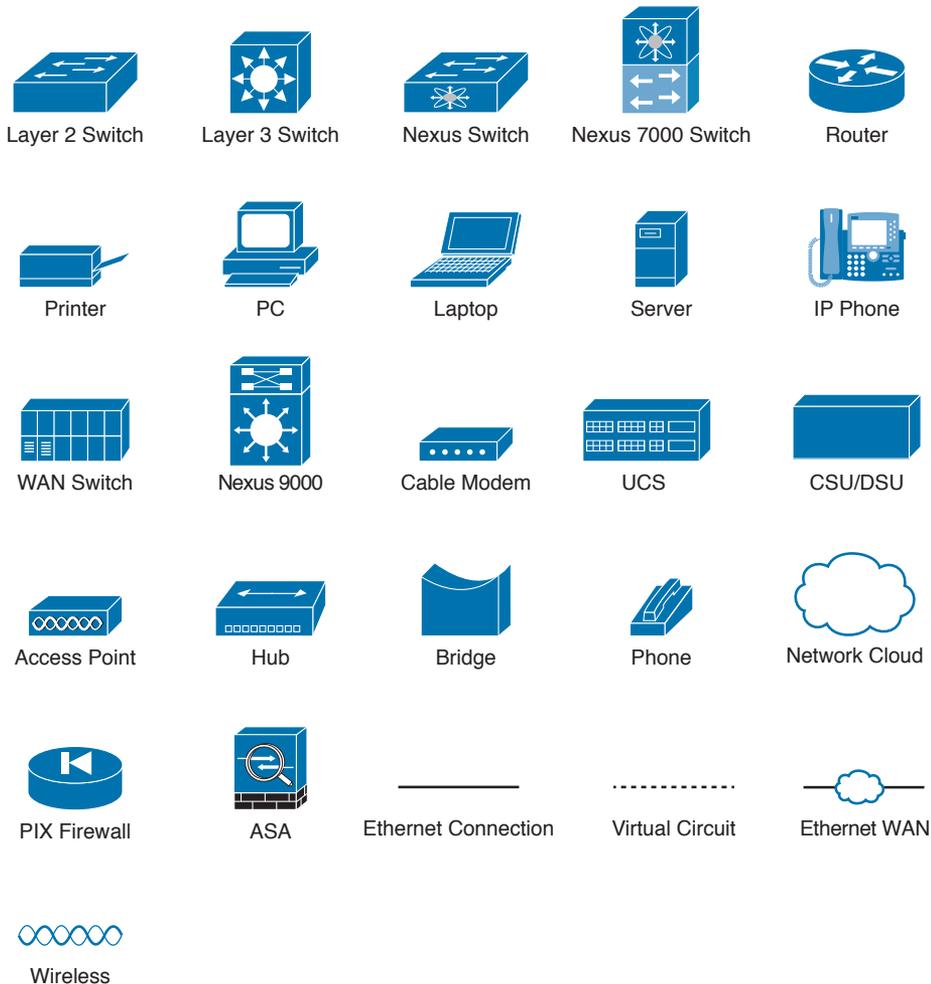
Practice Hands-On CLI Skills 675

Other Study Tasks 675

Final Thoughts 676

Appendix A Answers to the “Do I Know This Already?” Quizzes 679**Appendix B Memory Tables 697****Appendix C Memory Tables Answer Key 709****Glossary 721****Index 752****On the Companion Website:****Appendix D Study Planner**

Icons Used in This Book



Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italic* indicates arguments for which you supply actual values.

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ({ [] }) indicate a required choice within an optional element.

Introduction

About the Exam

Congratulations! If you are reading far enough to look at this book's Introduction, you've probably already decided to pursue your Cisco CCNA Data Center certification. Cisco dominates the networking marketplace, and after a few short years of entering the server marketplace, Cisco has achieved significant market share and has become one of the primary vendors for server hardware. If you want to succeed as a technical person in the networking industry in general, and in data centers in particular, you need to know Cisco. Getting your CCNA Data Center certification is a great first step in building your skills and becoming a recognized authority in the data center field.

Exams That Help You Achieve CCNA Data Center Certification

Cisco CCNA Data Center is an entry-level Cisco data center certification that is also a prerequisite for other Cisco Data Center certifications. CCNA Data Center itself has no other prerequisites. To achieve the CCNA Data Center certification, you must pass two exams: 200-150 Introduction to Cisco Data Center Networking (DCICN) and 200-155 Introduction to Cisco Data Center Technologies (DCICT), as shown in Figure I-1.

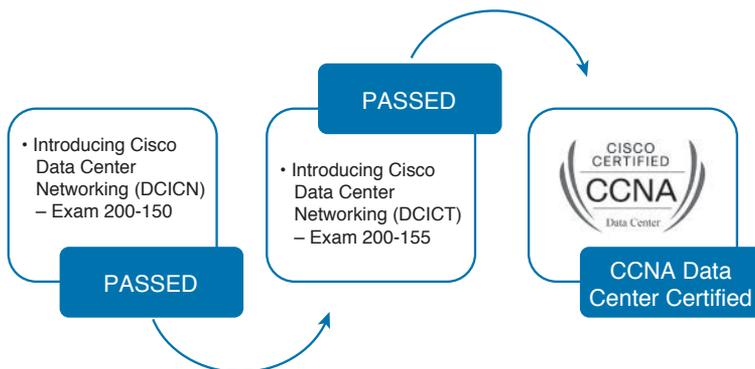


Figure I-1 Path to Cisco CCNA Data Center Certification

The DCICN and DCICT exams differ quite a bit in terms of the topics covered. DCICN focuses on networking technology. In fact, it overlaps quite a bit with the topics in the ICND1 100-105 exam, which leads to the Cisco Certified Entry Network Technician (CCENT) certification. DCICN explains the basics of networking, focusing on Ethernet switching and IP routing. The only data center focus on the DCICN exam is that all the configuration and verification examples use Cisco Nexus data center switches.

The DCICT exam instead focuses on technologies specific to the data center. These technologies include Unified Computing, Unified Fabric, network virtualization, Application Centric Infrastructure, cloud computing, automation and orchestration, as well as the other data networking features unique to the Cisco Nexus series of switches.

Types of Questions on the Exams

Cisco certification exams follow the same general format. At the testing center, you sit in a quiet room in front of the PC. Before the exam timer begins, you can complete a few other tasks on the PC; for example, you can take a sample quiz just to get accustomed to the PC and the testing engine. Anyone who has basic skills in getting around a PC should have no problems with the testing environment.

After the exam starts, you are presented with a series of questions, one at a time, on the PC screen. The questions typically fall into one of the following categories:

- Multiple choice, single answer
- Multiple choice, multiple answers
- Testlet
- Drag-and-drop (DND)
- Simulated lab (sim)
- Simlet

The first three items in the list are all multiple-choice questions. The multiple-choice format requires you to point to and click a circle or square beside the correct answer(s).

Cisco traditionally tells you how many answers you need to choose, and the testing software prevents you from choosing too many answers. The testlet asks you several multiple-choice questions, all based on the same larger scenario.

DND questions require you to move some items around in the graphical user interface (GUI). You left-click the mouse to hold the item, move it to another area, and release the mouse button to place the item in its destination (usually into a list). For some questions, to get the question correct, you might need to put a list of items in the proper order or sequence.

The last two types, sim and simlet questions, both use a network simulator to ask questions. Interestingly, the two types enable Cisco to assess two very different skills. First, sim questions generally describe a problem, while your task is to configure one or more routers and switches to fix it. The exam then grades the question based on the configuration you changed or added. Basically, these questions begin with a broken configuration, and you must fix it to answer the question correctly.

Simlet questions also use a network simulator, but instead of answering the question by changing or adding the configuration, they include one or more multiple-choice questions. These questions require you to use the simulator to examine network behavior by interpreting the output of **show** commands you decide to leverage to answer the question. Whereas sim questions require you to troubleshoot problems related to configuration, simlets require you to analyze both working and broken networks, correlating **show** command output with your knowledge of networking theory and configuration commands.

You can watch and even experiment with these command types using the Cisco Exam Tutorial. To find the Cisco Certification Exam Tutorial, go to www.cisco.com and search for “exam tutorial.”

What's on the DCICT Exam?

Everyone has always wanted to know what is on the test, for any test, since the early days of school. Cisco openly publishes the topics of each of its certification exams. Cisco wants the candidates to know the variety of topics and get an idea about the kinds of knowledge and skills required for each topic.

Exam topics are very specific, and the verb used in their description is very important. The verb tells us to what degree the topic must be understood and what skills are required. For example, one topic might begin with “Describe...,” another with “Configure...,” another with “Verify...,” and another with “Troubleshoot...” Questions beginning with “Troubleshoot” require the highest skills level, because to troubleshoot, you must understand the topic, be able to configure it (to see what's wrong with the configuration), and be able to verify it (to find the root cause of the problem). Pay attention to the question verbiage.

Cisco's posted exam topics, however, are only guidelines. Cisco's disclaimer language mentions that fact. Cisco makes an effort to keep the exam questions within the confines of the stated exam topics, and we know from talking to those involved that every question is analyzed for whether it fits the stated exam topic.

DCICT 200-155 Exam Topics

The exam topics for both the DCICN and the DCICT exams can be easily found at Cisco.com by searching. Alternatively, you can go to www.cisco.com/go/ccna, which gets you to the page for CCNA Routing and Switching, where you can easily navigate to the nearby CCNA Data Center page.

Over time, Cisco has begun making two stylistic improvements to the posted exam topics. In the past, the topics were simply listed as bullets with indentation to imply subtopics under a major topic. More often today, including for the DCICN and DCICT exam topics, Cisco also numbers the exam topics, making it easier to refer to specific topics. Additionally, Cisco lists the weighting for each of the major topic headings. The weighting tells the percentage of points from your exam, which should come from each major topic area. The DCICT contains five major headings with their respective weighting, shown in Table I-1.

Table I-1 Five Major Topic Areas in the DCICT 200-155 Exam

Number	Exam Topic	Weighting
1.0	Unified Computing	25%
2.0	Network Virtualization	17%
3.0	Cisco Data Center Networking Technologies	26%
4.0	Automation and Orchestration	15%
5.0	Application Centric Infrastructure	17%

Note that while the weighting of each topic area tells you something about the exam, in the authors' opinion, the weighting probably does not change how you study. All five topic areas hold enough weighting so that if you completely ignore an individual topic, you probably will not pass. Furthermore, data center technologies require you to put many concepts together, so you need all the pieces before you can understand the holistic view. The weighting might indicate where you should spend a little more time during the last days before taking the exam, but otherwise, plan to study all the exam topics.

Tables I-2 through I-6 list the details of the exam topics, with one table for each of the major topic areas listed in Table I-1. Note that these tables also list the book chapters that discuss each of the exam topics.

Table I-2 Exam Topics in the First Major DCICT Exam Topic Area

Number	Exam Topic	Chapter
1.0	Unified Computing	7, 8, 9, 10, 11
1.1	Describe common server types and connectivity found in a data center	7
1.2	Describe the physical components of the Cisco UCS	7
1.3	Describe the concepts and benefits of Cisco UCS hardware abstraction	9
1.4	Perform basic Cisco UCS configuration	8, 10
1.4.a	Cluster high availability	8, 10
1.4.b	Port roles	8, 10
1.4.c	Hardware discovery	8, 10
1.5	Describe server virtualization concepts and benefits	11
1.5.a	Hypervisors	11
1.5.b	Virtual switches	11
1.5.c	Shared storage	11
1.5.d	Virtual Machine components	11
1.5.e	Virtual Machine Manager	11

Table I-3 Exam Topics in the Second Major DCICT Exam Topic Area

Number	Exam Topic	Chapter
2.0	Network Virtualization	2, 4, 5, 6
2.1	Describe the components and operations of Cisco virtual switches	4
2.2	Describe the concepts of overlays	5
2.2.a	OTV	5

Number	Exam Topic	Chapter
2.2.b	NVGRE	5
2.2.c	VXLAN	5
2.3	Describe the benefits and perform simple troubleshooting of VDC STP	6
2.4	Compare and contrast the default and management VRFs	6
2.5	Differentiate between the data, control, and management planes	2

Table I-4 Exam Topics in the Third Major DCICT Exam Topic Area

Number	Exam Topic	Chapter
3.0	Cisco Data Center Networking Technologies	1, 2, 3, 6
3.1	Describe, configure, and verify FEX connectivity	6
3.2	Describe, configure, and verify basic vPC features	1
3.3	Describe, configure, and verify FabricPath	1
3.4	Describe, configure, and verify unified switch ports	3
3.5	Describe the features and benefits of Unified Fabric	3
3.6	Describe and explain the use of role-based access control within the data center infrastructure	2

Table I-5 Exam Topics in the Fourth Major DCICT Exam Topic Area

Number	Exam Topic	Chapter
4.0	Automation and Orchestration	15, 16, 17
4.1	Explain the purpose and value of using APIs	15
4.2	Describe the basic concepts of cloud computing	15
4.3	Describe the basic functions of a Cisco UCS Director	16
4.3.a	Management	16
4.3.b	Orchestration	17
4.3.c	Multitenancy	16
4.3.d	Chargeback	16
4.3.e	Service offerings	16
4.3.f	Catalogs	16
4.4	Interpret and troubleshoot a Cisco UCS Director workflow	17

Table I-6 Exam Topics in the Fifth Major DCICT Exam Topic Area

Number	Exam Topic	Chapter
5.0	Application Centric Infrastructure	12, 13, 14
5.1	Describe the architecture of an ACI environment	12
5.1.a	Basic policy resolution	12
5.1.b	APIC controller	12
5.1.c	Spine leaf	12
5.1.d	APIs	12
5.2	Describe the fabric discovery process	12
5.3	Describe the policy-driven, multitier application deployment model and its benefits	14
5.4	Describe the ACI logical model	13
5.4.a	Tenants	13
5.4.b	Context	13
5.4.c	Bridge domains	13
5.4.d	EPG	13
5.4.e	Contracts	13

NOTE Because it is possible for exam topics to change over time, it might be worth the time to double-check the exam topics as listed on the Cisco website (<http://learningnetwork.cisco.com> and navigate to the CCNA Data Center page).

About the Book

This book discusses the content and skills needed to pass the 200-155 DCICT certification exam, which is the second and final exam to achieve CCNA Data Center certification. This book's companion title, *CCNA Data Center DCICN 200-150 Official Cert Guide*, discusses the content needed to pass the 200-150 DCICN certification exam.

We strongly recommend that you plan and structure your learning to align with both exam requirements.

Book Features

The most important and somewhat obvious objective of this book is to help you pass the DCICT exam and help you achieve the CCNA Data Center certification. In fact, if the primary objective of this book were different, the book's title would be misleading! At the same time, the methods used in this book to help you pass the exam are also designed to

make you much more knowledgeable in the general field of the data center and help you in your daily job responsibilities.

This book uses several tools to help you discover your weak topic areas, to help you improve your knowledge and skills with those topics, and to prove that you have retained your knowledge of those topics. Importantly, this book does not try to help you pass the exams only by memorization, but by truly learning and understanding the topics. CCNA entry-level certification is the foundation for many of the Cisco professional-level certifications, and it would be a disservice to you if this book did not help you truly learn the material. This book helps you pass the CCNA exam by using the following methods:

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

Chapter Features

To help you customize study time using this book, the core chapters have several features that help you make the best use of your time:

- **“Do I Know This Already?” Quizzes:** Each chapter begins with a quiz that helps you determine the amount of time you need to spend studying the chapter.
- **Foundation Topics:** These are the core sections of each chapter. They explain the protocols, concepts, and configuration for the topics in the chapter.
- **Exam Preparation Tasks:** At the end of the “Foundation Topics” section of each chapter, the “Exam Preparation Tasks” section lists a series of study activities that should be completed at the end of the chapter. Each chapter includes the activities that make the most sense for studying the topics in that chapter. The activities include the following:
 - **Review All Key Topics:** The Key Topic icon is shown next to the most important items in the “Foundation Topics” section of the chapter. The “Review All Key Topics” activity lists the key topics from the chapter and their corresponding page numbers. Although the content of the entire chapter could appear on the exam, you should definitely know the information listed in each key topic.
 - **Complete Tables and Lists from Memory:** To help you exercise your memory and memorize certain lists of facts, many of the more important lists and tables from the chapter are included in Appendix B, “Memory Tables.” This document lists only partial information, allowing you to complete the table or list.
 - **Define Key Terms:** Although the exam might be unlikely to ask a question such as, “Define this term,” the CCNA exams require that you learn and know a lot of networking terminology. This section lists the most important terms from the chapter,

asking you to write a short definition and compare your answer to the Glossary at the end of this book.

- **References:** Some chapters contain a list of reference links for additional information and details on the topics discussed in that particular chapter.

Part Review

The part review tasks help you prepare to apply all the concepts you learned in that part of the book. Each book part contains several related chapters. The part review includes sample test questions that require you to apply the concepts from multiple chapters in that part, uncovering what you truly understood and what you did not quite yet understand.

The part reviews list tasks, along with checklists, so that you can track your progress. The following list explains the most common tasks you will see in the part review:

- **Repeat All “Do I Know This Already?” Questions:** Although you have already seen the “Do I Know This Already?” questions from the chapters in a part, answering those questions again can be a useful way to review facts. The “Part Review” section suggests that you repeat the “Do I Know This Already?” questions, but use the Pearson IT Certification Practice Test (PCPT) exam software that comes with the book for extra practice in answering multiple-choice questions on a computer.
- **Answer “Part Review” Questions:** The PCPT exam software includes several exam databases. One exam database holds “Part Review” questions, written specifically for part reviews. These questions purposefully include multiple concepts in each question, sometimes from multiple chapters, to help build the skills needed for the more challenging analysis questions on the exams.
- **Review Key Topics:** Yes, again! They are indeed the most important topics in each chapter.
- **Self-Assessment Questionnaire:** The exam is unlikely to ask a question such as, “Define this term,” but the CCNA exams require that you learn and know a lot of technology concepts and architectures. This section asks you some open questions that you should try to describe or explain in your own words. This will help you develop a thorough understanding of important exam topics pertaining to that part.

Final Prep Tasks

Chapter 18, “Final Preparation,” lists a series of tasks that you can use for your final preparation before taking the exam.

Other Features

In addition to the features in each of the core chapters, this book, as a whole, has additional study resources, including the following:

- **Practice Exam:** The companion website contains the powerful PCPT exam engine. You can answer the questions in study mode or take simulated DCICT exams with the website and activation code included in this book.
- **eBook:** If you are interested in obtaining an e-book version of this title, we have included a special offer on a coupon card inserted in the cardboard sleeve in the back of the book. This offer enables you to purchase the *CCNA Data Center DCICT 200-155 Official Cert Guide, Premium Edition* e-book and practice test at a 70 percent discount off the list price. In addition to three versions of the e-book—PDF (for reading on your computer), EPUB (for reading on your tablet, mobile device, or Nook or other e-reader), and Mobi (the native Kindle version)—you will receive additional practice test questions and enhanced practice test features.
- **Companion website:** The website www.ciscopress.com/title/9781587205910 posts up-to-the-minute material that further clarifies complex exam topics. Check this site regularly for new and updated postings written by the authors that provide further insight into the more troublesome topics on the exam.
- **PearsonITCertification.com:** The website www.pearsonitcertification.com is a great resource for all things IT-certification related. Check out the great CCNA articles, videos, blogs, and other certification preparation tools from the industry's best authors and trainers.

Book Organization, Chapters, and Appendixes

This book contains 17 core chapters—Chapters 1 through 17, with Chapter 18 including some suggestions for how to approach the actual exams. Each core chapter covers a subset of the topics on the 200-155 DCICT exam. The core chapters are organized into sections. The core chapters cover the following topics:

Part I: Data Center Networking Technologies

- **Chapter 1, “Data Center Networking:”** This chapter provides an overview of the data center networking architecture and design practices relevant to the exam. It goes into the detail of multilayer data center network design and technologies, such as port channel, virtual port channel, and Cisco FabricPath. Basic configuration and verification commands for these technologies are also included in this chapter.
- **Chapter 2, “Management and Monitoring of Cisco Nexus Devices:”** This chapter is an overview of operational planes of the Nexus platform—the data plane, control plane, and management plane. It explains the functions performed by each plane and provides an overview of out-of-band and in-band management interfaces. The Nexus platform provides several methods for device configuration and management. These methods are discussed, inclusive of important commands for initial setup, configuration, and verification. Readers will be introduced to Cisco NX-API, which allows using

HTTP/HTTPS as a transport to access and program the Cisco Nexus switches. This chapter also identifies the mechanism available in the Nexus platform to protect the control plane of the switch.

- **Chapter 3, “Unified Fabric Overview:”** This chapter offers an overview of challenges faced by today’s data centers. It focuses on how Cisco Unified Fabric architecture addresses those challenges by converging traditionally disparate network and storage environments while providing a platform for scalable, secure, and intelligent services. It also takes a look at some of the technologies allowing extension of the Unified Fabric environment beyond the boundary of a single data center.

Part II: Network Virtualization

- **Chapter 4, “Cisco Nexus 1000V and Virtual Switching:”** This chapter starts by describing the challenges of current virtual switching layers in data centers and then introduces the distributed virtual switches and, in particular, the Cisco vision—Cisco Nexus 1000V. The chapter explains installation options, commands to verify initial configuration of virtual Ethernet modules, virtual supervisor modules, and integration with VMware vCenter Server.
- **Chapter 5, “Data Center Overlay Networks:”** This chapter covers the latest Cisco innovations in the data center extension solutions and in the LAN extension in particular. Readers will learn about OTV, VXLAN, and NVGRE overlay protocols, which provide Layer 2 over Layer 3 connectivity in today’s data center networks.
- **Chapter 6, “Virtualizing Cisco Network Devices:”** This chapter covers the virtualization capabilities of the Nexus switches, using virtual device contexts (VDCs) and network interface virtualization (NIV). It details the VDC concept and the VDC deployment scenarios. The different VDC types and commands used to configure and verify the setup are also included in the chapter. Also covered in this chapter is the NIV—what it is, how it works, and how it is configured.

Part III: Cisco Unified Computing

- **Chapter 7, “Cisco UCS Architecture:”** This chapter begins with a quick fly-by on the evolution of server computing, followed by an introduction to the Cisco UCS value proposition, hardware, and software portfolio. Then the chapter explains UCS architecture in terms of component connectivity options and unification of blade and rack-mount server connectivity. It also details the Cisco Integrated Management Controller architecture and purpose.
- **Chapter 8, “Cisco UCS Manager:”** This chapter starts by describing how to set up, configure, and verify the Cisco UCS Fabric Interconnect cluster. It then describes the process of hardware and software discovery in Cisco UCS. It also explains how to monitor and verify this process.
- **Chapter 9, “Cisco Unified Computing System Pools, Policies, Templates, and Service Profiles:”** This chapter explains the hardware abstraction layer in more detail and how it relates to stateless computing. Then it explains logical and physical resource pools and the essentials to create templates and aid rapid deployment of service profiles. As a bonus, you also see notes, tips, and most relevant features.

- **Chapter 10, “Administration, Management, and Monitoring of Cisco UCS:”** This covers some of the important features used when administering and monitoring Cisco UCS. It also introduces Cisco UCS XML, goUCS automation toolkit, and the well-documented UCS XML API using the Python SDK.
- **Chapter 11, “Server Virtualization Solutions:”** This chapter takes a peek into the history of server virtualization and discusses fundamental principles behind different types of server virtualization technologies. It evaluates the benefits and challenges of server virtualization while offering approaches to mitigate performance and security concerns.

Part IV: Application Centric Infrastructure

- **Chapter 12, “ACI Architecture:”** This chapter discusses the ACI in four main areas: the ACI architecture and how it ties to software-defined networking (SDN), the Cisco application policy infrastructure controller (APIC) and ACI components, the application programming interface (API), and basic policy resolution in ACI. As a bonus, you will learn the latest ACI product portfolio, fundamental concepts, and advantages of ACI.
- **Chapter 13, “ACI Logical Model and Policy Framework:”** This chapter outlines the ACI policy model logical constructs. It introduces readers to a basic level of understanding about the model, what this policy model contains, and how to work with it. The complete object model contains a hierarchy of data center interactions. The most extensive information resource, which is available at this moment, is the APIC Management Information Model Reference packaged with the APIC itself.
- **Chapter 14, “Operating ACI:”** This chapter starts by explaining the three-tier Application model. It outlines the ACI hypervisor integration from multiple vendors, service integration, Cisco application virtual switch (AVS), and Openstack integration. Technologies that deliver telemetry for ACI fabric are also explained.

Part V: Automation and Orchestration

- **Chapter 15, “Cloud Computing:”** This chapter peeks into the history of cloud computing and provides an overview of various cloud computing deployment and services models for public, private, and hybrid environments. It discusses the use of application programming interfaces (APIs) to power a programmatic approach to provisioning and operating IT infrastructure elements for an unparalleled agility.
- **Chapter 16, “UCS Director:”** This chapter gives readers an overview of UCS Director features and functions that are relevant to the DCICT exam. It also provides configuration and troubleshooting guidelines and information about using UCS Director to simplify data center management.
- **Chapter 17, “Understanding and Troubleshooting UCS Director Workflows:”** This chapter offers readers a glimpse of UCS Director orchestration. It provides necessary information about workflow designer to create, validate, execute, and troubleshoot UCS Director workflows.

Part VI: Final Preparation

- **Chapter 18, “Final Preparation:”** This chapter suggests a plan for exam preparation after you have finished the core parts of the book, in particular explaining the many study options available in the book.
- **Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes:”** This includes the answers to all the questions from Chapters 1 through 17.
- **Appendix B, “Memory Tables:”** This holds the key tables and lists from each chapter, with some of the content removed. You can print this appendix and, as a memory exercise, complete the tables and lists. The goal is to help you memorize facts that can be useful on the exams.
- **Appendix C, “Memory Tables Answer Key:”** This contains the answer key for the exercises in Appendix B.
- The **Glossary** contains definitions for all the terms listed in the “Define Key Terms” section at the conclusion of Chapters 1 through 17.

Appendix On the Companion Website

- **Appendix D, “Study Planner:”** This is a spreadsheet with major study milestones enabling you to track your progress through your study.

Reference Information

This short section contains a few topics available for reference elsewhere in the book. You may read these when you first use the book, but you may also skip these topics and refer to them later. In particular, make sure to note the final page of this Introduction, which lists several contact details, including how to get in touch with Cisco Press.

Companion Website

Register this book to get access to the Pearson IT Certification test engine and other study materials plus additional bonus content. Check this site regularly for new and updated postings written by the authors that provide further insight into the more troublesome topics on the exam. Be sure to check the box that you would like to hear from us to receive updates and exclusive discounts on future editions of this product or related products.

To access this companion website, follow these steps:

1. Go to www.pearsonITcertification.com/register and log in or create a new account.
2. Enter the ISBN: 9781587205910.
3. Answer the challenge question as proof of purchase.

4. Click on the “Access Bonus Content” link in the Registered Products section of your account page, to be taken to the page where your downloadable content is available.

Please note that many of our companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following the steps above, please visit www.pearsonITcertification.com/contact and select the “Site Problems/Comments” option. Our customer service representatives will assist you.

Pearson IT Certification Practice Test Engine and Questions

The companion website includes the Pearson IT Certification Practice Test engine—software that displays and grades a set of exam-realistic multiple-choice questions. Using the Pearson IT Certification Practice Test engine, you can either study by going through the questions in Study Mode, or take a simulated exam that mimics real exam conditions. You can also serve up questions in a Flash Card Mode, which will display just the question and no answers, challenging you to state the answer in your own words before checking the actual answers to verify your work.

The installation process requires two major steps: installing the software and then activating the exam. The website has a recent copy of the Pearson IT Certification Practice Test engine. The practice exam (the database of exam questions) is not on this site.

NOTE The cardboard sleeve in the back of this book includes a piece of paper. The paper lists the activation code for the practice exam associated with this book. Do not lose the activation code. On the opposite side of the paper from the activation code is a unique, one-time-use coupon code for the purchase of the Premium Edition eBook and Practice Test.

Install the Software

The Pearson IT Certification Practice Test is a Windows-only desktop application. You can run it on a Mac using a Windows virtual machine, but it was built specifically for the PC platform. The minimum system requirements are as follows:

- Windows 10, Windows 8.1, or Windows 7
- Microsoft .NET Framework 4.0 Client
- Pentium-class 1GHz processor (or equivalent)
- 512 MB RAM
- 650 MB disk space plus 50 MB for each downloaded practice exam
- Access to the Internet to register and download exam databases

The software installation process is routine compared with other software installation processes. If you have already installed the Pearson IT Certification Practice Test software from another Pearson product, there is no need for you to reinstall the software. Simply launch the software on your desktop and proceed to activate the practice exam from this book by using the activation code included in the access code card sleeve in the back of the book.

The following steps outline the installation process:

1. Download the exam practice test engine from the companion site.
2. Respond to windows prompts as with any typical software installation process.

The installation process will give you the option to activate your exam with the activation code supplied on the paper in the cardboard sleeve. This process requires that you establish a Pearson website login. You need this login to activate the exam, so please do register when prompted. If you already have a Pearson website login, there is no need to register again. Just use your existing login.

Activate and Download the Practice Exam

Once the exam engine is installed, you should then activate the exam associated with this book (if you did not do so during the installation process) as follows:

1. Start the Pearson IT Certification Practice Test software from the Windows Start menu or from your desktop shortcut icon.
2. To activate and download the exam associated with this book, from the My Products or Tools tab, click the Activate Exam button.
3. At the next screen, enter the activation key from paper inside the cardboard sleeve in the back of the book. Once you've entered the key, click the Activate button.
4. The activation process will download the practice exam. Click **Next**, and then click **Finish**.

When the activation process completes, the My Products tab should list your new exam. If you do not see the exam, make sure that you have selected the **My Products** tab on the menu. At this point, the software and practice exam are ready to use. Simply select the exam and click the **Open Exam** button.

To update a particular exam you have already activated and downloaded, display the **Tools** tab and click the **Update Products** button. Updating your exams will ensure that you have the latest changes and updates to the exam data.

If you want to check for updates to the Pearson Cert Practice Test exam engine software, display the **Tools** tab and click the **Update Application** button. You can then ensure that you are running the latest version of the software engine.

Activating Other Exams

The exam software installation process, and the registration process, only has to happen once. Then, for each new exam, only a few steps are required. For instance, if you buy another Pearson IT Certification Cert Guide, extract the activation code from the cardboard sleeve in the back of that book; you do not even need the exam engine at this point. From there, all you have to do is start the exam engine (if not still up and running) and perform Steps 2 through 4 from the previous list.

Assessing Exam Readiness

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

Premium Edition eBook and Practice Tests

This book also includes an exclusive offer for 70 percent off the Premium Edition eBook and Practice Tests edition of this title. Please see the coupon code included with the cardboard sleeve for information on how to purchase the Premium Edition.

How to View Only “Do I Know This Already?” Questions by Part

Each “Part Review” section asks you to repeat the “Do I Know This Already?” quiz questions from the chapters in that part. Although you could simply scan the book pages to review these questions, it is slightly better to review these questions from inside the PCPT software, just to get a little more practice in how to read questions from the testing software.

To view these “Do I Know This Already?” (book) questions inside the PCPT software, follow these steps:

- Step 1.** Start the PCPT software.
- Step 2.** From the main (home) menu, select the item for this product, with a name like DCICT 200-155 Official Cert Guide, and click **Open Exam**.
- Step 3.** The top of the next window that appears should list some exams; select the box beside DCICT Book Questions, and deselect the other boxes. This selects the “book” questions (that is, the “Do I Know This Already?” questions from the beginning of each chapter).
- Step 4.** In this same window, click at the bottom of the screen to deselect all objectives (chapters), and then select the box beside each chapter in the part of the book you are reviewing.

Step 5. Select any other options on the right side of the window.

Step 6. Click **Start** to start reviewing the questions.

How to View Only Part Review Questions by Part

The exam databases you get with this book include a database of questions created solely for study during the part review process. “Do I Know This Already?” questions focus more on facts, with basic application. The part review questions instead focus more on application and look more like real exam questions.

To view these questions, follow the same process as you did with “Do I Know This Already?” book questions, but select the part review database instead of the book database, as follows:

Step 1. Start the PCPT software.

Step 2. From the main (home) menu, select the item for this product, with a name like DCICT 200-155 Official Cert Guide, and click **Open Exam**.

Step 3. The top of the next window should list some exams; select the box beside Part Review Questions, and deselect the other boxes. This selects the questions intended for part-ending review.

Step 4. On this same window, click at the bottom of the screen to deselect all objectives, and then select (check) the box beside the book part you want to review. This tells the PCPT software to give you part review questions from the selected part.

Step 5. Select any other options on the right side of the window.

Step 6. Click **Start** to start reviewing the questions.

For More Information

If you have any comments about the book, submit them via www.ciscopress.com. Just go to the website, select **Contact Us**, and type your message.

Cisco might make changes that affect the CCNA data center certification from time to time. You should always check www.cisco.com/go/certification for the latest details.

The *CCNA Data Center DCICT 200-155 Official Cert Guide* helps you attain the CCNA data center certification. This is the DCICT exam prep book from the only Cisco-authorized publisher. We at Cisco Press believe that this book certainly can help you achieve CCNA data center certification, but the real work is up to you. We trust that your time will be well spent.



Cisco Nexus 1000V and Virtual Switching

With server virtualization, a physical server can now run multiple workloads as virtual machines. Each virtual machine will be running its own guest operating system and invariably will require access to different networks. The virtual switching layers on hypervisors achieve this by giving you port groups (virtual patch cords), vSwitches (virtual switches in software), and associated uplinks (physical uplinks associated to those vSwitches), for example, on VMware ESX or ESXi hypervisors. Similar constructs are offered in various flavors of hypervisors available today to enable virtual switching, including Hyper-V virtual switch and Open Virtual Switch on KVM.

This chapter begins by describing the limitations of physical access layers. Then, as server virtualization has increased within data centers, the challenges being faced with virtual switching layers and static vSwitches are discussed. After you understand the limitations and challenges, you will start understanding distributed virtual switches and the Cisco virtual network vision. The chapter introduces you to the Cisco Nexus 1000V virtual distributed switch and explains how it solves the challenges in virtual data centers today.

As the chapter progresses, you learn about the Cisco Nexus 1000V virtual switching architecture and understand its integration with VMware ESX or ESXi and VMware vCenter. You also learn the installation methods and commands that enable you to verify the initial configuration and module status of the Cisco Nexus 1000V Series switch and validate connectivity between the virtual Ethernet module (VEM), virtual supervisor module (VSM), and VMware vCenter using VMware ESX command-line interface (CLI), Nexus 1000V CLI, and VMware vCenter vSphere Client.

NOTE The Cisco Nexus 1000V Series switch was developed with close cooperation between Cisco and VMware. This chapter focuses on the Cisco Nexus 1000V integration between VMware ESXi hypervisor and VMware vCenter. At the time of writing this chapter, the Cisco Nexus 1000V distributed virtual switch was also supported on other independent hypervisors, such as Microsoft Hyper-V and Linux/KVM.

Today, there are multiple hypervisor vendors and different implementations of distributed virtual switches. For all intents and purposes in this chapter, explanations and comparisons will be based on the VMware static vSwitch and VMware distributed virtual switch (DVS) technology.

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz enables you to assess whether you should read this entire chapter thoroughly or jump to the “Exam Preparation Tasks” section. If you are in doubt about your answers to these questions or your own assessment of your knowledge of the topics, read the entire chapter. Table 4-1 lists the major headings in this chapter and their corresponding “Do I Know This Already?” quiz questions. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes.”

Table 4-1 “Do I Know This Already?” Section-to-Question Mapping

Foundation Topics Section	Questions
“Evolution of Virtual Switching”	1, 2, 3
“Virtual Networking Component”	4
“VMware vDS Overview”	5
“Advantages of VMware vDS and Enhancements”	6
“Cisco Nexus 1000V Series Switch Salient Features and Benefits”	7
“Cisco Nexus 1000V Architecture”	8
“Cisco Nexus 1000V Component Communication”	9
“Cisco Nexus 1000V Port Profiles”	10
“Initial VSM Configuration Verification”	11
“Verifying VEM Agent”	12
“Validating VM Port Groups and Port Profiles”	13

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. What were key challenges with virtual machines and their hosts’ physical access layer prior to Nexus 1000V? (Choose two.)
 - a. VMs could not be moved from one host to another.
 - b. Enforcing policies at the VM vNIC level.
 - c. Virtual machines did not have enough bandwidth and could not scale.
 - d. Making these policies “sticky” and making sure they follow the VM.

2. True or false? You could control a VM's communication belonging to the same VLAN from the virtual access layer, similar to a physical switch port access control prior to Cisco Nexus 1000V Series switch.
 - a. True
 - b. False
3. What were key challenges with standard VMware static vSwitches? (Choose two.)
 - a. A compromised VM could be traced only as far as its host's physical switch port.
 - b. VM administrators were suddenly tasked with managing access layer networking, for which they did not have extensive experience around data center networking.
 - c. The VMware static vSwitches needed to be upgraded on a regular basis.
 - d. The VMware static vSwitch had very limited ports and could connect only a certain number of VMs.
4. What are valid virtual networking components?
 - a. Port groups, vNICs, VMNICs, and guest operating systems
 - b. IP addresses, NICs, switches, routers, and load balancers
 - c. vNICs, VMNICs, port groups, physical networks, and virtual networks
 - d. VLANs and vSwitches
5. Which statement best describes a VMware vDS?
 - a. A VMware vDS is a distributed switch that virtually enables you to manage one switch across multiple VMware ESX/ESXi hosts.
 - b. The VMware vDS is the latest version of the VMware vSwitch.
 - c. The VMware vDS is an API that enables you to programmatically alter the configuration of VMware vSwitches.
 - d. The VMware vDS is a centralized switch management platform that enables you to manage geographically distributed vSwitches from a single location.
6. What are two main advantages of a VMware vDS when compared with VMware vSwitch?
 - a. VM LAN, SAN, Admin Pane managed from a single point.
 - b. VMware cluster (span across multiple VMware ESX or ESXi hosts)–level network management and configuration capability.
 - c. Policies now have the capability to follow the VM as it traverses hosts within a VMware cluster.
 - d. Equipment, LAN, SAN, and Admin Panes managed from a single point.
7. True or false? The Cisco Nexus 1000V Series switch is based on a standard Cisco iOS software command-line interface (CLI).
 - a. True
 - b. False

8. What are the main “components” of a Cisco Nexus 1000V Series switch?
 - a. Nexus 1000, 2000, 3000, 4000, 5000, 6000, 7000, and 9000 Series switches
 - b. The Virtual Supervisory Module and Virtual Ethernet Module
 - c. The Virtual Chassis and Virtual Line Cards
 - d. The Nexus 1010/11000 Series virtual services appliance
9. True or false? The Cisco Nexus 1000V Series switch can be managed and controlled over Layer 2 and Layer 3 networks.
 - a. True
 - b. False
10. Which statement best describes a Cisco Nexus 1000V port profile?
 - a. You can define two types of port profiles in a Cisco Nexus 1000V Series switch (Ethernet, vEthernet). These port profiles are equivalent to VMware port groups and can be assigned to VMs or uplinks on VMware ESX/ESXi hosts.
 - b. Port profiles contain a VLAN ID, which can be assigned to a trunk interface.
 - c. A port profile is the Cisco terminology to refer to a VM as seen from the Cisco Nexus 1000V Series switch.
 - d. Port profiles are defined on VMware vCenter server and can be accessed by all Cisco networking devices within a virtualized data center.
11. Which is the valid CLI command to verify successful VSM connectivity to VMware vCenter?
 - a. **show running-config**
 - b. **show svcs domain**
 - c. **show svcs connectivity**
 - d. **show vcenter connectivity**
12. True or false? The command to verify whether the VEM agent is running on a VMware ESX or ESXi host is the **vem status** command to be executed on ESX CLI.
 - a. True
 - b. False
13. Which command do you use to validate and verify a Cisco Nexus 1000V port profile that already exists?
 - a. **show port-profile name <name>**
 - b. **show run port-profile name <name>**
 - c. **port-profile list <name>**
 - d. **show running-config | port-profile name <name>**

Foundation Topics

Evolution of Virtual Switching

The effects of virtualization today are felt across the IT spectrum, from the physical layer in a data center, up to the application layer. Particularly, server virtualization has blurred the traditional demarcations between networks, security, servers, application technologies and their respective administration responsibilities. Furthermore, server virtualization has made the enforcement of security policies, rules, and networking fuzzy and hard to pinpoint because a server workload or application can now run from anywhere, any server, and any port. It is dynamic, and that poses a significant challenge in the data center today.

Before Server Virtualization

Traditionally, you were used to a single application running its own operating system on an individual piece of server hardware. That server consisted of network interface cards (NICs) and host bus adapters (HBAs); their ports would connect to redundant network access layer and Fibre Channel switches. Network control security policies, quality of service (QoS), and so on were created at this access layer switch port level, related to that application's usage. These policy enforcements were static and rarely needed modifications.

If that server required maintenance or upgrades and needed to be offline, but the application needed to be available, redundancy was built with a cold or hot standby server or host-based clustering, at a cost. Figure 4-1 shows typical server connectivity in the physical world, before server virtualization.

Key
Topic

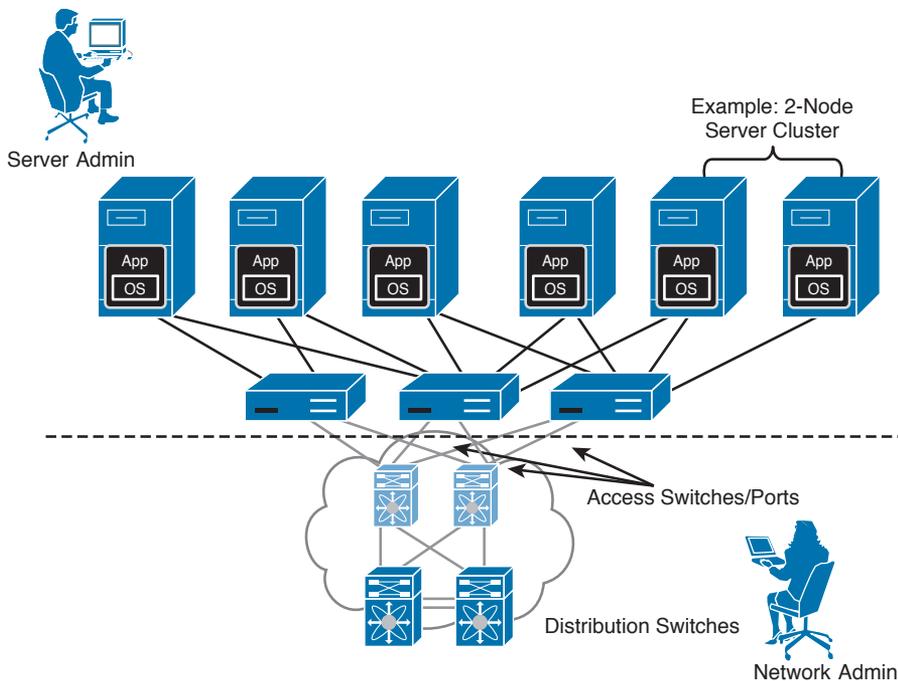


Figure 4-1 Server Connectivity Before Server Virtualization

With the advent and adoption of server virtualization technologies, multiple workloads or applications with their own guest operating systems can be hosted on a single server. This increases server utilization, improves efficiency, and keeps server hardware proliferation and costs under control in data centers. All that is achieved without compromising enterprise-class service levels. Features such as VMware vMotion, high availability (HA), and dynamic resource scheduling can move a workload or application running on one host to another in minutes, while in operation. This has essentially driven the need to enforce network policies, security controls, QoS, and the like at the virtual machine (VM) level rather than at the physical host level.

Key Topic

The following are some challenges with VMs and their hosts' physical access layer:

- Enforcing policies at the VM vNIC level
- Making those policies “sticky” and allowing them to follow the VM-in-motion
- Ensuring that a VM that moves to another host is not affected by that host's physical network switch port configuration
- Knowing that most configurations on a physical switch port that the ESX or ESXi host connects to will affect all VMs
- Viewing and controlling the vNIC at VM level from an access switch
- Redirecting traffic to network services, such as load balancers, firewalls, and so on, based on VM vNIC (service chaining)

Given each of these VM workloads or applications, a physical network adapter is not practical and seldom possible. Therefore, VMware developed the vSwitch to help overcome this limitation. Providing each VM one or more virtual network connections (vNIC), they connect to a vSwitch(s), providing access to the physical network (through its uplinks). Figure 4-2 shows typical host server connectivity with server virtualization.

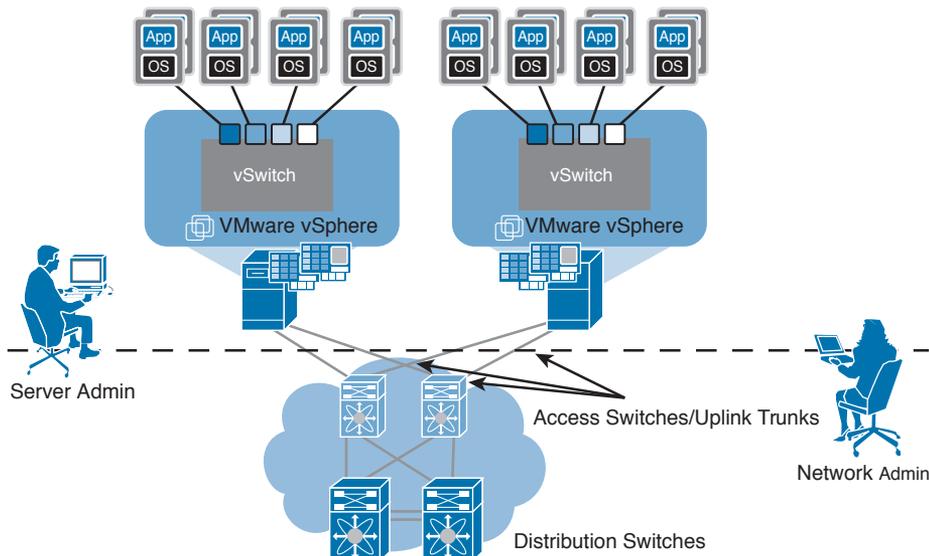


Figure 4-2 *Server Virtualization Effect*

Key Topic

Server Virtualization with Static VMware vSwitch

Server virtualization technologies like VMware ESX and ESXi enable you to run multiple workloads or applications as full virtual machines on physical servers; therefore, a traditional 1:1 relationship between a server and application is seen less in today's data centers. This highlights the requirement to have Layer 2 capabilities closely embedded into the ESX hypervisor to switch packets between VMs and the external world. Therefore, Layer 2 software switches are embedded within the ESX hypervisor to switch packets between VMs and the outside world.

A single host consisting of multiple VMs would utilize the same Ethernet uplinks (VMNICs) to access the network. So how would you differentiate between VM1 and VM2 traffic? This is currently achieved through IEEE 802.1Q VLANs. Therefore, the Ethernet uplinks would be defined as “trunks” for server connectivity into the external network, allowing multiple VLANs to be carried on a single wire, as shown earlier in Figure 4-2.

Interestingly, now server virtualization administrators suddenly owned and managed the virtual network configurations through, for example, VMware vCenter. Figure 4-3 shows a graphical view of standard static virtual switches being utilized on a couple of physical hosts running VMs.

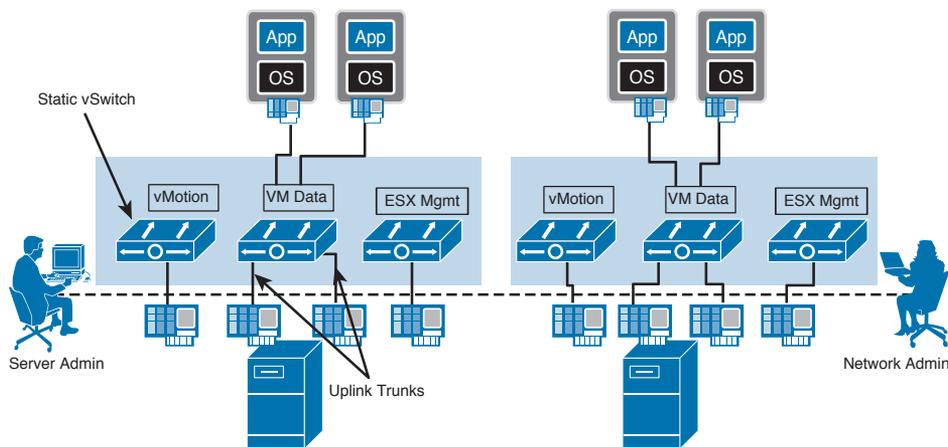


Figure 4-3 VMware Static vSwitch Configuration at Host Level

Key Topic

Here are some challenges with standard VMware vSwitch:

- When troubleshooting, the network administration team does not have visibility into the virtual switch layer. The internal vSwitch hides all VMs from the physical network.
- A compromised VM can be traced only as far as its physical switch port to which its host is connected.
- Shutting down a physical network switch port to which an ESX or ESXi host is connected can have severe consequences on production environments, because it would affect all VMs running on that server, and not just the compromised VM.
- In a VMware ESX or ESXi infrastructure, to avail features such as network vMotion and DRS (within a VMware cluster), all vSwitches must have the same port groups available.

Creating a port group on a vSwitch does not automatically propagate to all vSwitches; it needs to be manually created on each concerned ESX or ESXi host.

- As the virtual access layer extends into the ESX or ESXi host (that is, the network closest to the VM), the boundary of server versus network administration responsibilities starts to blur, requiring both functions to up-skill on each other's technology, which can cause some resistance between these functions. Table 4-2 compares the visibilities of physical and virtual network components that contribute to this resistance.

Table 4-2 Network Visibility and Configuration Comparison

Features	Physical Network	Virtual Network
Network visibility	Individual server	Physical server
Port configuration	Individual server	Physical server
Network configuration	Network administrator	VM and network administrator
Security policies	Individual server	Physical server

- As virtualized data centers evolved, more complex virtual network access layer requirements arose, such as spanning multiple disaster recovery sites and stretched active/active data centers. These vSwitches have limited functionalities.
- Many security- and compliance-related issues stem from the fact that traffic between VMs on the same host does not leave the server to run over the physical network. This makes it difficult for networking teams to monitor or manage this traffic. The lack of visibility means the network firewalls, QoS, access control lists (ACLs), and intrusion detection system/intrusion prevention system (IDS/IPS) cannot see this and differentiate individual data transfer activities over the physical network.

Key Topic

Virtual Network Components

As explained earlier, server virtualization extends the network access layer into the VMware ESX or ESXi hypervisor software. To provide virtual networking functions to virtual machines, multiple virtual network components work in union to deliver them. Figure 4-4 shows the virtual network components that work together to deliver networking functionality.

- **Physical NIC:** A VMware ESX or ESXi host requires physical NICs to connect to the external network. These physical NICs will be used as Ethernet uplinks, trunked. Typically, multiple NICs are defined as Ethernet uplinks; they are also known as VMNICs.
- **Port groups:** A subset of ports defined on a vSwitch for connectivity. Virtual machines connect to a virtual Ethernet switch (vSwitch) via port groups; they correspond to patch cords.
- **Virtual Ethernet switch (vSwitch):** Similar to a physical Layer 2 access switch, a vSwitch maintains a table of all connected VMs, similar to a MAC:Port address table and uses it to forward Layer 2 frames. This is a switch and not a hub, whereas it avoids unnecessary deliveries. This vSwitch can be connected to an external network via a physical NIC (VMNICs).

- **vNIC (virtual network interface card):** The virtual equivalent of a physical network interface card or port. It is used by virtual machines to essentially communicate over Layer 2 or Layer 3 protocols.
- **Virtual network:** A network defined in software, where multiple virtual machines share the same system resource.
- **Physical network:** Defined on physical switches that are used to connect VMware ESX or ESXi hosts to physical networks, the same as in regular LANs.

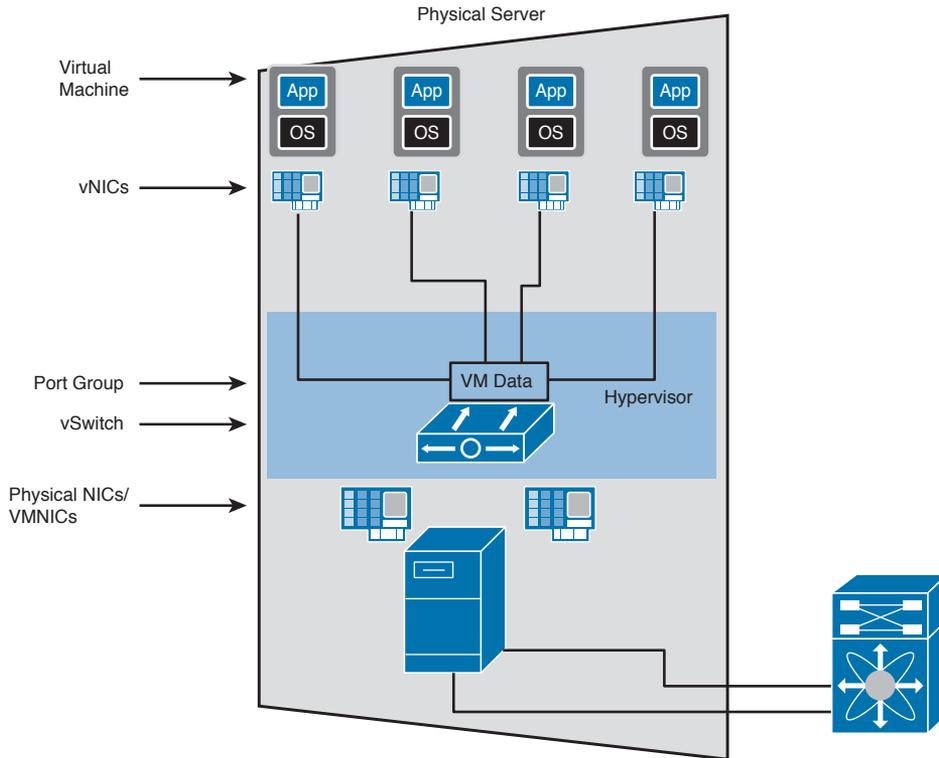


Figure 4-4 *Virtual Network Components*

Virtual Access Layer

In a modern virtualized data center with hundreds and thousands of VMware ESX or ESXi hosts running a large number of virtual machines, you would have multiples of these virtual network components that form a wide virtual access layer. That makes two significant access layers that need to be managed. While the server or virtualization administrator is managing the virtual access layer, your network administrator will manage the physical access layer. Figure 4-5 shows virtual and physical access layers in a data center with multiple VMware ESX or ESXi hosts.

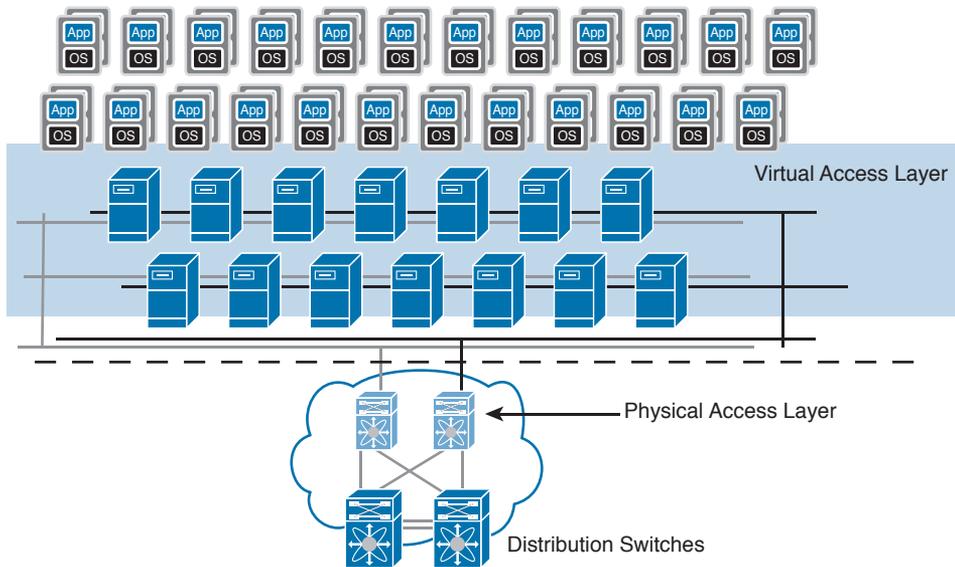


Figure 4-5 *The Virtual and Physical Access Layer*

Due to differences in mode of operation and challenges that the virtual access layer presents, certain tasks that were carried out by network administrators traditionally can move to administrators of the VM environment within a modern data center. Bear in mind, this can have an impact on SLAs and the responsibilities within your IT department.

While Cisco's Nexus 1000V series switches address the challenges of virtual access layers, further explained in this section, Cisco has revolutionized the software defined data center network architecture by introducing Cisco Application Centric Infrastructure (Cisco ACI) that helps to bring out further operational efficiencies. Please refer to Chapters 12, 13, and 14 for some information.

Standard VMware vSwitch Overview

A virtual switch (vSwitch) is a software-based implementation of a physical Layer 2 access switch that appears as a virtual construct. On a host, it enables networking between VMs and the external LAN. It brings together VMNICs and vNICs, facilitating internal and external communications, thus allowing the following:

- VM communications within and between ESX hosts
 - ESX Management (Service Console) communication
 - Port assigned to a VLAN
- VMKernel for vMotion, Fault Tolerance
 - Port assigned to a VLAN
- VMs assigned to port groups
 - Port groups assigned to a VLAN

- Uplinks
 - Used for external network connectivity
 - A VMNIC being associated with a vSwitch only (cannot be shared between vSwitches)

Because it's a software construct, you can create many ports. But note that there are configuration maximums defined by VMware that introduce an upper ceiling limit. Figure 4-6 shows an overview of a VMware ESX or ESXi Virtual Standard Switch (VSS, or vSwitch).

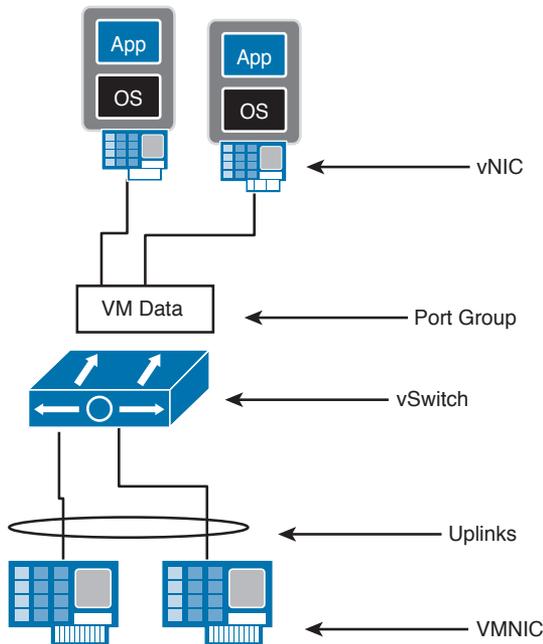


Figure 4-6 *Virtual Standard Switch Overview*

NOTE You can have multiple vSwitches in a VMware ESX or ESXi host, but the total number of ports per host is 4096. Furthermore, there are limits on the maximum active ports per host: 1016. These are known limits at the time of writing this chapter, defined by VMware. It is recommended that you consult the VMware-defined maximums for the particular VMware version being implemented before designing virtual infrastructure solutions.

Standard VMware vSwitch Operations

A single VMware ESX or ESXi host can consist of multiple vSwitches, such as to separate management and data traffic at the Layer 2 software switch level. As a general rule, there is no means for communication between vSwitches. As mentioned earlier, a vSwitch is a Layer 2 device, which means that only switching of traffic between VMs on the same host or Layer 2 forwarding over the uplinks takes place. These vSwitches cannot perform Layer 3 routing—that is, knowing and communicating with different IP networks. Hence, traffic that

does not belong to VMs on the host is forwarded to the uplink port. Figure 4-7 shows some of the operations graphically.

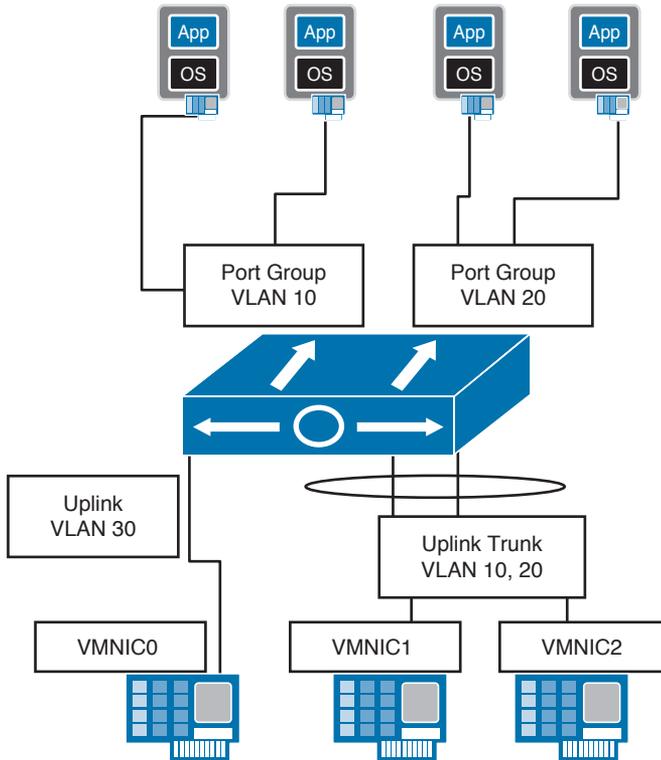


Figure 4-7 *Virtual Standard Switch Operations*

The vSwitch supports the following operations:

- Trunking functionality
- Port channels
- Cisco Discovery Protocol (CDP) for discovering and responding to neighboring network devices
- No participation in Spanning Tree Protocol (STP), Dynamic Trunking Protocol (DTP), or Port Aggregation Protocol (PAgP)
- Virtual guest tagging (reserved VLAN 4095)
 - Tagged traffic passed to guest operating system
- Outbound load balancing only
 - vSwitch port-based
 - Source MAC-based
 - IP Hash-based

Standard VMware vSwitch Configuration

Standard vSwitches on VMware are configured on each VMware ESX or ESXi host via the VMware vCenter Server or by directly connecting to the host using the VMware vSphere client utility. A single ESX or ESXi host can consist of multiple vSwitches, depending on your intended configuration and available physical NICs (VMNICs) available to be designated as uplinks. Figure 4-8 shows two standard vSwitches configured at the host level, via VMware vCenter, and their logical representation, showing the physical and virtual components.

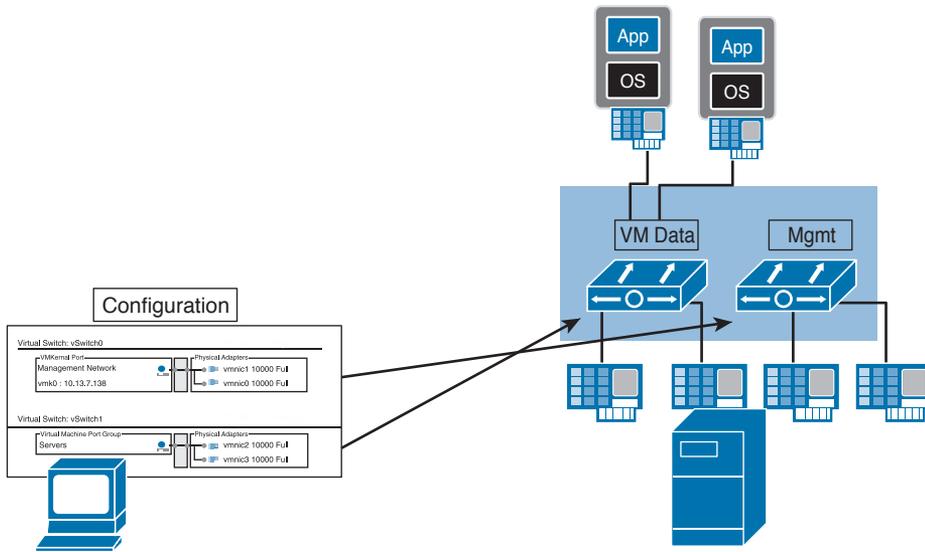


Figure 4-8 Virtual Standard Switch Configuration

NOTE A VMware ESX or ESXi host can have a total of 1016 active ports and as many as 32 VMNICs per vSwitch. Refer to the VMware configuration maximums document corresponding to the VMware version for current information.

A VMware standard vSwitch can consist of different port types on a VMware ESX or ESXi host. Figure 4-9 shows the different port types available on VMware ESX/ESXi hosts.

- VM port group types are the most common, used by VMs for their data access.
- VMKernel port types are used for advanced functions, such as vMotion and NIFS/CIFS/iSCSI storage connectivity and for management communication on ESXi hosts (for example, for console CLI access and vCenter server communication).

These port types usually belong to different logical networks, which allows coexistence of multiple logical networks (VLANs) within the same vSwitch. Or you can segregate networks into different vSwitches. It depends on your intended configuration and circumstances.

Figure 4-10 graphically depicts single or multinetworks on vSwitches.

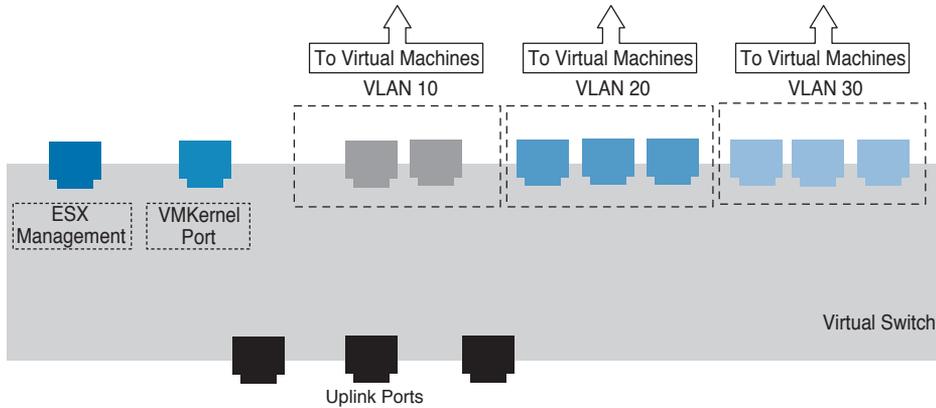


Figure 4-9 *Virtual Standard Switch Port Types*

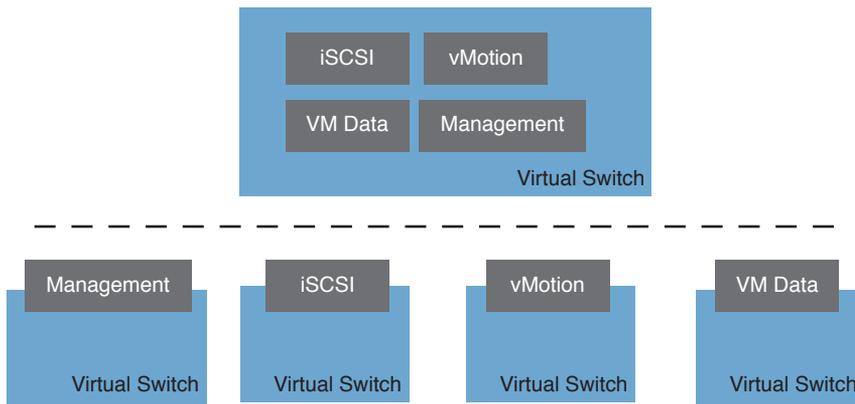


Figure 4-10 *Virtual Standard Switch Single—Multinetwork Examples*

Key Topic

VMware vDS Overview

The VMware Virtual Distributed Switch (vDS) was introduced with VMware vSphere version 4. As opposed to managing static vSwitches on every VMware ESX or ESXi host, the vDS brings the capability to configure and manage a collection of vSwitches in a VMware ESX or ESXi host or multiple hosts within a VMware Cluster. This can be achieved from a central location, via the VMware vCenter Server. Figure 4-11 shows a comparison of a VMware vDS versus standard vSwitches architecture.

NOTE The VMware vDS feature is licensed and cannot be managed on individual hosts; it requires a VMware vCenter Server.

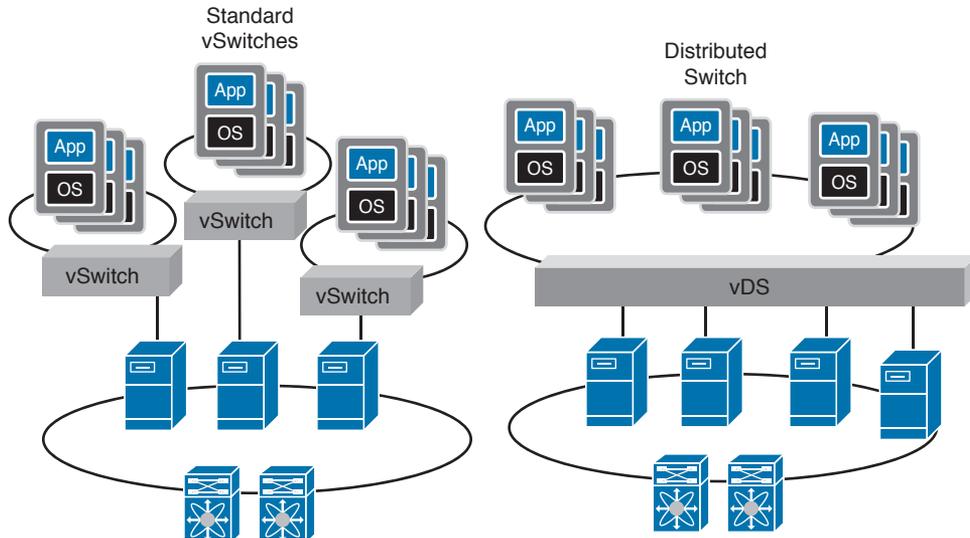
Key
Topic

Figure 4-11 VMware Virtual Distributed Switch (vDS) Overview

The VMware vDS was an important enhancement to virtual networking or the virtual access layer in data centers. It simplified management while introducing some enhancements, such as private VLAN (PVLAN) support, inbound traffic rate limiting, and tracking VM port state with migrations, to name a few. You will read some explanations of these features in the next sections in this chapter.

The introduction of vDS and enhancement with VMware vSphere version 4 also paved the way to a VMware vNetwork third-party vSwitch API that, as a matter of fact, was used by the Cisco Nexus 1000V Series switch to develop a fully VMware-compatible implementation of a virtual distributed switch. Figure 4-12 shows you an overview of the VMware vDS topology with multiple ESXi hosts.

VMware vDS Configuration

Management of a vDS differs from management of a vSwitch. A vDS is managed from a different management panel. This panel can be opened from any host in that cluster (under network configurations for that host) using the VMware vCenter server, and changes made are reflected on all hosts (within a VMware Cluster) rather than configuring each vSwitch on every host manually, as practiced with standard vSwitches. Figure 4-13 shows the two management panels from a host network configuration window.

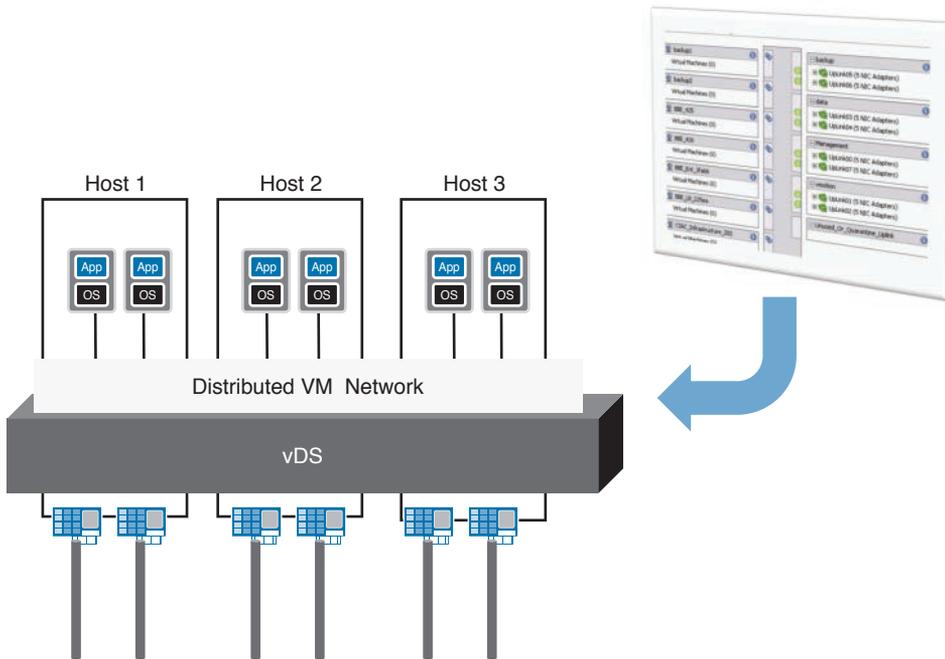


Figure 4-12 VMware Virtual Distributed Switch (vDS) Topology

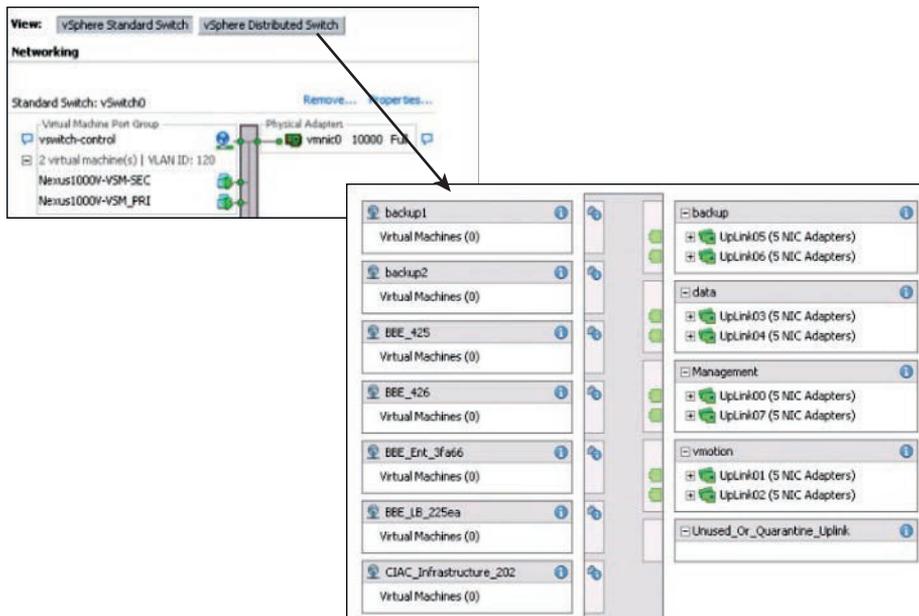


Figure 4-13 VMware vDS and vSwitch Management Panels

VMware vDS Enhancements

Apart from the procedural and operational benefits, several feature advancements were introduced with VMware vDS. The following are some of these enhancements, such as PVLANS support, Rx rate limiting, and Port state migration:

Key Topic

- **Port state migration:** When network vMotion is executed on a VM, its network state in terms of counters or port statistics is now tracked or is stateful because the VM moves from host to host on a vDS. So regardless of the location of the VM, its network interface on the vDS is consistent, which helps troubleshooting and network monitoring capabilities immensely.
- **Rx rate limiting:** The vDS builds upon the VM-to-network (Egress) traffic-shaping capabilities by enabling Ingress (from network to VM) as well. Now traffic-shaping policies can be applied on the port group definitions, and these definitions will be applied no matter which host the VM will reside on. Traffic shaping becomes a useful option when you want to prevent oversubscription to and from a VM or other traffic, when faced with limited resources. These policies are defined by average, peak bandwidth, or burst size.
- **PVLAN support:** You must know the concept of VLANs by now, which includes a broadcast domain, a network segment, and implementation over Layer 2 principles. PVLANS enable users to restrict communication between VMs on the same VLAN or network segment. This feature helps reduce the need for a number of subnets, depending upon your network configurations. Within a subnet, PVLANS are implemented on the vDS with allocations made to the following:
 - **Promiscuous PVLAN:** VMs on the promiscuous VLAN can communicate with all VMs.
 - **Community PVLAN:** VMs can communicate among themselves and with VMs on the promiscuous PVLAN.
 - **Isolated PVLAN:** VMs can communicate only with the VMs on promiscuous PVLANS.

NOTE Adjacent physical network switches must also support PVLANS and should be configured to support the vDS allocations.

VMware vSwitch and vDS

The VMware standard vSwitch and vDS can be used simultaneously on the same VMware ESX or ESXi host and managed by the same VMware vCenter Server, based on your design requirements and circumstances. You can also move the ESX Management and VMKernel ports assigned by default to vSwitch0 to your vDS (dvSwitch0), or you can opt to maintain them on the vSwitch0.

For example, you might want to separate ESX management traffic away from VM data traffic and opt to configure a vSwitch, assign separate uplinks (VMNICs), and connect a VMKernel port group to that vSwitch while you utilize the vDS for all VM data traffic with separate uplinks (VMNICs) assignments. Figure 4-14 shows a graphical view of vDS and vSwitch working in tandem.

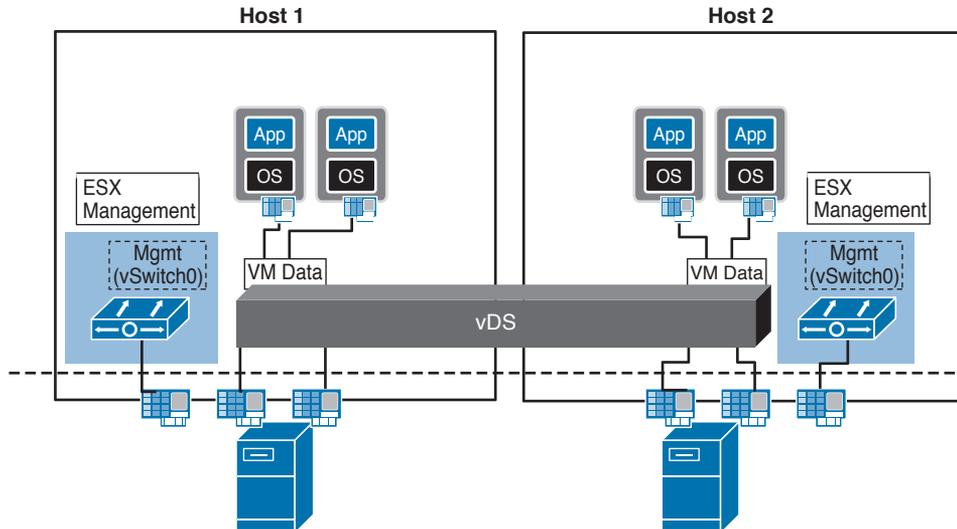


Figure 4-14 VMware vDS and vSwitch on the Same Host

The following are some of the advantages and enhancements of VMware vDS:

- The VMware vDS simplifies and unifies the virtual access layer management function by allowing central management and access.
- It moves away from the VMware ESX or ESXi host-level management requirement with standard vSwitches, into a VMware Cluster-level network management and configuration capability.
- With VMware vDS, policies now have the capability to follow the VM as it traverses hosts within a VMware Cluster, enhancing security; it is less prone to errors and has better debugging and troubleshooting capabilities.
- It forms the foundation for network resource pools and can determine the bandwidth that different network traffic types are given on a VMware virtual distributed switch (vDS).

Cisco Nexus 1000V Virtual Networking Solution

The Cisco virtualized network access layer solution for VM environments uses a technology that Cisco and VMware developed jointly. Primarily, the network access layer is moved further down into the virtual environment to overcome the challenges with the virtual access layers, improving its visibility, management, and enhanced features at the VM level. Figure 4-15 shows the evolution of the VMware vNetwork, virtual access layer.

The Cisco Nexus 1000V solution was introduced with VMware vSphere version 4. It consists of an Ethernet module and a supervisory (management) module. This management module can be deployed on specialized hardware or as a software appliance, depending on your data center design requirements and circumstances. Both these deployment models offer the same improvements and enhancements in terms of VM visibility, policy-based VM connectivity, policy mobility, and a highly available nondisruptive operational model. Figure 4-16 shows you an overview of the Cisco Nexus 1000V distributed virtual software switch architecture.

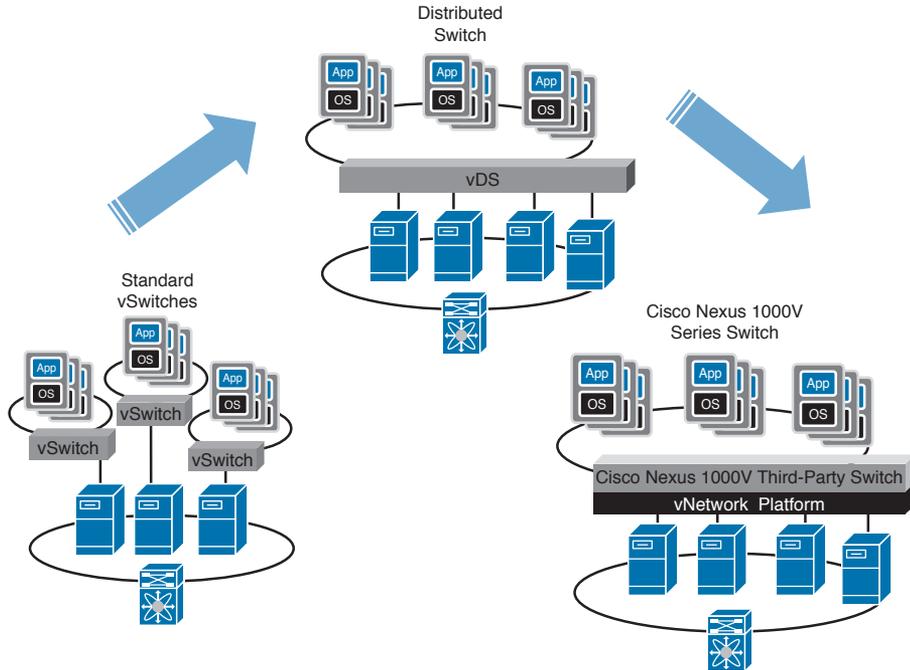


Figure 4-15 Evolution of VMware vNetwork

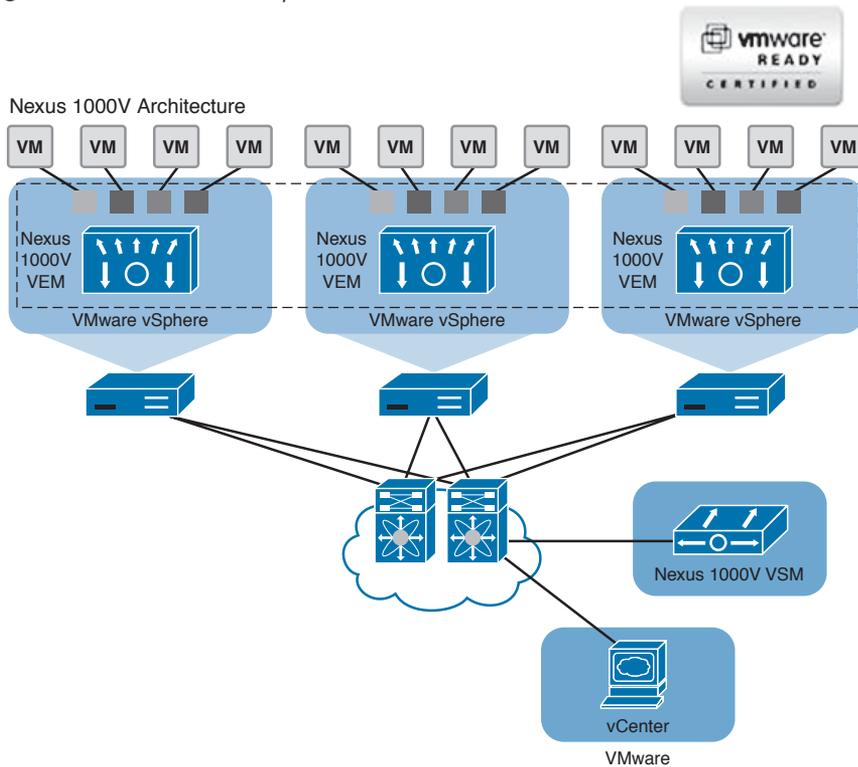


Figure 4-16 Cisco Nexus 1000V Distributed Virtual Software Switch Architecture

NOTE The release of VMware vSphere version 4 and the vNetwork vSwitch APIs that were made available by VMware enabled third parties to develop vSwitch implementations. The Cisco Nexus 1000V is the first third-party vDS implementation that is fully supported by VMware.

Key Topic

Cisco Nexus 1000V System Overview

The Cisco Nexus 1000V is a software-based solution that works with any upstream physical switching system to provide standard networking functionality and controls to the virtual machine environment. For your clarity and understanding, Figure 4-17 shows a comparison of a Cisco Nexus 1000V solution to a physical switch.

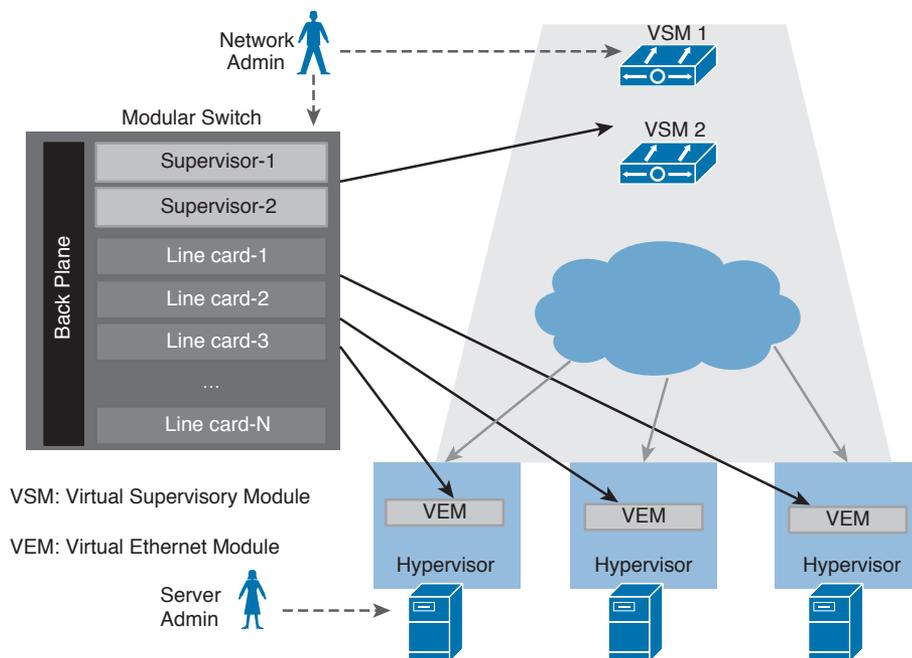


Figure 4-17 Cisco Nexus 1000V Analogy to a Modular Physical Switch

NOTE With the introduction of VMware vSphere version 4, you had the option of three virtual networking solutions: VMware vSwitch, vDS, and the Cisco Nexus 1000V Series switch. At the time of writing this chapter, VMware vSphere 6.0 was at general availability. The VMware NSX network virtualization solution was also an option.

Cisco Nexus 1000V Salient Features and Benefits

The Cisco Nexus 1000V Series switch bypasses the standard vSwitch by using a Cisco software switch. This model provides a single point of configuration for the networking

environment of multiple ESX or ESXi hosts. Some of the salient features of the Cisco Nexus 1000V Series switch are listed here:

Key Topic

- Policy-based VM connectivity:** The network administrators, rather than the VM administrators, define these policies, allowing network administrators control of VM-level virtual access layer connectivity. These policies are defined in the form of a port profile containing many characteristics and pushed to VMware vCenter automatically; the VM administrators can then utilize these policies by assigning them to the respective VMs. (Port profiles are discussed later in this chapter.) Figure 4-18 shows you how policy-based VM connectivity is assigned.

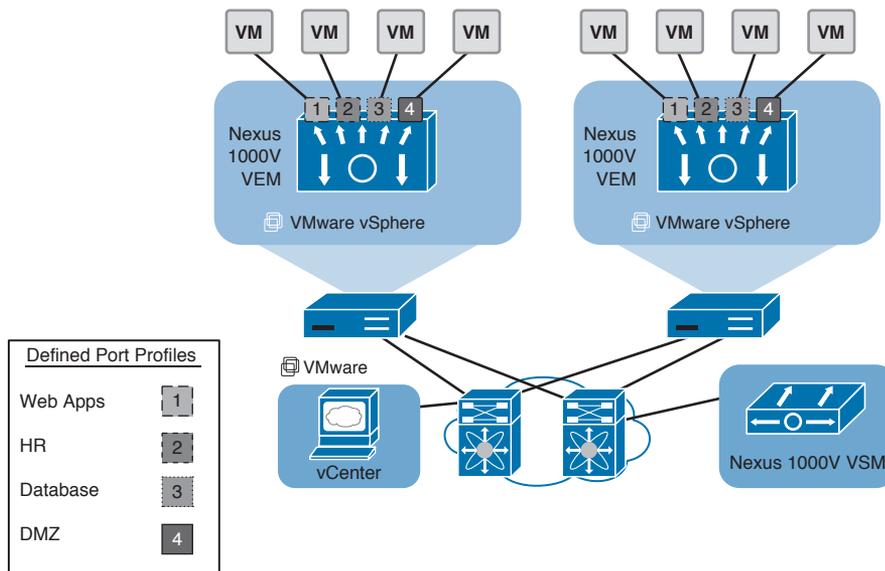


Figure 4-18 Cisco Nexus 1000V Policy-Based VM Connectivity

- Mobility of networking and security properties:** All policies defined in Cisco Nexus 1000V and applied to VMs fully support VMware mobility features, such as VMware vMotion and High Availability (HA). These policies remain attached to the VM, even when the VM moves from host to host. Figure 4-19 shows the mobility of policies along with movement of VMs.
- Nondisruptive operational model:** The Cisco Nexus 1000V Series switch can be introduced into existing virtual environments nondisruptively by utilizing best practices to migrate from VMware standard virtual networking to the Cisco Nexus 1000V Series switch. Furthermore, this switch unifies the management plane with the rest of the IP network by using the same familiar Cisco NX-OS operating system, commands, and techniques that are available on other Cisco network platforms. Figure 4-20 summarizes the Cisco Nexus 1000V Series switch feature set; for more detailed information, see the links in the “Reference List” section of this chapter.

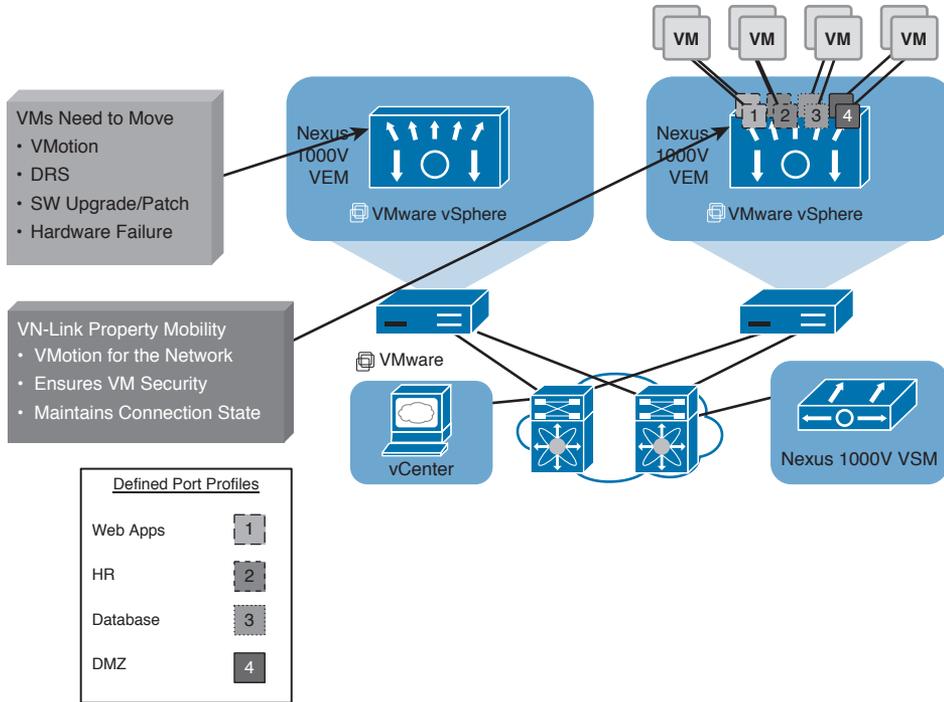


Figure 4-19 Cisco Nexus 1000V Policy Mobility

Switching	<ul style="list-style-type: none"> • L2 switching, 802.1Q tagging, VLAN, rate limiting (TX), VXLAN • IGMP snooping, QoS marking (COS and DSCP), class-based WFQ
Security	<ul style="list-style-type: none"> • Policy Mobility, Private VLANs w/local PVLAN Enforcement • Access control lists, port security, Cisco TrustSec support • Dynamic ARP inspection, IP Source Guard, DHCP snooping
Network Services	<ul style="list-style-type: none"> • Virtual Services Datapath (vPath) support for traffic steering and fast-path off-load [leveraged by Virtual Security Gateway (VSG), vWAAS, ASA1000V]
Provisioning	<ul style="list-style-type: none"> • Port profiles, integration with vC, vCD, SCVMM*, BMC CLM • Optimised NIC teaming with virtual port channel – host mode
Visibility	<ul style="list-style-type: none"> • VM migration tracking, VC plugin, NetFlow v.9 w/NDE, CDP v.2 • VM-level interface statistics, vTracker • SPAN and ERSPAN (policy-based)
Management	<ul style="list-style-type: none"> • Virtual Centre VM provisioning, vCenter plugin, Cisco LMS, DCNM • Cisco CLI, Radius, TACACs, Syslog, SNMP (v1, 2, 3) • Hitless upgrade, SW Installer

Figure 4-20 Cisco Nexus 1000V Series Switch Feature Set Summarized

Some of the benefits at the server and network levels are summarized here:

- Server Level
 - VM management preserved
 - Reduced operational workload
 - VM-level visibility
 - Compatible with Network vMotion, Storage vMotion, DRS, HA, and FT (Fault Tolerance)
- Network Level
 - Unified management and operations
 - Improved operational security (L2, L3, L4 access lists, port security)*
 - Enhanced network features (VLANs, PVLANS 802.1q, LACP, vPC host mode, QoS)*
 - Policy persistence
 - VM-level visibility
 - SPAN and ERSPAN, NetFlow*

(*Note that some of these additional functionalities are not available on basic vSwitches and vDS; you may require advanced editions.)



Cisco Nexus 1000V Series Virtual Switch Architecture

As mentioned in the overview section, the Cisco Nexus 1000V Series switch has two major components: the VEM, which runs inside the hypervisor, and the VSM, which manages these VEMs.

Cisco Nexus 1000V Virtual Supervisory Module

The VSM is the virtual equivalent of the supervisor module that exists in physical Cisco modular devices running Cisco Nexus Operating System (NX-OS), such as Nexus 7000 Series switches. The VSM is running NX-OS; it manages and controls all other components (VEMs) that are part of the Cisco Nexus 1000V Series switch. All VEMs, equivalent to physical line cards, connect to the VSM and behave as a single virtual modular device. Furthermore, the VSM is responsible for communicating with vCenter and programming, managing the VEMs.

The VSM can be managed via the NX-OS CLI; the CLI has the same syntax and behavior as the CLI on other Cisco Nexus devices. The VSM can reside as a virtual appliance or be hosted on Cisco Nexus 1010/1100 Series virtual services appliances. Figure 4-21 shows the various deployment scenarios for the Cisco Nexus 1000V VSM.

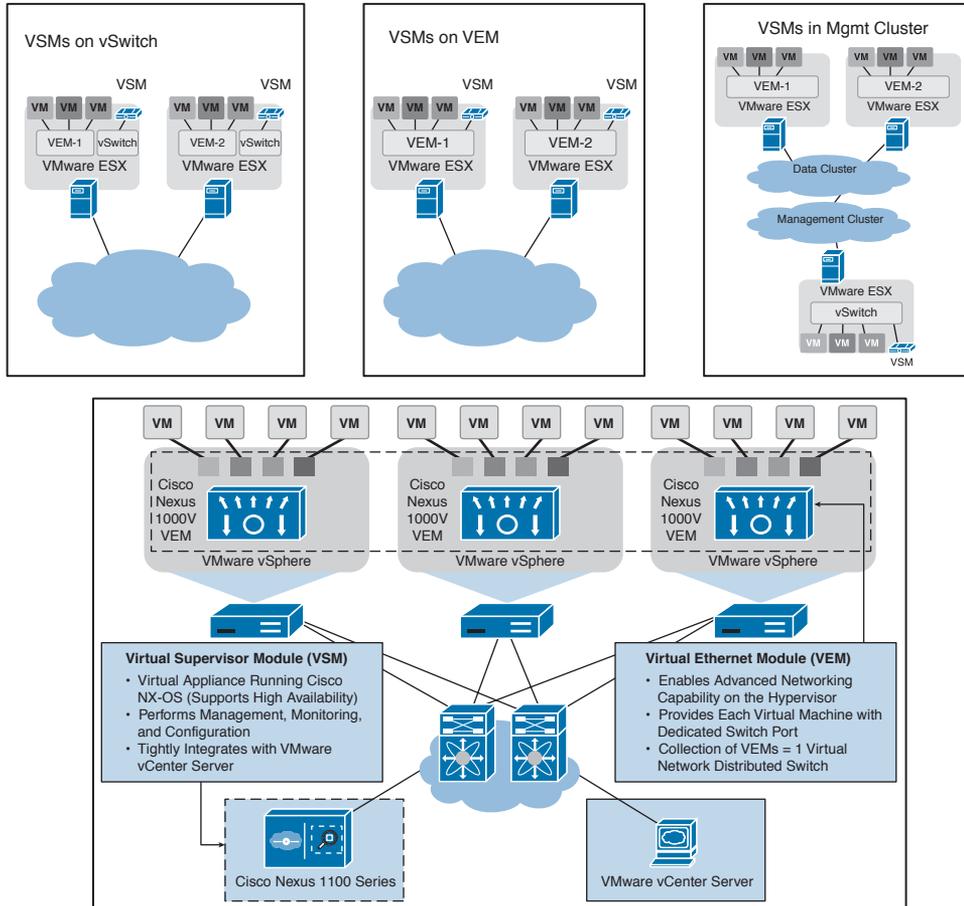


Figure 4-21 Cisco Nexus 1000V VSM Deployment Scenarios

Cisco Nexus 1000V Virtual Ethernet Module

The Virtual Ethernet Module (VEM) is the virtual equivalent of a line card that exists in a physical Cisco modular switch. The VEM resides on every VMware ESX or ESXi host; on the hypervisor. The VEM provides all the VMs their connectivity, through physical NICs on the host, and forms that virtual access layer. Multiple VEMs communicating with a VSM or multiple VSMs construct one logical switch.

VEMs on different hosts don't have a direct line of communication with each other; they require an external switch to link them together, such as a physical access layer switch of the VMware ESX or ESXi host. Figure 4-22 graphically represents the Cisco Nexus 1000V VEM.

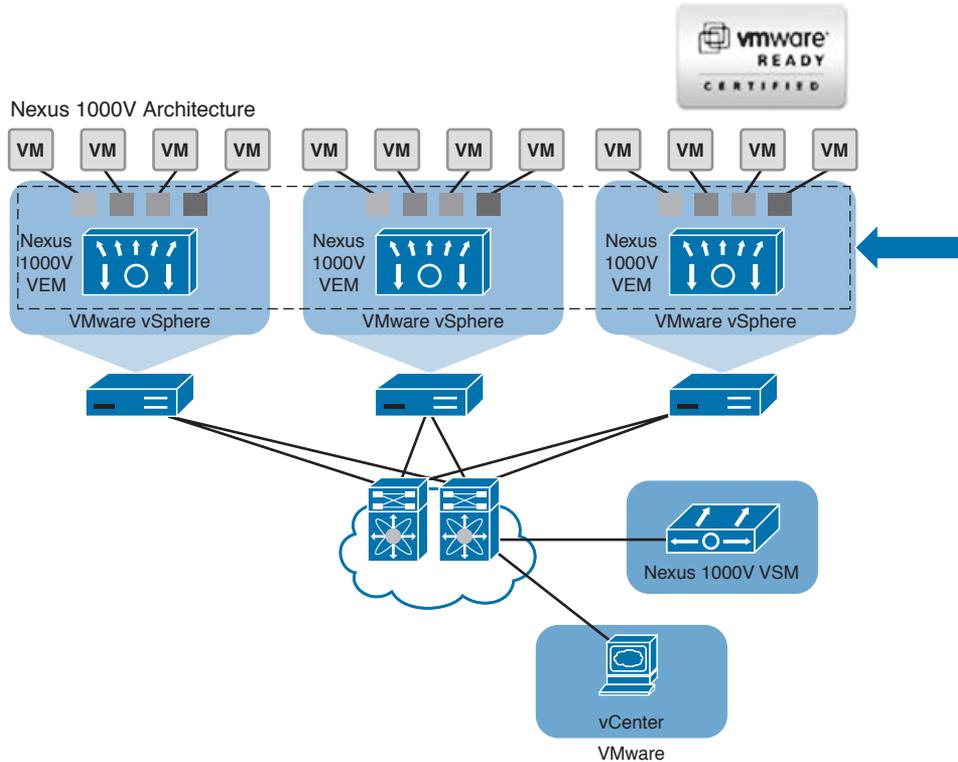


Figure 4-22 Cisco Nexus 1000V VEM Architecture

NOTE The VEM-to-VSM communication carries only control traffic. Traditionally, this communication took place over Layer 2 mode. With the later releases of Cisco Nexus 1000V Series switch, Layer 3 mode is valid and is the default option. Every VEM on each VMware ESX or ESXi host requires a VMKernel interface/NIC with an IP address. This VMK interface must be moved to the VEM. The management modes are briefly explained in the subsequent sections.

As described earlier, the Cisco Nexus 1000V can be logically viewed as a virtual modular switch chassis. This virtual chassis contains the Cisco Nexus 1000V Series components, such as redundant VSMS and installed VEMs. Figure 4-23 shows the **show module** command output from the primary VSM, where it displays the VSMS and the VEMs in the same manner as it would display supervisors and line cards on a Cisco Nexus 7000 Series switch.

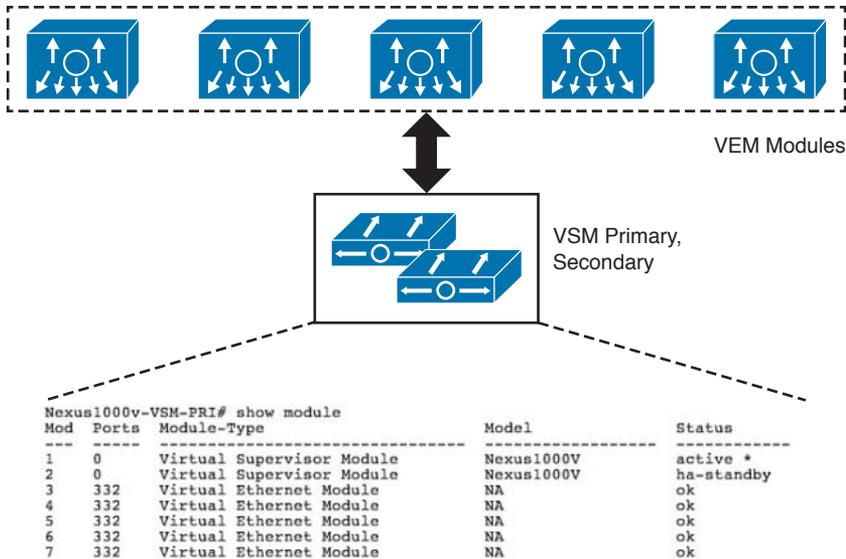


Figure 4-23 Cisco Nexus 1000V Virtual Chassis Architecture

Cisco Nexus 1000V Component Communication

The communication between VEMs and VSMs is facilitated by two separate virtual interfaces known as *control* and *packet* interfaces.

Key Topic

- **Control:** The control interface carries “heartbeat” traffic from VEMs to VSM; it’s communicated every 2 seconds, with a 6-second timeout to confirm its presence to the VSM. Furthermore, the control interface maintains synchronization between primary and secondary VSMs over this interface.
- **Packet:** The packet interface carries packets such as Cisco Discovery Protocol (CDP) or Internet Group Management Protocol (IGMP) control messages, from VEM to VSM.

NOTE You can use one or separate VLANs for the control and packet interfaces. Usually, two separate VLANs are used for these interfaces; therefore, Layer 2 connectivity between VEMs and VSMs is a prerequisite, unless Layer 3 mode is utilized.

Cisco Nexus 1000V Management Communication

The communication between the VSM and the VMware vCenter server is conducted through this management interface. The VSM uses the VMware virtual infrastructure methodology (VIM) application programming interface (API) over Secure Socket Layer (SSL) to communicate with VMware vCenter server. The connection can be manually set up on the VSM or established during the installation process. After communication between the VSM and VMware vCenter server is established, the Cisco Nexus 1000V is created in the vCenter server.

The management interface is known as the out-of-band management interface. The best practice is to have this interface and the VMware vCenter server ESX management interface on the same VLAN.

NOTE Cisco Nexus 1000V introduced Layer 3 mode for control, packet, and management traffic since version 1.2. At that time, Layer 3 connection mode was optional; at the time of writing this book, the default when installing the VSM is to establish Layer 3 connectivity instead of Layer 2. With Layer 3 mode, every VEM requires an IP address and all control, packet, and management traffic is carried over this network. Figure 4-24 shows the interaction during VSM installation and the control mode selection.

```

Configure the default domain name? (yes/no) [n]: y
  Default domain name : cisco.com
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]: y
  Type of ssh key you would like to generate (dsa/rsa) [rsa]:
  Number of rsa key bits <768-2048> [1024]:
Enable the http-server? (yes/no) [y]:
Configure the ntp server? (yes/no) [n]: y
  NTP server IPv4 address :
Uem feature level will be set to 4.2(1)SU1(5.1). Do you want to reconfigure? (yes/no) [n]:
Configure svcs domain parameters? (yes/no) [y]:
  Enter SUS Control mode (L2 / L3) : L3_

```

Figure 4-24 Cisco Nexus 1000V VSM Installation Layer 3 Mode

Key Topic

Cisco Nexus 1000V Port Profiles

Port profiles are the Cisco Nexus 1000V Series switch equivalent of port groups/dvPortgroups in VMware. When a port profile is created, you will see a corresponding port group being created in VMware vCenter server.

These port profiles are used to configure interfaces on the Cisco Nexus 1000V Series switch. A port profile can be assigned to multiple interfaces (physical or virtual). All changes to a port profile are automatically propagated across all interfaces.

In the VMware vCenter server, port profiles are represented as port groups/dvPortgroups. Both physical and virtual interfaces are assigned in VMware vCenter server to their respective port groups. These port groups perform the following functions:

- Define a port configuration by policy.
- Apply a single policy across all concerned ports.
- Support both virtual and physical Ethernet ports.

When a VMware ESX or ESXi host port (VMNIC) is assigned to the Cisco Nexus 1000V Series switch, an uplink port group/dvPortgroup is assigned to it and its settings are applied to that VMNIC. Similarly, when a NIC is added to a VM (vNIC), an available VM port

group/dvPortgroup is assigned, and those network settings associated with that port profile are inherited.

NOTE When implementing the Cisco Nexus 1000V Series switch, manual configuration of port groups/dvPortgroups via VMware vCenter server is not recommended in general, because it overrides the Cisco Nexus 1000V port profile. Instead, all port profile creation is performed via the VSM.

Types of Port Profiles in Cisco Nexus 1000V

Key Topic

There are two main types of port profile constructs in the Cisco Nexus 1000V Series switch.

- **Type Ethernet:** The Ethernet type port profiles are usually uplink port profiles, normally associated with VMNICs (physical interfaces on ESX or ESXi hosts). All “system VLANs” must be allowed on this uplink.
- **Type vEthernet:** The vEthernet type port profiles are usually assigned to virtual interfaces on VMs (vNICs) and represented as port groups/dvPortgroups in the VMware vCenter server. These port profiles are typically a collection of NX-OS CLI commands, consisting of the port profile name, switch port mode, VLAN number, and so on.

NOTE System VLANs are Control, Packet, Management VLANs; they are classified as system VLANs in a Cisco Nexus 1000V Series switch implementation.

The Cisco Nexus 1000V Series software switch provides a model in which network administrators define network-related policies that a server administrator can consume as new VMs are provisioned. These port profiles automatically sync with VMware vCenter server and are available as port groups to be assigned by a server administrator.

Cisco Nexus 1000V Administrator View and Roles

This admin model is implemented on Cisco Nexus 1000V Series switches with the port profiles feature. With this port profile feature, server administrators don't need to create and maintain vSwitches or port group configurations on any of their VMware ESX or ESXi hosts.

Port profiles allow separation of the network and server administrator functions. Network administrators now have the ability to define port profiles with the familiar syntax as existing physical Cisco switches and ensure consistent policy enforcement both at physical and virtual machine level.

The server administrator can now focus on server administration tasks; when connectivity is required, the administrator can consume the policies defined through port profiles. Figure 4-25 shows the VMware vCenter view of choosing port groups, which are populated when port profiles are created on a Cisco Nexus 1000V Series switch and assigned to a virtual machine interface (vNIC).

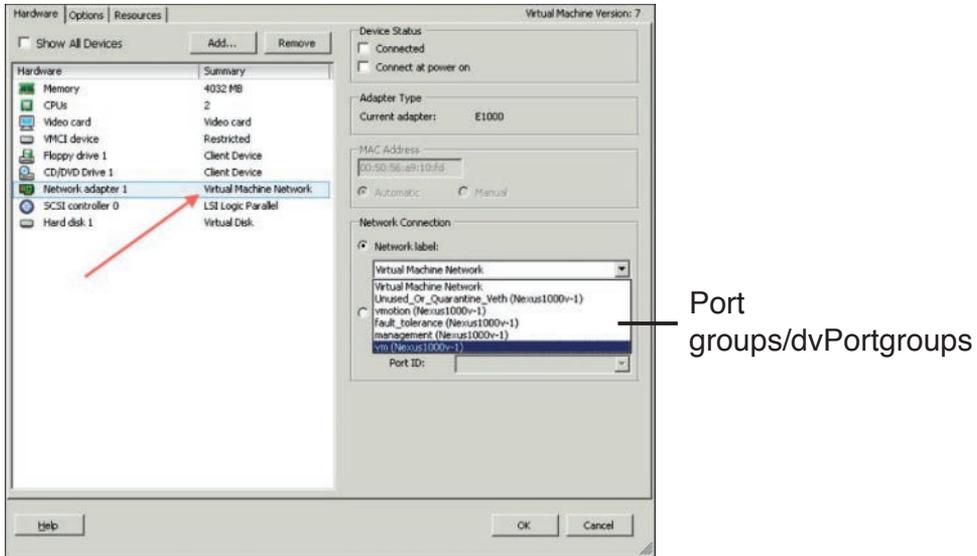


Figure 4-25 VMware vCenter Port Groups/dvPortgroups

Table 4-3 compares the typical administrative tasks before and after the Cisco Nexus 1000V Series switch or vDS availability. The network administrators’ participation facilitates more focus, control, and manageability around the virtual access layer.

Table 4-3 Administrator Tasks Before and After Cisco Nexus 1000V Series Switch

Tasks	VMware Administrator Before	VMware Administrator After	Network Administrator Before	Network Administrator After
vSwitch config	Per ESX host	Automated	—	Same as physical network
Port group config	Per ESX host	Automated	—	Policy-based
Add ESX host	vCenter-based	vCenter-based	—	—
NIC teaming config	Per vSwitch	Automated	—	Port channel optimized
Virtual machine creation	vCenter-based	vCenter-based	—	—
Security	—	Policy-based	—	ACL, PVLAN, port security, TrustSec
vm visibility	vCenter	VM-specific	—	VM-specific
Management	vCenter	vCenter-based	—	Cisco CLI, XML API, SNMP, DCNM

Cisco Nexus 1000V Verifying Initial Configuration

This section focuses on explaining the initial configuration and validation steps you need to understand when using the VMware ESX and Cisco Nexus 1000V Series switch and the CLI commands to validate connectivity to the VSM, VEM, and VMware vCenter.

Cisco Nexus 1000V VSM Installation Methods

You can use multiple methods to deploy and install your Cisco Nexus 1000V VSM; the preferred installation method is to use an open virtualization appliance (OVA) file. When you download the Cisco Nexus 1000V zipped package from software.cisco.com, it contains a number of folders that include VSM, VEM, XML-APIs, and the like, and a README.txt file.

In the VSM folder, under Install, you will find the required OVA file, which you can use to deploy the VSM directly on to your ESX or ESXi host.

From the .iso file, follow these steps:

- Step 1.** Create a VSM VM in vCenter.
- Step 2.** Configure VSM networking.
- Step 3.** Perform initial VSM setup via the VSM console.
- Step 4.** Install the VSM plug-in in vCenter.
- Step 5.** Configure the SVS connection in the VSM console.
- Step 6.** Add hosts to the virtual distributed switch in vCenter.

From the .ovf file:

- Step 1.** Use the wizard to deploy an OVF from template (Steps 1 and 2)
- Step 2.** All other steps are manual as above.

From the .ova file:

- Step 1.** Use the wizard to deploy an OVF from template (Steps 1–4).
- Step 2.** All other steps are identical and manual.

NOTE Open virtualization format (OVF) and open virtualization appliance (OVA) are similar. The main difference is that the OVF folder contains (metadata) structured files—.ovf, .mf, .vmdk/.vhd, and so on—whereas the OVA file is a single zipped file. Figure 4-26 shows the OVA file deployment process.

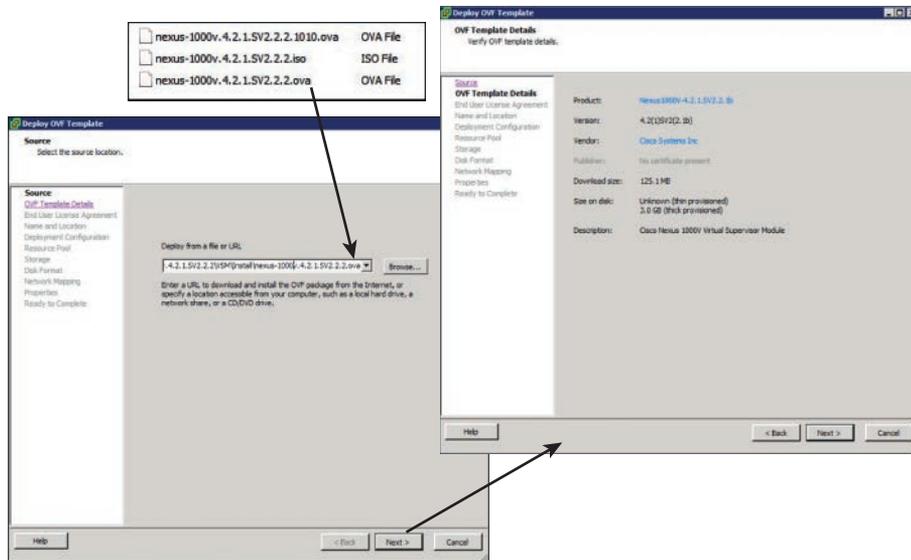


Figure 4-26 Cisco Nexus 1000V VSM Installation



Initial VSM Configuration Verification

When installing the Cisco Nexus 1000V Series VSM, as a VM or instantiated in the Cisco Nexus 1010/1100 appliance, the network administrator performs the initial configuration to provide the basic parameters for the Cisco Nexus 1000V Series switch. To verify this initial configuration and subsequent verifications to the configuration, use the **show running-config** command at the Cisco Nexus 1000V CLI. Figure 4-27 shows sample output of this command.

Verifying VMware vCenter Connectivity

During initial configuration, to establish a connection between the Cisco Nexus 1000V Series switch and the VMware vCenter, the network administrator must configure a Software Virtual Switch (SVS) connection. This is required for the Cisco Nexus 1000V Series switch to push configurations and policies such as port profiles to VMware vCenter server.

To verify whether the SVS connection is in place explicitly, use the **show svcs connections** command on the Cisco Nexus 1000V Series switch CLI. Figure 4-28 shows sample output of the **show svcs connections** command.

```

Nexus1000v-VSM-PRI# show running-config
!Command: show running-config
!Time: Sat Sep 20 16:00:58 2014

version 4.2(1)SV2(2.2)
svs switch edition essential

no feature telnet
feature lisp

username admin password 5 $!$enMjuth8k263InUe6p5X5NA9NM01Q0 role network-admin
banner motd #Nexus 1000v Switch#

ip domain-lookup
ip host Nexus1000v-VSM-PRI 10.113.97.202
switchname Nexus1000v-VSM-PRI
errdisable recovery cause failed-port-state
ip access-list Demo_Mob_Quarantine
10 deny lsmg any any
20 permit ip any any

vsm 3
 host id 8895c837-7794-e311-0001-0000000000af
vsm 4
 host id 8895c837-7794-e311-0001-0000000000af
vsm 5
 host id 8895c837-7794-e311-0001-0000000000bf
vsm 6
 host id 8895c837-7794-e311-0001-0000000000cf
vsm 7
 host id 8895c837-7794-e311-0001-0000000000df
snmp-server user admin network-admin auth md5 0x53c76bb82205f81aa2612f6b25e1ed priv 0x53c76bb8
2205f81aa2612f6b25e1ed localizedKey

vrf context management
 ip route 0.0.0.0/0 10.113.97.193
vlan 1,90-92,99,101-105,120-122,198-191,201-202,301-303,400-427,650,900-901
 name ESX_Management
vlan 91
 name UCS_Management
vlan 92
 name VMware_Management
vlan 99
 name vmotion
vlan 101
 name pord1
vlan 102
 name pord2
vlan 103
 name pord3
vlan 104
 name pord4
vlan 105
 name pord5
vlan 120
 name control
vlan 121
 name packet
vlan 122
 name data
vlan 192
 name VHG_Data
vlan 193
 name VHG_BA
vlan 201
 name Infra_Transit_Vlan
vlan 202
 name CIAC_Infrastr
vlan 301
 name VFX_Trans
vlan 302
 name CIAC
%

boot kickstart bootflash:/nexus-1000v-kickstart.4.2.1.SV2.2.2.bin sup-1
boot system bootflash:/nexus-1000v.4.2.1.SV2.2.2.bin sup-1
boot kickstart bootflash:/nexus-1000v-kickstart.4.2.1.SV2.2.2.bin sup-2
boot system bootflash:/nexus-1000v.4.2.1.SV2.2.2.bin sup-2

svs-domain
 domain id 1
 control vlan 1
 packet vlan 1
 sv mode L3 l3interface sgm0
 svs connection vcenter-vsalab
 protocol vmware-vim
 remote ip address 10.113.97.216 port 80
 vmware dvs uuid "4d 64 20 50 74 83 09 8e-40 91 13 1a 3b 3d fb 52"
 max-ports 8192
 connect
 vservice global type vsag
 top state-checks invalid-ack
 top state-checks seq-past-window
 no top state-checks window-variation
 no bypass sas-traffic
 vsm-policy-agent
 registration-ip 0.0.0.0
 shared-secret *****
 log-level

```

Figure 4-27 Cisco Nexus 1000V `show running-config` Command Output

```

Nexus1000v-VSM-PRI# show svcs connections
connection vcenter-vsalab:
 ip address:
 remote port: 80
 protocol: vmware-vim https
 certificate: default
 datacenter name: vsalab
 admin:
 max-ports: 8192
 DVS uuid: 4d 64 20 50 74 83 09 8e-40 91 13 1a 3b 3d fb 52
 config status: Enabled
 operational status: Connected
 sync status: Complete
 version: VMware vCenter Server 5.1.0 build-799731
 vc-uuid: EA62B6C7-0BB4-44B3-ABAF-11E79848C55F
 ssl-cert: self-signed or not authenticated

```

Figure 4-28 Cisco Nexus 1000V `show svcs connections` Command Output

You can perform further verifications of the connection between the Cisco Nexus 1000V Series switch and VMware vCenter server by using the `show svcs domain` command. It is important to understand that each Cisco Nexus 1000V Series switch uses a domain ID. All ESX or ESXi hosts with a VEM installed are members of a domain and listen to updates from that domain, which demarcates the virtual chassis where they reside. Figure 4-29 shows a sample output of the `show svcs domain` command.

NOTE In the sample output, see the Status field that confirms an operationally successful connection to VMware vCenter server and the use of Cisco-recommended Layer 3 mode for VSM-to-VEM communications, not Layer 2 mode.

```

Nexus1000v-VSM-PRI# show svcs domain
SVS domain config:
Domain id: 1
Control vlan: NA
Packet vlan: NA
L2/L3 Control mode: L3
L3 control interface: mgmt0
Status: Config push to VC successful.
Control type multicast: No

Note: Control VLAN and Packet VLAN are not used in L3 mode
Nexus1000v-VSM-PRI#

```

Figure 4-29 Cisco Nexus 1000V show svcs domain Command Output

Key Topic

Verifying Nexus 1000V Module Status

After an ESX or ESXi host has been added to the distributed switch and the VEM has been installed on it successfully, the VEM appears as a module on the VSM as part of the virtual chassis, similar to modules that are added to a physical module switch chassis.

Figure 4-30 shows a sample output of the **show module** command, demonstrating the primary and standby supervisor modules in slots 1 and 2, respectively, and the VMware ESX or ESXi hosts/modules that have been added to the Cisco Nexus 1000V Series switch instance starting from slot 3.

Slot 1 and 2 Reserved for VSM: Primary and Standby

```

Nexus1000v-VSM-PRI# show module
-----
Mod  Ports  Module-Type          Model          Status
-----
1    0      Virtual Supervisor   Nexus1000V     active *
2    0      Virtual Supervisor   Nexus1000V     ha-standby
3    332    Virtual Ethernet     NA             ok
4    332    Virtual Ethernet     NA             ok
5    332    Virtual Ethernet     NA             ok
6    332    Virtual Ethernet     NA             ok
7    332    Virtual Ethernet     NA             ok

Mod  Sw          Hw
-----
1    4.2(1)SV2(2.2)  0.0
2    4.2(1)SV2(2.2)  0.0
3    4.2(1)SV2(2.2)  VMware ESXi 5.1.0 Releasebuild-1065491 (3.1)
4    4.2(1)SV2(2.2)  VMware ESXi 5.1.0 Releasebuild-1065491 (3.1)
5    4.2(1)SV2(2.2)  VMware ESXi 5.1.0 Releasebuild-1065491 (3.1)
6    4.2(1)SV2(2.2)  VMware ESXi 5.1.0 Releasebuild-1065491 (3.1)
7    4.2(1)SV2(2.2)  VMware ESXi 5.1.0 Releasebuild-799733 (3.1)

Mod  Server-IP      Server-UUID          Server-Name
-----
1    10.113.97.202  NA                   NA
2    10.113.97.202  NA                   NA
3    10.113.97.201  8895c837-7794-e311-0001-0000000000af
4    10.113.97.198  8895c837-7794-e311-0001-0000000000ef
5    10.113.97.199  8895c837-7794-e311-0001-0000000000bf
6    10.113.97.200  8895c837-7794-e311-0001-0000000000cf
7    10.113.97.197  8895c837-7794-e311-0001-0000000000df

* this terminal session
Nexus1000v-VSM-PRI#

```

New VEM Modules
Start from Slot 3

Figure 4-30 Cisco Nexus 1000V show module Command Output

NOTE Slots 1 and 2 are reserved for VSM (primary and standby) supervisory modules; the VEM module slots start from slot 3.

You can further verify module status with the **show module vem map** command, which shows the status of all visible VEMs, as well as their respective hosts (VMware ESX or ESXi), Universally Unique Identifier (UUID), and license status. Figure 4-31 shows a sample output of the **show module vem map** command.

```

Nexus1000v-VSM-PRI# show module vem map
Mod      Status      UUID                               License Status
-----
3        powered-up  8895c837-7794-e311-0001-0000000000af  licensed
4        powered-up  8895c837-7794-e311-0001-0000000000ef  licensed
5        powered-up  8895c837-7794-e311-0001-0000000000bf  licensed
6        powered-up  8895c837-7794-e311-0001-0000000000cf  licensed
7        powered-up  8895c837-7794-e311-0001-0000000000df  licensed
Nexus1000v-VSM-PRI#

```

Figure 4-31 Cisco Nexus 1000V *show module vem map* Command Output

Cisco Nexus 1000V VEM Installation Methods

There are multiple methods to install your Cisco Nexus 1000V VEM. Based on the chosen method, refer to the installation and configuration guides for detailed steps and prerequisites.

- If you are using the VUM (VMware Update Manager) to install or upgrade, you will have to create a host patch baseline and include the appropriate VMware patch or update bulletin and the corresponding Cisco Nexus 1000V VEM bulletin in the baseline.
- If you are using the vCLI on the ESX or ESXi host, use the **esxupdate**, **esxcli** command using the downloaded .vib file.
- You can also prepare ESXi images with VEMs installed and deploy them along with the ESX installation.

Initial VEM Status on ESX or ESXi Host

This topic will show you how to verify the status of a VEM on the VMware ESX or ESXi host by using VMware vCenter server or the Cisco Nexus 1000V Series switch CLI.



Verifying VSM Connectivity from vCenter Server

After VSM connectivity is established via the SVS connection, you will see the vNetwork Distributed Switch (vDS) in the vCenter networking inventory panel.

You will initially see the port profiles that you created, such as control, packet, and management. Typically, other port profiles are also created by default, such as the Unused_Or_Quarantined DVUplinks port group, which connects to physical NICs. Figure 4-32 shows a sample output of the initial port profiles. This output also shows physical uplinks already attached to the Cisco Nexus 1000V Series switch and note; it does not show “control” and “packet” port profiles because this environment is using L3 mode.

Verifying VEM Agent Running

If you have SSH enabled on your ESX or ESXi host, you can log in to your host by opening an SSH connection, log in with valid credentials, and at the CLI of the host, you can execute the **vem status** command. This command verifies whether the VEM module is loaded and shows whether the VEM Agent is running on this host. The command also confirms the physical interfaces (VMNICs) that are used as uplinks on this host. Figure 4-33 shows a sample output of the **vem status** command.

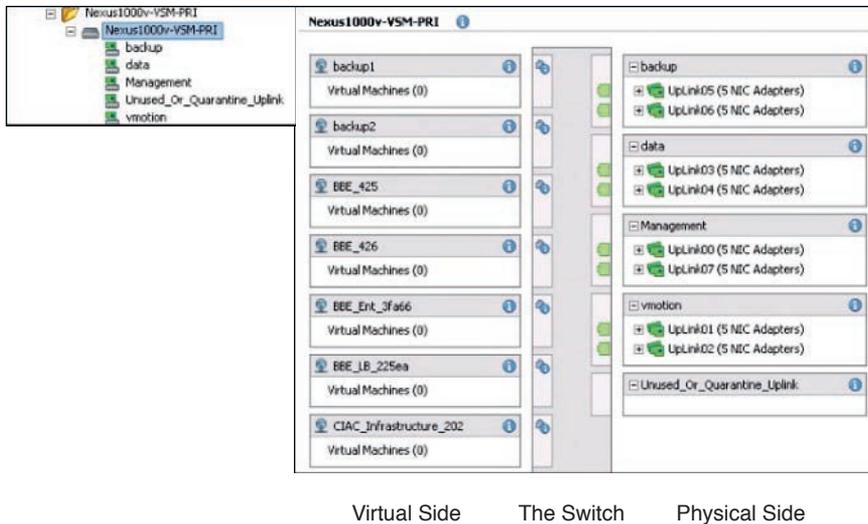


Figure 4-32 Cisco Nexus 1000V Series Switch from vCenter Server

Cisco Nexus 1000V Series VEM Module

```

- # vem status
VEM modules are loaded
Switch Name      Num Ports  Used Ports  Configured Ports  MTU      Uplinks
vSwitch1        128        1           128               1500
DVS Name         Num Ports  Used Ports  Configured Ports  MTU      Uplinks
Nexus1000v-VSM-PRI 1024       21          1024              1500     vmnic0,vmnic7,vmnic6,vmnic5,vmnic4,vmnic3,vmnic2,vmnic1
VEM Agent (vemdp) is running
  
```

Cisco Nexus 1000V
Series VEM Agent Status

VMNICs Used as Cisco
Nexus 1000V Series Uplinks

Figure 4-33 VEM Agent Verification on Host



Verifying VEM Uplink Interface Status

You can also verify the uplink interface status on the VMware ESX or ESXi hosts by using these commands: **vemcmd show port** and **vemcmd show port vlans**.

The **vemcmd show port** command verifies the VEM port that is used on the host (uplink) and the Cisco Nexus 1000V Series switch. It provides details of the port state, link status, and so on. As you can see in Figure 4-34, the port state (F/B*) indicates that it's blocking some VLANs.

The previous command (**vemcmd show port**) showed that certain VLANs were blocked on some uplinks. You can use the **vemcmd show port vlans** command to verify which VLANs are carried across the uplink and whether any VLANs are missing or blocked intentionally. In Figure 4-34, the uplink VMNIC3 on VEM allows only VLAN 99.

```

- # vemcmd show port
LTL VSM Port Admin Link State PC-LTL SCID Vem Port Type
17 Eth7/1 UP UP F/B* 561 0 vmnic0
18 Eth7/2 UP UP F/B* 561 1 vmnic1
19 Eth7/3 UP UP F/B* 562 2 vmnic2
20 Eth7/4 UP UP F/B* 562 3 vmnic3
21 Eth7/5 UP UP F/B* 563 4 vmnic4
22 Eth7/6 UP UP F/B* 563 5 vmnic5
23 Eth7/7 UP UP F/B* 564 6 vmnic6
24 Eth7/8 UP UP F/B* 564 7 vmnic7
49 Veth1 UP UP FWD 0 0 vmk0 L3VNSER
50 Veth2 UP UP FWD 0 2 vmk1
561 Po4 UP UP F/B* 0
562 Po3 UP UP F/B* 0
563 Po2 UP UP F/B* 0
564 Po1 UP UP F/B* 0

* F/B: Port is BLOCKED on some of the vlans.
One or more vlans are either not created or
not in the list of allowed vlans for this port.
Please run "vemcmd show port vlans" to see the details.

```

VEM Uplink

```

- # vemcmd show port vlans
LTL VSM Port Mode Native VLAN State* Vlans
17 Eth7/1 T 1 FWD 90-92,120-122,190-191,202,302
18 Eth7/2 T 1 FWD 90-92,120-122,190-191,202,302
19 Eth7/3 T 1 FWD 99
20 Eth7/4 T 1 FWD 99
21 Eth7/5 T 1 FWD 101-105,190,201,301
22 Eth7/6 T 1 FWD 101-105,190,201,301
23 Eth7/7 T 1 FWD 900-901
24 Eth7/8 T 1 FWD 900-901
49 Veth1 A 92 FWD 92
50 Veth2 A 99 FWD 99
561 Po4 T 1 FWD 90-92,120-122,190-191,202,302
562 Po3 T 1 FWD 99
563 Po2 T 1 FWD 101-105,190,201,301
564 Po1 T 1 FWD 900-901

* VLAN State: VLAN State represents the state of allowed vlans.

```

VLANs Allowed

Figure 4-34 VEM Uplink Interface Status

Verifying VEM Parameters with VSM Configuration

To ensure that your VEM's parameters match the VSM, the `vemcmd show card` command executed on the VMware ESX or ESXi host can be very useful. Using the command verifies the following components. Figure 4-35 shows the output of the `vemcmd show card` command and identifies the following configurations.

- Card name
- Card domain ID
- Card slot
- Control traffic connectivity mode between VSM and VEM
- VEM control agent MAC ID
- Used VSM MAC address
- Used control and packet VLANs

Key Topic

Validating VM Port Groups and Port Profiles

Now that you have verified the VSM and the VEM configurations and cross-checked the configurations from both perspectives, you will understand how to validate that VMs are using the correct vDS port groups by using VMware vCenter and Cisco Nexus 1000V Series switch CLI.

First, use the `show port-profile name <name>` command to verify the details of that port profile configuration and parameters. With this command, you can check which switchport mode this port profile uses, which VLANs are associated with this port profile, and which virtual interfaces are assigned to this port profile. Figure 4-36 shows a sample port profile output using the `show port-profile name <name>` command.

```

- # vemcmd show card
Card UUID type 2: 8895c837-7794-e311-0001-0000000000df
Card name: esxi-vsamgmt2
Switch name: Nexus1000v-VSM-PRI
Switch alias: DvsPortset-0
Switch uuid: 4d 64 20 50 74 83 09 8e-40 91 13 1a 3b 3d fb 52
Card domain: 1
Card slot: 7
VEM Tunnel Mode: L3 Mode
L3 Ctrl Index: 49
L3 Ctrl VLAN: 92
VEM Control (AIPC) MAC: 00:02:3d:10:01:06
VEM Packet (Inband) MAC: 00:02:3d:20:01:06
VEM Control Agent (DPA) MAC: 00:02:3d:40:01:06
VEM SPAN MAC: 00:02:3d:30:01:06
Primary VSM MAC : 00:50:56:a0:6b:97
Primary VSM PKT MAC : 00:50:56:a0:4f:48
Primary VSM MGMT MAC : 00:50:56:a0:60:5f
Standby VSM CTRL MAC : 00:50:56:a0:3c:ce
Management IPv4 address: 10.113.97.197
Management IPv6 address: 0000:0000:0000:0000:0000:0000:0000:0000
Primary L3 Control IPv4 address: 10.113.97.202
Secondary VSM MAC : 00:00:00:00:00:00
Secondary L3 Control IPv4 address: 0.0.0.0
Upgrade : Default
Max physical ports: 32
Max virtual ports: 300
Card control VLAN: 1
Card packet VLAN: 1
Control type multicast: No
Card Headless Mode : No
Processors: 32
Processor Cores: 16
Processor Sockets: 2
Kernel Memory: 100628640
Port link-up delay: 5s
Global UFB: DISABLED
Heartbeat Set: True
PC LB Algo: source-mac
Datapath portset event in progress : no
Licensed: Yes

```

Card Name

Card Domain

Card Slot

VSM to VEM
Connectivity Mode

VEM Agent MAC ID

VSM MAC IDs and
L3 Control IP Address

Control/Packet VLAN IDs
(Not Used Since L3 Mode)

The 2-Second Heartbeat
from VEM to VSM Enabled

Figure 4-35 Verifying VEM Parameters with the `vemcmd show card` Command

```

Nexus1000v-VSM-PRI# show port-profile name Demo_Web_Network

port-profile Demo_Web_Network
type: Vethernet
description:
status: enabled
max-ports: 32
min-ports: 1
inherit:
config attributes:
  switchport mode access
  switchport access vlan 190
no shutdown
evaluated config attributes:
  switchport mode access
  switchport access vlan 190
no shutdown
assigned interfaces:
  Vethernet12
  Vethernet13
  Vethernet16
  Vethernet17
port-group: Demo_Web_Network
system vlans: 190
capability l3control: no
capability iscsi-multipath: no
capability vxlan: no
capability l3-vservice: no
port-profile role: none
port-binding: static

```

Port Profile Type
= Vethernet

Switchport Mode and
Allowed VLANs

Port Profiles Being
Used and Assigned
Interfaces

Figure 4-36 Cisco UCS Nexus 1000V Sample `show port-profile <name>` Command Output

After this port profile (Demo_Web_Network) is created on the Cisco Nexus 1000V Series switch and automatically pushed to VMware vCenter server, it will be available for the VMware/server administrator to use and connect a VM to that port profile.

Now let us see that same port profile (Demo_Web_Network) from VMware vCenter server; by right-clicking the VM and choosing Edit Settings, you can add a VM to a port group or port profile. After choosing Edit Settings, you will be presented with the screen shown

in Figure 4-37, where you can see Demo_Web_Network as a port group, and now you can assign this port group Demo_Web_Network to a VM's vNIC.

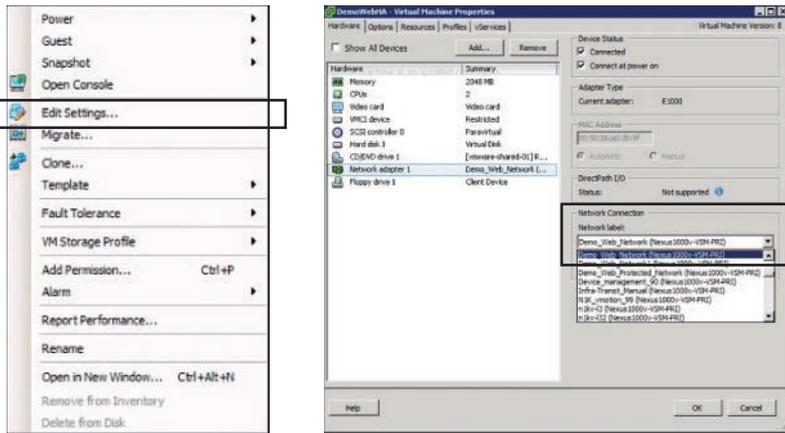


Figure 4-37 VMware vCenter Sample Port Profile Assignment to VM

Verifying Port Profile and Groups

From within VMware vCenter, choose Inventory, Networking, and in the navigation pane you will see the networking inventory objects, including the Demo_Web_Networking port profile. This port profile is a VM port profile, meaning it can be assigned to VMs, VMK (they get tied to a vEth number), and not an uplink port profile (note the different symbols); they get tied to an Eth x/x number.

In the same window, you can see the uplink port profiles created on the Cisco Nexus 1000V Series switch and already pushed to VMware vCenter server so that the virtual switch can provide external connectivity for the VMs that reside on that host. Figure 4-38 shows both the Demo_Web_Networking port profile and the uplink port profiles, such as data, backup, vmotion, and the like.

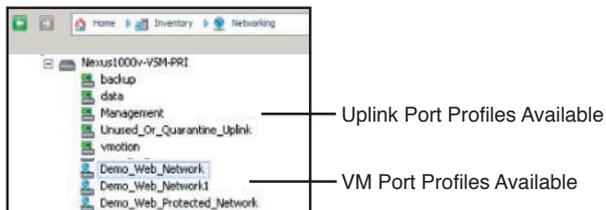


Figure 4-38 Verifying Port Profiles from VMware vCenter Inventory Networking Pane

Key New Technologies Integrated with Cisco Nexus 1000V

In addition, here are some brief clarifications and explanations on some concepts and new technologies that have been integrated to work with the Cisco Nexus 1000V Series software switch. In-depth details of these concepts and technologies are beyond the scope of this certification guide.

What Is the Difference Between VN-Link and VN-Tag?

VN-Link (Virtual Network—Link) is not a technology; it is a marketing term that refers to network visibility of nonphysical, nondirect attached devices. This could include virtual machines or virtual interfaces, or it could mean physical interfaces on nonswitching remote devices like Nexus 2000 Series devices. Two approaches offer VN-Link capabilities:

- The Nexus 1000V is VN-Link capable because every virtual machine connected to this software switch receives a virtual Ethernet port that can be configured and controlled just like a physical Ethernet switch port on a standard physical switch.
- The Nexus 5000/2000 Series combination and Cisco UCS Fabric Interconnect/I/O module combination both use an additional header in the Ethernet frame called VN-Tag (Virtual Network—Tag), which identifies a remote port (on Nexus 2000 / UCS I/O module) that will be assigned a virtual port on the Cisco Nexus 5000 Series or Cisco UCS 6100/6200/6300 Series Interconnects. This enables the Nexus 2000/UCS I/O module to be managed as a line card on the Cisco Nexus 5000/UCS Fabric 6100/6200/6300 Interconnect. All switching takes place on the Cisco Nexus 5000 or UCS Fabric Interconnects. This same VN-Tag is used to identify a virtual interface on a Cisco Virtual Interface Card (Cisco VIC); the host device (Cisco Nexus 5000 or UCS Fabric Interconnect) can now manage that port as if it were a physical port on the host device.

What Is VXLAN?

Virtual extensible LAN (VXLAN) is an overlay network that is superimposed on top of a traditional network. It is a Layer 2 overlay scheme utilizing Layer 3 networks. Cisco VXLAN is a Layer 2 network isolation technology that uses a 24-bit segment identifier to scale beyond the 4K limitations of VLANs. VXLAN technology creates LAN segments by using an overlay approach with MAC in IP encapsulation.

The Virtual Ethernet Module (VEM) encapsulates the original Layer 2 frame from the VM. Each VEM is assigned an IP address, which is used as the source IP address when encapsulating MAC frames to be sent on the network. This is accomplished by creating virtual network adapters (VMNIC) on each VEM. You can have multiple VMNICs per VEM and use them for this encapsulated traffic. The encapsulation carries the VXLAN identifier, which is used to scope the MAC address of the payload frame. Figure 4-39 shows a graphical view of a host overlay network with virtual endpoints.

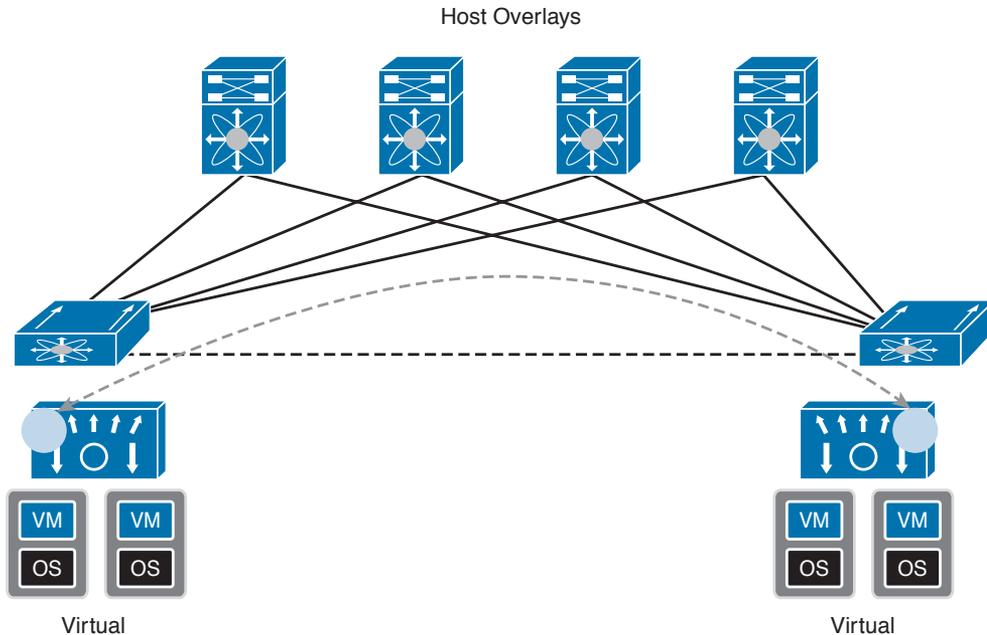


Figure 4-39 *Host Overlay with Virtual Endpoints*

NOTE Cisco Nexus 1000V Series switches fully support VXLAN technology starting from version 1.5 onward.

What Is vPath Technology?

In virtualized data centers, services such as firewalls, server load balancers, security gateways, and the like are often offered and implemented as virtual services. Given that, the traditional physical connections for these services do not exist, and retrofitting VMs to avail these services causes significant inefficiencies and overheads.

Cisco vPath provides embedded intelligence for virtual services and offers traffic steering capabilities to direct traffic to these virtual services nodes. These topology-agnostic service-chaining capabilities can be built in to the respective Nexus 1000V port profile, for vPath interception. Cisco vPath (Virtual Services Datapath) technology is fully integrated with Cisco Nexus 1000V Series software switches. Figure 4-40 shows an overview of vPath—virtual network services architecture.

What Is Cisco Application Virtual Switch (AVS)?

Cisco AVS is underpinned by Cisco Nexus 1000V Series virtual switch. It is fully integrated with Cisco Application Centric Infrastructure (ACI); thus, the VSM is integrated into the application policy infrastructure controller (APIC), allowing enforcement of application-centric policies automatically, all the way up to the virtual edge of the data center. Please refer to Chapters 12 through 14 for some information.

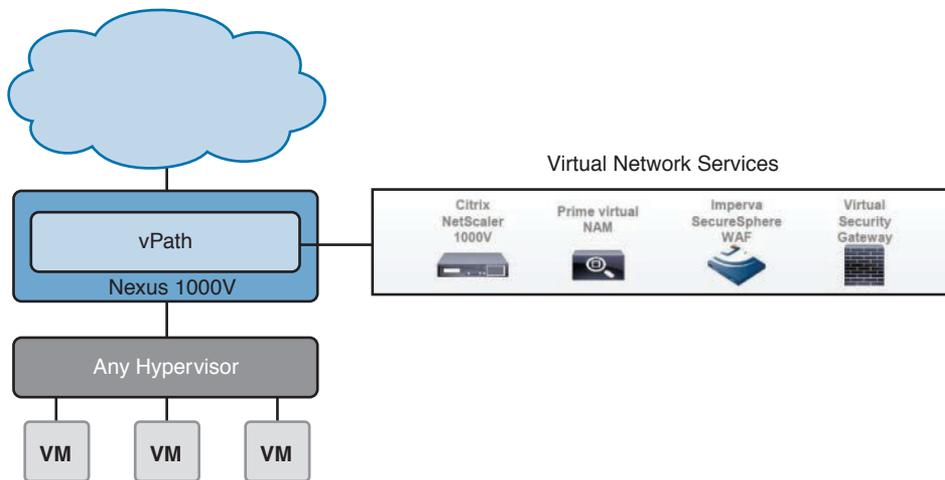


Figure 4-40 Cisco Virtual Network Services with vPath Technology

Reference List

VMware Virtual Networking Concepts: http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf

Cisco Nexus 1000V Configuration Guides: <http://www.cisco.com/c/en/us/support/switches/nexus-1000v-switch-vmware-vsphere/products-installation-and-configuration-guides-list.html>

Cisco Nexus 1000V Installation Guides: <http://www.cisco.com/c/en/us/support/switches/nexus-1000v-switch-vmware-vsphere/products-installation-guides-list.html>

VMware Configuration Maximums: <http://www.vmware.com/pdf/vsphere5/r55/vsphere-55-configuration-maximums.pdf>

VMware vSphere 5.1 Documentation: <http://pubs.vmware.com/vsphere-51/index.jsp#com.vmware.vsphere.doc/GUID-1B959D6B-41CA-4E23-A7DB-E9165D5A0E80.html>

Virtual Networking Features of VMware vDS and Cisco Nexus 1000V: http://www.cisco.com/c/dam/en/us/products/collateral/switches/nexus-1000v-switch-vmware-vsphere/solution_overview_c22-526262.pdf

Cisco Nexus 1100 Series Virtual Services Appliance: http://www.cisco.com/c/en/us/products/collateral/switches/nexus-1100-series-cloud-services-platforms/data_sheet_c78-297641.html

Cisco VXLAN on Nexus 1000V: http://www.cisco.com/c/en/us/products/collateral/switches/nexus-1000v-switch-vmware-vsphere/guide_c07-702975.html

Exam Preparation Tasks

Review All Key Topics

Review the most important topics in the chapter, noted with the key topic icon in the outer margin of the page. Table 4-4 lists a reference for these key topics and the page numbers on which each is found.

Table 4-4 Key Topics for Chapter 4

Key Topic Element	Description	Page
Figure 4-1	“Server Connectivity Before Server Virtualization”	119
List	“Physical Access Layer Challenges”	120
Section	“Server Virtualization with Static VMware vSwitch”	121
List	“Challenges with Standard VMware vSwitch”	121
Section	“Virtual Network Components”	122
Section	“VMware vDS Overview”	128
Figure 4-11	“VMware Virtual Distributed Switch (vDS) Overview”	129
List	“Advantages and Enhancements of VMware vDS”	131
Section	“Cisco Nexus 1000V System Overview”	134
List	“Cisco Nexus 1000V Salient Features and Benefits”	135
Section	“Cisco Nexus 1000V Series Virtual Switch Architecture”	137
List	“Cisco Nexus 1000V Component Communication”	140
Section	“Cisco Nexus 1000V Port Profiles”	141
List	“Types of Port Profiles in Cisco Nexus 1000V”	142
Section	“Initial VSM Configuration Verification”	145
Section	“Verifying Nexus 1000V Module Status”	147
Section	“Verifying VSM Connectivity from vCenter Server”	148
Section	“Verifying VEM Uplink Interface Status”	149
Section	“Validating VM Port Groups and Port Profiles”	150

Complete Tables and Lists from Memory

Print a copy of Appendix B, “Memory Tables” or at least the section for this chapter, and complete the tables and lists from memory. Appendix C, “Memory Tables Answer Key” includes completed tables and lists to check your work.

Define Key Terms

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

Virtual LAN (VLAN), Virtual Extensible LAN (VXLAN), Virtual Services Datapath (vPath), Virtual Network Interface Card (vNIC), VMNIC, Uplink Trunk, VN-Link, VN-Tag, Cisco Virtual Interface Cards (VIC), VMware Update Manager (VUM), vEthernet Ethernet Application Programming Interface (API), Private VLAN (PVLAN), 802.1q, Link Aggregation Control Protocol (LACP), vPC Host Mode, Quality of Service (QoS), Internet Small Computer Interface (iSCSI), VMKernel, Application Centric Infrastructure (ACI), Application Policy Infrastructure Controller (APIC), Application Virtual Switch (AVS)