

Official Cert Guide

Learn, prepare, and practice for exam success

- ▶ Master the #2V0-641 exam with this official guide
- ▶ Assess your knowledge with chapter-opening quizzes
- ▶ Review key concepts with Exam Preparation Tasks
- ▶ Practice with realistic exam questions

VCP6-NV

(Exam #2V0-641)

ELVER SENA SOSA

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



VCP6-NV
Official Cert Guide
(Exam #2V0-641)

VMware Press is the official publisher of VMware books and training materials, which provide guidance on the critical topics facing today's technology professionals and students. Enterprises, as well as small- and medium-sized organizations, adopt virtualization as a more agile way of scaling IT to meet business needs. VMware Press provides proven, technically accurate information that helps them meet their goals for customizing, building, and maintaining their virtual environment.

VMware Press provides proven, technically accurate information that will help you achieve your goals for customizing, building, and maintaining a virtual environment—from the data center to mobile devices to the public, private, and hybrid cloud.

With books, certification and study guides, video training, and learning tools produced by world-class architects and IT experts, VMware Press helps IT professionals master a diverse range of topics on virtualization and cloud computing and is the official source of reference materials for preparing for the VMware Certified Professional certification.

VMware Press is also pleased to have localization partners that can publish its products in more than 42 languages, including, but not limited to, Chinese (Simplified), Chinese (Traditional), French, German, Greek, Hindi, Japanese, Korean, Polish, Russian, and Spanish.

For more information about VMware Press, please visit **vmwarepress.com**.

This page intentionally left blank

VCP6-NV

Official Cert Guide

(Exam #2V0-641)

Elver Sena Sosa

vmware® PRESS

Hoboken, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris • Madrid
Cape Town • Sydney • Tokyo • Singapore • Mexico City

VCP6-NV Official Cert Guide (Exam #2V0-641)

Copyright © 2017 VMware, Inc.

Published by Pearson Education, Inc.

Publishing as VMware Press

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise.

ISBN-10: 0-7897-5480-0

ISBN-13: 978-0-7897-5480-6

Library of Congress Control Number is on file.

Printed in the United States of America

First Printing: August 2016

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. The publisher 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.

VMware terms are trademarks or registered trademarks of VMware in the United States, other countries, or both.

Warning and Disclaimer

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 provided is on an “as is” basis. The authors, VMware Press, VMware, and the publisher 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 author and are not necessarily those of VMware.

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 United States, please contact intlcs@pearson.com.

EDITOR IN CHIEF

Mark Taub

PRODUCT LINE MANAGER

Brett Bartow

EXECUTIVE EDITOR

Mary Beth Ray

VMWARE PRESS PROGRAM MANAGER

Karl Childs

DEVELOPMENT EDITOR

Christopher Cleveland

MANAGING EDITOR

Sandra Schroeder

PROJECT EDITOR

Mandie Frank

TECHNICAL EDITORS

William Grismore, Richard Hackman, Jon Hall

COPY EDITOR

Geneil Breeze

PROOFREADER

The Wordsmithery LLC

INDEXER

Ken Johnson

EDITORIAL ASSISTANT

Vanessa Evans

DESIGNER

Chuti Prasertsith

COMPOSITOR

TnT Design

Contents at a Glance

INTRODUCTION xx

CHAPTER 1	Introduction to VMware NSX	3
CHAPTER 2	Network and VMware vSphere Requirements for NSX	21
CHAPTER 3	NSX Architecture and NSX Manager	61
CHAPTER 4	VXLAN, NSX Controllers, and NSX Preparation	87
CHAPTER 5	NSX Switches	127
CHAPTER 6	Logical Switch Packet Walks	161
CHAPTER 7	Logical Router	195
CHAPTER 8	Logical Router Packet Walks	227
CHAPTER 9	NSX Edge Services Gateway	253
CHAPTER 10	Layer 2 Extensions	275
CHAPTER 11	Layer 3 Connectivity Between Virtual and Physical Networks	315
CHAPTER 12	Routing Protocols	343
CHAPTER 13	NSX Edge VPN Services	379
CHAPTER 14	NSX Edge Network Services and Security	413
CHAPTER 15	Distributed Logical Firewall	445
CHAPTER 16	Security Services	477
CHAPTER 17	Additional NSX Features	503
CHAPTER 18	NSX Automation	527
CHAPTER 19	Upgrade to NSX for vSphere 6.2	551
CHAPTER 20	Final Preparation	571
APPENDIX A	Answers to the “Do I Know This Already?” Quizzes	581
APPENDIX B	VCP6-NV Exam 2V0-641 Updates	585
GLOSSARY		588
INDEX		596

ONLINE ELEMENTS

APPENDIX C	Memory Tables
APPENDIX D	Memory Tables Answer Key
APPENDIX E	Study Planner

Contents

Introduction xx

- About This Book xx
- Who Should Read This Book xxi
- Book Features xxii
- How to Use This Book xxiii
- Certification Exam and This Preparation Guide xxv
- Book Content Updates xxvi
- Companion Website xxviii
- Pearson IT Certification Practice Test Engine and Questions xxviii
 - Install the Software xxix
 - Activate and Download the Practice Exam xxix
 - Activating Other Exams xxx
 - Assessing Exam Readiness xxx
 - Premium Edition eBook and Practice Tests xxxi

Chapter 1 Introduction to VMware NSX 3

- Do I Know This Already? 3
- Foundation Topics 6**
- Physical Network Challenges 6
- Ethernet Challenges 7
- IP Network Challenges 10
- Security Challenges 12
- VMware NSX 12
- Exam Preparation Tasks 17**
- Review All the Key Topics 17
- Complete Tables and Lists from Memory 18
- Define Key Terms 18

Chapter 2 Network and VMware vSphere Requirements for NSX 21

- Do I Know This Already? 21
- Foundation Topics 24**
- Physical Network Infrastructure 24
- POD Design 28
- Collapsed Access Layer 30
- Spine and Leaf Design 31
- NSX and Physical Network Infrastructure 33
- NSX and vSphere 34
 - ESXi Host Network Connectivity 35
 - vSphere Standard Switch 36
 - vSS Configuration 39

vSphere Distributed Switch	43
Create vSphere Distributed Switch	45
Migrate to vSphere Distributed Switch	49
Configure LACP	52
Configure QoS Marking	55
Exam Preparation Tasks	58
Review All the Key Topics	58
Complete Tables and Lists from Memory	59
Define Key Terms	59

Chapter 3 NSX Architecture and NSX Manager 61

Do I Know This Already?	61
Foundation Topics	64
Network Planes	64
NSX Architecture	66
NSX Manager	68
NSX Manager Base Configuration	73
Cross vCenter NSX	80
Exam Preparation Tasks	84
Review All the Key Topics	84
Complete Tables and Lists from Memory	85
Define Key Terms	85

Chapter 4 VXLAN, NSX Controllers, and NSX Preparation 87

Do I Know This Already?	87
Foundation Topics	90
VXLAN Introduction	90
VXLAN	90
NSX Controllers	96
Deploying NSX Controllers	97
Verifying NSX Controllers	101
Creating an NSX Controller Cluster	104
NSX Controller Master and Recovery	106
IP Pools	107
Host Preparation	109
Host Configuration	113
VNI Pools, Multicast Pools, and Transport Zones	120
Exam Preparation Tasks	124
Review All the Key Topics	124
Complete Tables and Lists from Memory	125
Define Key Terms	125

Chapter 5 NSX Switches 127

Do I Know This Already?	127
Logical Switches	130

Foundation Topics 130

Creating a Logical Switch	131
Verifying Logical Switches	135
Adding Virtual Machines to Logical Switches	137
Logical Switch Tables	138
VTEP Table	139
Example: Populating the VTEP Table	140
Example: Updating the VTEP Table	143
MAC Table	144
ARP Table	147
Logical Switch Table Verification	149
Unknown Unicast or ARP Request	152
Replication Mode	152
Multicast Replication Mode	154
Unicast Replication Mode and Proxy VTEP	155
Hybrid Replication Mode	156
Exam Preparation Tasks 158	
Review All the Key Topics	158
Complete Tables and Lists from Memory	159
Define Key Terms	159

Chapter 6 Logical Switch Packet Walks 161

Do I Know This Already?	161
Foundation Topics 165	
Logical Switches Packet Walks	165
Logical Switch Packet Walk Example 1	169
Logical Switch Packet Walk Example 2	170
Logical Switch Packet Walk Example 3	171
Logical Switch Packet Walk Example 4	177
Logical Switch Packet Walk Example 5	189
Exam Preparation Tasks 193	
Review All the Key Topics	193
Define Key Terms	193

Chapter 7 Logical Router 195

Do I Know This Already?	195
Foundation Topics 198	
NSX Logical Router	198
Logical Router Control VM	201
Creating and Deploying the Logical Router	204
Logical Router Verification	210
Connectivity Testing	216
Locale ID	221

Exam Preparation Tasks 225

Review All the Key Topics 225

Define Key Terms 225

Chapter 8 Logical Router Packet Walks 227

Do I Know This Already? 227

Foundation Topics 232

Logical Router Packet Walks 232

Logical Router Packet Walk Example 1 235

Logical Router Packet Walk Example 2 240

Logical Router Packet Walk Example 3 246

Exam Preparation Tasks 250

Review All the Key Topics 250

Define Key Terms 250

Chapter 9 NSX Edge Services Gateway 253

Do I Know This Already? 253

Foundation Topics 256

NSX Edge 256

NSX Edge Size 259

Edge HA 260

Creating and Deploying an NSX Edge 262

Undeployed NSX Edge 267

Monitoring and Verifying the NSX Edge 269

Exam Preparation Tasks 273

Review All the Key Topics 273

Complete Tables and Lists from Memory 273

Define Key Terms 273

Chapter 10 Layer 2 Extensions 275

Do I Know This Already? 275

Foundation Topics 278

Layer 2 VPN 278

Configuring Layer 2 VPN 280

Verifying Layer 2 VPN 289

Layer 2 VPN Packet Walk 290

Layer 2 Bridging 296

Configuring Layer 2 Bridging 299

Verifying Layer 2 Bridging 301

Layer 2 Bridging Packet Walk 303

Hardware VTEPs 307

Exam Preparation Tasks 312

Review All the Key Topics 312

Complete Tables and Lists from Memory 313

Define Key Terms 313

Chapter 11 Layer 3 Connectivity Between Virtual and Physical Networks 315

Do I Know This Already? 315

Foundation Topics 318

Logical Router VLAN LIF 318

Designated Instance 324

NSX Edge Gateway 330

Equal Cost Multipathing 336

Exam Preparation Tasks 340

Review All the Key Topics 340

Complete Tables and Lists from Memory 340

Define Key Terms 341

Chapter 12 Routing Protocols 343

Do I Know This Already? 343

Foundation Topics 346

Routing 346

Administrative Distance and Cost 347

Static Routes 349

OSPF 351

OSPF Areas 352

OSPF Neighbor Adjacencies 354

LSA Types 355

Configuring OSPF 356

Verifying OSPF 360

BGP 361

Configuring BGP 366

Verifying BGP 368

IS-IS 369

IS-IS Areas and IS Types 369

Configuring IS-IS 370

Verifying IS-IS 373

Route Redistribution 373

Exam Preparation Tasks 376

Review All the Key Topics 376

Complete Tables and Lists from Memory 377

Define Key Terms 377

Chapter 13 NSX Edge VPN Services 379

Do I Know This Already? 379

Foundation Topics 382

IPsec VPNs 382

IPsec VPN Establishment 384

Configuring IPsec VPNs 386

Verifying IPsec VPNs 392

SSL VPN-Plus	394
Configure SSL VPN-Plus	395
SSL VPN-Plus Server Settings	396
Creating a Web Resource	397
Configuring Authentication	398
Enable SSL VPN-Plus Service	402
Adding the Installation Package	403
Adding an IP Pool	405
Adding Private Networks	406
Verifying SSL VPN-Plus	408
Exam Preparation Tasks	411
Review All the Key Topics	411
Complete Tables and Lists from Memory	411
Define Key Terms	411

Chapter 14 NSX Edge Network Services and Security 413

Do I Know This Already?	413
Foundation Topics	416
Network Address Translation	416
NSX Edge Load Balancer	420
Configuring the Edge Load Balancer	427
Application Profile	427
Server Pools	430
Virtual Server	431
Enable Load Balancer	433
NSX Edge Protocol and Port Groupings	433
Configure NSX Edge DHCP and DNS	434
NSX Edge Logical Firewall	436
Configuring an Edge Firewall	439
Exam Preparation Tasks	443
Review All the Key Topics	443
Complete Tables and Lists from Memory	443
Define Key Terms	443

Chapter 15 Distributed Logical Firewall 445

Do I Know This Already?	445
Foundation Topics	449
Traditional Firewall Design Compromises	449
Distributed Logical Firewall	453
DFW Thresholds and Limits	458
Exclusion List	460
Logical Firewall Rules	460
Creating Firewall Sections and Rules	462
Firewall Rules Saved Configurations	468
NSX Manager and Domains	469

- Verifying DFW Functionality 470
- SpoofGuard 471
- Exam Preparation Tasks 474**
- Review All the Key Topics 474
- Complete Tables and Lists from Memory 475
- Define Key Terms 475

Chapter 16 Security Services 477

- Do I Know This Already? 477
- Foundation Topics 480**
- Security Services for NSX 480
 - Registering Service with NSX 482
 - Deploying the Security Service Appliance 484
- Service Composer 486
 - Security Groups 487
 - Security Policies 491
- Logical Firewall Service Redirection 496
- Security Tags 497
- IP Sets and MAC Sets 499
- Exam Preparation Tasks 501**
- Review All the Key Topics 501
- Complete Tables and Lists from Memory 501
- Define Key Terms 501

Chapter 17 Additional NSX Features 503

- Do I Know This Already? 503
- Foundation Topics 506**
- VMware Data Security 506
- Activity Monitoring 509
 - VM Activity 511
 - Inbound Activity 512
 - Outbound Activity 513
 - Inter Container Interaction 513
 - Outbound AD Group Activity 514
 - Viewing Activity Report 514
- Flow Monitoring 514
- Traceflow 519
- Role Based Access Control 521
- Exam Preparation Tasks 524**
- Review All the Key Topics 524
- Complete Tables and Lists from Memory 525
- Define Key Terms 525

Chapter 18 NSX Automation 527

Do I Know This Already? 527

Foundation Topics 530

REST 530

NSX API Calls for Logical Switch 532

NSX API Calls for Logical Router 536

NSX API Calls for NSX Edge 540

vRealize Automation 542

External Network Profile 544

Routed Network Profile 544

Private Network Profile 545

NAT Network Profile 546

Exam Preparation Tasks 548

Review All the Key Topics 548

Complete Tables and Lists from Memory 548

Define Key Terms 549

Chapter 19 Upgrade to NSX for vSphere 6.2 551

Do I Know This Already? 551

Foundation Topics 555

Upgrade vCloud Network and Security to NSX for vSphere 555

Upgrade to NSX Manager 555

Upgrade to NSX VIBs 558

Upgrade to NSX DFW 559

Upgrade to NSX Edge 559

Upgrade to USVM 560

Upgrade NSX for vSphere to NSX for vSphere 6.2 561

Upgrade to NSX Manager 6.2 561

Upgrade NSX Controllers to 6.2 563

Upgrade Host Clusters to 6.2 565

Upgrade NSX Edges to 6.2 566

Exam Preparation Tasks 568

Review All the Key Topics 568

Define Key Terms 568

CHAPTER 20 Final Preparation 571

Getting Ready 571

Taking the Exam 574

Tools for Final Preparation 575

Review Tools on the Companion Website 575

Pearson Cert Practice Test Engine and Questions 576

Using the Exam Engine 578

Appendix A Answers to the “Do I Know This Already?” Quizzes 581

Appendix B VCP6-NV Exam 2V0-641 Updates 585

Always Get the Latest at the Book’s Product Page 585

Technical Content 586

GLOSSARY 588

Index 596

ONLINE ELEMENTS

APPENDIX C Memory Tables

APPENDIX D Memory Tables Answer Key

APPENDIX E Study Planner

Dedication

I am dedicating this book to my father, who told me when I was still in high school to learn as much about computers as I could. He convinced me to take a Lotus 1-2-3 class and later an A+ class! Thanks, Dad!

About the Author

Elver Sena Sosa, CCIE 7321 Emeritus (R&S), VCDX-NV (#154), CCSI, VCI. Elver has been working in IT since the late 1990s. Elver started his IT career as an intern network engineer in Appleton, Wisconsin, later moving to Columbus, Ohio, to work with AT&T Solutions. Over the years Elver continued to learn more about different technologies and how these technologies could help solve business problems. Feeling constrained and limited working in a siloed environment, Elver decided to become an independent contractor so that he could help provide technical solutions for as many different clients as possible. Elver currently is the data center infrastructure architect at Hydra 1303, Inc. You can follow Elver on Twitter @ElverS_Opinion, or his blog, <http://blog.senasosa.com>.

Acknowledgments

I have a lot of people to thank for making this first book a reality. The biggest and most important are my wife, Katy, and son, Danilo. Katy served as my non-technical English editor, reading chapters while having no idea about what she was reading and somehow translating and fixing what she read from Elver to English. They both endured my physical and emotional absence throughout this process, at times encouraging me to keep going when I wanted to quit (did I allude to how hard it is to write a book?). Although at times it looked as if they were more pleased than not that I was locked in my office, their support is what made this project possible.

I also want to thank those at VMware (Chris McCain, Jenny Lawrence, Quang Nguyen) who provided the opportunities that put me on the path to writing this book. I want to thank those at Pearson (Mary Beth Ray, Chris Cleveland) who took a chance on me and provided me guidance along the way to get this done.

Special thanks goes to my editors (Richard Hackman, William Grismore, Jon Hall) for going through the pain of reading my drafts. I know it wasn't easy, but your feedback was very valuable (well, most of the feedback ☺).

I want to save the last thanks to those who kept asking me “when is the book coming out?” Every few weeks someone would ask me this, and although I didn't say it, it was encouraging that someone out there was interested in reading what I wrote. Muchas gracias.

We Want to Hear from You!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our way.

We welcome your comments. You can email or write us directly to let us know what you did or didn't like about this book—as well as what we can do to make our books better.

Please note that we cannot help you with technical problems related to the topic of this book.

When you write, please be sure to include this book's title and author as well as your name, email address, and phone number. We will carefully review your comments and share them with the author and editors who worked on the book.

Email: VMwarePress@vmware.com

Mail: VMware Press
 ATTN: Reader Feedback
 800 East 96th Street
 Indianapolis, IN 46240 USA

Reader Services

Register your copy of *VCP6-NV Official Cert Guide (Exam #2V0-641)* at www.pearsonitcertification.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.pearsonitcertification.com/register and log in or create an account*. Enter the product ISBN, 9780789754806, and click Submit. Once the process is complete, you will find any available bonus content under Registered Products.

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

Introduction

Hola y bienvenidos. I'm grateful that you have decided to pick up a copy of the *VCP6-NV Official Cert Guide (Exam #2V0-641)* and read it. Or if the book was given to you, I'm grateful that you decided to keep the book and read it instead of donating it to someone else. Why am I grateful? I'm grateful because I understand that your time is valuable, and out of all the available sources of information about NSX for vSphere, you chose my book as one of your study sources. Thank you.

About This Book

I was lucky to be in the right place at the right time when NSX for vSphere came out. I was one of the few folks around who knew vSphere, vRealize Automation (formerly VCAC), and vCloud Director well enough, and also had a better than average understanding of networking and network security. Being one of the few folks who fit the mold, it was a natural progression for me to get involved with NSX and thus I took the plunge. Over the last three years I have been traveling the world educating about NSX for vSphere and software defined networks, including delivering the first week of training to the first group of NSX Ninja candidates. I have also served as a mentor to many of the current NSX professionals and instructors, some of whom have grown to be way more competent than me in the subject. Before plunging into NSX, I was already working as an independent consultant as well as a VMware and Cisco instructor (and at one time a high school math teacher in the Bronx). I have delivered many courses over the years and have met many people.

Before writing the book I was heavily involved in writing the *NSX for vSphere: Install, Configure, Manage* and *NSX for vSphere: Fast Track* courses, one of which must be attended before you can be certified as a VMware Certified Professional 6 - Network Virtualization (VCP6-NV) (if you don't already have a VCP from another VMware solution track). Having now done both I can attest that writing a course is a cakewalk compared to writing a certification book. From time to time I try to pen some stuff in my blog, <http://blog.senasosa.com/>, as well as give talks at VMUGs, which I greatly enjoy.

This is my first book, so I'm really hoping you like it and find it useful. Although VMware puts out an exam blueprint to help students prepare for the exam, located in this site www.vmware.com/go/vcp6nv, this book does not follow the layout of the blueprint. The book's layout is designed to help the student fully understand what NSX is, the problems it solves, and the different features it provides. You will notice

the book starts with a short trip down memory lane on how data center networking used to be and how it evolved to what it is today, followed by the introduction to NSX and its components. In Chapters 6 and 8 I opted for walking the reader through different packet walks so as to better illustrate how logical switches and logical routers work. While the book covers all the objects in the blueprint (as of January 2016), it is possible that the blueprint could be modified at VMware's discretion at any time.

In writing the book I assumed that you know what a virtual machine is and not much more. I assumed that your knowledge of the vSphere switches and basic networking is limited, thus I spent some time covering those basics where needed in the book. If you feel that you are above average in those topics, feel free to skip over them. If you are not sure how to rate yourself in those topics, the material is here for you to read; it should be a quick read anyway.

I also strongly advise you to get your hands on an NSX lab as part of your studies. There is nothing like having practical experience beyond reading and memorizing. If you can't get yourself your own lab, you can try the ones provided by VMware (for free) at the Hands On Labs, <http://labs.hol.vmware.com>.

And with that said, I wish you best of luck in your studies, and let's set sail.

Who Should Read This Book

If you work in the data center as a network administrator, storage administrator or vSphere administrator, this book is for you. By now you should have noticed that infrastructure components you work with in the data center have been prepended with a "Software Defined" term in front of it. The days of having a strict silo where you only knew one aspect of the data center infrastructure are numbered as all those Software Defined *whatever* have a strong co-dependency with each other. In the data center, infrastructure will be automated but to get us there (and for you to have a job in the data center) you must understand how each of those silos work. This book is one of the steps in the ladder to get you there by helping you become VCP6-NV certified.

Book Features

To help you customize your 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?” quiz:** Each chapter begins with a quiz that helps you determine how much time you need to spend studying that chapter.
- **Foundation Topics:** These are the core sections of each chapter. They explain the concepts for the topics in that chapter.
- **Exam Preparation Tasks:** After the “Foundation Topics” section of each chapter, the “Exam Preparation Tasks” section lists a series of study activities that you should do at the end of the chapter. Each chapter includes the activities that make the most sense for studying the topics in that chapter:
 - **Review All the Key Topics:** The Key Topic icon appears next to the most important items in the “Foundation Topics” section of the chapter. The “Review All the Key Topics” section lists the key topics from the chapter, along with their page numbers. Although the contents of the entire chapter could be on the exam, you should definitely know the information listed in each key topic, so you should review these.
 - **Complete Tables and Lists from Memory:** To help you memorize some lists of facts, many of the more important lists and tables from the chapter are included in a document on the book’s website. This document lists only partial information, allowing you to complete the table or list.
 - **Define Key Terms:** Although the exam may be unlikely to ask a question such as “Define this term,” the VCP-NV exam does require that you learn and know a lot of 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 the book.
 - **Web-based practice exam:** The companion website includes the Pearson Cert Practice Test engine that allows you to take practice exam questions. Use these to prepare with a sample exam and to pinpoint topics where you need more study.

How to Use This Book

The book is organized by chapters that cover a topic that I believe is needed to fully understand NSX. Some chapters should be read sequentially, such as Chapters 4, 5, and 6, while other chapters can be read in any order, such as Chapters 15 and 18. Be aware that I do make references throughout the book to previously covered chapters.

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

- **Chapter 1, “Introduction to VMware NSX:”** This chapter covers some of the history behind the data center network infrastructure, the challenges (Ethernet, IP, and security) that must be designed for, and how VMware NSX attempts to handle these challenges by eliminating them outright.
- **Chapter 2, “Network and VMware vSphere Requirements for NSX:”** This chapter covers the different types of data center infrastructure designs, the NSX underlay requirements, and the vSphere requirements for NSX.
- **Chapter 3, “NSX Architecture and NSX Manager:”** This chapter introduces the architecture of NSX and NSX Manager, describing its functions as well as how to deploy it.
- **Chapter 4, “VXLAN, NSX Controllers, and NSX Preparation:”** This chapter introduces VXLAN, one of the control planes of NSX, NSX Controllers, and how to prepare the vSphere environment for NSX.
- **Chapter 5, “NSX Switches:”** This chapter introduces logical switches, both global logical switches and universal logical switches.
- **Chapter 6, “Logical Switch Packet Walks:”** This chapter describes multiple step-by-step scenarios of the flow of virtual machine frames over logical switches.
- **Chapter 7, “Logical Router:”** This chapter introduces logical routers, including distributed logical routers and universal logical routers.
- **Chapter 8, “Logical Router Packet Walks:”** This chapter describes multiple step-by-step scenarios of the flow of virtual machine frames over logical routers.
- **Chapter 9, “NSX Edge Services Gateway:”** This chapter introduces the NSX Edge Services Gateway, describes its characteristics, and lists the features it supports.

- **Chapter 10, “Layer 2 Extensions:”** This chapter explains the ways in which NSX allows for a broadcast domain to be extended between a logical switch and a VLAN.
- **Chapter 11, “Layer 3 Connectivity Between Virtual and Physical Networks:”** This chapter explains how traffic between a virtual machine and a physical entity can take place when the virtual machine is connected to a logical switch.
- **Chapter 12, “Routing Protocols:”** This chapter describes the routing protocols supported by NSX: OSPF, BGP, and ISIS.
- **Chapter 13, “NSX Edge VPN Services:”** This chapter explains the virtual private network features supported by the NSX Edge.
- **Chapter 14, “NSX Edge Network Services and Security:”** This chapter explains the NSX Edge features of Network Address Translation, load balancer, and logical firewall.
- **Chapter 15, “Distributed Logical Firewall:”** This chapter introduces the distributed logical firewall (as well as the universal logical firewall), integration with LDAP/AD, and SpoofGuard.
- **Chapter 16, “Security Services:”** This chapter covers Security Composer, its components (security groups, security services), and the types of security services that can be offered by NSX.
- **Chapter 17, “Additional NSX Features:”** This chapter covers Layer 7 and Application security services, and troubleshooting tools native to NSX such as VMware Data Security, Activity Monitoring, and Traceflow.
- **Chapter 18, “NSX Automation:”** This chapter introduces RESTful APIs and how NSX APIs are used to create various NSX objects. There is a discussion of integration between NSX and vRealize Automation.
- **Chapter 19, “Upgrade to NSX for vSphere 6.2:”** This chapter covers how to upgrade a vCloud network and security or pre-NSX 6.2 installation to NSX 6.2.
- **Chapter 20, “Final Preparation:”** This chapter identifies tools for final exam preparation and helps you develop an effective study plan. It contains tips on how to best use the web-based material to study.

Certification Exam and This Preparation Guide

As mentioned earlier, this book is written in a way that best helps you understand NSX, which doesn't always make it clear as to which blueprint objectives are being covered in a particular chapter. Some objectives are covered over multiple chapters. Table I-1 lists the VCP6-NV Exam Blueprint Objectives and the chapters in the book that covers them.

Table I-1 VCP6-NV Exam Topics and Chapter References

Exam Section/Objective	Chapter Where Covered
Section 1—Understand VMware NSX Technology and Architecture	
Objective 1.1—Compare and Contrast the Benefits of a VMware NSX Implementation	Chapters 1, 15
Objective 1.2—Understand VMware NSX Architecture	Chapter 3
Objective 1.3—Differentiate Physical and Virtual Network Technologies	Chapters 2, 15
Objective 1.4—Understand VMware NSX Integration with Third-Party Products and Services	Chapter 10
Objective 1.5—Understand VMware NSX Integration with vRealize Automation (vRA)	Chapter 18
Section 2—Understand VMware NSX Physical Infrastructure Requirements	
Objective 2.1—Compare and Contrast the Benefits of Running VMware NSX on Physical Network Fabrics	Chapter 2
Objective 2.2—Determine Physical Infrastructure Requirements for a VMware NSX Implementation	Chapter 2
Section 3—Configure and Manage vSphere Networking	
Objective 3.1—Configure and Manage vSphere Distributed Switches (vDS)	Chapter 2
Objective 3.2—Configure and Manage vDS Policies	Chapter 2
Section 4—Install and Upgrade VMware NSX	
Objective 4.1—Configure Environment for Network Virtualization	Chapter 2
Objective 4.2—Deploy VMware NSX Components	Chapters 3, 4
Objective 4.3—Upgrade Existing vCNS/NSX Implementation	Chapter 19
Objective 4.4—Expand Transport Zone to Include New Cluster(s)	Chapter 4

Exam Section/Objective	Chapter Where Covered
Section 5—Configure VMware NSX Virtual Networks	
Objective 5.1—Create and Administer Logical Switches	Chapters 5, 6
Objective 5.2—Configure VXLAN	Chapters 4, 5
Objective 5.3—Configure and Manage Layer 2 Bridging	Chapter 10
Objective 5.4—Configure and Manage Logical Routers	Chapters 7, 8, 9, 11, 12
Section 6—Configure and Manage NSX Network Services	
Objective 6.1—Configure and Manage Logical Load Balancing	Chapter 14
Objective 6.2—Configure and Manage Logical Virtual Private Networks (VPN)	Chapters 4, 10, 13
Objective 6.3—Configure and Manage DHCP/DNS/NAT	Chapter 14
Objective 6.4—Configure and Manage Edge Services High Availability	Chapter 9
Section 7—Configure and Administer Network Security	
Objective 7.1—Configure and Administer Logical Firewall Services	Chapter 14
Objective 7.2—Configure Distributed Firewall Services	Chapter 15
Objective 7.3—Configure and Manage Service Composer	Chapter 16
Section 8—Deploy a Cross-vCenter NSX Environment	
Objective 8.1—Differentiate Single and Cross-vCenter NSX Deployments	Chapters 3, 5, 7, 15
Objective 8.2—Determine Cross-vCenter Requirements and Configurations	Chapters 3, 4, 5, 7
Section 9—Perform Operations Tasks in a VMware NSX Environment	
Objective 9.1—Configure Roles, Permissions, and Scopes	Chapter 17
Objective 9.2—Understand NSX Automation	Chapter 18
Objective 9.3—Monitor a VMware NSX Implementation	Chapter 17
Objective 9.4—Perform Auditing and Compliance	Chapter 17
Objective 9.5—Administer Logging	Chapters 12, 13, 14, 15
Objective 9.6—Backup and Recover Configurations	Coming Soon. Check Appendix B on the book website

Exam Section/Objective	Chapter Where Covered
Section 10—Troubleshoot a VMware Network Virtualization Implementation	
Objective 10.1—Compare and Contrast Tools Available for Troubleshooting	Chapter 17
Objective 10.2—Troubleshoot Common NSX Installation/Configuration Issues	Coming Soon. Check Appendix B on the book website
Objective 10.3—Troubleshoot Common NSX Component Issues	Chapters 3, 4, 7, 10
Objective 10.4—Troubleshoot Common Connectivity Issues	Chapter 4
Objective 10.5—Troubleshoot Common vSphere Networking Issues	Coming Soon. Check Appendix B on the book website

Book Content Updates

Since VMware occasionally updates exam topics without notice, VMware Press might post additional preparatory content on the web page associated with this book at <http://www.pearsonitcertification.com/title/9780789754806>. It is a good idea to check the website a couple of weeks before taking your exam, to review any updated content that might be posted online. We also recommend that you periodically check back to this page on the Pearson IT Certification website to view any errata or supporting book files that may be available.

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 author 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:

- Step 1.** Go to www.pearsonITcertification.com/register and log in or create a new account.
- Step 2.** Enter the ISBN: **9780789754806**
- Step 3.** Answer the challenge question as proof of purchase.
- Step 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 preceding steps, 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 displays 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 4.5 Client
- Pentium-class 1 GHz 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 as 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:

- Step 1.** Download the exam practice test engine from the companion site.
- Step 2.** Respond to Windows' prompts as with any typical software installation process.

The installation process gives 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:

- Step 1.** Start the Pearson IT Certification Practice Test software from the Windows Start menu or from your desktop shortcut icon.
- Step 2.** To activate and download the exam associated with this book, from the My Products or Tools tab, click the **Activate Exam** button.

- Step 3.** At the next screen, enter the activation key from the paper inside the cardboard sleeve in the back of the book. Once entered, click the **Activate** button.
- Step 4.** The activation process downloads 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 ensures 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% off the Premium Edition eBook and Practice Tests edition of this title. See the coupon code included with the cardboard sleeve for information on how to purchase the Premium Edition.

VXLAN, NSX Controllers, and NSX Preparation

Deploying NSX Manager and attaching it to vCenter are the first steps in allowing you to deploy your software defined network. Your goal is to have logical switches, distributed logical routers, and create and enforce security policies with the distributed firewall and service composer.

Before you can reach your goal, you need to deploy our NSX Controllers and tell NSX Manager which ESXi hosts will be part of the NSX domain. The steps to tell NSX Manager which ESXi hosts will be part of the NSX Domain are

- Install NSX modules.
- Configure VXLAN networking in each ESXi host.
- Create VNI pools.
- Create transport zones.

This chapter covers all the steps needed to prepare your NSX domain. The chapter begins with a proper introduction of what VXLAN is.

Do I Know This Already?

The “Do I Know This Already?” quiz allows you to assess whether you should read this entire chapter or simply jump to the “Exam Preparation Tasks” section for review. If you are in doubt, read the entire chapter. Table 4-1 outlines the major headings in this chapter and the 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 Headings and Questions

Foundation Topic Section	Questions Covered in This Section
VXLAN	1-2
NSX Controllers	3-4
IP Pools	5
Host Preparation	6-7

Foundation Topic Section	Questions Covered in This Section
Host Configuration	8-9
VNI Pools, Multicast Pools, and Transport Zones	10

1. What is the source Layer 4 port number of a VXLAN frame?
 - a. It is statically configured to TCP 8472.
 - b. It is statically configured to UDP 8472.
 - c. It is randomly generated by the VTEP.
 - d. It is derived from the encapsulated frame.

2. At least how many bytes does the VXLAN encapsulation add to the encapsulated frame?
 - a. 50
 - b. 100
 - c. 1500
 - d. 9000

3. How many NSX universal controllers are required to be deployed in a production NSX environment?
 - a. 1
 - b. 2
 - c. 3
 - d. 4

4. What NSX entity is responsible for slicing the distributed logical router?
 - a. The NSX Manager
 - b. The distributed router control virtual machine
 - c. The API provider NSX Controller Master
 - d. The Layer 3 NSX Controller Master

5. What are two use cases of IP pools by NSX Manager? (Choose two.)
 - a. To assign IPs to virtual machines in the virtual network.
 - b. To assign the default gateway for VTEPs.
 - c. To assign IPs to NSX Manager.
 - d. To assign IPs to NSX Controllers.

6. Which of the following is an action that takes place during host preparation?
 - a. The NSX Manager tells vCenter to add the selected hosts in the NSX host clusters.
 - b. The NSX Manager installs NSX modules on the ESXi hosts.
 - c. vCenter adds the VXLAN VMkernel port to the ESXi hosts.
 - d. The NSX Controller Master uploads the NSX configuration data to the ESXi hosts.

7. Which NSX feature does not require logical networking preparation to be completed before it can be used?
 - a. VXLAN
 - b. Logical switches
 - c. Distributed firewall
 - d. Distributed logical routers

8. How many vDS switches does NSX Manager support in a single host cluster?
 - a. 1
 - b. 2
 - c. 32
 - d. 128

9. During host configuration you select a VMKNic teaming policy of enhanced LACP. How many VTEPs does NSX Manager create per ESXi host?
 - a. 1
 - b. 2
 - c. As many dvUplinks as are configured on the vDS
 - d. As many VMNICs as are installed on the ESXi hosts

10. How many universal transport zones are supported in a cross vCenter NSX domain?
 - a. 1
 - b. 1 per NSX Manager in the cross vCenter NSX domain
 - c. Up to the number of VNIs in the segment ID pool
 - d. 1 per NSX universal controller

Foundation Topics

VXLAN Introduction

Multitier applications have long been designed to use separate Ethernet broadcast domains or virtual local area networks (VLANs) to separate tiers within the application. In a vSphere environment, the number of multitier applications can be quite large, which eats up the number of available VLANs and makes it challenging to scale the virtual environment. For example, if a client has 100 four-tier applications, the client may need 400 separate Ethernet broadcast domains or VLANs to support these applications. Now multiply that by 10 clients. You are basically hitting the limit on how many Ethernet broadcast domains you can support using VLANs. As the virtual machines (VMs) for these applications are distributed among multiple vSphere clusters or even different data centers, the Ethernet broadcast domains must be spanned across the physical network, necessitating the configuration of Spanning Tree Protocol to prevent Ethernet loops.

Virtual Extensible LAN (VXLAN) addresses the Layer 2 scaling challenges in today's data centers by natively allowing for the transparent spanning of millions of distinct Ethernet broadcast domains over any IP physical network or IP transport, reducing VLAN sprawl and thus eliminating the need to enable Ethernet loop-preventing solutions such as Spanning Tree.

VXLAN

Key Topic

VXLAN is an open standard supported by many of the key data center technology companies, such as VMware. VXLAN is a Layer 2 encapsulation technology that substitutes the usage of VLAN numbers to label Ethernet broadcast domains with VXLAN numbers. A traditional Ethernet switch can support up to 2^{12} (4096) Ethernet broadcast domains or VLAN numbers. VXLAN supports 2^{24} Ethernet broadcast domains or VXLAN numbers. That is 16,777,216 Ethernet broadcast domains. A VXLAN number ID is referred to as VNI. There is a one-to-one relationship between an Ethernet broadcast domain and a VNI. A single Ethernet broadcast domain can't have more than one VNI. Two distinct Ethernet broadcast domains can't have the same VNI.

Figure 4-1 shows a traditional design with two ESXi hosts in different racks, each one with a powered on VM. If both VMs need to be in the same Ethernet broadcast domain, the broadcast domain must be spanned, or extended, across all the Ethernet switches shown in the diagram. This makes it necessary for either the Spanning Tree Protocol to be configured in all the Ethernet switches or a more expensive loop-preventing solution such as Transparent Interconnection of Lots of Links

(TRILL) to be deployed. With VXLAN deployed, the ESXi hosts can encapsulate the VM traffic in a VXLAN frame and send it over the physical network, which can be IP-based rather than Ethernet-based, thus removing the need to configure Spanning Tree or deploy solutions such as TRILL.

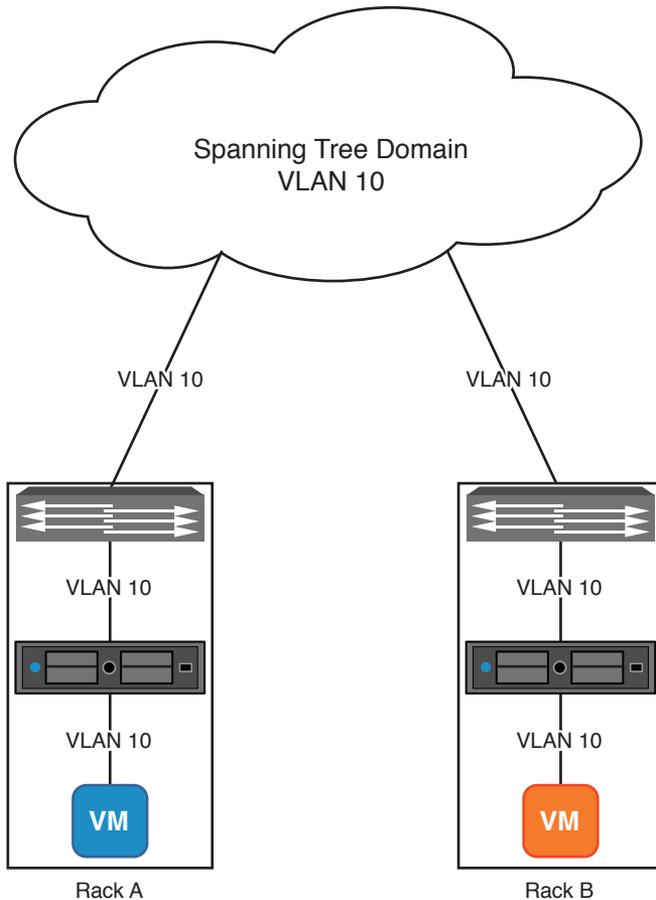


Figure 4-1 Spanning broadcast domain across multiple ESXi racks

Traditionally, any network technology that encapsulates traffic the way VXLAN does is called a *tunnel*. A tunnel hides the original frame's network information from the IP physical network. A good example of a tunnel is Genetic Routing Encapsulation (GRE), which hides Layer 3 and Layer 4 information from IP network devices, although GRE could be set up to also hide Layer 2 information. VXLAN tunnels hide Layer 2, Layer 3, and Layer 4 information. It is possible to deploy a new IP network topology by just using tunnels, without having to do major reconfiguration of the IP physical network. Such a network topology is called an *overlay*, whereas the

IP physical network that switches and routes the tunnels that make up the overlay is called the *underlay*.

Just as GRE requires two devices to create and terminate the tunnel, VXLAN requires two devices to create and terminate VXLAN tunnels. A device that can create or terminate the VXLAN tunnel is called the VXLAN Tunnel Endpoint (VTEP). NSX enables ESXi hosts to have VTEPs. A VTEP performs these two roles:

- Receive Layer 2 traffic from a source, such as a VM, in an Ethernet broadcast domain, encapsulating it within a VXLAN frame and sending it to the destination VTEP.
- Receive the VXLAN frame, stripping the encapsulation to reveal the encapsulated Ethernet frame, and forwarding the frame toward the destination included in the encapsulated Ethernet frame.

Figure 4-2 shows an Ethernet frame from a VM encapsulated in a VXLAN frame. The source VTEP of the VXLAN frame is a VMkernel port in the ESXi host. You can see the encapsulated Ethernet frame, or original frame, and the new header, thus creating the VXLAN overlay.

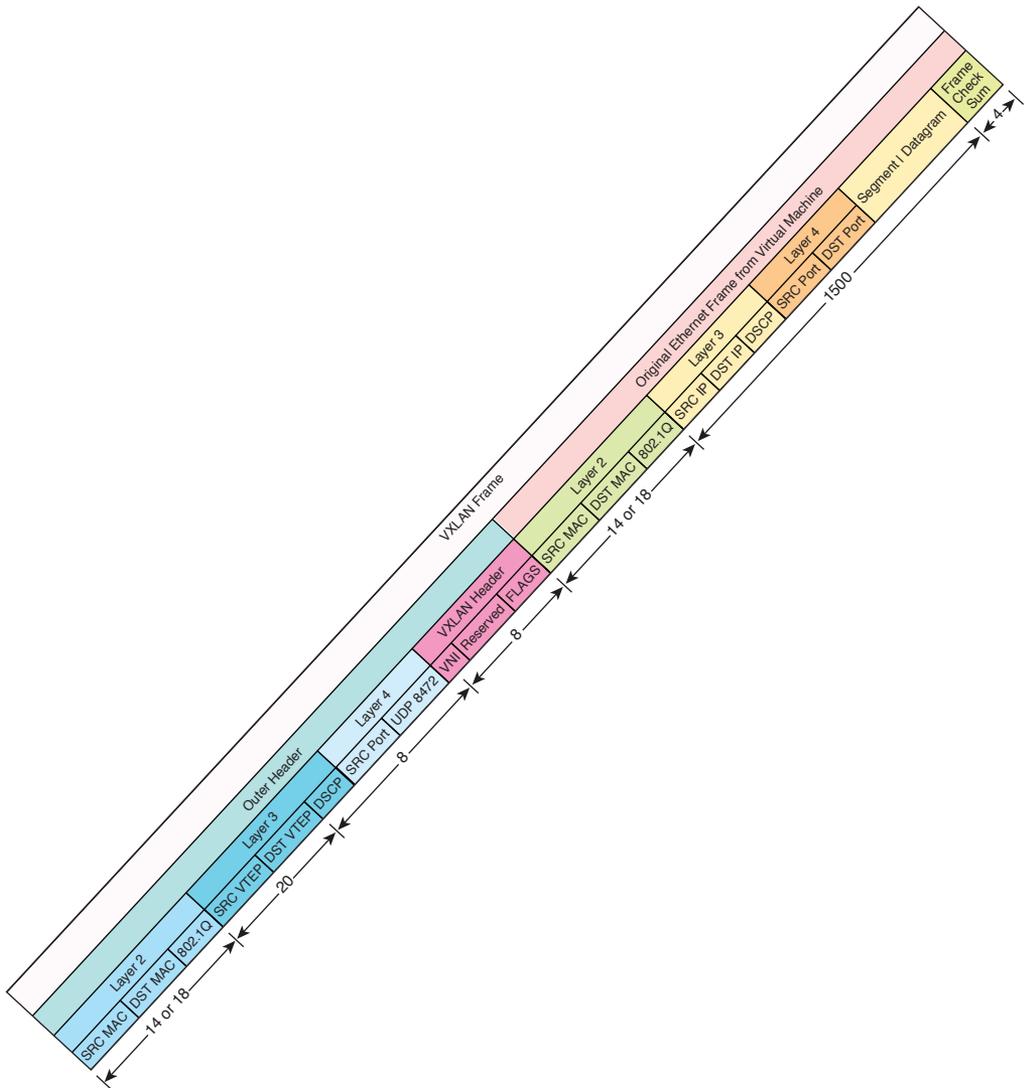


Figure 4-2 VXLAN encapsulation

Key Topic

The VXLAN frame contains the following components:

- New Layer 2 header distinct from the encapsulated Layer 2 header. This header has new source and destination MAC addresses and a new 802.1Q field.

- This header is 14 bytes long if not using 802.1Q.
- If using 802.1Q, this header is 18 bytes long.
- Class of Service (CoS) markings copied from the original frame's 802.1Q field in the Layer 2 header, if any.
- New Layer 3 header distinct from the encapsulated Layer 3 header. This header has new source and destination IP addresses, including
 - The source and destination IPs are VTEPs. In some cases the destination IP could be a multicast group (we expand further on this during Chapter 5, "NSX Switches").
 - This header is 20 bytes long, with no extensions.
 - DSCP markings, if any, are copied from the encapsulated DSCP files in the Layer 3 header.
 - The *do not fragment (DF) bit* is set to 1.
- New Layer 4 header distinct from the encapsulated Layer 4 header. This header is always UDP.
 - This header is be 8 bytes long.
 - NSX VTEPs use a destination port of 8472. As of April 2013, the standard VXLAN UDP port is 4789. NSX supports changing the UDP port number via the NSX APIs. We cover the NSX APIs in Chapter 18, "NSX Automation."
 - The source port is derived from the encapsulated Layer 4 header.
- New VXLAN header.
 - This header is 8 bytes long.
 - 3 bytes are dedicated for VNI labeling of the tunnel.
 - 4 bytes are reserved for future use.
 - 1 byte is dedicated for flags.

To aggregate a few things stated in the preceding content about VXLAN: Any QoS markings, such as DSCP and CoS from the VM Ethernet frame being encapsulated, are copied to the VXLAN frame, and the destination UDP port of the VXLAN frame is derived from the header information from the encapsulated frame. For this to work, VXLAN has to support virtual guest VLAN Tagging (VGT). Without VGT support, the VM's guest OS couldn't do QoS markings. If the encapsulated frame does not have any QoS markings, none would be copied to the VXLAN frame; however, there is nothing stopping you from adding QoS markings directly to the VXLAN frame.

Then there is the part where the VXLAN frame traverses the physical network, called the *VXLAN underlay* or simply *underlay*. The underlay uses VLANs. It is almost certain that the VXLAN underlay will place the VXLAN frames in their own Ethernet broadcast domain, thus requiring its own VLAN. The VLAN used by the underlay for VXLAN frames is referred to as the *VXLAN VLAN*. If the ESXi host with the source VTEP is connected to a physical switch via a trunk port, the ESXi host could be configured to add a VLAN tag, 802.1Q, to the VXLAN frame or send the VXLAN frame without a VLAN tag, in which case the physical switch's trunk needs to be configured with a native VLAN.

**Key
Topic**

All this means that VXLAN encapsulation adds 50+ bytes to the original frame from the VM. The 50+ bytes come from the following addition:

Original Layer 2 (minus Frame Check Sum) + VXLAN Header + Outer Layer 4 Header + Outer Layer 3 Header

Without Original Frame 802.1Q field: $14 + 8 + 8 + 20 = 50$

With Original Frame 802.1Q field: $18 + 8 + 8 + 20 = 54$

VMware recommends that the underlay for VXLAN support jumbo frames with an MTU of at least 1600 bytes to support VMs sending frames with the standard 1500 bytes MTU. This includes any routers that are part of the underlay; otherwise, the routers will discard the VXLAN frames when they realize they can't fragment the VXLAN frames with more than 1500 bytes payload. ESXi hosts with VTEPs also configure the VXLAN tunnel with the Do Not Fragment bit, DF, in the IP header of the VXLAN overlay to 1.

Figure 4-3 shows two VMs on the same Ethernet broadcast domain communicating with each other. The two VMs are connected to the same VNI, and the two ESXi hosts have the VTEPs. This diagram does not show the nuances of how the VTEPs know about each other's existence or how they determine where to forward the VXLAN frame. Chapter 5 covers these details in more depth.

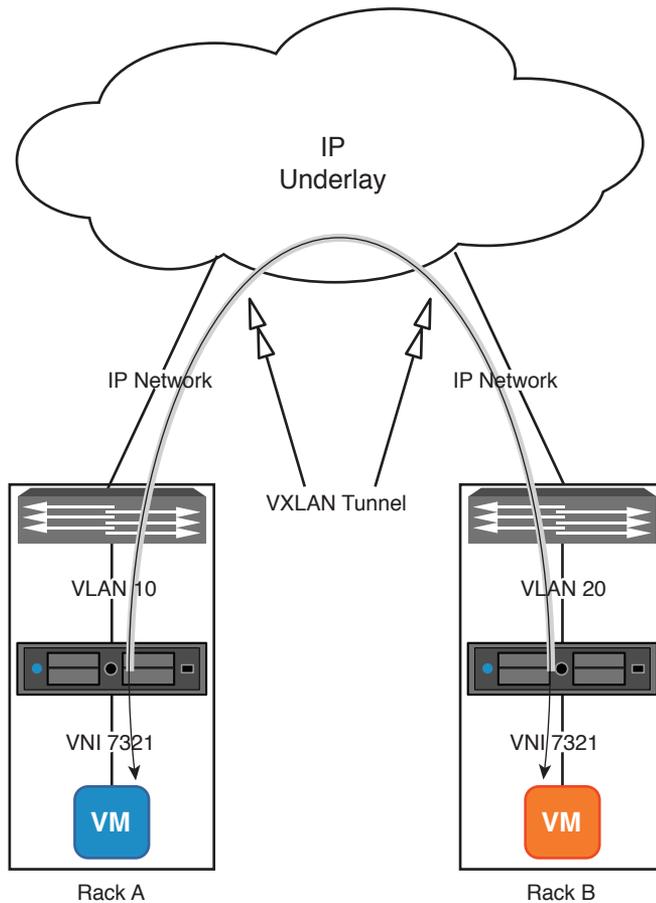


Figure 4-3 Virtual machine communication via VXLAN

NSX Controllers

The NSX Controllers are responsible for most of the control plane. The NSX Controllers handle the Layer 2 control plane for the logical switches, and together with the distributed logical router control virtual machine, the NSX Controllers handle the Layer 3 control plane. We review the role of the Layer 3 control plane and the distributed logical router control virtual machine in Chapter 7, “Logical Router.”

Key Topic

For Layer 2, the NSX Controllers have the principal copy of three tables per logical switch, which are used to facilitate control plane decisions by the ESXi host. The three tables are

- VTEP table:** Principal table that lists all VTEPs that have at least one VM connected to the logical switch. There is one VTEP table per logical switch.

- **MAC table:** Principal table containing the MAC addresses for VMs connected to logical switches as well as any physical end system in the same broadcast domain as the logical switch.
- **ARP table:** Principal table containing the ARP entries for VMs connected to logical switches as well as any physical end system in the same broadcast domain as the logical switch.

For Layer 3, the NSX Controllers have the routing table for each distributed logical router as well as the list of all hosts running a copy of each distributed logical router.



NSX Controllers do not play any role in security, such as the distributed firewall, nor do they provide control plane services to the NSX Edge Service Gateway.

Deploying NSX Controllers

The NSX Controllers are virtual appliances deployed by the NSX Manager. The NSX Controllers must be deployed in the same vCenter associated with NSX Manager. In our examples from the figures, that would be vCenter-A if the NSX Controller is from NSXMGR-A. At least one NSX Controller must be deployed before logical switches and distributed logical routers can be deployed in an NSX Manager with a Standalone role.

Deploying NSX Controllers might be the most infuriating thing about setting up an NSX environment. I restate some of this in context a little later, but in short if NSX Manager can't establish communication with the NSX Controller after it is deployed, it has the NSX Controller appliance deleted. The process of deploying the NSX Controller can take a few minutes or more, depending on the available resources in the ESXi host where you deploy it and the datastore. If the NSX Controller deployment fails for whatever reason, NSX Manager doesn't attempt to deploy a new one. You can view the NSX Manager's log to find the reason to why the deployment failed and then try again. But you won't be doing much networking with NSX until you get at least one NSX Controller deployed.

Let's now cover the steps to deploying the NSX Controllers, but I wanted to point out this *little* annoyance first. A single NSX Controller is all that is needed to deploy logical switches and distributed logical routers; however for redundancy and failover capability, VMware supports only production environments with three NSX Controllers per standalone NSX Manager. The NSX Controllers can be deployed in separate ESXi clusters as long as

- Each NSX Controller has IP connectivity with NSX Manager, over TCP port 443.
- Each NSX Controller has IP connectivity with each other, over TCP port 443.

- Each NSX Controller has IP connectivity with the management VMkernel port of ESXi hosts that will be part of the NSX domain over TCP port 1234.

The following steps guide you in how to deploy NSX Controllers via the vSphere Web Client. You can also deploy NSX Controllers using the NSX APIs.

Key Topic

You must be an NSX administrator or enterprise administrator to be allowed to deploy NSX Controllers. We cover Role Based Access Control (RBAC), in Chapter 17, “Additional NSX Features.”

- Step 1.** From the Networking and Security home page, select the **Installation** field.
- Step 2.** Select the **Management** tab.
- Step 3.** In the NSX Controller Nodes section click the green + icon, as shown in Figure 4-4.

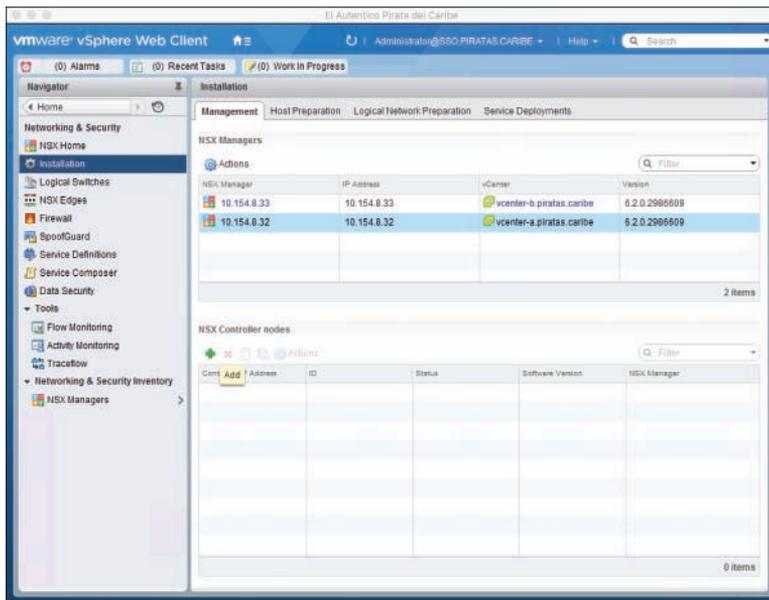


Figure 4-4 Add NSX Controller

- Step 4.** In the NSX Controller Wizard, select the NSX Manager that would deploy the NSX Controller.

The vSphere Web Client supports multiple vCenters, and thus multiple NSX Managers.

- Step 5.** Select the data center on which you are adding the NSX Controller.

- Step 6.** Select the datastore where the NSX Controller will be deployed.
- Step 7.** Select the ESXi cluster or resource pool where the NSX Controller will be deployed.
- Step 8.** Optionally, select the ESXi host and folder where the NSX Controller will be deployed. If the ESXi cluster selected in step 5 is configured with DRS with automatic virtual machine placement, you can skip the host selection.
- Step 9.** Select the standard portgroup or vDS portgroup where the NSX Controller's management interface will be connected. All communication from the NSX Controller to NSX Manager, other NSX Controllers, and the ESXi hosts will take place over this connection.
- Step 10.** Select the pool of IPs from which the NSX Controller will be assigned an IP by the NSX Manager.
- If no IP pool exists, you have the option to create one now. We review the creation of an IP pool later in this chapter.
- Step 11.** If this is your first NSX Controller, you need to provide a CLI password, as shown in Figure 4-5. You do not need to provide a password for subsequent NSX Controllers as the NSX Manager automatically assigns them all the same password from the first deployed NSX Controller. The default username of the CLI prompt is **admin**.

The screenshot shows a dialog box titled "Add Controller" with a help icon in the top right corner. The dialog contains the following fields and controls:

- NSX Manager:** A text box containing "10.154.8.32" with a dropdown arrow on the right.
- Datacenter:** A dropdown menu showing "Santo Domingo".
- Cluster/Resource Pool:** A dropdown menu showing "MGT-A1".
- Datastore:** A dropdown menu showing "MGT_A1".
- Host:** An empty dropdown menu.
- Folder:** An empty dropdown menu.
- Connected To:** A text box containing "MGT_A1-VMMGT" with "Change" and "Remove" buttons to its right.
- IP Pool:** A text box containing "Controllers" with a "Select" button to its right.
- Password:** A text box containing "*****".
- Confirm password:** A text box containing "*****".

At the bottom of the dialog are "OK" and "Cancel" buttons.

Figure 4-5 Adding first NSX Controller

When NSX Controllers get deployed, they automatically form a cluster among themselves. The first NSX Controller needs to be deployed and have joined the NSX Controller cluster by itself before the other NSX Controllers can be deployed. If you try to deploy a second NSX Controller before the first one is deployed, you get an error message.

When NSX Manager receives the request to deploy an NSX Controller from vCenter, who got it from the vSphere Web Client, or when NSX Manager receives the request via the NSX APIs, the following workflow takes place:

- Step 1.** NSX Manager gives the NSX Controller off to vCenter to deploy, per your configurations during the Add NSX Controller Wizard. This includes
- The data center, datastore, and cluster/resource pool to place the NSX Controller
 - The ovf import specifications, which includes the IP from the IP pool, the private and public certificates for communication back to NSX Manager, and the cluster IP, which is the IP of the first NSX Controller
 - A request to place the NSX Controller in the Automatic Startup of the ESXi host

- Step 2.** vCenter deploys the NSX Controller, powers it on, and then tells NSX Manager the Controllers are powered on.
- Step 3.** NSX Manager makes contact with the NSX Controller.

If NSX Manager cannot establish an IP connection to the NSX Controller to complete its configuration, the NSX Manager has vCenter power off the NSX Controller and delete it.

Verifying NSX Controllers

You can verify the status of the NSX Controller installation by selecting the Installation view from the Networking and Security page, as shown in Figure 4-6.

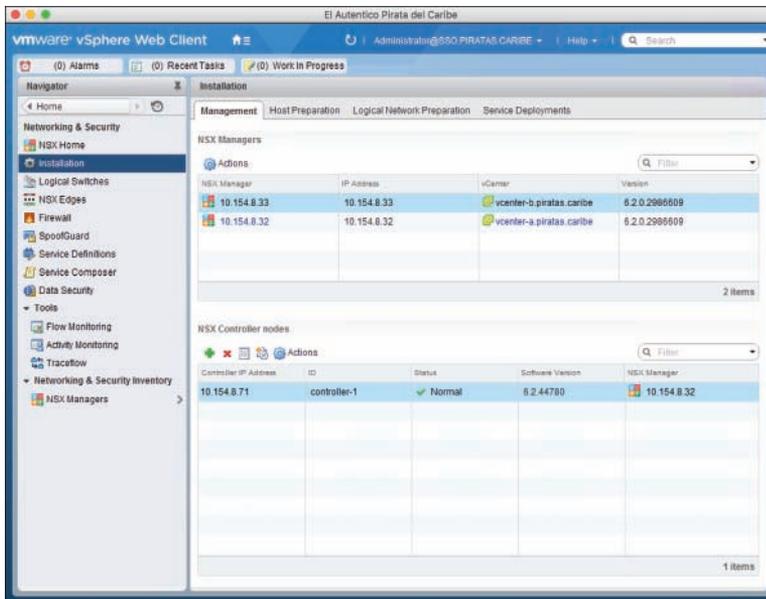


Figure 4-6 An NSX Controller successfully deployed

In this view you can verify the following:

- **Controller IP Address:** The IP address of the NSX Controller. This is one of the IP addresses from the IP pool. Clicking on the controller IP address brings up information about the ESXi host and datastore the NSX Controller is in, as shown in Figure 4-7.
- **ID:** The ID of the NSX Controller. This ID is assigned by the NSX Manager that is communicating with the NSX Controller and has no impact on the role or function of the NSX Controller.

- **Status:** This is the status of the NSX Controller. The statuses we care about are Deploying and Normal.
 - Deploying is self-explanatory.
 - Disconnected means the NSX Manager lost connectivity to the NSX Controller.
 - Normal means the NSX Controller is powered up and NSX Manager has normal operation communication with it.
- **Software Version:** The version of NSX software running in the NSX Controller. The version number is independent of the NSX Manager's version.
- **NSX Manager:** The NSX Manager that is communicating with this NSX Controller. Yes, this is here because a single vSphere Web Client supports multiple vCenters and thus Multiple NSX Managers. If one of the NSX Managers is participating in cross vCenter NSX, a sixth column becomes visible:
- **Managed By:** The IP of the Primary NSX Manager that deployed the NSX Controller.

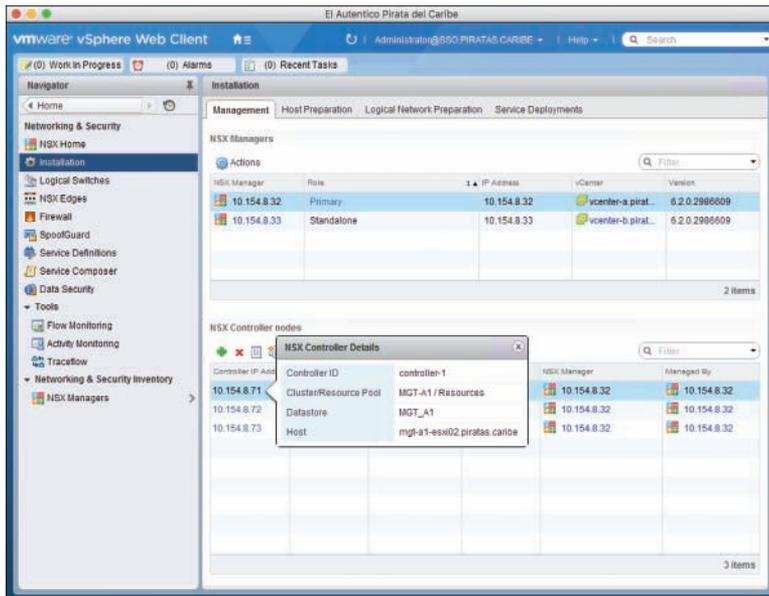


Figure 4-7 NSX Controller details

If you assign a role of Primary to an NSX Manager, the NSX Manager's three NSX Controllers become NSX universal controllers. NSX universal controllers can communicate with Secondary NSX Managers in the same cross vCenter NSX domain as

well as Secondary NSX Manager's participating entities such as ESXi hosts. Before you add Secondary NSX Managers, their existing NSX Controllers, if any, must be deleted.

You can also verify the deployment of the NSX Controllers by viewing the NSX Controller virtual machine in the Host and Clusters or VM and Templates view. The NSX Controller is deployed using the name **NSX_Controller_** followed by the NSX Controller's UUID. Figure 4-8 shows the first NSX Controller in the Host and Clusters view. Notice in Figure 4-8 the number of vCPUs, memory, memory reservation, and HDD configured in the NSX Controller.

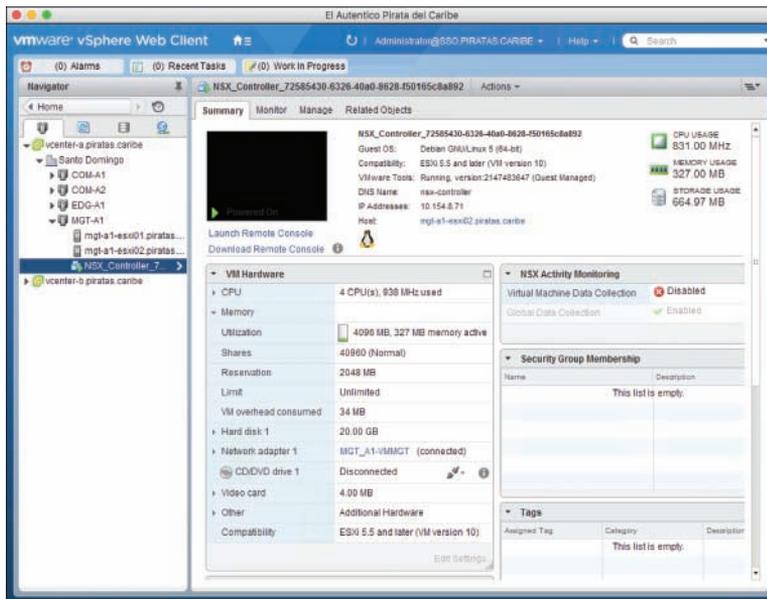


Figure 4-8 NSX Controller's virtual machine Summary view

Each NSX Controller gets deployed with these settings:

- 4 vCPUs
- 4 GB vRAM, with 2 GB reservation
- 20 GB HDD
- 1 vNIC
- VM hardware version 10



VMware does not support changing the hardware settings of the NSX Controllers.

**Key
Topic**

If the NSX Manager is participating in a Secondary role in cross vCenter NSX, the NSX Manager will not have any NSX Controllers of its own. Instead the Secondary NSX Managers create a logical connection to the existing NSX universal controllers from the Primary NSX Manager in the same cross vCenter NSX domain.

Creating an NSX Controller Cluster

When more than one NSX Controller is deployed, the NSX Controllers automatically form a cluster. They know how to find each other because NSX Manager makes them aware of each other's presence. To verify that the NSX Controller has joined the cluster successfully, connect to the NSX Controllers via SSH or console using the username of **admin** and the password you configured during the first NSX Controller deployment. Once logged in the NSX Controller, issue the CLI command **show control-cluster status** to view the NSX Controller's cluster status. You need to do this for each NSX Controller to verify its cluster status. Figure 4-9 shows the output of the command for an NSX Controller that has joined the cluster successfully.

TIP You can use the Tab key to autocomplete CLI commands in NSX Manager and the NSX Controllers.

```

nsx-controller # show control-cluster status
-----
Type                Status                Since
-----
Join status:        Join complete         10/21 04:29:19
Majority status:    Connected to cluster majority 10/21 04:58:16
Restart status:     This controller can be safely restarted 10/21 04:58:05
Cluster ID:         b3f7d956-e177-43b1-aeb8-e5ae8532e67e
Node UUID:          b3f7d956-e177-43b1-aeb8-e5ae8532e67e

Role                Configured status    Active status
-----
api_provider        enabled              activated
persistence_server enabled              activated
switch_manager      enabled              activated
logical_manager     enabled              activated
directory_server    enabled              activated
nsx-controller #

```

Figure 4-9 Output of **show control-cluster status**

Figure 4-9 depicts the following cluster messages:

- **Join status:** Join complete. This message indicates this NSX Controller has joined the cluster.
- **Majority status:** Connected to cluster majority. This message indicates that this NSX Controller can see the majority of NSX Controllers (counting itself). If this NSX Controller were not connected to the cluster majority, it would

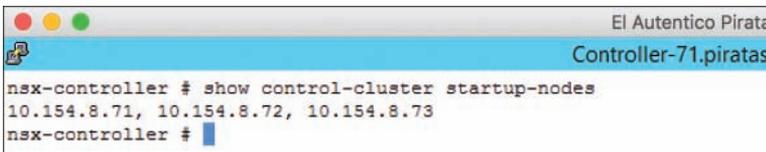
remove itself from participation in the control plane until it can see the majority of NSX Controllers again.

- **Restart status:** This controller can be safely restarted.
- **Cluster ID:** {UUID}. This is the Universal Unique ID of the cluster.
- **Node UUID:** {UUID}. This is the Universal Unique ID of this NSX Controller.

Key Topic

The clustering algorithm used by the NSX Controllers depends on each NSX Controller having IP communication with a majority of the NSX Controllers, counting itself. If the NSX Controller does not belong to the majority, or quorum, it removes itself from control plane participation. To avoid a split-brain situation where no NSX Controller is connected to the cluster majority and potentially each one removing itself from control plane participation, VMware requires that three of the NSX Controllers be deployed in production environments.

Figure 4-10 shows the output of the command **show control-cluster startup-nodes**, which shows the NSX Controllers that are known to be cluster members. All NSX Controllers should provide the same output. You can also issue the NSX Manager basic mode command **show controller list all** to list all the NSX Controllers the NSX Manager is communicating with plus their running status.



```

nsx-controller # show control-cluster startup-nodes
10.154.8.71, 10.154.8.72, 10.154.8.73
nsx-controller #
  
```

Figure 4-10 Output of **show control-cluster startup-nodes**

Additional CLI commands that could be used in the NSX Controllers to verify cluster functionality and availability are as follows:

- **show control-cluster roles:** Displays which NSX Controller is the master for different roles. We cover roles in the next section.
- **show control-cluster connections:** Displays the port number for the different roles and the number of established connections.
- **show control-cluster management-address:** Displays the IP used by the NSX Controller for management.

We review additional CLI commands in NSX Manager and NSX Controllers related to logical switches and distributed logical routers in Chapter 5 and Chapter 7.

NSX Controller Master and Recovery

When deploying multiple NSX Controllers, the control plane responsibilities for Layer 2 and Layer 3 are shared among all controllers. To determine which portions each NSX Controller handles, the NSX Controllers cluster elects an API provider, Layer 2 and Layer 3 NSX Controller Master. The masters are selected after the cluster is formed. The API provider master receives internal NSX API calls from NSX Manager. The Layer 2 NSX Controller Master assigns Layer 2 control plane responsibility on a per logical switch basis to each NSX Controller in the cluster, including the master. The Layer 3 NSX Controller Master assigns the Layer 3 forwarding table, on a per distributed logical router basis, to each NSX Controller in the cluster, including the master.

The process of assigning logical switches to different NSX Controllers and distributed logical routers to different NSX Controllers is called *slicing*. By doing slicing, the NSX Controller Master for Layer 2 and Layer 3 distributes the load of managing the control plane for logical switches and distributed routers among all the NSX Controllers. No two NSX Controllers share the Layer 2 control plane for a logical switch nor share the Layer 3 control plane for a distributed logical router. Slicing also makes the NSX Layer 2 and Layer 3 control planes more robust and tolerant of NSX Controller failures.

Once the master has assigned Layer 2 and Layer 3 control plane responsibilities, it tells all NSX Controllers about it so all NSX Controllers know what each NSX Controller is responsible for. This information is also used by the NSX Controllers in case the NSX Controller Master becomes unresponsive or fails.

If your NSX environment has only a single distributed logical router and three NSX Controllers, only one of the NSX Controllers would be responsible for the distributed logical router while the other two would serve as backups. No two NSX Controllers are responsible for the Layer 2 control plane of the same logical switch. No two NSX Controllers are responsible for the Layer 3 forwarding table of the same logical router.

When an NSX Controller goes down or becomes unresponsive, the data plane continues to operate; however, the Layer 2 NSX Controller Master splits among the surviving NSX Controllers Layer 2 control plane responsibilities for all the impacted logical switches. The Layer 3 NSX Controller Master splits among all the surviving NSX Controllers Layer 3 control plane responsibilities for all the affected distributed logical routers.

What if the NSX Controller that fails was the master? In this case, the surviving NSX Controllers elect a new master, and the new master then proceeds to recover the control plane of the affected logical switches and/or distributed logical routers. How does the new master determine which logical switches and/or distributed logical routers were affected and need to have their control plane responsibilities re-assigned? The new master uses the assignment information distributed to the cluster by the old master.

For Layer 2 control plane, the newly responsible NSX Controller queries the hosts in the transport zone so it can repopulate the logical switch's control plane information. We learn about transport zones later in this chapter. For Layer 3, the newly responsible NSX Controller queries the logical router control virtual machine. We learn about the logical router control virtual machine in Chapter 7.

IP Pools

IP pools are the only means to provide an IP address to the NSX Controllers. IP pools may also be used to provide an IP address to the ESXi hosts during NSX host preparations. We review NSX host preparation later in this chapter in the section "Host Preparation." IP pools are created by an NSX administrator and are managed by NSX Manager. Each NSX Manager manages its own set of IP pools. NSX Manager selects an IP from the IP pool whenever it needs one, such as when deploying an NSX Controller. If the entity using the IP from the IP pool is removed or deleted, NSX Manager places the IP back into the pool. The IPs in the IP pool should be unique in the entire IP network (both physical and virtual).

There are two ways to start the creation of an IP pool. The first method we mentioned during the deployment of the NSX Controllers. This option to create an IP pool is also available during NSX host preparation, which we discuss later in this chapter.

The second method involves the following steps:

- Step 1.** Select the **NSX Managers** field in the Networking and Security page.
- Step 2.** Select the NSX Manager you want to create an IP pool in.
- Step 3.** Select the **Manage** tab.
- Step 4.** Select the **Grouping Objects** button.
- Step 5.** Select **IP Pools**.
- Step 6.** Click the green + icon, as shown in Figure 4-11.

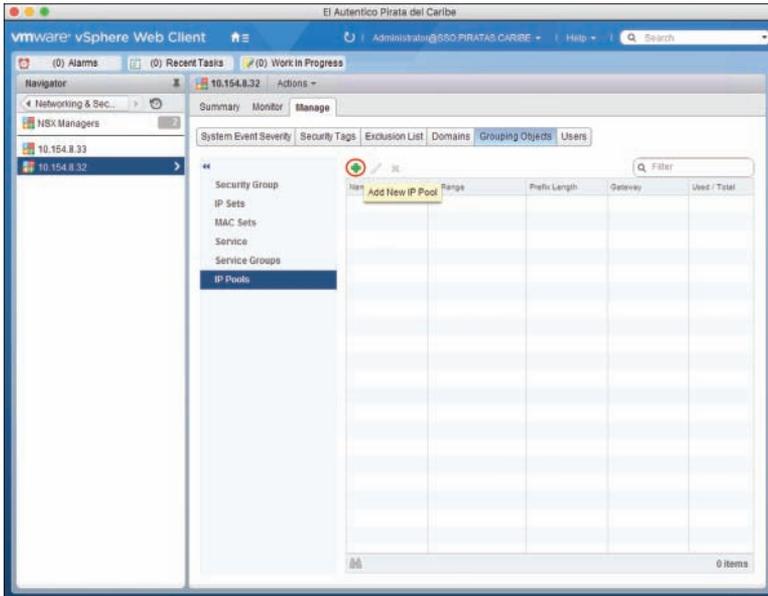


Figure 4-11 Create an IP pool

Regardless of how you choose to create an IP pool, the same IP Pool Wizard comes up, as shown in Figure 4-12.

Add Static IP Pool

Name: *

Gateway: *
A gateway can be any IPv4 or IPv6 address.

Prefix Length: *

Primary DNS:

Secondary DNS:

DNS Suffix:

Static IP Pool: *
for example 192.168.1.2-192.168.1.100 or
 abcd:87:87::10-abcd:87:87::20

Figure 4-12 IP Pool Wizard

In the IP Pool Wizard, populate the following information:

- Step 1.** Give the IP pool a unique name.
- Step 2.** Enter the default gateway for this IP pool. This entry cannot be changed once the IP pool is created.
- Step 3.** Enter the subnet prefix for the IP pool. For example, enter 24 for a mask for 255.255.255.0.
- Step 4.** Optionally, enter the IP of the primary and secondary DNS servers.
- Step 5.** Optionally, enter a DNS suffix.
- Step 6.** Enter the range of IPs that will be part of this IP pool.

Once an IP pool is created, you can modify or delete it. To make changes to an IP pool, follow these steps:

- Step 1.** Return to Object Groupings for the NSX Manager that owns the IP pool.
- Step 2.** Select **IP Pools**.
- Step 3.** Select the IP pool you want to modify.
- Step 4.** Click the **Edit IP Pools** icon.
- Step 5.** You can change almost all fields desired, including adding IPs to the pool, except the name and the default gateway fields.

The IP pool's IP range can't be shrunk if at least one IP has already been assigned. An IP pool can't be deleted if at least one IP has been already assigned.

Host Preparation

Now that you deployed your NSX Controllers, it's time to focus on the next steps that must take place before you can start deploying your virtual network and deploying security services. The NSX Controllers can also be deployed *after* host preparation.

The next step is to install the NSX vSphere Infrastructure Bundles (VIBs) in the ESXi hosts that will be in the NSX domain. The VIBs give the ESXi hosts the capability to participate in NSX's data plane and in kernel security. We do this by selecting the Host Preparation tab from the Installation view in the Networking and Security page, as shown in Figure 4-13. An alternative would be to use vSphere ESXi Image Builder to create an image with the NSX VIBs installed.

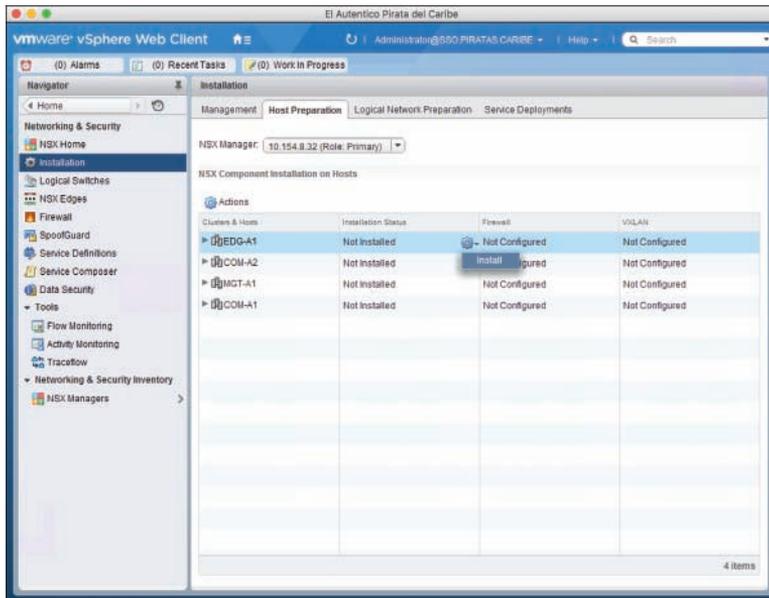


Figure 4-13 Host Preparation tab

In the Host Preparation tab you see a list of all the ESXi host clusters configured in vCenter. Under the Installation Status column, hover toward the right until the mouse is over the cog, click it and select **Install**. That's it. NSX Manager pushes the VIBs to each ESXi host that is in the cluster. Successfully adding the VIBs is non-disruptive, and there is no need to place the ESXi host in maintenance. Yes, I wrote “successfully” because if the VIB installation fails you might need to reboot the ESXi host(s) to complete it, as shown in Figure 4-14. The good thing is that NSX Manager tries to reboot the ESXi host for you, first putting in Maintenance mode. The moral of this: Don't execute any type of infrastructure changes or upgrades outside of a maintenance window. You would also need to reboot the ESXi host if you wanted to remove the NSX VIBs.

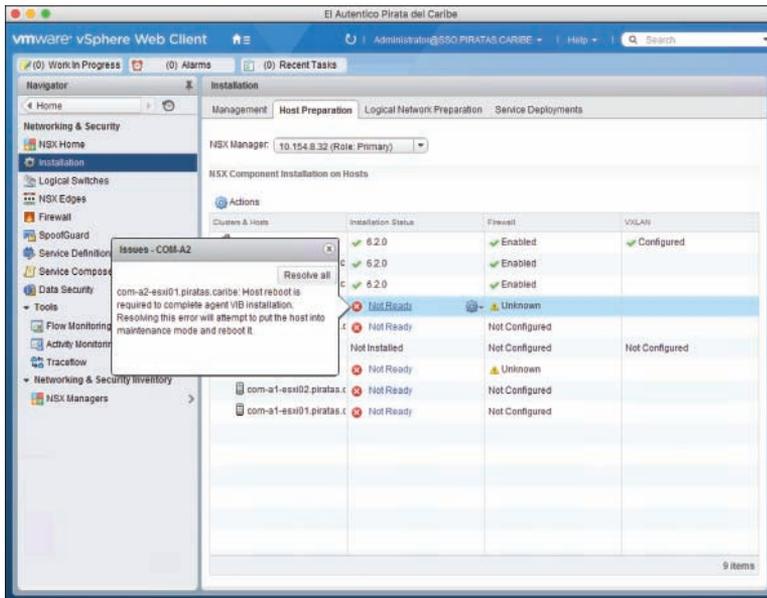


Figure 4-14 Incomplete NSX VIB installation

So what superpowers exactly are these VIBs giving the ESXi hosts? The modules and the over-and-above human capabilities they give the ESXi hosts are as follows:

- **The VXLAN module:** Enables the ESXi host to have logical switches. We discuss logical switching in Chapter 5.
- **The Switch Security (SwSec) module:** It is the logical switch’s assistant. It is a dvFilter that sits in Slot 1 of the IOChain and helps with Layer 2 broadcast suppression.
- **The Routing module:** Enables the ESXi host to run distributed logical routers. We review distributed logical routers in Chapter 7.
- **The distributed firewall:** Enables the ESXi host to do Layer 2, Layer 3, and Layer 4 security in kernel. It also allows the ESXi host to leverage, out of network, additional security services. We start the conversation about the distributed firewall and security in Chapter 15, “Distributed Logical Firewall.”

Any other superpowers? Well, maybe this can be considered as a superpower: If you add an ESXi host to a cluster that has already been prepared, the ESXi host gets the NSX VIBs automatically. How about that for cool?! And before I forget, installing the VIBs takes minimal time. Even in my nested-ESXi-hosts running lab with scant available CPU, memory, and an NFS share that is slower at delivering I/O than a delivery pigeon, the VIBs install quickly.

Figure 4-15 shows the ESXi host clusters that have been prepared with version 6.2.0 of the NSX VIBs by NSMGR-A, 10.154.8.32. Have a look at the two columns to the right, the Firewall and VXLAN columns. The Firewall module has its own column because it can be installed independently from the other modules. The VIB that has the Firewall module is called VSFWD. If the Firewall status reads Enabled, with the green check mark, you could go over to the Firewall view of Networking and Security, where the distributed firewall policies get created and applied, or the Service Composer view of Networking and Security, where service chaining is configured, to start creating and applying security rules for VMs. The distributed firewall VIB for NSX 6.0 can be installed with ESXi hosts running version 5.1 or higher. For NSX 6.1 and higher, the ESXi hosts must run 5.5 or higher.

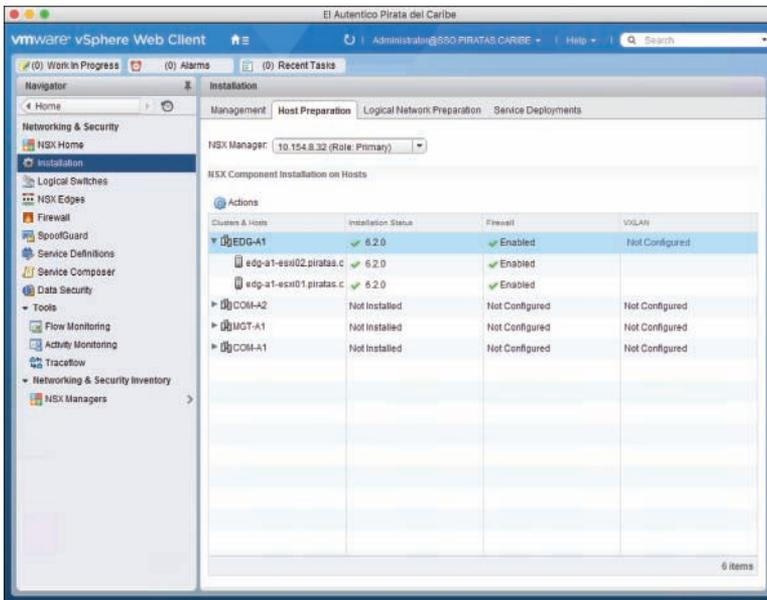


Figure 4-15 Host Preparation tab after NSX modules have been installed

The VXLAN column confirms the installation of the VXLAN VIB. The VXLAN VIB has the VXLAN module, the Security module, and the Routing module. If the column reads **Not Configured** with a hyperlink, the VXLAN VIB is installed. The VXLAN VIB can be installed with ESXi hosts running version 5.1 or higher; however, with version 5.1 ESXi hosts logical switches can only be deployed in Multicast Replication Mode. We cover Replication Mode in Chapter 5. For NSX 6.1 and higher, the ESXi hosts must run 5.5 or higher. The Routing module only works in ESXi hosts running vSphere 5.5 or higher. Table 4-2 shows the vSphere and vCenter version supported by each module.

**Table 4-2** vSphere Versions Supported by the NSX Modules

NSX Modules	vSphere Version
Security	5.1 or later
VXLAN	5.1 (only for Multicast Replication Mode) and later
Routing	5.5 or later

Host Configuration

If you want to deploy logical switches, you must complete the Logical Network Preparation tab in the Installation view. In this section you set up an NSX domain with the variables needed to create VXLAN overlays. Three sections need to be configured. If you skip any of them, you are not going to be deploying logical switches.

First, you need to tell NSX Manager how to configure the ESXi hosts. Oddly enough, you don't start the logical network configuration from the Logical Network Preparation tab. Rather, click the **Configure** hyperlink in the VXLAN column in the Host Preparation tab to open the Configure VXLAN Networking Wizard. Optionally, hover toward the right and click on the cog to see a menu list and choose **Configure VXLAN**, as shown in Figure 4-16.

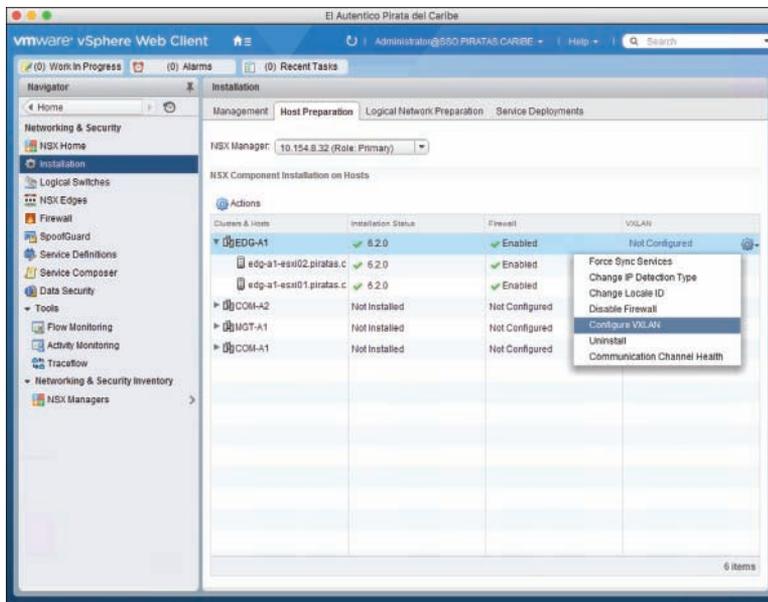
**Figure 4-16** VXLAN host configuration

Figure 4-17 shows the Configure VXLAN Networking window. Here we can configure the following:

- The vDS where the new VXLAN VMkernel portgroup will be created.
- The IP ESXi hosts will use as their VTEP. A new VMkernel port gets created for this, typically referred to as the VXLAN VMkernel port, and it is this VMkernel port that is the VTEP. Since the ESXi host owns the VXLAN VMkernel port, it is common practice to refer to the ESXi host as the VTEP itself. Moving forward, from time to time I refer to both the ESXi hosts and the VXLAN VMkernel ports as VTEPs.
- The number of VXLAN VMkernel ports, per ESXi hosts, that will be configured. Each VXLAN VMkernel port will have a different IP.

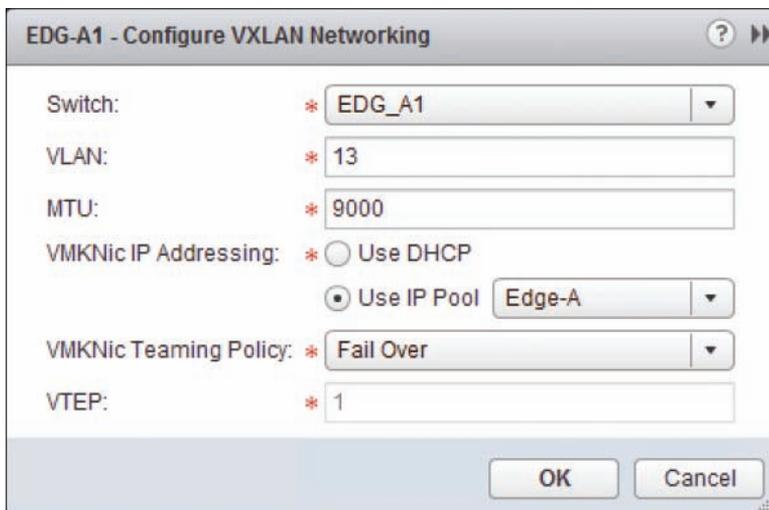


Figure 4-17 Configure VXLAN Networking Wizard

Key Topic

All ESXi hosts, per host cluster, must be in the same vDS that will be used by NSX for host configuration. NSX can work with different clusters having different vDSes. This has zero impact on the performance of VMs in the NSX domain. If running a vSphere version before 6.0, not using the same vDS across multiple clusters may impact the capability of vMotion virtual machines connecting to logical switches. We touch on this topic in Chapter 5.

The VLAN in Figure 4-17 is the VXLAN VLAN. The vDS switch selected in Figure 4-17 will be used by NSX Manager to create a portgroup for the VXLAN VMkernel port and portgroups to back the logical switches, which we cover in Chapter 5. All these portgroups will be configured by NSX Manager with the VXLAN VLAN. If the MTU configured is larger than the MTU already configured in the

vDS, the vDS's MTU will be updated. The vDS that gets assigned to the cluster for VXLAN may also continue to be used for other non-NSX connectivity, such as a portgroup for vMotion.

You can assign an IP address to the VXLAN VMkernel port by using DHCP or an IP pool. In both cases, the VXLAN VMkernel port would be getting a default gateway. This would typically present a problem for the ESXi host since it already has a default gateway, most likely pointing out of the management VMkernel port. Luckily for NSX, vSphere has supported multiple TCP/IP stacks since version 5.1. In other words, the ESXi host can now have multiple default gateways. The original default gateway, oddly enough referred to as *default*, would still point out of the management VMkernel port, or wherever you originally had it configured for. The new default gateway, which you probably correctly guessed is referred to as VXLAN, would point out of the VXLAN VMkernel port. The VXLAN TCP/IP stack default gateway and the VXLAN VMkernel port will only be used for the creation and termination of VXLAN overlays. Figure 4-18 shows the VMkernel ports of an ESXi host, with only the VXLAN VMkernel port using the VXLAN TCP/IP stack.

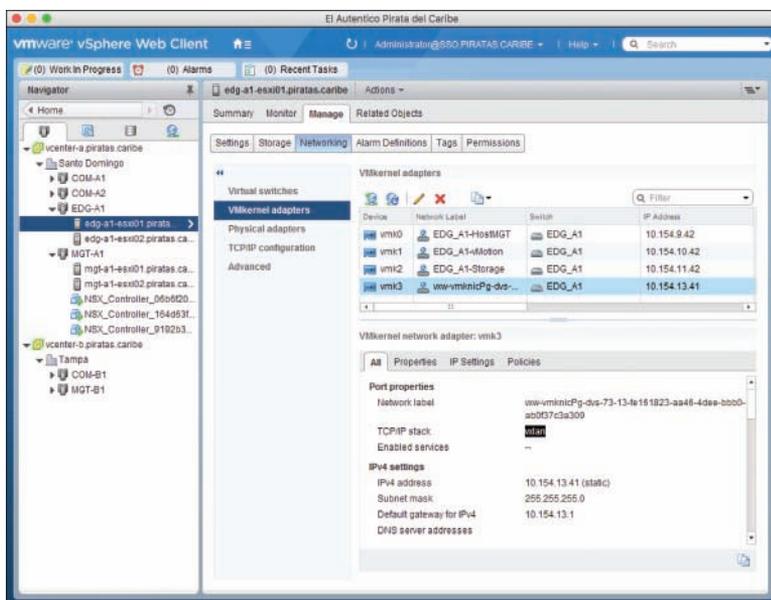


Figure 4-18 VXLAN VMkernel port with VXLAN TCP/IP stack

One final thing you can configure here is the VMKNic Teaming Policy, a name I'm not too fond of. Why couldn't they name it VXLAN Load Share Policy? After all, this is how the vDS load shares egress traffic from the VXLAN VMkernel port. Anyhow, the selection you make here has great implications for the behavior of your VXLAN overlays. For one, the policy must match the configuration of the physical

switches to which the vDS uplinks connect, which means the vDS must also be configured to match the selected policy, such as enhanced LACP.

These are the VMKNic Teaming Policy options available:

- Fail over
- Static EtherChannel
- Enhanced LACP
- Load Balance – SRCID
- Load Balance – SRCMAC

Go back and have a look at Figure 4-17. Do you see the VTEP field at the bottom? It says 1, meaning 1 VXLAN VMkernel port is created for each ESXi host in the cluster being configured. Where did the 1 come from? NSX Manager put it there. Notice the text box for the 1 is grayed out, which means you can't edit it. And how did NSX Manager know to put a 1 in there? Go back to the VMKNic Teaming Policy selection. If you choose anything other than Load Balance – SRCID or Load Balance – SRCMAC, NSX Manager puts a 1 in the VTEP text box.

**Key
Topic**

If, on the other hand, you choose VMKNic Teaming Policy of Load Balance – SRCID or Load Balance – SRCMAC, NSX Manager creates multiple VXLAN VMkernel ports, one per dvUplink in the vDS. Now that the ESXi hosts have multiple VXLAN VMkernel ports, load sharing can be achieved on a per VM basis by pinning each VM to a different VXLAN VMkernel port and mapping each VXLAN VMkernel port to a single dvUplink in the vDS. Figure 4-19 shows the configured ESXi hosts with multiple VXLAN VMkernel ports.

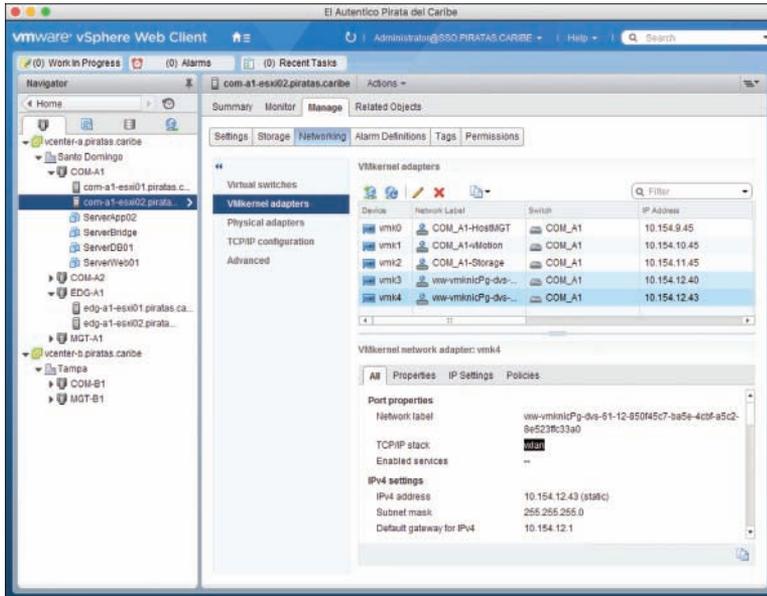


Figure 4-19 ESXi hosts with multiple VTEPs

Figure 4-20 shows the logical/physical view of two ESXi hosts, each with two dvUplinks, two VMs, and two VTEPs. The VMs are connected to logical switches.

Key Topic Element	Multi-VTEP Support	vDS Teaming Mode	vDS Version
Load Balance - SRCID	Yes	Source Port	5.5 and later
Load Balance - SRCMAC	Yes	Source MAC (MAC Hash)	5.5 and later

Now why would NSX Manager allow the option of multiple VTEPs in the same ESXi host? It allows the option because there is no other good way to load share, yes *load share*, egress traffic sourced from an ESXi host if the load sharing hash is using the source interface (SRCID) or the source MAC (SRCMAC). I won't spend too long explaining why NSX Manager achieves the load sharing the way it does. I would just say, think of how the physical network would react if the source MAC in egress frames from the ESXi host were seen in more than one discrete dvUplink from the same ESXi host.

After you finish the Configure VXLAN Networking Wizard, you can go over to the Logical Network Preparation tab to verify the configuration. Figure 4-21 shows the VXLAN Transport section listing the ESXi hosts that have been configured and the details of their configuration.

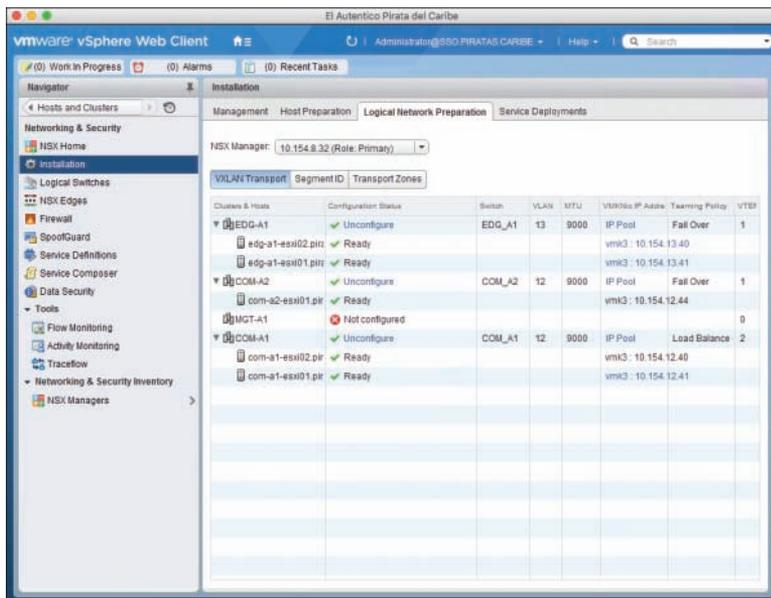


Figure 4-21 ESXi host clusters that have been configured for VXLAN

In the Network view of vCenter, you can verify that the portgroup was created for connecting the VXLAN VMkernel port. Figure 4-22 shows the VXLAN VLAN for the EDG-A1 host cluster, 13, is configured in the portgroup. Notice that there

are other portgroups in the same vDS. If you were to look at the vDS configuration, you would see the MTU is set to at least the size you configured in Configure VX-LAN networking.

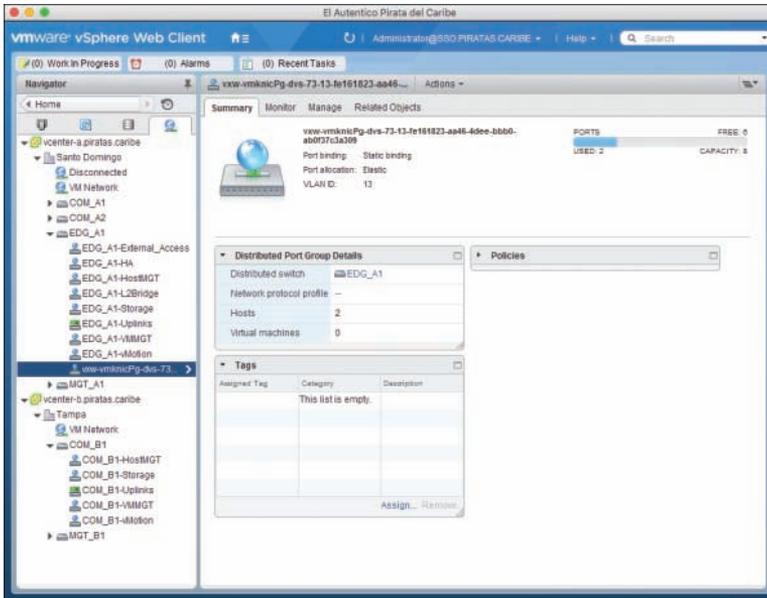


Figure 4-22 VXLAN vDS

VNI Pools, Multicast Pools, and Transport Zones

You need to undertake two more preparations for the NSX networks.

The first thing you should do is provide the range or pool of VNIs and multicast groups that NSX Manager would be using for its local use as well as do the same for cross vCenter NSX use. Local VNI pools and universal VNI pools shouldn't overlap. Local multicast groups and universal multicast groups shouldn't overlap either. The VNI pool can start at 5000. To create the VNI pools, go to the Segment ID section of the Logical Network Preparation tab and select the Primary NSX Manager. If you require multicast support, you can enter the multicast group pools for NSX Manager to use in the same place. We discuss multicast in the "Replication Mode" section of Chapter 5. Secondary NSX Managers can only configure local VNI and multicast group pools.

The second thing you should do is create global transport zones, at least one per NSX Manager, and a universal transport zone. When a logical switch is created, NSX Manager needs to know which ESXi hosts in the NSX domain have to be informed about the logical switch. The global transport zone is a group of ESXi host clusters under the same NSX domain that would be told about the creation of logi-

cal switches. Global transport zone only includes ESXi host clusters local to a vCenter. The universal transport zone is a group of ESXi host clusters under the same cross vCenter NSX domain that would be told about the creation of universal logical switches. Universal transport zones may include ESXi host clusters in all vCenters in the same cross vCenter NSX domain. The logical switch's global transport zone assignment and a universal logical switch's universal transport zone assignment are done during the creation of the switches.

NOTE For the rest of the book, when I refer to *transport zone*, my comment applies to both the global transport zone and the universal transport zone.

A transport zone can contain as many clusters as you want. An ESXi host cluster can be in as many transport zones as you want, and it can belong to both types of transport zones at the same time. And yes, you can have as many global transport zones as your heart desires, although you typically don't deploy more than one or two per NSX Manager. However, you can only have a single universal transport zone. More importantly, both types of transport zones can have ESXi host clusters each with a different vDS selected during Configure VXLAN networking. Again, transport zones matter only for the purpose of letting the NSX Manager know which ESXi hosts should be told about a particular logical switch or universal logical switch.

To create a transport zone, head over to the Logical Network Preparation tab, select the NSX Manager that will own the transport zone, and go to the Transport Zones section. There, click the green + sign. There you can assign the transport zone a name, select its Replication Mode, and choose the ESXi host clusters that will be part of the transport zone. If the NSX Manager is the Primary NSX Manager, you have a check box to turn this transport zone into a universal transport zone, as shown in Figure 4-23.

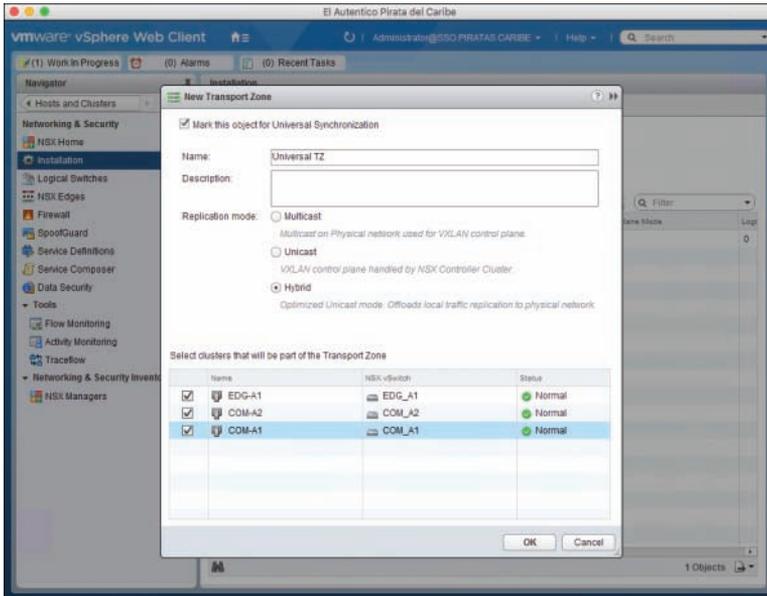


Figure 4-23 Creating a transport zone

As mentioned, Chapter 5 discusses what Replication Mode is. For now, you should know that if you select **Multicast** or **Hybrid** you need to create a multicast group pool in the Segment ID section mentioned previously. Finally, after a transport zone is created, you can't change the transport zone type. However, you can modify it by adding or removing ESXi host clusters from the NSX Manager that owns the association to the vCenter that owns those clusters. If an NSX switch (a logical switch or a universal logical switch) has already been created before the ESXi host cluster is added to the transport zone, NSX Manager automatically updates the newly added ESXi hosts in the ESXi host cluster with the NSX switch information.

**Key
Topic**

To add an ESXi host cluster in a transport zone, return to the Transport Zone section of the Logical Network Preparation tab and select the NSX Manager that prepared the ESXi host cluster that will be added. Select the Transport Zone and click the Connect Clusters icon. Select the ESXi host clusters you want to add and click **OK**.

To remove an ESXi host cluster from a transport zone, select the transport zone in the Transport Zones section and select the Disconnect Clusters icon. Select the ESXi host clusters you want to remove and click **OK**. For the operation to succeed, all VMs (powered on or not) in the ESXi host you want to remove must be disconnected from all logical switches that belong to the transport zone. We cover how to disconnect a VM from a logical switch in Chapter 5.

A transport zone that has any logical switches can't be deleted. The logical switches must be deleted first. We cover how to delete logical switches in Chapter 5. To delete a transport zone, select the transport zone, then select **Actions, All NSX User Interface Plugin Actions**, and then select **Remove**.



One more note on this section. It should be clear by now that NSX Manager *loves* ESXi host clusters. If you add an ESXi host to an already prepared and configured ESXi host cluster, NSX Manager would make sure that the ESXi host gets the NSX VIBs, the VXLAN VMkernel ports get created with the right IP and subnets, and make the new ESXi host aware of any logical switches, and so forth. On the reverse, if you remove an ESXi host from an already prepared and configured ESXi host cluster, the ESXi host would lose its VXLAN VMkernel ports and IPs, and lose knowledge of any logical switches.

That wraps up all the prep work that needs to be done to get your NSX network and security going. The next chapter begins the coverage of the process of actually building stuff that you can put virtual machines on.

Exam Preparation Tasks

Review All the Key Topics

Review the most important topics from inside the chapter, noted with the Key Topic icon in the outer margin of the page. Table 4-4 lists these key topics and the page numbers where each is found.

Table 4-4 Key Topics for Chapter 4

**Key
Topic**

Key Topic Element	Description	Page Number
Paragraph	Define what VXLAN is and the scaling capabilities native to the protocol.	90
List	VXLAN frame inherits Layer 2 and Layer 3 QoS from the encapsulated frame. The source UDP port is derived from the encapsulated frame.	93
Paragraph	VXLAN requires jumbo frame support from the underlay.	95
Paragraph	The NSX Controllers maintain the principle copies of the VTEP, MAC, and ARP tables.	96
Paragraph	NSX Controllers have no role in Network Security	97
Paragraph	The user must have the correct administrator account to deploy NSX Controllers.	98
Paragraph	Changing the NSX Controllers's hardware settings is not supported by VMware.	103
Paragraph	Secondary NSX Managers do not deploy NSX Controllers	104
Paragraph	VMware requires three NSX Controllers in a production deployment of NSX.	105
Paragraph	The NSX Controller taking over for a failed one queries the ESXi hosts in the VTEP table.	105
Table 4-2	The versions of vSphere supported by the NSX modules.	113
Paragraph	All members of the ESXi host cluster must belong to the same vDS for NSX host preparation.	114
Paragraph	NSX supports multiple VTEPs per ESXi host.	116
Paragraph	Each NSX Manager in the cross vCenter NSX domain is responsible for adding clusters to the universal transport zone.	122
Paragraph	NSX Manager only interacts with ESXi hosts that are members of clusters.	123

Complete Tables and Lists from Memory

Download and print a copy of Appendix C, “Memory Tables” (found on the book’s website), or at least the section for this chapter, and complete the tables and lists from memory. Appendix D, “Memory Tables Answer Key,” also on the website, includes the completed tables and lists so you can check your work.

Define Key Terms

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

VXLAN, VTEP, VNI, NSX Controller Master, VXLAN module, Routing module, Distributed Firewall module, transport zone, slicing



VCP6-NV Exam 2V0-641 Updates

Over time, reader feedback allows Pearson to gauge which topics give our readers the most problems when taking the exams. To assist readers with those topics, the authors create new materials clarifying and expanding on those troublesome exam topics. As mentioned in the Introduction, the additional content about the exam is contained in a PDF on this book's companion website at <http://www.ciscopress.com/title/9780789754806>.

This appendix provides you with updated information if VMware makes minor modifications to the exam upon which this book is based. When VMware releases an entirely new exam, the changes are usually too extensive to provide in a simple updated appendix. In those cases, you might need to consult the new edition of the book for the updated content. This appendix attempts to fill the void that occurs with any print book. In particular, this appendix does the following:

- Mentions technical items that might not have been mentioned elsewhere in the book
- Covers new topics if VMware adds new content to the exam over time
- Provides a way to get up-to-the-minute current information about content for the exam

Always Get the Latest at the Book's Product Page

You are reading the version of this appendix that was available when your book was printed. However, given that the main purpose of this appendix is to be a living, changing document, it is important that you look for the latest version online at the book's companion website. To do so, follow these steps:

- Step 1.** Browse to <http://www.ciscopress.com/title/9780789754806>.
- Step 2.** Click the **Updates** tab.
- Step 3.** If there is a new Appendix B document on the page, download the latest Appendix B document.

NOTE The downloaded document has a version number. Comparing the version of the print Appendix B (Version 1.0) with the latest online version of this appendix, you should do the following:

- **Same version:** Ignore the PDF that you downloaded from the companion website.
- **Website has a later version:** Ignore this Appendix B in your book and read only the latest version that you downloaded from the companion website.

Technical Content

The current Version 1.0 of this appendix does not contain additional technical coverage.

This page intentionally left blank



Index

Numbers

802.1Q standard, 9

A

ABR (Area Border Routers), OSPF routing, 352-353, 356

access

compliance. *See* Activity Monitoring
RBAC, 521

assigning roles, 523
NSX Manager/SSO integration, 522
NSX roles, 521-522
user access control, 522
user roles, 522

Access Layer (networks), 25

collapsed Access Layer design, 30
Spine and Leaf design, 31

access routers, physical networks, 25

access switches. *See* ToR switches

activation codes, practice exams (Pearson Cert Practice Test engine), 577

Active mode (LACP), 28

Activity Monitoring, 509

activity reports
Inbound Activity reports, 512

Outbound Activity reports, 513

Outbound AD Group Activity reports, 514

running, 511

viewing, 514

VM Activity reports, 511

data collection, enabling in VM, 510

Inter Container Interaction reports, 513

admin credentials, NSX Manager, 73

administrative distance (routing), 348

AES-NI (Advanced Encryption Standard-New Instructions), 278

Allowed Flows option (Flow Monitoring), 516-517

anti-replay feature, IPsec VPN, 382

API (NSX)

calls for

logical routers, 536-540

logical switches, 532-536

NSX Edge, 540-542

client headers, 531

HTTP responses, 531

supported HTTP methods, 530

application profiles (load balancers), 421, 427-428

areas

IS-IS areas, 369

OSPF routing

- backbone areas, 352*
- inter-reas, 352*
- intra-areas, 352*
- normal areas, 354*
- NSSA, 354*
- stubby areas, 354*

ARP (Address Resolution Protocol)

ARP requests

- Designated Instances, 325-327*
- logical switches, 147, 152*
- recursive routing, 347*
- VLAN LIF, 322-323*

ARP tables

- logical switches, 147-152*
- NSX Controllers, 97*

ASBR (Autonomous System Border Routers), OSPF routing, 352-353

ASN (Autonomous System Numbers), BGP ASN, 361-364

audit logs, verifying DFW, 471

authentication

- data origin authentication, IPsec VPN, 382

- MD 5 authentication, 355

- OSPF routing, 355

- SSL VPN-Plus, 398-402

authorization requests, exam preparation, 573

automation

- consumption plane, defining, 530

NSX API

- calls for logical routers, 536-540*
- calls for logical switches, 532-536*
- calls for NSX Edge, 540-542*
- client headers, 531*

- HTTP responses, 531*

- supported HTTP methods, 530*

REST

- API, 530*

- client headers, 531-532*

- HTTP responses, 531*

- HTTP URI, 530*

- NSX API calls for logical routers, 536-540*

- NSX API calls for logical switches, 532-536*

- NSX API calls for NSX Edge, 540-542*

- NSX API HTTP supported methods, 530*

- NSX Manager, 531*

vRA, 542

- external network profiles, 543-544*

- NAT network profiles, 543, 546*

- private network profiles, 543-545*

- routed network profiles, 543-545*

- vRO, 543*

B

backbone areas, OSPF routing, 352

BDR (Backup Designated Routers), OSPF neighbor adjacencies, 355

BGP (Border Gateway Protocol)

- ASN, 361-364

- best path selection, 365

- configuring, 366-367

- eBGP peers, 362-364

- iBGP peers, 362

- route advertisements, 363-364

- synchronization, 364

- verifying, 368-369

Blocked Flows option (Flow Monitoring), 516-517

blueprint (exam), exam preparation, 572

bridges (Layer 2), 296

Bridge Instances, 297-298, 308

configuring, 299-300

packet walks, 303-307

verifying, 301-303

broadcast domains, Ethernet, 10

BUM (Broadcast, Unknown unicast, Multicast), logical switches

Multicast Replication Mode, 154-155

Replication Mode, 153

C

certification requirements, exam preparation, 572

chapter resources (exam preparation), 573

CLI commands, NSX Manager, 71-72

Client Error API responses, 531

client headers, REST, 531-532

cloud-computing

automation via vRA, 542

external network profiles, 543-544

NAT network profiles, 543, 546

private network profiles, 543-545

routed network profiles, 543-545

vRO, 543

physical networks, 7

clusters (NSX Controller), 104-105

collapsed Access Layer network design, 30

confidentiality (data), IPsec VPN, 382

configuring

BGP routing, 366-367

DFW, 460-468

dvPortgroups (vDS), 48, 54-57

ESXi hosts, 113-119

IPsec VPN, 386-390

IS-IS routing, 370-372

LACP, 52-54

Layer 2 bridging, 299-300

Layer 2 VPN, 280-283

server settings, 284-285

user details and certificate, 285

VPN Client details, 288

NAT

DNAT rules, 416-418

SNAT rules, 418-420

NSX Edge, 259, 262-266

DHCP Relay, 435

DHCP servers, 434-435

DNS servers, 435

firewalls, 439-442

load balancers, 427-432

NSX Manager, 71-76

Network & Security access, 78-80

vSphere Web Client, 77

OSPF routing, 356-360

QoS marking, 55-57

SSL VPN-Plus, 395

adding installation packages, 403-404

adding IP pools, 405-406

adding private networks, 406-408

authentication, 398-402

enabling SSL VPN-Plus service, 402

security, 398-402

server settings, 396-397

Web Resources, 397

static routes, 350

Traceflow, 519-520

virtual servers for load balancers, 431-432

VMNIC*SR-IOV configuration, 40**TSO, 41**vSS configuration, 39***vSS***VMkernel portgroup configuration,
41-42**VMNIC configuration, 39***connectivity testing, 216-220****consumption plane, defining, 530****control plane, 65****Control VM**

logical routers, 201-203, 213

privileged mode, 215

user mode, 215

Controllers (NSX), 16, 67

ARP tables, 97

clusters, 104-105

Controller Masters, 106

deploying, 97-103

ESXi hosts, 109-111

host preparation, 109-111

IP pools, 107-109

logical switches, 138-139, 142-157

MAC tables, 97

NSX Controllers 6.2, 563-564

recovery, 107

responsibilities of, 96

slicing, 106

universal controllers, 102

verifying, 101-103

VTEP tables, 96

cost, routing, 348**cross vCenter NSX, 80, 83-84**

features of, 81-82

use cases, 81

D

**data center networks, infrastructure
of, 24-28****data confidentiality, IPsec VPN, 382****data integrity, IPsec VPN, 382****data origin authentication, IPsec
VPN, 382****data plane, 65****data security**

NSX Data Security, 509

PCI Data Security Standard, 506

vCNS Data Security, 555

VMware Data Security, 506-509

DELETE API calls, 531**deleting**

dvPortgroups (vDS), 48

vSS portgroups, 42

delta exams, exam preparation, 572**deploying**

logical routers, 204-210

NSX Controllers, 97-103

NSX Edge, 259, 262-266

*Edge HA mode, 261**monitoring deployments, 269**undeployed NSX Edges, 267**verifying deployments, 269-270*

NSX Manager, 69, 72

security service appliances, 484-485

Designated Instances, 324-327**DFW (Distributed Firewalls), 449,
453-455**

configuring, 460-468

domain groups, adding to NSX
Manager, 469-470

Exclusion Lists, 460

memory, 459

- microsegmentation, 457
- packet walk, 456-458
- rules
 - creating*, 462-468
 - default rules*, 458-460
 - enforcing*, 458
 - local rules*, 462-464
 - security policy rules*, 493-494
 - universal rules*, 462-465
 - updating*, 467-468
- saving configurations, 468
- security groups, 490
- service redirection, 496-497
- thresholds/limits, 458-460
- vCNS upgrades to NSX, 559
- verifying, 470-471
- DH (Diffie-Hellman) keys, IPsec VPN, 385**
- DHCP (Dynamic Host Configuration Protocol), NSX Edge**
 - DHCP Relay, 435
 - DHCP servers, 434-435
- directly connected subnets/subnet masks, 346**
- distributed firewalls, NSX, 16**
- distributed logical routers, 16, 201**
- DLR (Device-Level Rings), Layer 2 bridging, 296-298**
 - configuring, 299-300
 - verifying, 301-303
- DNAT (Destination NAT)**
 - configuring rules, 416-418
 - load balancers, 421
- DNS servers, NSX Edge, 435**
- domain groups, NSX Manager, 469-470**

- DR (Designated Routers), OSPF neighbor adjacencies, 355**
- dvPortgroups, 43**
 - configuring, 48, 54-57
 - creating, 47
 - deleting, 48
 - LACP, 54
 - logical routers, 199
 - logical switches, 132-135
 - QoS marking, 55-57
 - verifying, 135
 - VLAN LIF, 200
- dvUplinks, vDS, 43**
- dynamic memberships, Service Composer security groups, 488-489**

E

- eBGP peers (BGP routing), 362-364**
- ECMP (Equal Cost Multipathing), 336-339, 348, 365**
- Edge. *See* NSX Edge**
- Edge Clusters, Edge VLAN, 257**
- Edge HA mode (NSX Edge), 260, 266**
- Edge VLAN, 256**
- EGP (External Gateway Protocols), BGP, 361**
- encryption**
 - AES-NI, 278
 - IPsec VPN, 389-390
 - Layer 2 VPN, 278
- ESXi hosts**
 - Bridge Instances, 297-298, 308
 - configuring, VXLAN, 113-114
 - NSX Manager*, 113-119
 - VMKNic Teaming Policies*, 116-118
 - VXLAN*, 115-119

- Designated Instances, 324
- DFW, 458-460, 471
- DW, updating rules, 468
- host clusters
 - NSX, 35
 - upgrading to NSX 6.2, 565*
- Layer 2 bridging, 297-298, 308
- Locale ID, 224
- logical routers, 198-200
- logical switches, 130, 135
 - ARP tables, 147-152*
 - MAC tables, 144-147, 150*
 - VTEP tables, 139-143, 150, 153*
- network connectivity, 35
- NSX, 35
- NSX and physical networks, 33
- NSX Edge Gateways, 331
- NSX Manager, 113-119, 123
- transport zones, 122
- troubleshooting, 151
- vDS, 43, 49-51
- VIB, 110-111
- VM security, 452
- VMkernel ports, 35
- VMNIC, 35
- VNI, 151
- vSS, 36, 43
- VXLAN, 90

Ethernet

- 802.1Q standard, 9
- broadcast domains, 10
- BUM, 153
- challenges of, 7-10
- collapsed Access Layer network design, 30
- MAC addresses, 7

- MAC tables, 8
- MTU, NSX and physical networks, 33
- reconvergence, 10
- trunks, 9
- unknown unicasts, 8-9
- VLAN, 9

exam preparation

- authorization requests, 573
- certification requirements, 572
- chapter resources, 573
- delta exams, 572
- exam blueprint, 572
- exam day strategies, 574-575
- exam details, 572-573
- exam format, 571
- exam requirements, 571
- hands-on experience, 571
- Pearson Cert Practice Test engine, 573
 - downloading/installing practice exams, 576*
 - downloading/installing software, 576*
 - exam activation codes, 577*
 - Practice Exam Mode, 578-579*
 - Premium Edition, 577*
 - Study Mode, 578-579*
- practice labs, 571
- product pages, 585
- recommended training details, 573
- registering for the exam, 573
- review tools, 575
- updates, 585
- web resources, 571-579, 585

Exclusion Lists (DFW), 460

experience (hands-on), exam preparation, 571

external network profiles (vRA), 543-544

F

final exam preparation

- authorization requests, 573
- certification requirements, 572
- chapter resources, 573
- delta exams, 572
- exam blueprint, 572
- exam day strategies, 574-575
- exam details, 572-573
- exam format, 571
- exam requirements, 571
- hands-on experience, 571
- Pearson Cert Practice Test engine, 573
 - downloading/installing practice exams*, 576
 - downloading/installing software*, 576
 - exam activation codes*, 577
 - Practice Exam Mode*, 578-579
 - Premium Edition*, 577
 - Study Mode*, 578-579
- practice labs, 571
- recommended training details, 573
- registering for the exam, 573
- review tools, 575
- web resources, 571-579

firewalls

- design compromises (traditional), 449, 453
- DFW, 449, 453-455
 - configuring*, 460-468
 - creating rules*, 462-468
 - default rules*, 458-460
 - domain groups, adding to NSX Manager*, 469-470
 - enforcing rules*, 458
 - Exclusion Lists*, 460

- local rules*, 462-464
- memory*, 459
- microsegmentation*, 457
- packet walk*, 456-458
- saving configurations*, 468
- security groups*, 490
- security policies*, 493-494
- service redirection*, 496-497
- thresholds/limits*, 458-460
- universal rules*, 462-465
- updating rules*, 467-468
- vCNS upgrades to NSX*, 559
- verifying*, 470-471

- distributed firewalls, NSX, 16
- inter-tier security, 450, 453
- Layer 2 firewalls, 12
- network security, 12
- NSX Edge firewall, 436-442
- physical networks, updating, 6
- ToR switches, 449, 453
- VLAN, 451-453
- VM, 451-452
- vMotion, 453

flooding

- MAC addresses, 8
- unknown unicasts, 9

Flow Monitoring, 514

- Allowed Flows option, 516-517
- Blocked Flows option, 516-517
- collected flows
 - exporting*, 517
 - viewing*, 516
- enabling, 515
- Flow Collectors, 518
- IPFix, 517

Live Flow Monitoring, 518

VM restrictions, 515-516

format of exam, exam preparation, 571

G

GET API calls, 531-532

GET transport zones, 532

global logical switches, 130

assigning, 131

creating, 132-134

Guest Introspection

Activity Monitoring, 509

data collection, enabling in VM, 510

Inbound Activity reports, 512

Inter Container Interaction reports, 513

Outbound Activity reports, 513

Outbound AD Group Activity reports, 514

running activity reports, 511

viewing activity reports, 514

VM Activity reports, 511

VMware Data Security, 506-509

Guest Introspection (security services), 480, 491

H

hairpinning, virtual networks, 11

hands-on experience, exam preparation, 571

hardware VTEP, 308, 311-312

Health Check (Service Monitor), NSX Edge load balancers, 425

HTTP (Hypertext Transfer Protocol), automation and REST

HTTP responses, 531

HTTP URI, 530

NSX API HTTP supported methods, 530

Hybrid Replication Mode (logical switches), 156-157

I

iBGP peers (BGP routing), 362

IGP (Interior Gateway Protocol), OSPF routing, 351

IKE (Internet Key Exchange), IPsec VPN, 384-385

In-Line (Transparent Mode) NSX Edge load balancers, 422

Inbound Activity reports, 512

installation packages, SSL VPN-Plus, 403-404

integrity (data), IPec VPN, 382

Inter Container Interaction reports, 513

inter-rea, OSPF routing, 352

Internal LIF, 200

internal segments, 200

intra-areas, OSPF routing, 352

intra-tier security and firewalls, 450, 453

IP (Internet Protocol)

IP networks

challenges of, 10-12

IPAM, 11-12

NAT, 11

overlays, 91

TCP/IP, 10

underlays, 91

IP pools

creating, 107-109

SSL VPN-Plus, 405-406

- IP sets, 499-500
- SpoofGuard, 471-473
- VIP, load balancers, 421
- IPAM (IP Address Management) platforms, 11-12**
- IPFix, Flow Monitoring, 517**
- IPsec VPN**
 - configuring, 386-390
 - creating, 391
 - encryption algorithms, 389-390
 - establishing, 384-385
 - security, 389-390
 - security features, 382
 - site-site IPsec VPN, 382-385, 394
 - support for, 386
 - verifying, 392-394
- IS-IS (Interior System to Interior System)**
 - areas, 369
 - configuring, 370-372
 - Mesh Groups, 370
 - route redistribution, 373-375
 - router types, 369-370
 - verifying, 373

J-K-L

- LACP (Link Aggregation Control Protocol)**
 - Active mode, 28
 - configuring, 52-54
 - dvPortgroups, 54
 - NIC teaming, 28
 - Passive mode, 28
 - physical networks, 27
 - Static mode, 28

- vDS, 52
- VMNIC, 53-54
- LAN (Local Area Networks)**
 - TCP/IP, 10
 - VLAN
 - vDS, 44-45*
 - vSS portgroups, 38-39*
 - VXLAN
 - encapsulation, 92, 95*
 - ESXi hosts, 90*
 - host configuration, 113-119*
 - tunnels, 91*
 - underlays, 95*
 - VIB, 112*
 - VTEP, 92*
- Layer 2 bridging, 296**
 - Bridge Instances, 297-298, 308
 - configuring, 299-300
 - packet walks, 303-307
 - verifying, 301-303
- Layer 2 firewalls, 12**
- Layer 2 VPN**
 - configuring, 280-283
 - server settings, 284-285*
 - user details and certificate, 285*
 - VPN Client details, 288*
 - encryption, 278
 - NSX Edge, 278-282, 286-287
 - packet walks, 290-296
 - security, 278
 - verifying, 289
- Layer 3 connectivity between virtual/physical networks**
 - Designated Instances, 324-327
 - ECMP, 336-339

logical router VLAN LIF, 318,
321-324, 330, 336

NSX Edge Gateways, 330-333

Layer 3 separation, network security, 12

Layer 3 switches, physical networks, 26

LDAP (Lightweight Directory Access Protocol), adding domain groups to NSX Manager, 469-470

LIF (Logical Interfaces)

logical routers, 198-200

logical switches, 200

VLAN LIF, 200, 318, 321-324,
330, 336

link state, OSPF routing, 351-352

Live Flow Monitoring, 518

load balancers (NSX Edge), 420

application profiles, 421, 427-428

configuring

application profiles, 427-428

server pools, 430

virtual servers, 431-432

enabling, 433

persistence states, 426

Proxy (One-Arm) Mode, 422-425

server pools, 430

Service Monitor (Health Check), 425

support for, 421

throttling, 426

Transparent (In-Line) Mode, 422

VIP, 421

virtual servers, configuring, 431-432

local DFW rules, 462-464

Locale ID

ESXi hosts, 224

ULR, 221-224

logical routers

connectivity testing, 216-220

Control VM, 201-203

creating, 204

deploying, 204-210

distributed logical routers, 16, 201

distributed placement of, 199

dvPortgroups, 199

ESXi hosts, 200

LIF, 198-200

NSX API calls (automation), 536-540

NSX Edge versus, 258

packet walks, 232-234

example 1, 235-239

example 2, 240-246

example 3, 246-249

pMAC, 200

routing, 346-347

OSPF, 351-360

static routes, 349-350

transport zones, 201

ULR, 201, 221-224

verifying, 210-214

VLAN LIF, 318, 321-324, 330, 336

vMAC, 200

XML tags, 536-537

logical switches

ARP requests, 147, 152

creating, 131-134

dvPortgroups, 132-135

ESXi hosts, 130, 135

ARP tables, 147-152

MAC tables, 144-147, 150

VTEP tables, 139-143, 150, 153

global logical switches, 130

assigning, 131

creating, 132-134

- Hybrid Replication Mode, 156-157
 - LIF, 200
 - MAC learning, 131
 - Multicast Replication Mode, 154-155
 - NSX, 16
 - NSX API calls (automation), 532-536
 - packet walks, 165-168
 - example 1*, 169-170
 - example 2*, 170-171
 - example 3*, 171-173, 176-177
 - example 4*, 177-179, 185-188
 - example 5*, 189-190
 - Replication Mode, 152-154
 - tables, 138
 - ARP tables*, 147-152
 - MAC tables*, 144-147, 150
 - verifying*, 149-151
 - VTEP tables*, 139-143, 150, 153
 - transport zones, 201
 - Unicast Replication Mode, 155-156
 - universal logical switches, 130
 - assigning*, 131
 - creating*, 132-134
 - unknown unicasts, 152
 - verifying, 135
 - VM migration, 137
 - VNI, 130
 - VTEP, 130
 - Hybrid Replication Mode*, 156-157
 - Multicast Replication Mode*, 154-155
 - Unicast Replication Mode*, 155-156
 - VTEP tables*, 139-143, 150, 153
- LSA (Link State Advertisements), OSPF routing, 355-356**
- LSDB (Link State Database), OSPF routing, 354**
- ## M
-
- MAC addresses**
 - Ethernet, 7
 - flooding, 8
 - pMAC, logical routers, 200
 - SpoofGuard, 471-473
 - vMAC, logical routers, 200
 - vSS, 37
 - MAC learning, 8**
 - defining, 130
 - logical switches, 131
 - MAC sets, 500**
 - MAC tables**
 - Ethernet, 8
 - logical switches, 144-147, 150
 - NSX Controllers, 97
 - management plane, 64-65**
 - Masters (NSX Controllers), 106**
 - MD 5 authentication, OSPF routing, 355**
 - memoryDFW, 459**
 - Mesh Groups, 370**
 - MIB (Management Information Bases), 530**
 - microsegmentation and DFW, 457**
 - MOID (Managed Object ID), 135**
 - monitoring NSX Edge, 269**
 - MTU (Maximum Transmission Units)**
 - NSX and physical networks, 33
 - OSPF neighbor adjacencies, 355
 - multicast pools, 120**
 - Multicast Replication Mode (logical switches), 154-155**
 - multipathing, ECMP, 336-339, 348, 365**

N

NAT (Network Address Translation)

DNAT

configuring rules, 416-418

load balancers, 421

IP networks, 11

network profiles (vRA), 543, 546

SNAT, 416-420

neighbor adjacencies, OSPF routing, 354-355

nested security groups (Service Composer), 489

Network Introspection (security services), 480-481, 491

networks

collapsed Access Layer design, 30

data center networks, infrastructure of, 24-28

Ethernet

802.1Q standard, 9

broadcast domains, 10

challenges of, 7-10

collapsed Access Layer network design, 30

MAC addresses, 7

MAC tables, 8

reconvergence, 10

trunks, 9

unknown unicasts, 8-9

VLAN, 9

functions

NSX, 14-16

NV, 14, 17

IP networks, 91

challenges of, 10-12

IPAM, 11-12

NAT, 11

TCP/IP, 10

overlays, 91

physical networks

Access Layer, 25

access routers, 25

challenges of, 6-7

cloud-computing, 7

collapsed Active Layer design, 30

infrastructure of, 24-28

LACP, 27

Layer 3 connectivity between virtual/physical networks, 318, 321-327, 330-333, 336-339

Layer 3 switches, 26

NSX, 14, 33-34

POD design, 28-29

scaling, 26

Spine and Leaf design, 31

STP, 25

SVI, 26

ToR switches, 24-26

updating firewalls, 6

virtualization, 7, 26

virtual workloads, 7

planes

control plane, 65

data plane, 65

dependencies of, 64

management plane, 64-65

POD design, 28-29

private networks, SSL VPN-Plus, 406-408

routers, 349

security, 12

Spine and Leaf design, 31

spokes, 349

- stubs, 349
- underlays, 91, 95
- virtual networks, Layer 3 connectivity
 - between virtual/physical networks, 318, 321-327, 330-333, 336-339
- NFV (Network Function Virtualization), 14, 17**
- NIC (Network Interface Cards)**
 - teaming, 28
 - vNIC, vSS portgroups, 37
- normal areas, OSPF routing, 354**
- NSSA (Not-So-Stubby Areas), OSPF routing, 354**
- NSX, 12, 17**
 - architecture of, 66-68
 - Data Security, 509
 - distributed firewalls, 16
 - distributed logical routers, 16
 - ESXi host clusters, 35
 - ESXi hosts, 33-35
 - logical switches, 16
 - MTU, 33
 - multicast pools, 120
 - network functions, 14-16
 - NSX 6.2, upgrading to
 - host clusters, 565*
 - NSX Controllers, 563-564*
 - NSX Edge, 566-567*
 - NSX Manager 6.2, 561-563*
 - NSX-MH, 13
 - NSX-V, 13
 - NSX vSwitch, 67
 - physical networks, 14, 33-34
 - security
 - Data Security, 509*
 - functions, 15-16*
 - registering services, 482-484*
 - service composer, 16
 - transport zones, 120
 - creating, 121*
 - ESXi hosts, 122*
 - modifying, 122*
 - use cases, 13
 - vCenter, 13, 34
 - cross vCenter NSX, 80-84*
 - integration, 73-77*
 - vCNS upgrades to NSX
 - DFW, 559*
 - NSX Edge, 559*
 - NSX Manager, 555-557*
 - NSX VIB, 558*
 - USVM, 560*
 - vCNS Data Security, 555*
 - vCNS partner integration, 555*
 - vDS, 45
 - VIB, 109
 - ESXi hosts, 110-111*
 - vCNS upgrades to NSX, 558*
 - VXLAN, 112*
 - VNI pools, 120
 - vSphere
 - ESXi host clusters, 35*
 - ESXi hosts, 35*
 - vCenter, 34*
 - vDS, 35*
 - vSS, 35*
- NSX API**
 - calls for
 - logical routers, 536-540*
 - logical switches, 532-536*
 - NSX Edge, 540-542*
 - client headers, 531

- HTTP responses, 531
- supported HTTP methods, 530
- NSX Controllers, 16, 67**
 - ARP tables, 97
 - clusters, 104-105
 - Controller Masters, 106
 - deploying, 97-103
 - ESXi hosts, 109-111
 - host preparation, 109-111
 - IP pools, 107-109
 - logical switches, 138-139, 142-157
 - MAC tables, 97
 - NSX Controllers 6.2, 563-564
 - recovery, 107
 - responsibilities of, 96
 - slicing, 106
 - universal controllers, 102
 - verifying, 101-103
 - VTEP tables, 96
- NSX Edge, 16, 66-67, 256**
 - BGP routing, 366-367
 - configuring, 259, 262-266
 - DHCP Relay*, 435
 - DHCP servers*, 434-435
 - DNS servers*, 435
 - firewalls*, 439-442
 - deploying, 259, 262-266
 - monitoring deployments*, 269
 - NSX Edges in Edge HA mode*, 261
 - undeployed NSX Edges*, 267
 - verifying deployments*, 269-270
 - DHCP
 - DHCP Relay*, 435
 - DHCP servers*, 434-435
 - DNS servers, 435
 - Edge HA mode, 260, 266
 - Edge VLAN, 256
 - firewalls, 436-442
 - Gateways, 330-333
 - IPsec VPN
 - configuring*, 386-390
 - creating*, 391
 - encryption algorithms*, 389-390
 - establishing*, 384-385
 - security*, 389-390
 - security features*, 382
 - site-site IPsec VPN*, 382-385, 394
 - support for*, 386
 - verifying*, 392-394
 - IS-IS routing
 - configuring*, 370-372
 - route redistribution*, 373-375
 - Layer 2 VPN, 278-282, 286-287
 - limits of, 260
 - load balancers, 420
 - application profiles*, 421, 427-428
 - configuring*, 427-432
 - enabling*, 433
 - persistence states*, 426
 - Proxy (One-Arm) Mode*, 422-425
 - server pools*, 430
 - Service Monitor (Health Check)*, 425
 - support for*, 421
 - throttling*, 426
 - Transparent (In-Line) Mode*, 422
 - VIP*, 421
 - virtual servers*, 431-432
 - logical routers, 210-211, 214, 258
 - monitoring, 269
 - NAT, 416-418
 - DNAT*, 421
 - SNAT*, 420

- NSX 6.2, upgrading to, 566-567
 - NSX API calls (automation), 540-542
 - Perimeter Edge, 256
 - port groupings, 433
 - privileged mode commands, 272
 - protocols, 433
 - routing, 346-347
 - OSPF*, 351, 354-360
 - static routes*, 349-350
 - security, 269
 - Services, creating, 433
 - services of, 260
 - services provided by, 258
 - sizes of, 259
 - SSL VPN-Plus
 - adding installation packages*, 403-404
 - adding IP pools*, 405-406
 - adding private networks*, 406-408
 - authentication*, 398-402
 - components of*, 395-396
 - configuring*, 395-408
 - enabling SSL VPN-Plus service*, 402
 - security*, 398-402
 - server settings*, 396-397
 - verifying*, 408-410
 - Web Resources*, 397
 - Standby Edges becoming Active Edges, 261-262
 - undeployed NSX Edges, 267
 - user mode commands, 272
 - vCNS upgrades to NSX, 559
 - verifying, 269, 270
 - VLAN, 278
 - VM, NSX Edge as, 259
- NSX Manager**
- access control, SSO integration via RBAC, 522
 - admin credentials, 73
 - CLI commands, 71-72
 - configuring, 71-80
 - deploying, 69, 72
 - DFW, 460
 - updating rules*, 468
 - verifying*, 471
 - domain groups, 469-470
 - ESXi hosts, 113-119, 123
 - home page, 75
 - host preparation, 113-119
 - installing, 68
 - IP pools, 107-109
 - IP sets, 499-500
 - Layer 2 bridging, verifying, 303
 - logical routers
 - creating*, 204
 - verifying*, 211
 - logical switches
 - creating*, 131-134
 - Multicast Replication Mode*, 155
 - login screen, 73
 - MAC sets, 500
 - Network & Security access, 78-80
 - NSX Controller deployments, 100-101
 - NSX Controller verification, 102
 - NSX Edge, 256, 259-261
 - NSX Manager 6.2, upgrading to, 561-563
 - Primary role, 82
 - responsibilities of, 68, 82
 - REST and automation, 531
 - role assignments, 82
 - Secondary role, 82
 - security
 - services, deploying appliances*, 484-485
 - tags*, 497-499

- SpoofGuard, 471-473
- Transit role, 84
- vCNS upgrades to NSX, 555-557
- vPostgres, 74
- vSphere Web Client, 77

O

One-Arm (Proxy Mode) NSX Edge load balancers, 422-425

OSPF (Open Shortest Path First)

- ABR, 352-353, 356

- areas

- backbone areas, 352*

- inter-areas, 352*

- intra-areas, 352*

- normal areas, 354*

- NSSA, 354*

- stubby areas, 354*

- ASBR, 352-353

- authentication, 355

- configuring, 356-360

- IGP, 351

- link state, 351-352

- LSA, 355-356

- LSDB, 354

- neighbor adjacencies, 354-355

- route redistribution, 352

- security, 355

- verifying, 360-361

Outbound Activity reports, 513

Outbound AD Group Activity reports, 514

overlays, IP networks, 91

P

packet walks

- DFW, 456-458

- Layer 2 bridging, 303-307

- Layer 2 VPN, 290-296

- logical routers, 232-234

- example 1, 235-239*

- example 2, 240-246*

- example 3, 246-249*

- logical switches, 165-168

- example 1, 169-170*

- example 2, 170-171*

- example 3, 171-173, 176-177*

- example 4, 177-179, 185-188*

- example 5, 189-190*

Passive mode (LACP), 28

passwords, OSPF authentication, 355

PCI (Payment Card Industry) Data Security Standard, 506

Pearson Cert Practice Test engine, exam preparation, 573

- downloading/installing

- practice exams, 576*

- software, 576*

- exam activation codes, 577

- Practice Exam Mode, 578-579

- Premium Edition, 577

- Study Mode, 578-579

Perimeter Edge, 256

persistence states, NSX Edge load balancers, 426

physical networks

- Access Layer, 25

- access routers, 25

- Active Layer
 - collapsed Active Layer design, 30*
 - Spine and Leaf design, 31*
- challenges of, 6-7
- cloud-computing, 7
- firewalls, updating, 6
- infrastructure of, 24-28
- LACP, 27
- Layer 3 connectivity between virtual/physical networks
 - Designated Instances, 324-327*
 - ECMP, 336-339*
 - logical router VLAN LIF, 318, 321-324, 330, 336*
 - NSX Edge Gateways, 330-333*
- Layer 3 switches, 26
- NSX, 14, 33-34
- POD design, 28-29
- scaling, 26
- STP, 25
- SVI, 26
- ToR switches, 24-26
- virtualization, 7, 26
- virtual workloads, 7
- pinning virtual networks, 11**
- planes**
 - control plane, 65
 - data plane, 65
 - dependencies of, 64
 - management plane, 64-65
- pMAC**
 - Designated Instances, 324, 327
 - logical routers, 200
- POD (Point of Delivery) network design, 28-29**
- portgroups (vSS), 37-39, 42**
- ports**
 - dvUplinks, vDS, 43
 - NSX Edge port groupings, 433
 - sinkports, 296
 - VMkernel ports
 - ESXi hosts, 35*
 - vSS portgroups, 37, 41-42*
- POST API calls, 531**
- Practice Exam Mode (Pearson Cert Practice Test engine), 578-579**
- practice labs, exam preparation, 571**
- practice tests, exam preparation, 573**
 - downloading/installing
 - practice exams, 576*
 - software, 576*
 - exam activation codes, 577
- preparing for exams**
 - authorization requests, 573
 - certification requirements, 572
 - chapter resources, 573
 - delta exams, 572
 - exam blueprint, 572
 - exam day strategies, 574-575
 - exam details, 572-573
 - exam format, 571
 - exam requirements, 571
 - hands-on experience, 571
 - Pearson Cert Practice Test engine, 573
 - downloading/installing practice exams, 576*
 - downloading/installing software, 576*
 - exam activation codes, 577*
 - Practice Exam Mode, 578-579*
 - Premium Edition, 577*
 - Study Mode, 578-579*
 - practice labs, 571

- product pages, 585
- recommended training details, 573
- registering for the exam, 573
- review tools, 575
- updates, 585
- web resources, 571-579, 585

private networks

- profiles (vRA), 543-545
- SSL VPN-Plus, 406-408

privileged mode (Control VM), 215

product page updates, exam preparation, 585

Proxy (One-Arm) Mode (NSX Edge load balancers), 422-425

proxy VTEP, 156

PUT API calls, 531

Q-R

QoS (Quality of Service) marking, 55-57

ranking security policies (Service Composer), 496

RBAC (Role-Based Access Control)

- assigning roles, 523
- NSX Manager/SSO integration, 522
- NSX roles, 521-522
- user access control, 522
- user roles, 522

reconvergence, Ethernet, 10

recovery, NSX Controllers, 107

recursive routing, 347

Redirection API responses, 531

redistributing routes, OSPF routing, 352

registering

- for the exam (exam preparation), 573
- security services with NSX, 482-484

Replication Mode (logical switches), 152-154

requirements for exams, exam preparation, 571

resource pools, security services, 485

REST (Representative State Transfer)

- API, 530

- client headers, 531-532

- HTTP

- HTTP responses, 531*

- HTTP URI, 530*

- NSX API HTTP supported methods, 530*

- NSX API calls

- logical routers, 536-540*

- logical switches, 532-536*

- NSX Edge, 540-542*

- NSX Manager, 531

review tools, exam preparation, 575

routed network profiles (vRA), 543-545

routing

- ABR, 352-353, 356

- access routers, physical networks, 25

- administrative distance, 348

- ARP requests, recursive routing, 347

- ASBR, 352-353

- BDR, OSPF neighbor adjacencies, 355

- BGP

- ASN, 361-364*

- best path selection, 365*

- configuring, 366-367*

- eBGP peers, 362-364*

- iBGP peers, 362*

- route advertisements, 363-364*
 - synchronization, 364*
 - verifying, 368-369*
 - costs, 348
 - directly connected subnets/subnet masks, 346
 - distributed logical routers, NSX, 16
 - DR, OSPF neighbor adjacencies, 355
 - IS-IS
 - areas, 369*
 - configuring, 370-372*
 - Mesh Groups, 370*
 - route redistribution, 373-375*
 - router types, 369-370*
 - verifying, 373*
 - logical routers, 346-347
 - connectivity testing, 216-220*
 - Control VM, 201-203*
 - creating, 204*
 - deploying, 204-210*
 - distributed logical routers, 201*
 - distributed placement of, 199*
 - dvPortgroups, 199*
 - ESXi hosts, 200*
 - LIF, 198-200*
 - NSX API calls (automation), 536-540*
 - NSX Edge versus, 258*
 - OSPF, 351-360*
 - packet walks, 232-235*
 - packet walks, example 1, 235-239*
 - packet walks, example 2, 240-246*
 - packet walks, example 3, 246-249*
 - pMAC, 200*
 - routing, 346-347*
 - routing, OSPF, 351-360*
 - routing, static routes, 349-350*
 - static routes, 349-350*
 - transport zones, 201*
 - ULR, 201, 221-224*
 - verifying, 210-214*
 - VLAN LIF, 318, 321-324, 330, 336*
 - vMAC, 200*
 - XML tags, 536-537*
 - logical switches, transport zones, 201
 - NSX Edge, 346-347
 - OSPF, 351, 354-360*
 - static routes, 349-350*
 - OSPF, 351
 - ABR, 352-353, 356*
 - areas, 352-361*
 - recursive routing, 347
 - route redistribution, OSPF routing, 352
 - routing tables, 346
 - spokes, 349
 - static routes, 349-350
 - stubs, 349
 - subnets/subnet masks, 346
 - rules (DFW)**
 - creating, 462-468
 - default rules, 458-460
 - enforcing, 458
 - local rules, 462-464
 - memory, 459
 - thresholds/limits, 458-460
 - universal rules, 462-465
-
- S**
- saving DFW configurations, 468**
 - scaling physical networks, 26**
 - security**
 - Activity Monitoring, 509
 - data collection, enabling in VM, 510*
 - Inbound Activity reports, 512*

- Inter Container Interaction reports*, 513
- Outbound Activity reports*, 513
- Outbound AD Group Activity reports*, 514
- running activity reports*, 511
- viewing activity reports*, 514
- VM Activity reports*, 511
- AES-NI, 278
- data security
 - NSX Data Security*, 509
 - PCI Data Security Standard*, 506
 - vCNS Data Security*, 555
 - VMware Data Security*, 506-509
- DFW, 449, 453-455
 - configuring*, 460-468
 - creating rules*, 462-468
 - default rules*, 458-460
 - domain groups, adding to NSX Manager*, 469-470
 - enforcing rules*, 458
 - Exclusion Lists*, 460
 - local rules*, 462-464
 - memory*, 459
 - microsegmentation*, 457
 - packet walk*, 456-458
 - saving configurations*, 468
 - security groups*, 490
 - security policies*, 493-494
 - service redirection*, 496-497
 - thresholds/limits*, 458-460
 - universal rules*, 462-465
 - updating rules*, 467-468
 - vCNS upgrades to DFW*, 559
 - verifying*, 470-471
- firewalls
 - design compromises (traditional)*, 449, 453
 - DFW*, 449, 453-471, 490, 493-497, 559
 - inter-tier security*, 450, 453
 - Layer 2 firewalls*, 12
 - network security*, 12
 - NSX Edge firewall*, 436-442
 - physical networks*, 6
 - ToR switches*, 449, 453
 - updating*, 6
 - VLAN*, 451-453
 - VM*, 451-452
 - vMotion*, 453
- Flow Monitoring, 514
 - Allowed Flows option*, 516-517
 - Blocked Flows option*, 516-517
 - enabling*, 515
 - exporting collected flows*, 517
 - Flow Collectors*, 518
 - IPFix*, 517
 - Live Flow Monitoring*, 518
 - viewing collected flows*, 516
 - VM restrictions*, 515-516
- inta-tier security and firewalls, 450, 453
- IPsec VPN, 382, 389-390
- Layer 2 VPN, 278
- Layer 3 separation, 12
- network security, 12
- NSX, 15-16
- NSX Data Security, 509
- NSX Edge, 269, 436-442
- NSX Manager, Network & Security access, 78-80

- OSPF routing, 355
- passwords, 355
- PCI Data Security Standard, 506
- security groups (Service Composer)
 - creating*, 487-488
 - dynamic memberships*, 488-489
 - nested security groups*, 489
 - universal security groups*, 489-490
- security policies (Service Composer)
 - associating security policies with*, 495-496
 - associating with security groups*, 495-496
 - changing rank of*, 496
 - creating*, 491-494
 - DFW rules*, 493-494
- security tags, 497-499
- services
 - deploying appliances*, 484-485
 - DFW service redirection*, 496-497
 - Guest Introspection*, 480, 491
 - IP sets*, 499-500
 - MAC sets*, 500
 - Network Introspection*, 480-481, 491
 - registering with NSX*, 482-484
 - resource pools*, 485
 - security tags*, 497-499
 - Service Composer*, 486-496
 - VM*, 486
- SpoofGuard, 471-473
- SSL VPN-Plus, Web Resources, 398-402
- VM, 451-452
- VMware Data Security, 506-509
- Server Error API responses, 531**
- server pools (load balancers), 430**
- Service Composer, 16, 486**
 - security groups
 - associating security policies with*, 495-496
 - creating*, 487-488
 - dynamic memberships*, 488-489
 - nested security groups*, 489
 - universal security groups*, 489-490
 - security policies
 - associating with security groups*, 495-496
 - changing rank of*, 496
 - creating*, 491-494
 - DFW rules*, 493-494
- Service Monitor (Health Check), NSX Edge load balancers, 425**
- sinkports, 296**
- site-site IPsec VPN, 382-385, 394**
- slicing NSX Controllers, 106**
- SNAT (Source NAT), 416-420**
- SNMP (Simple Network Management Protocol), MIB, 530**
- Spine and Leaf network design, 31**
- spokes (routers), 349**
- SpoofGuard, 471-473**
- SR-IOV, VMNIC configuration, 40**
- SSL VPN-Plus**
 - components of, 395-396
 - configuring, 395
 - adding installation packages*, 403-404
 - adding IP pools*, 405-406
 - adding private networks*, 406-408
 - authentication*, 398-402
 - enabling SSL VPN-Plus service*, 402
 - server settings*, 396-397
 - Web Resources*, 397
 - verifying, 408-410

**SSO (Single Sign-On), NSX Manager/
SSO integration via RBAC, 522**

**standard portgroups. See VSS,
portgroups**

**Standby NSX Edges becoming Active
Edges, 261-262**

Static mode (LACP), 28

static routes, 349-350

**STP (Spanning Tree Protocol),
physical networks, 25**

stubby areas, OSPF routing, 354

stubs (routers), 349

**Study Mode (Pearson Cert Practice
Test engine), 578-579**

subnets/subnet masks, routing, 346

Success API responses, 531

**SVI (Switched Virtual Interface), phys-
ical networks, 26**

switches

access switches. *See* ToR switches

Layer 3 switches, physical networks, 26

logical switches

ARP requests, 147, 152

ARP tables, 147-152

creating, 131-134

dvPortgroups, 132-135

ESXi hosts, 130, 135, 139-153

global logical switches, 130-134

Hybrid Replication Mode, 156-157

LIF, 200

MAC learning, 131

MAC tables, 144-147, 150

Multicast Replication Mode, 154-155

NSX, 16

NSX API calls (automation), 532-536

packet walks, 165-168

packet walks, example 1, 169-170

packet walks, example 2, 170-171

*packet walks, example 3, 171-173,
176-177*

*packet walks, example 4, 177-179,
185-188*

packet walks, example 5, 189-190

Replication Mode, 152-154

tables, 138-139, 149-151

Unicast Replication Mode, 155-156

universal logical switches, 130-134

unknown unicasts, 152

verifying, 135

VM migration, 137

VNI, 130

VTEP, 130, 139-143, 150, 153-157

VTEP tables, 139-143, 150, 153

NSX vSwitch, 67

ToR switches

firewalls, 449, 453

NSX and physical networks, 33

vDS, 35

creating, 45

dvPortgroups, 43, 47-48, 54-57

dvUplinks, 43

ESXi hosts, 43, 49-51

LACP, 52

migrating to, 49-51

NSX, 45

topology view, 51

VLAN, 44-45

vSS, 35

ESXi hosts, 36, 43

forwarding decision rules, 38

MAC addresses, 37

portgroups, 37-39, 42

VMKernel portgroup configuration,
41-42

VMNIC configuration, 39

synchronization, BGP routing, 364

system event logs, verifying DFW, 471

T

**TCP/IP (Transport Control Protocol/
Internet Protocol), 10**

throttling, NSX Edge load
balancers, 426

ToR (Top of Rack) switches

firewalls, 449, 453

NSX and physical networks, 33

physical networks, 24-26

Traceflow, 519-520

traffic (data),

Flow Monitoring, 514

Allowed Flows option, 516-517

Blocked Flows option, 516-517

collected flows, 516-517

enabling, 515

Flow Collectors, 518

IPFix, 517

Live Flow Monitoring, 518

VM restrictions, 515-516

Traceflow, 519-520

**Transparent (In-Line) Mode (NSX
Edge load balancers), 422**

transport zones, 120

creating, 121

ESXi hosts, adding/removing, 122

GET transport zones, 532

logical routers, 201

logical switches, 201

modifying, 122

XML tags, 532-534

**TRILL (Transparent Interconnection
if Lots of Links), Spine and Leaf
network design, 31**

trunks, 9

**TSO (TCP Segmentation Offload),
VMNIC, 41**

tunnels (VXLAN), 91

U

**ULR (Universal Logical Routers), 201,
221-224**

undeployed NSX Edges, 267

underlays

IP networks, 91

VXLAN, 95

**Unicast Replication Mode (logical
switches), 155-156**

universal DFW rules, 462-465

universal logical switches, 130

assigning, 131

creating, 132-134

universal NSX Controllers, 102

**universal security groups (Service
Composer), 489-490**

unknown unicasts. *See also* BUM

Ethernet, 8-9

logical switches, 152

updating

DFW rules, 467-468

exam preparation, 585

firewalls, physical networks, 6

VTEP tables, 143

upgrading

NSX 6.1/6.1 to NSX 6.2

host clusters, 565

NSX Controllers, 563-564
NSX Edge, 566-567
NSX Manager 6.2, 561-563

vCNS to NSX

Data Security, 555
DFW, 559
NSX Edge, 559
NSX Manager, 555-557
NSX VIB, 558
partner integration, 555
USVM, 560

Uplink LIF, 200

uplink segments, 200

user mode (Control VM), 215

USVM (Universal Services Virtual Machine)

vCNS upgrades to NSX, 560
 VMware Data Security, 506-508

UTEF (Unicast proxy VTEF), 156

V

vCenter

MOID, 135
 NSX, 13, 34
cross vCenter NSX, 80-84
integration, 73-77
NSX Controller deployments, 101
NSX Controller verification, 104

vDS, 135

vMotion, 147

vCNS (vCloud Network and Security), upgrading to NSX

Data Security, 555
 DFW, 559
 NSX Edge, 559
 NSX Manager, 555-557

NSX VIB, 558
 partner integration, 555
 USVM, 560

vDS (vSphere Distributed Switches), 35

creating, 45
 dvPortgroups, 43
configuring, 48, 54-57
creating, 47
deleting, 48
LACP, 54
QoS marking, 55-57

dvUplinks, 43

ESXi hosts, 43, 49-51

LACP, 52

migrating to, 49-51

MOID, 135

NSX, 45

sinkports, 296

topology of, 51

VLAN, 44-45

verifying

BGP routing, 368-369
 DFW, 470-471
 dvPortgroups, 135
 IPsec VPN, 392-394
 IS-IS routing, 373
 Layer 2 bridging, 301-303
 Layer 2 VPN, 289
 logical routers, 210-214
 logical switches, 135, 149-151
 NSX Controllers, 101-103
 NSX Edge, 269-270
 OSPF routing, 360-361
 SSL VPN-Plus, 408-410

versions, exam updates, 585

VIB (vSphere Infrastructure Bubbles), 109

- ESXi hosts, 110-111
- host clusters, upgrading to NSX 6.2, 565
- vCNS upgrades to NSX, 558
- VXLAN, 112

VIP (Virtual IP), load balancers, 421**virtual networks**

- Layer 3 connectivity between virtual/physical networks
 - Designated Instances, 324-327*
 - ECMP, 336-339*
 - logical router VLAN LIF, 318, 321-324, 330, 336*
 - NSX Edge Gateways, 330-333*
- pinning, 11

virtual servers, load balancers, 431-432**virtual workloads, physical networks, 7****virtualization**

- NFV, 14, 17
- physical networks, 7, 26

VLAN (Virtual LAN)

- Edge VLAN, 256
- Ethernet, 9
- firewalls, 451-453
- hardware VTEP, 308, 311-312
- Layer 2 bridging, 296-300
- NSX Edge, 278
- vDS, 44-45
- vSS portgroups, 38-39

VLAN LIF, 200, 318, 321-324, 330, 336**VM (Virtual Machines)**

- Activity Monitoring
 - data collection, enabling in VM, 510*
 - VM Activity reports, 511*

connectivity testing, 216-220

Control VM

- logical routers, 201-203, 213*
- privileged mode, 215*
- user mode, 215*

DFW, 454-460

firewalls, 451-452

Flow Monitoring, VM restrictions, 515-516

logical switches, VM migration, 137

NSX Edge as, 259

security groups (Service Composer), 487

security services, 486

security tags, 497-499

SpoofGuard, 471-473

USVM

- vCNS upgrades to NSX, 560*
- VMware Data Security, 506-508*

vMAC (Virtual MAC)

- Designated Instances, 328
- logical routers, 200

VMkernel ports

- ESXi hosts, 35
- vSS portgroups, 37, 41-42

VMKNic Teaming Policies, ESXi host configuration, 116-118**VMNIC (Virtual Machine Network Interface Cards)**

- ESXi hosts, 35
- LACP, 53-54
- SR-IOV, 40
- TSO, 41
- vSS, 39

vMotion, 147

- firewalls, 453
- NSX and physical networks, 33

VMware Certification website, 571

VMware Data Security, 506-509

VMware NSX. *See* NSX

VNI (VXLAN Network Identifiers)

ESXi hosts, 151

logical switches, 130

pools, 120

vNIC (virtual Network Interface Cards)

DFW, 455, 458, 466-468

security groups (Service Composer), 487

SpoofGuard, 471-473

vSS portgroups, 37

VPN (Virtual Private Networks)

IPsec VPN

configuring, 386-390

creating, 391

encryption algorithms, 389-390

establishing, 384-385

security, 389-390

security features, 382

site-site IPsec VPN, 382-385, 394

support for, 386

verifying, 392-394

Layer 2 VPN

configuring, 280-285, 288

encryption, 278

NSX Edge, 278-282, 286-287

packet walks, 290-296

security, 278

verifying, 289

SSL VPN-Plus

adding installation packages, 403-404

adding IP pools, 405-406

adding private networks, 406-408

authentication, 398-402

components of, 395-396

configuring, 395-408

enabling SSL VPN-Plus service, 402

security, 398-402

server settings, 396-397

verifying, 408-410

Web Resources, 397

vPostgres, 74

vRA (vRealize Automation), 542

external network profiles, 543-544

NAT network profiles, 543-546

private network profiles, 543-545

routed network profiles, 543-545

vRO, 543

vSphere

ESXi hosts, network connectivity, 35

NSX, 34-35

vDS, 35

creating, 45

dvPortgroups, 43, 47-48, 54-57

dvUplinks, 43

ESXi hosts, 43, 49-51

LACP, 52

migrating to, 49-51

NSX, 45

topology view, 51

VLAN, 44-45

VIB, 109

ESXi hosts, 110-111

VXLAN, 112

vSS, 35

ESXi hosts, 36, 43

forwarding decision rules, 38

MAC addresses, 37

portgroups, 37-39, 42

VMkernel portgroup configuration,
41-42

VMNIC configuration, 39

vSphere Web Client

logical routers, deploying, 204-210

logical switches, 132

NSX Manager, configuring, 77

vSS (vSphere Standard Switches), 35

configuring

VMkernel portgroup configuration,
41-42

VMNIC configuration, 39

ESXi hosts, 36, 43

forwarding decision rules, 38

MAC addresses, 37

portgroups, 37-39, 42

VTEP (VXLAN Tunnel

Endpoints), 92

hardware VTEP, 308, 311-312

logical switches, 130

Hybrid Replication Mode, 156-157

Multicast Replication Mode, 154-155

Unicast Replication Mode, 155-156

VTEP tables, 139-143, 150, 153

proxy VTEP, 156

tables, NSX Controllers, 96

UTEF, 156

VXLAN (Virtual Extensible LAN)

encapsulation, 92, 95

ESXi hosts, 90

hardware VTEP, 308, 311-312

host configuration, 113-119

Layer 2 bridging, 296-300

tunnels, 91

underlays, 95

VIB, 112

VTEP, 92

W

web resources

exam preparation, 571, 585

Pearson Cert Practice Test engine,
573, 576-579

review tools, 575

VMware Certification website, 572

SSL VPN-Plus, 397

VMware Certification website, 571

X-Y-Z

XML (Extensible Markup Language)

logical router XML tags, 536-537

transport zone XML tags, 532-534