

PEARSON IT  
CERTIFICATION

vmware®



Practice  
Tests



Flash  
Cards



Review  
Exercises



Study  
Planner

# Official Cert Guide

Advance your IT career with hands-on learning

# VCP-DCV for vSphere 7.x

Exam 2V0-21.20



JOHN A. DAVIS  
STEVE BACA  
OWEN THOMAS

## About This eBook

ePUB is an open, industry-standard format for eBooks. However, support of ePUB and its many features varies across reading devices and applications. Use your device or app settings to customize the presentation to your liking. Settings that you can customize often include font, font size, single or double column, landscape or portrait mode, and figures that you can click or tap to enlarge. For additional information about the settings and features on your reading device or app, visit the device manufacturer's Web site.

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a "Click here to view code image" link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.

# **VCP-DCV for vSphere 7.x (Exam 2V0-21.20) Official Cert Guide**

**John A. Davis, Steve Baca, Owen Thomas**



# **VCP-DCV for vSphere 7.x (Exam 2V0-21.20) Official Cert Guide**

Copyright © 2021 by Pearson Education, Inc.

Published by Pearson Education, Inc.

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-10: 0-13-589819-6

ISBN-13: 978-0-13-589819-2

ScoutAutomatedPrintCode

Library of Congress Control Number: 2020923071

## **Trademarks**

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

## **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 author 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.

## **Special Sales**



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

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

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

**EDITOR-IN-CHIEF**

Mark Taub

**DIRECTOR, ITP PRODUCT MANAGEMENT**

Brett Bartow

**EXECUTIVE EDITOR**

Nancy Davis

**TECHNICAL EDITOR**

Joseph Cooper

**DEVELOPMENT EDITOR**

Ellie Bru

**MANAGING EDITOR**

Sandra Schroeder

**PROJECT EDITOR**

Mandie Frank

**COPY EDITOR**

Kitty Wilson

**PROOFREADER**

Betty Pessagno

**INDEXER**

Erika Millen

**EDITORIAL ASSISTANT**

Cindy Teeters

**DESIGNER**

Chuti Prasertsith

**COMPOSITOR**

codeMantra

# Credits

## Figure

## Attribution/Credit Line

Chapter Opener

Charlie  
Edwards/Photodisc/G  
etty Images

Figure 3-2, Figure 3-4

vSphere Networking  
Guide

Figure 5-1, Figure 5-2, Figure 5-3, Figure 5-4, Figure 8-1,  
Figure 10-1, Figure 10-2, Figure 10-3, Figure 10-4

VMware Hands on  
Lab

Figure 13-1, Figure 13-2

Screenshot from  
VMware Hands on  
Labs

## **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: [www.informit.com](http://www.informit.com)

# Reader Services

Visit our website at [www.pearsonitcertification.com/title/9780135898192](http://www.pearsonitcertification.com/title/9780135898192) and register this book for convenient access to any updates, downloads, or errata that might be available for this book.

# Dedications

*Dedicated to Madison, Emma, Jaxon, Ethan, Eli, and Robbie, the six wonderful children to whom I am blessed to be known as “Grampy.” They fill my days with joy and fun, especially after a hard day of writing or working for their namesake, MEJEER, LLC.*

—John Davis

*First and foremost, I would like to dedicate this book to my loving wife, Sharyl. Without your support, I would not be able to commit the time necessary to co-author a book. Thank you for believing in me and allowing me to have the time for my many endeavors. I would also like to dedicate this book to my children: Zachary, Brianna, Eileen, Susan, Keenan, and Maura.*

—Steve Baca

*I would like to dedicate this book to my wife, Angela, and our daughter, Emma. May it be a reminder of pushing for bigger and brighter things in life. I love you both with all of my heart.*

—Owen Thomas

## About the Authors

**John A. Davis**, now an independent contractor and senior integration architect at MEJEER, LLC, became a VMware Certified Instructor (VCI) and VMware Certified Professional (VCP) in 2004. Since then, all of his work has focused on VMware-based technologies. He has experience in teaching official VMware curriculum in five countries and delivering VMware professional services throughout the United States. Recently, his work has involved designing and implementing solutions for hybrid clouds, cloud automation, disaster recovery, and virtual desktop infrastructure (VDI). He has authored several white papers and co-authored *VCP6-DCV Cert Guide* and *VCAP5-DCA Cert Guide* (VMware Press). He holds several advanced certifications, including VCIX6-DCV, VCAP5-DTD, VCAP5-CID, and VCIX6-NV. He has been a vExpert since 2014. He is the author of the vLoreBlog.com and can be found on Twitter @johnnyadavis.

**Steve Baca**, VCAP, VCI, VCP, and NCDA, has been in the computer industry for more than 20 years. Originally a computer programmer and a system administrator working on Unix and Windows systems, he migrated over to technical training and wrote a course for Sun Microsystems. After teaching various courses for Sun, he eventually transitioned to VMware about 10 years ago, to do technical training. Currently he is a badged employee for VMware and lives in Omaha, Nebraska. He thoroughly enjoys teaching and writing and believes that the constant evolution of the computer industry requires continuously learning to stay ahead. Steve can be found on Twitter @scbaca1.

**Owen Thomas** holds a number of VMware certifications and has taught more than 400 authorized VMware classes for vSphere, Horizon, vCloud, and vRealize products. He has operated as a VMware Solutions Provider and performed a number of VMware partner services for customers across the United States.

## About the Reviewer

**Joseph Cooper** is a Principal Instructor and a member of America's Tech Lead Team with VMware's Education Department. Joe has spoken at several VMworld conferences, VMUG events, and vForum events, and is a featured instructor in the VMware Learning Zone. Prior to joining VMware, Joe was an instructor at the State University of New York, College at Cortland, where he taught technology courses to Sport Management and Kinesiology students.

You can find him on Twitter @joeicooper and on the newly launched YouTube channel <https://youtube.com/channel/UCYrPi0AqS8f8QxChAgZa5Sg>.



# Acknowledgments

Thanks to my wife and best friend, Delores, who tolerates my late-night writing, supports my recent business venture, and makes me happy every day. Thanks to my parents, Monica and Norman Davis, who provided me with a great education and taught me the importance of hard work. Thanks to God for placing me in an environment with unmeasurable blessings and opportunities.

I would like to thank my co-authors and partners, Steve Baca and Owen Thomas. Thanks to our technical editor, Joe Cooper, for his hard work and dedication. Special thanks to Nancy Davis (executive editor) and Ellie Bru (development editor) for coordinating everything and keeping this project moving.

—John Davis

There are so many people to acknowledge and thank for making this book possible. First, thanks to my wife and family for supporting me while writing this book. I would also like to thank my fellow co-authors, John Davis and Owen Thomas, who deserve much of the credit for this book. Thank you to the production team and editors at Pearson, who do a tremendous amount of work from the initial planning of the book to the final printing.

—Steve Baca

Thank you to my wife, Angela, and our daughter, Emma, for your patience with me while I worked on this book. Thank you, John Davis, for working with me yet again. It is always a pleasure, and I hope to work with you more in the future. Thank you, Pearson, for letting us write another version of this book and for all of the awesome support. And thank you, VMware, for continuing to trailblaze.

—Owen Thomas

# Contents at a Glance

Foreword

Introduction

## **Part I: vSphere Architecture, Integration, and Requirements**

**Chapter 1** vSphere Overview, Components, and Requirements

**Chapter 2** Storage Infrastructure

**Chapter 3** Network Infrastructure

**Chapter 4** Clusters and High Availability

**Chapter 5** vCenter Server Features and Virtual Machines

**Chapter 6** VMware Product Integration

**Chapter 7** vSphere Security

## **Part II: vSphere Installation/Configuration**

**Chapter 8** vSphere Installation

**Chapter 9** Configuring and Managing Virtual Networks

## **Part III: vSphere Management and Optimization**

**Chapter 10** Managing and Monitoring Clusters and Resources

**Chapter 11** Managing Storage

**Chapter 12** Managing vSphere Security

**Chapter 13** Managing vSphere and vCenter Server

**Chapter 14** Managing Virtual Machines

**Chapter 15** Final Preparation

**Appendix A** Answers to the “Do I Know This Already?” Quizzes and Review Questions

Index

**Online Elements:**

**Appendix B** Memory Tables

**Appendix C** Memory Tables Answer Key

**Appendix D** Study Planner

Glossary

# Table of Contents

**Foreword**

**Introduction**

**Part I: vSphere Architecture, Integration, and Requirements**

**Chapter 1 vSphere Overview, Components, and Requirements**

“Do I Know This Already?” Quiz

**Foundation Topics**

vSphere Components and Editions

    vSphere Components

    Editions and Licenses

vCenter Server Topology

    Single Sign-On (SSO) Domain

    Enhanced Linked Mode

    vCenter HA

Infrastructure Requirements

    Compute and System Requirements

    Storage Requirements

    Network Requirements

    Infrastructure Services

Other Requirements

    Additional Requirements

    vSphere Replication Requirements

    vCenter High Availability Requirements

    SDDC Requirements

VMware Cloud vs. VMware Virtualization

- Server Virtualization
- VMware SDDC
- vCloud Suite and Private Clouds
- VCF and Hybrid Clouds
- VMC on AWS
- VMware vCloud Director
- Cloud Automation

### **Exam Preparation Tasks**

- Review All the Key Topics
- Complete Tables and Lists from Memory
- Define Key Terms
- Answer Review Questions

## **Chapter 2 Storage Infrastructure**

“Do I Know This Already?” Quiz

### **Foundation Topics**

- Storage Models and Datastore Types
  - How Virtual Machines Access Storage
  - Storage Virtualization: The Traditional Model
  - Software-Defined Storage Models
  - Datastore Types
  - Storage in vSphere with Kubernetes
  - VMware NVMe
- vSAN Concepts
  - vSAN Characteristics
  - vSAN Terminology
  - What Is New in vSAN 7.0
  - vSAN Deployment Options
  - vSAN Limitations
  - vSAN Space Efficiency
  - vSAN Encryption
  - vSAN File Services

- vSAN Requirements
- Other vSAN Considerations
- vSphere Storage Integration
  - VASA
  - VAAI
  - Virtual Volumes (vVols)
- Storage Multipathing and Failover
  - Multipathing Overview
  - Pluggable Storage Architecture (PSA)
- Storage Policies
  - Storage Policy Based Management (SPBM)
  - Virtual Disk Types
  - vSAN-Specific Storage Policies
- Storage DRS (SDRS)
  - Initial Placement and Ongoing Balancing
  - Space Utilization Load Balancing
  - I/O Latency Load Balancing
  - SDRS Automation Level
  - SDRS Thresholds and Behavior
  - SDRS Recommendations
  - Anti-affinity Rules
  - Datastore Cluster Requirements
  - NIOC, SIOC, and SDRS
- Exam Preparation Tasks**
  - Review All Key Topics
  - Complete Tables and Lists from Memory
  - Define Key Terms
  - Review Questions

## **Chapter 3 Network Infrastructure**

“Do I Know This Already?” Quiz

## **Foundation Topics**

### Networking Terms and Concepts

- Traditional Networking Terminology

- Virtual NICs

- Virtual Switch Concepts

- VLANs

### vSphere Standard Switch (vSS)

- MTU

- vSS Network Policies

- NIC Teaming Policies

- Network Security Policies

- Traffic Shaping Policy

- VLAN Policies

### Distributed Virtual Switch (vDS)

- Distributed Port Groups

- Uplink Port Groups

- vSS and vDS Comparison

- vDS Network Policies

- Inbound Traffic Shaping

- Port-Blocking Policies

- Load-Based NIC Teaming

- Resource Allocation Policy

- NetFlow and Monitoring Policy

- Traffic Filtering and Marking Policy

### vDS Settings and Features

- Private VLANs

- Data Center–Level Management

- Port State Monitoring

- Port State with vMotion

- Port Mirroring

- Port Binding and Allocation

LACP Support

vDS Health Check

Other vSphere Networking Features

Multicast Filtering Mode

Discovery Protocol

TCP Segmentation Offload

DirectPath I/O

Single Root I/O Virtualization (SR-IOV)

VMkernel Networking and TCP/IP Stacks

### **Exam Preparation Tasks**

Review All Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## **Chapter 4 Clusters and High Availability**

“Do I Know This Already?” Quiz

### **Foundation Topics**

Cluster Concepts and Overview

Enhanced vMotion Compatibility (EVC)

vSAN Services

Distributed Resource Scheduler (DRS)

Recent DRS Enhancements

DRS Rules

DRS Migration Sensitivity

Resource Pools

vSphere High Availability (HA)

vSphere HA Requirements

vSphere HA Response to Failures

Heartbeats

vSphere HA Admission Control

vSphere HA Advanced Options



- Virtual Machine Settings
- VM Component Protection (VMCP)
- Virtual Machine and Application Monitoring
- vSphere HA Best Practices
- Proactive HA

Other Resource Management and Availability Features

- Predictive DRS
- Distributed Power Management (DPM)
- Fault Tolerance (FT)
- vCenter Server High Availability
- VMware Service Lifecycle Manager

### **Exam Preparation Tasks**

- Review All Key Topics
- Complete Tables and Lists from Memory
- Define Key Terms
- Review Questions

## **Chapter 5 vCenter Server Features and Virtual Machines**

“Do I Know This Already?” Quiz

### **Foundation Topics**

vCenter Server and vSphere

- vSphere Managed Inventory Objects

- Host Profiles

- Content Libraries

- vSphere with Tanzu

Virtual Machine File Structure

- Configuration File

- Virtual Disk Files

- Snapshot Files

Virtual Machine Snapshots

- Snapshot Use Cases

- What a Snapshot Preserves

- Parent Snapshots
- Snapshot Behavior
- Limitations
- Virtual Machine Settings
  - VM Hardware/Compatibility
  - Virtual Disk Provisioning
  - VMware Tools
  - Virtual Machine Options
  - Virtual Machine Advanced Settings
- Virtual Machine Migration
  - Virtual Machine Migration
  - vMotion Details
  - Storage vMotion Details
- Virtual Machine Cloning
  - Clones
  - Rapid Provisioning with Templates
  - Instant Clones
- Exam Preparation Tasks**
  - Review All Key Topics
  - Complete Tables and Lists from Memory
  - Define Key Terms
  - Review Questions

## **Chapter 6 VMware Product Integration**

“Do I Know This Already?” Quiz

### **Foundation Topics**

vSphere Add-ons

- vSphere with Tanzu

- vCenter Converter

- VMware vSphere Replication

- VMware SkyLine

vRealize Suite

vRealize Operations Manager (vROps)

vRealize Log Insight (vRLI)

vRealize Automation (vRA)

vRealize Orchestrator (vRO)

vRealize Network Insight (vRNI)

Desktop and Application Virtualization

VMware Horizon

App Volumes

Replication and Disaster Recovery

vSphere Replication

Site Recovery Manager (SRM)

Private, Public, and Hybrid Clouds

VMware Cloud Foundation (VCF)

VMware Hybrid Cloud Extension (HCX)

VMware on AWS

Azure VMware Solution

Networking and Security

AppDefense

NSX

**Exam Preparation Tasks**

Review All Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## **Chapter 7 vSphere Security**

“Do I Know This Already?” Quiz

**Foundation Topics**

vSphere Certificates

vSphere Certificates Overview

Certificate Requirements

ESXi Host Certificates

## vSphere Permissions

- Authentication and Authorization

- Inventory Hierarchy and Objects

- Privileges and Roles

- Permissions

- Global Permissions

- Best Practices for Roles and Permissions

- Required Privileges for Common Tasks

- How Permissions Are Applied by vCenter Server

## ESXi and vCenter Server Security

- Built-in Security Features

- Security Profiles

- ESXi Password Hardening

- Joining an ESXi Host to a Directory Service

- vSphere Authentication Proxy

- ESXi Host Access

- Control MOB Access

- ESXi Secure Boot and TPM

- vSphere Trust Authority (vTA)

- vCenter Server Security

## vSphere Network Security

### Virtual Machine Security

- Virtual Machine Hardening Best Practices

- Configuring UEFI Boot

- Disabling Unexposed Features

- Other Common Settings

- Virtual Machine Risk Profiles

- Protecting Virtual Machines Against Denial-of-Service Attacks

- Controlling VM Device Connections

- Virtual Machine Encryption

- Encrypted vSphere vMotion

virtual Trusted Platform Module (vTPM)  
virtual Intel Software Guard Extension (vSGX)

Available Add-on Security

Compliance Using vRealize Operations Manager  
VMware NSX  
AppDefense

### **Exam Preparation Tasks**

Review All the Key Topics  
Complete Tables and Lists from Memory  
Define Key Terms  
Review Questions

## **Part II: vSphere Installation/Configuration**

### **Chapter 8 vSphere Installation**

“Do I Know This Already?” Quiz

#### **Foundation Topics**

Installing ESXi Hosts

Installing ESXi Interactively  
Scripted ESXi Installation  
Using Auto Deploy

Deploying vCenter Server Components

vCenter Server Database  
Platform Services Controller (PSC)  
vCenter Server Appliance  
Configuring and Managing VMware Certificate Authority (VMCA)

Configuring Single Sign-On (SSO)

SSO and Identity Sources Overview  
Adding, Editing, and Removing SSO Identity Sources  
Adding an Active Directory Identity Source  
Adding an LDAP Authentication Source  
Enabling and Disabling Single Sign-On (SSO) Users

- Configuring SSO Policies
- Configuring Identity Federation

#### Initial vSphere Configuration

- Implementing vSphere Client
- Implementing VMware vSphere Lifecycle Manager
- Configuring the vCenter Server Inventory
- Implementing vCenter HA
- Using Host Profiles
- VMware Tools
- Advanced ESXi Host Options

#### **Exam Preparation Tasks**

- Review All the Key Topics
- Complete Tables and Lists from Memory
- Define Key Terms
- Review Questions

### **Chapter 9 Configuring and Managing Virtual Networks**

- “Do I Know This Already?” Quiz

#### **Foundation Topics**

##### vSphere Standard Switches (vSS)

- Creating and Configuring vSphere Standard Switches
- Creating and Configuring Standard Port Groups

##### vSphere Distributed Switches (vDS)

- Creating and Configuring vSphere Distributed Switches
- Creating and Configuring Distributed Port Groups

##### VMkernel Networking

- Configuring and Managing VMkernel Adapters
- Configuring TCP/IP Stacks

##### Configuring and Managing Networking Features

- Configuring Network I/O Control (NIOC)
- Creating a Network Resource Pool
- Using Private VLANs

- Using DirectPath I/O
- Single Root I/O Virtualization (SR-IOV)
- Configuring and Managing Port Mirroring
- Configuring and Managing Link Aggregation Groups (LAGs)
- Managing Host Networking with vDS
  - Adding Hosts to a vDS
  - Managing Host Physical Network Adapters on a vDS
  - Migrating VMkernel Network Adapters to a vDS
  - Removing Hosts from a vDS
  - Migrating Virtual Machines to a vDS
  - Monitoring the State of Ports in a Distributed Port Group
  - Using the vDS Health Check
  - Networking Policies and Advanced Features

### **Exam Preparation Tasks**

- Review All the Key Topics
- Complete Tables and Lists from Memory
- Define Key Terms
- Review Questions

## **Part III: vSphere Management and Optimization**

### **Chapter 10 Managing and Monitoring Clusters and Resources**

“Do I Know This Already?” Quiz

#### **Foundation Topics**

Creating and Configuring a vSphere Cluster

- Creating a Cluster

- Configuring a Cluster with Quickstart

- EVC Mode

Creating and Configuring a vSphere DRS Cluster

- Creating a vSphere DRS Cluster

- Creating a Resource Pool

- Configuring Advanced DRS Options

## Creating and Configuring a vSphere HA Cluster

- Creating a vSphere HA Cluster

- Configuring Advanced vSphere HA Options

- Configuring vSphere HA Admission Control

- Configuring VMCP

- Configuring Virtual Machine and Application Monitoring

- Configuring Proactive HA

- Configuring vSphere Fault Tolerance

## Monitoring and Managing vSphere Resources

- Metrics

- vSphere Client Performance Charts

- Troubleshooting and Optimizing Performance

- Monitoring and Managing Cluster Resources

- Monitoring and Managing Resource Pool Resources

- Monitoring and Managing Host Resources and Health

- Monitoring and Managing Virtual Machine Resources

- ESXTOP

- VIMTOP

- vCenter Server Appliance Management Interface (VAMI)

## Events, Alarms, and Automated Actions

- Events

- Viewing Events in the vSphere Client

- Viewing the System Event Log

- Streaming Events to a Remote Syslog Server

- Alarms

- Viewing and Acknowledging Triggered Alarms

- Creating Alarm Definitions

- Alarm Actions

- Advanced Use Cases for Alarms

## Logging in vSphere

- ESXi Logs



- vCenter Server Logs
- Uploading System Logs to VMware
- Log Levels
- Configuring Syslog on ESXi Hosts
- vRealize Log Insight (vRLI)

### **Exam Preparation Tasks**

- Review All the Key Topics
- Complete Tables and Lists from Memory
- Define Key Terms
- Review Questions

## **Chapter 11 Managing Storage**

- “Do I Know This Already?” Quiz

### **Foundation Topics**

- Configuring and Managing vSAN
  - Preparing for vSAN
  - Creating a vSAN Cluster with Quickstart
  - Manually Enabling vSAN
  - Editing vSAN Settings
  - Licensing vSAN
  - Viewing a vSAN Datastore
  - Configuring vSAN and vSphere HA
  - Disabling vSAN
  - Shutting Down and Restarting vSAN
  - Deploying vSAN with vCenter Server
  - Expanding a vSAN Cluster
  - Working with Maintenance Mode
  - Managing vSAN Fault Domains
  - Extending a vSAN Datastore Across Two Sites
  - Managing Devices in a vSAN Cluster
  - Increasing Space Efficiency in a vSAN Cluster
  - Using Encryption in a vSAN Cluster

- Using vSAN Policies
- Viewing vSAN Storage Providers
- Using vSAN File Service

#### Managing Datastores

- Managing VMFS Datastores
- Managing Raw Device Mappings (RDMs)
- Managing NFS Datastores

#### Storage DRS and SIOC

- Configuring and Managing Storage DRS
- Configuring and Managing SIOC

#### NVMe and PMem

- Managing VMware NVMe
- Managing PMem

#### Multipathing, Storage Policies, and vVols

- Managing Multipathing
- Managing Storage Policies
- Configuring and Managing vVols

### **Exam Preparation Tasks**

- Review All the Key Topics
- Complete Tables and Lists from Memory
- Define Key Terms
- Review Questions

## **Chapter 12 Managing vSphere Security**

- “Do I Know This Already?” Quiz

### **Foundation Topics**

- Configuring and Managing Authentication and Authorization

- Managing SSO
- Users and Groups
- Privileges and Roles
- Permissions
- Global Permissions

- Editing Permissions
- Configuring and Managing vSphere Certificates
  - Managing vSphere Client Certificates
  - Using Custom Certificates
  - Managing ESXi Certificates
- General ESXi Security Recommendations
  - Configuring ESXi Using Host Profiles
  - Using Scripts to Manage Host Configuration Settings
  - ESXi Passwords and Account Lockout
  - SSH and ESXi Shell Security
  - PCI and PCIe Devices and ESXi
  - Disabling the Managed Object Browser
  - ESXi Networking Security Recommendations
  - ESXi Web Proxy Settings
  - vSphere Auto Deploy Security Considerations
  - Controlling CIM Access
- Configuring and Managing ESXi Security
  - Configuring the ESXi Firewall
  - Customizing ESXi Services
  - Using Lockdown Mode
  - Managing the Acceptance Levels of Hosts and VIBs
  - Assigning Privileges for ESXi Hosts
  - Using Active Directory to Manage ESXi Users
  - Configuring vSphere Authentication Proxy
  - Configuring Smart Card Authentication for ESXi
  - Configuring UEFI Secure Boot for ESXi Hosts
  - Securing ESXi Hosts with Trusted Platform Module
  - Securing ESXi Log Files
- Additional Security Management
  - Key Management Server
  - Changing Permission Validation Settings

Configuring and Managing vSphere Trust Authority (vTA)  
Securing Virtual Machines with Intel Software Guard Extensions (SGX)  
Encrypting a Virtual Machine

### **Exam Preparation Tasks**

Review All the Key Topics  
Complete Tables and Lists from Memory  
Define Key Terms  
Review Questions

## **Chapter 13 Managing vSphere and vCenter Server**

“Do I Know This Already?” Quiz

### **Foundation Topics**

vCenter Server Backup

Upgrading to vSphere 7.0

- vCenter Server Data Transfer

- Upgrading vCenter Server Appliance

- Migrating vCenter Server for Windows to vCenter Server Appliance

- Upgrading ESXi and Virtual Machines

- Using Update Planner

Using vSphere Lifecycle Manager

- About VMware Update Manager

- VMware Update Manager Download Service (UMDS)

- Baselines and Images

- ESXi Quick Boot

- ESXi Firmware Updates

- Hardware Compatibility Checks

- Exporting and Importing Cluster Images

- Backup and Restore Scenarios

- Upgrading Virtual Machines

Managing ESXi Hosts

Monitoring and Managing vCenter Server

- Monitoring and Managing vCenter Server with the VAMI
- Monitoring and Managing vCenter Server with the vSphere Client
- Updating the vCenter Server
- Managing a vCenter HA Cluster
- Repointing a vCenter Server to Another Domain

### **Exam Preparation Tasks**

- Review All the Key Topics
- Complete Tables and Lists from Memory
- Define Key Terms
- Review Questions

## **Chapter 14 Managing Virtual Machines**

“Do I Know This Already?” Quiz

### **Foundation Topics**

Creating and Configuring Virtual Machines

- Creating a New Virtual Machine
- Powering on a VM
- Opening a Console to a VM
- Installing and Upgrading VMware Tools
- Shutting Down a Guest
- Cloning a Virtual Machine
- Converting Between a VM and a Template
- Deploying a Virtual Machine from a Template
- Customizing the Guest OS
- Deploying OVF/OVA Templates

Managing Virtual Machines

- Configuring Virtual Machine Hardware
- Editing Virtual Machine Options
- Configuring Guest User Mappings
- Editing OVF Details
- Creating and Managing Virtual Machine Snapshots
- Migrating Virtual Machines

## Advanced Virtual Machine Management

Managing OVF Templates

Virtualization-Based Security

Managing VMs by Using PowerCLI

Configuring VMs to Support vGPUs

## Content Libraries

Introduction to Content Libraries

Creating a Content Library

Publishing a Content Library

Subscribing to a Content Library

Content Library Permissions

Content Library Synchronization Options

Adding Items to a Content Library

Deploying VMs by Using a Content Library

## **Exam Preparation Tasks**

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## **Chapter 15 Final Preparation**

Getting Ready

Taking the Exam

## **Appendix A Answers to the “Do I Know This Already?” Quizzes and Review Questions**

## **Index**

## **Online Elements:**

## **Appendix B Memory Tables**

## **Appendix C Memory Table Answer Key**

## **Appendix D Study Planner**

### **Glossary**

# Foreword

Certification exams aren't easy. They're not supposed to be. If they were, they wouldn't mean much.

Certifications validate a specific minimum level of expertise of an individual, and in the case of VMware certifications, that means we, as a company, stand behind those individuals earning the certification. We create exams that are aimed at real job skills, that align to actual job roles that exist in the industry, and that properly test the baseline expertise required to perform those roles.

The authors of this book have multiple certifications among themselves, equaling decades of hands-on experience. They are teachers, learners, administrators, and architects of key IT technologies. Their combined knowledge provides them the ability to share their expertise through this book, which in turn allows you, as the reader and prospective certification holder, to be better prepared to pass that hard certification exam.

This study guide is a great asset and tool for you to use in your study and preparation. Take advantage of the practice exams, the suggestions and tips, and the content in the book. By using this guide and thoroughly preparing yourself, getting hands-on practice through labs and sandbox or production environments, and paying close attention to the objectives the exam will cover, you will be prepared to not only pass the exams on your way to getting certified but add real value to your organizations through a software-defined approach to business and IT.

I wish you the best of luck in your study and subsequent exam attempt. And when you earn that certification, remember: You've just done something hard. And that means something to VMware, to the industry, and to you. Well done!

Karl Childs

Senior Manager, VMware Certification



# Introduction

This book focuses on one major goal: helping you prepare to pass the Professional VMware vSphere 7.0 (2V0-21.20) exam, which is a key requirement for earning the VCP-DCV 2021 certification. This book may be useful for secondary purposes, such as learning how to implement, configure, and manage a vSphere environment or preparing to take other VCP-DCV qualifying exams.

The rest of this introduction provides details on the VCP-DCV certification, the 2V0-21.20 exam, and this book.

## VCP-DCV Requirements

The primary objective of the VCP-DCV 2021 certification is to demonstrate that you have mastered the skills to successfully install, configure, and manage VMware vSphere 7 environments. You can find the exam requirements, objectives, and other details on the certification web portal, at <http://mylearn.vmware.com/portals/certification/>. On the website, navigate to the Data Center Virtualization track and to the VCP-DCV certification. Examine the VCP-DCV 2021 requirements based on your qualifications. For example, if you select that you currently hold no VCP certifications, then the website indicates that your path to certification is to gain experience with vSphere 7.0, attend one of the following required training courses, and pass the Professional vSphere 7.0 (2V0-21.20) exam:

- VMware vSphere: Install, Configure, Manage [V7]
- VMware vSphere: Optimize and Scale [V7]
- VMware vSphere: Troubleshooting [V7]
- VMware vSphere: Fast Track [V7]

If you select that you currently hold a VCP6-DCV certification, the website indicates that your path includes a recommendation, but not a requirement, to take a training course.

VMware updates the VCP-DCV certification requirements each year. So, the requirements for the VCP-DCV 2021 certification may differ slightly from VCP-DCV 2020 certification. Likewise, VMware updates the qualifying exams. Each year, as VMware updates the Professional VMware vSphere 7.x exam, the authors of this book will create an appendix to supplement the original book. To prepare for a future version of the exam, download the corresponding online appendix from the book's companion website and use it to supplement the original book.

After you identify your path to certification, you can select the Professional VMware vSphere 7.x (2V0-21.20) exam to closely examine its details and to download the Exam Preparation Guide (also known as the exam blueprint).

## **Details on the 2V0-21.20 Exam**

The 2V0-21.20 exam blueprint provides details on exam delivery, minimum qualifications for candidates, exam objectives, recommended courses, and references to supporting VMware documentation. It also contains 10 sample exam questions. The 2V0-21.20 exam is a proctored exam delivered through Pearson VUE. See [Chapter 15, “Final Preparation,”](#) for details on registering and taking the exam.

A minimally qualified candidate (MQC) has 6 to 12 months of hands-on experience implementing, managing, and supporting a vSphere environment. The MQC has knowledge of storage, networking, hardware, security, business continuity, and disaster recovery concepts.

The exam characteristics are as follows:

- Format: Proctored exam
- Question type: Multiple choice
- Number of questions: 70
- Duration: 130 minutes
- Passing score: 300
- Cost: \$250 (in the United States)

## **2V0-21.20 Exam Objectives**

The 2V0-21.20 exam blueprint lists the exam objectives, which are summarized

here:

## Section 1: Architectures and Technologies

- Objective 1.1: Identify the prerequisites and components for a vSphere implementation
- Objective 1.2: Describe vCenter Server topology
- Objective 1.3: Identify and differentiate storage access protocols for vSphere (NFS, iSCSI, SAN, etc.)
  - 1.3.1: Describe storage datastore types for vSphere
  - 1.3.2: Explain the importance of advanced storage configuration (vSphere Storage APIs for Storage Awareness (VASA), vSphere Storage APIs Array Integration (VAAI), etc.)
  - 1.3.3: Describe storage policies
  - 1.3.4: Describe basic storage concepts in K8s, vSAN and vSphere Virtual Volumes (vVols)
- Objective 1.4: Differentiate between vSphere Network I/O Control (NIOC) and vSphere Storage I/O Control (SIOC)
- Objective 1.5: Describe instant clone architecture and use cases
- Objective 1.6: Describe ESXi cluster concepts
  - 1.6.1: Describe Distributed Resource Scheduler (DRS)
  - 1.6.2: Describe vSphere Enhanced vMotion Compatibility (EVC)
  - 1.6.3: Describe how Distributed Resource Scheduler (DRS) scores virtual machines
  - 1.6.4: Describe vSphere High Availability
  - 1.6.5: Describe datastore clusters
- Objective 1.7: Identify vSphere distributed switch and vSphere standard switch capabilities
  - 1.7.1: Describe VMkernel networking
  - 1.7.2: Manage networking on multiple hosts with vSphere distributed switch
  - 1.7.3: Describe networking policies

- 1.7.4: Manage Network I/O Control (NIOC) on a vSphere distributed switch
- Objective 1.8: Describe vSphere Lifecycle Manager concepts (baselines, cluster images, etc.)
- Objective 1.9: Describe the basics of vSAN as primary storage
  - 1.9.1: Identify basic vSAN requirements (networking, disk count + type)
- Objective 1.10: Describe the vSphere Trust Authority architecture
- Objective 1.11: Explain Software Guard Extensions (SGX)

## Section 2: VMware Products and Solutions

- Objective 2.1: Describe the role of vSphere in the software-defined data center (SDDC)
- Objective 2.2: Identify use cases for vCloud Foundation
- Objective 2.3: Identify migration options
- Objective 2.4: Identify DR use cases
- Objective 2.5: Describe vSphere integration with VMware Skyline

Section 3: Planning and Designing (There are no testable objectives for this section.)

## Section 4: Installing, Configuring, and Setup

- Objective 4.1: Describe single sign-on (SSO) deployment topology
  - 4.1.1: Configure a single sign-on (SSO) domain
  - 4.1.2: Join an existing single sign-on (SSO) domain
- Objective 4.2: Configure VSS advanced virtual networking options
- Objective 4.3: Set up identity sources
  - 4.3.1: Configure Identity Federation
  - 4.3.2: Configure Lightweight Directory Access Protocol (LDAP) integration

- 4.3.3: Configure Active Directory integration
- Objective 4.4: Deploy and configure vCenter Server Appliance
- Objective 4.5: Create and configure VMware High Availability and advanced options (Admission Control, Proactive High Availability, etc.)
- Objective 4.6: Deploy and configure vCenter Server High Availability
- Objective 4.7: Set up content library
- Objective 4.8: Configure vCenter Server file-based backup
- Objective 4.9: Analyze basic log output from vSphere products
- Objective 4.10: Configure vSphere Trust Authority
- Objective 4.11: Configure vSphere certificates
  - 4.11.1: Describe Enterprise PKIs role for SSL certificates
- Objective 4.12: Configure vSphere Lifecycle Manager/VMware Update Manager (VUM)
- Objective 4.13: Securely Boot ESXi hosts
- Objective 4.14: Configure different network stacks
- Objective 4.15: Configure Host Profiles
- Objective 4.16: Identify boot options
  - 4.16.1: Configure Quick Boot

## Section 5: Performance-tuning, Optimization, Upgrades

- Objective 5.1: Identify resource pools use cases
  - 5.1.1: Explain shares, limits, and reservations (resource management)
- Objective 5.2: Monitor resources of vCenter Server Appliance and vSphere environment
- Objective 5.3: Identify and use tools for performance monitoring
- Objective 5.4: Configure Network I/O Control (NIOC)
- Objective 5.5: Configure Storage I/O Control (SIOC)
- Objective 5.6: Explain the performance impact of maintaining virtual machine snapshots

- Objective 5.7: Plan for upgrading various vSphere components

Section 6: Troubleshooting and Repairing (There are no testable objectives for this section.)

Section 7: Administrative and Operational Tasks

- Objective 7.1: Create and manage virtual machine snapshots
- Objective 7.2: Create virtual machines using different methods (Open Virtual Machine Format (OVF) templates, content library, etc.)
- Objective 7.3: Manage virtual machines
- Objective 7.4: Manage storage (datastores, storage policies, etc.)
  - 7.4.1: Configure and modify datastores (expand/upgrade existing datastore, etc.)
  - 7.4.2: Create virtual machine storage policies
  - 7.4.3: Configure storage cluster options
- Objective 7.5: Create Distributed Resource Scheduler (DRS) affinity and anti-affinity rules for common use cases
- Objective 7.6: Configure and perform different types of migrations
- Objective 7.7: Configure role-based user management
- Objective 7.8: Configure and manage the options for securing a vSphere environment (certificates, virtual machine encryption, virtual Trusted Platform Module, lock-down mode, virtualization-based security, etc.)
- Objective 7.9: Configure and manage host profiles
- Objective 7.10: Utilize baselines to perform updates and upgrades
- Objective 7.11: Utilize vSphere Lifecycle Manager
  - 7.11.1: Describe Firmware upgrades for ESXi
  - 7.11.2: Describe ESXi updates
  - 7.11.3: Describe component and driver updates for ESXi
  - 7.11.4: Describe hardware compatibility check
  - 7.11.5: Describe ESXi cluster image export functionality

- Objective 7.12: Configure alarms

### **Note**

Sections 3 and 6 currently do not apply to the 2V0-21.20 exam, but they may be used for other exams.

### **Note**

For future exams, download and examine the objectives in the updated exam blueprint. Be sure to use the future Pearson-provided online appendix specific to the updated exam.

## **Who Should Take This Exam and Read This Book?**

The VCP-DCV certification is the most popular certification at VMware; more than 100,000 professionals around the world hold this certification. This book is intended for anyone who wants to prepare for the 2V0-21.20 exam, which is a required exam for VCP-DCV 2021 certification. The audience includes current and prospective IT professionals such as system administrators, infrastructure administrators, and virtualization engineers.

## **Book Features and Exam Preparation Methods**

This book uses several key methodologies to help you discover the exam topics on which you need more review, to help you fully understand and remember those details, and to help you prove to yourself that you have retained your knowledge of those topics. This book does not try to help you pass the exam only by memorization but by truly learning and understanding the topics.

The book includes many features that provide different ways to study so you can be ready for the exam. If you understand a topic when you read it but do not study it any further, you probably will not be ready to pass the exam with confidence. The features included in this book give you tools that help you determine what you know, review what you know, better learn what you don't know, and be well prepared for the exam. These tools include:

- **“Do I Know This Already?” Quizzes:** Each chapter begins with a quiz that helps you determine the amount of time you need to spend studying that chapter.
- **Foundation Topics:** These are the core sections of each chapter. They explain the protocols, concepts, and configuration for the topics in that chapter.
- **Exam Preparation Tasks:** This section of each chapter lists a series of study activities that should be done after reading the “Foundation Topics” section. Each chapter includes the activities that make the most sense for studying the topics in that chapter. The activities include the following:
  - **Key Topics Review:** The Key Topics icon appears next to the most important items in the “Foundation Topics” section of the chapter. The “Key Topics Review” section lists the key topics from the chapter and their page numbers. Although the contents of the entire chapter could be on the exam, you should definitely know the information listed for each key topic. Review these topics carefully.
  - **Memory Tables:** To help you exercise your memory and memorize some important facts, memory tables are provided. The memory tables contain only portions of key tables provided previously in the chapter, enabling you to complete the table or list. [Appendix B, “Memory Tables,”](#) provides the incomplete tables, and [Appendix C, “Memory Tables Answer Key,”](#) includes the completed tables (answer keys). These appendixes are also provided on the companion website that is provided with your book.
  - **Define Key Terms:** The VCP-DCV exam requires you to learn and know a lot of related terminology. This section lists some of the most important terms from the chapter and asks you to write a short definition and compare your answer to the Glossary.
- **Practice Exams:** The companion website contains an exam engine.

## Book Organization

The chapters in this book are organized such that [Chapters 1](#) through [7](#) provide in-depth material on vSphere concepts, and [Chapters 8](#) through [14](#) describe procedures for the installation, configuration, and management of vSphere components and features. The authors recommend that you read the entire book



from cover to cover at least once. As you read about any topic in [Chapters 1 to 7](#), keep in mind that you can find corresponding “how to” steps in [Chapters 8 to 14](#). As you read about any specific procedure in [Chapters 8 to 14](#), keep in mind that you can find associated details (concepts) in [Chapters 1 to 7](#).

Optionally, you can prepare for the exam by studying for the exam objectives in order, using [Table I-1](#) as your guide. As you prepare for each exam objective, you can focus on the most appropriate chapter and section. You can also refer to related chapters and sections. For example, as you prepare for Objective 1.2 (Describe vCenter Server topology), you should focus on the “vCenter Server Topology” section in [Chapter 1](#), but you may also want to review the “Deploying vCenter Server Components” section in [Chapter 8](#) and the “vSphere Managed Inventory Objects” section in [Chapter 5](#).

When preparing for a specific exam objective, you can use [Table I-1](#) to identify the sections in the book that directly address the objective and the sections that provide related information.

**Table I-1** Mapping of Exam Objectives to Book Chapters and Sections

<b>Objective</b>	<b>Description</b>	<b>Chapter/Section</b>	<b>Related (Supporting) Chapters/Sections</b>
<b>1 Architectures and Technologies</b>			
1.1	Identify the pre-requisites and components for a vSphere implementation	1—vSphere Overview, Components, and Requirements	8—vSphere Installation
		<ul style="list-style-type: none"> <li>■ Installing ESXi Hosts</li> <li>■ Infrastructure Requirements</li> <li>■ Other Requirements</li> </ul>	<ul style="list-style-type: none"> <li>■ Deploying vCenter Server Components</li> </ul>

---

1 . 2	Describe vCenter Server topology	1—vSphere Overview, Components, and Requirements	8—vSphere Installation  Deploying vCenter Server Components
		<ul style="list-style-type: none"> <li>■ vCenter Server Topology</li> </ul>	5—vCenter Server Features and Virtual Machines
			<ul style="list-style-type: none"> <li>■ vSphere Managed Inventory Objects</li> </ul>

---

1 . 3	Identify and differentiate storage access protocols for vSphere (NFS, iSCSI, SAN, etc.)	2—Storage Infrastructure	
		<ul style="list-style-type: none"> <li>■ Storage Virtualization—Traditional Model</li> </ul>	

---

1 . 3 . 1	Describe storage datastore types for vSphere	2—Storage Infrastructure	11—Managing Storage
-----------------------	--	--------------------------	---------------------

- Software-Defined Storage Models
- Manage Datastores

- Datastore Types

---

1 Explain the importance of  
 . advanced storage configuration  
 3 (VAAI + VASA, multipathing)  
 .  
 2

2—Storage Infrastructure

11—Managing Storage

- VASA
- VAAI

- VASA: Register a Storage Provider
- VASA: Manage Storage Providers

---

1 Describe storage policies  
 .  
 3  
 .  
 3

2—Storage Infrastructure

11—Managing Storage

- Storage Policies
- Managing Storage Policies

---

1 Describe basic storage concepts  
 . in K8s, vSAN and vVOLs  
 3  
 .  
 4

2—Storage Infrastructure

11—Managing Storage

- Storage Virtualization—Traditional Model
- Managing vSAN
- Software-Defined Storage Models
- Managing Datastore
- Datastore Types
- Configuring and Managing vVols
- Storage in vSphere with Kubernetes

---

1	Identify use cases for RDMs, PMEMs, VVOLs, and NVMe	2—Storage Infrastructure	11—Managing Storage
.			
3			
.			
5		<ul style="list-style-type: none"> <li>■ Raw Device Mappings (RDMs)</li> <li>■ vVols</li> <li>■ VMware NVMe</li> </ul>	<ul style="list-style-type: none"> <li>■ Managing RDMs</li> <li>■ Managing Storage Policies</li> <li>■ Managing VMware NVMe</li> <li>■ Managing PMEM</li> </ul>

---

1 Differentiate between NIOC and  
. SIOC  
4

2—Storage  
Infrastructure

3—Network  
Infrastructure

- NIOC, SIOC, and SDRS
- Network I/O Control

9—Configuring  
and Managing  
Virtual  
Networks

- Configuring  
Network I/O  
Control (NIOC)

11—Managing  
Storage

- Configuring and  
Managing SIOC

---

1 Describe instant clone  
. architecture and use cases  
5

5—vCenter  
Server Features  
and Virtual

## Machines

- Instant Clone

---

1 . 6	Describe ESXi cluster concepts	4—Clusters and High Availability	10—Managing and Monitoring Clusters and Resources
		<ul style="list-style-type: none"><li>■ Cluster Concepts and Overview</li><li>■ Distributed Resources Scheduler (DRS)</li><li>■ High Availability (HA)</li></ul>	<ul style="list-style-type: none"><li>■ Creating and Configuring a vSphere Cluster</li><li>■ Creating and Configuring a vSphere DRS Cluster</li><li>■ Creating and Configuring a vSphere HA cluster</li></ul>
1 . 6 . 1	Describe VMware Distributed Resource Scheduler (DRS)	4—Clusters and High Availability	10—Managing and Monitoring Clusters and Resources

---

- Cluster Concepts and Overview
- Distributed Resources Scheduler (DRS)
- Creating and Configuring a vSphere DRS Cluster

---

1 . 6 . 2	Describe Enhanced vMotion Compatibility (EVC)	4—Clusters and High Availability	10— Managing and Monitoring Clusters and Resources
-----------------------	---	----------------------------------	--

- Enhanced vMotion Compatibility (EVC)
- EVC Mode

---

1 . 6 . 3	Describe how DRS scores VMs	4—Clusters and High Availability	10—Managing and Monitoring Clusters and Resources
-----------------------	-----------------------------	----------------------------------	---

- How DRS Scores VMs
- Creating and Configuring a vSphere DRS Cluster

---

1 . 6	Describe vSphere HA	4—Clusters and High Availability	10—Managing and Monitoring Clusters and
-------------	---------------------	----------------------------------	---

.  
4

## Resources

- vSphere High Availability (HA)
  - Creating and Configuring a vSphere HA Cluster

---

1 Describe how vSphere HA  
. calculates slot size  
6  
. 4  
. 1

4—Clusters and High Availability

10—Managing and Monitoring Clusters and Resources

- vSphere HA Admission Control

- Creating and Configuring a vSphere HA Cluster

---

1 Describe datastore clusters  
. 6  
. 5

2—Storage Infrastructure

11—Managing Storage

- Storage DRS (SDRS)

- Configuring and Managing SDRS

---

1 Identify vSphere Distributed  
. Switch (VDS) and vSphere

3—Network Infrastructure

9—Configuring and Managing



7 Standard Switch (VSS) capabilities

Virtual Networks

- vSphere Standard Switch (vSS)
  - vSphere Distributed Switch (vDS)
  - vDS Settings and Features
- Creating and Configuring vSphere Standard Switches
  - Creating and Configuring vSphere Distributed Switches

---

1 Describe VMkernel Networking  
. 7  
. 1

3—Network Infrastructure

9—Configuring and Managing Virtual Networks

- VMkernel Networking and TCP/IP Stacks

- Configuring and Managing VMkernel Adapters

- Configuring TCP/IP Stacks
-

1 . 7 . 2	Managing Networking on multiple hosts with vSphere Distributed Switch (vDS)	9—Configuring and Managing Virtual Networks	3—Network Infrastructure
		<ul style="list-style-type: none"> <li>■ Managing Host Networking with vDS</li> </ul>	<ul style="list-style-type: none"> <li>■ vSphere Distributed Switch (vDS)</li> </ul>

---

1 . 7 . 3	Describe networking policies	3—Network Infrastructure	9—Configuring and Managing Virtual Networks
		<ul style="list-style-type: none"> <li>■ vSS Networking Policies</li> </ul>	<ul style="list-style-type: none"> <li>■ Networking Policies and Advanced Features</li> </ul>

---

1 . 7 . 4	Managing Network I/O Control (NIOC) on a Distributed Switch (vDS)	9—Configuring and Managing Virtual Networks	3—Network Infrastructure
		<ul style="list-style-type: none"> <li>■ Configuring Network I/O Control (NIOC)</li> </ul>	<ul style="list-style-type: none"> <li>■ Network I/O Control</li> </ul>

---

1	Describe vSphere Lifecycle	13—Managing	8—vSphere
---	----------------------------	-------------	-----------

. 8	Manager concepts (baselines, cluster images, etc)	vSphere and vCenter Server	Installation
		■ Using vSphere Lifecycle Manager	■ VMware vSphere Lifecycle Manager Implementation

---

1 . 9	Describe the basics of vSAN as primary storage	2—Storage Infrastructure	
		■ vSAN Concepts	

---

1 . 9 . 1	Identify basic vSAN requirements (networking, disk count + type)	2—Storage Infrastructure	11—Managing Storage
		■ vSAN Requirements	■ Configuring and Managing vSAN

---

1 . 1 0	Describe the vSphere Trust Authority architecture	7—vSphere Security	12—Managing vSphere Security
------------------	--	-----------------------	------------------------------------

- vSphere Trust Authority (vTA)
- Configuring and Managing vSphere Trust Authority (vTA)

---

1 . 1 1	Explain Virtual SGX—Software Guard Extensions	7—vSphere Security	12—Managing vSphere Security
------------------	---	--------------------	------------------------------

- Securing Virtual Machines with Virtual Intel Software Guard Extension (vSGX)
- Securing Virtual Machines with Intel Software Guard Extensions (SGX)

---

## 2 VMware Products and Solutions

---

2 . 1	Describe the role of vSphere in the software-defined data center	1—vSphere Overview, Components, and Requirements
-------------	--	--

- VMware SDDC

---

2 . 2	Identify use cases for vCloud Foundation	6—VMware Product Integration
-------------	--	------------------------------

- VMware Cloud Foundation (VCF)

---

2 Identify migration options  
.  
3

6—VMware  
Product  
Integration

5—vCenter  
Server Features  
and Virtual  
Machine

- Inbound and  
Outbound  
vSphere  
Migration

- Virtual Machine  
Migration

---

2 Identify DR use cases  
.  
4

6—VMware  
Product  
Integration

- vSphere  
Replication

- Site Recovery  
Manager (SRM)

---

2 Describe vSphere integration  
.  
5 with VMware Skyline

6—VMware  
Product  
Integration

- VMware Skyline Integration

---

### 3 Planning and Designing

---

### 4 Installing, Configuring, and Setup

---

4	Plan SSO deployment topology	1—vSphere Overview, Components, and Requirements	12—Managing vSphere Security
---	------------------------------	--	------------------------------

- Managing SSO
- vCenter Server Topology

---

4	Configure an SSO domain	8—vSphere Installation	1—vSphere Overview, Components, and Requirements
---	-------------------------	------------------------	--

- Deploying vCenter Server Components
- Configuring Single Sign-On (SSO)
- vCenter Server Topology

12—Managing vSphere

## Security

- Managing SSO

---

4	Join an existing SSO domain	8—vSphere Installation	1—vSphere Overview, Components, and Requirements
.			
1			
.			
2			
		■ Deploying vCenter Server Components	■ vCenter Server Topology
		■ Configuring Single Sign-On (SSO)	
			12—Managing vSphere Security
			■ Managing SSO

---

4	Configure VSS advanced virtual networking options	9—Configuring and Managing Virtual Networks	3—Network Infrastructure
.			
2			

- Creating and Configuring vSphere Standard Switches
- vSphere Standard Switch (vSS)
- Creating and Configuring Standard Port Groups

---

4	Set up identity sources	8—vSphere Installation	12—Managing vSphere Security
.			
3			

- Adding, Editing, and Removing SSO Identity Sources
- Managing SSO

---

4	Configure Identity Federation	8—vSphere Installation	12—Managing vSphere Security
.			
3			
.			
1			

- Configuring Identity Federation
- Managing SSO

---

4	Configure LDAP integration	8—vSphere Installation	12—Managing vSphere
.			



3  
.  
2

## Security

- Adding, Editing, and Removing SSO Identity Sources
- Managing SSO
- Adding an LDAP Authentication Source

---

4 Configure Active Directory  
integration  
3  
.  
3

8—vSphere  
Installation

12—Managing  
vSphere  
Security

- Adding an Active Directory Identity Source
- Managing SSO

12—Managing  
vSphere Security

- Using Active Directory to Manage ESXi Users

4 Deploy and configure vCenter  
Server Appliance (VCSA)  
4

8—vSphere  
Installation

1—vSphere  
Overview,  
Components,  
and  
Requirements

- vCenter Server  
Appliance

- vCenter Server  
Topology

13—Managing  
vSphere and  
vCenter Server

- Upgrading to  
vSphere 7.0

- Repointing a  
vCenter Server  
to Another  
Domain

---

4 Create and configure VMware  
HA and DRS advanced options  
(Admission Control, Proactive  
HA, etc.)  
5

10—Managing  
and Monitoring  
Clusters and  
Resources

4—Clusters and  
High  
Availability

- Distributed

- Creating and Configuring a vSphere DRS Cluster

Resource Scheduler (DRS)

- Creating and Configuring a vSphere HA Cluster

- vSphere High Availability (HA)

---

4 . 6 Deploy and configure vCenter Server High Availability

8—vSphere Installation

1—vSphere Overview, Components, and Requirements

- Implementing VCSA HA

- vCenter Server Topology

- vCenter High Availability Requirements

4—Clusters and High Availability

- vCenter Server

High  
Availability

13—Managing  
vSphere and  
vCenter Server

- Managing the  
vCenter HA  
Cluster

---

4 Set up content library  
. 7

14—Managing  
Virtual Machines

5—vCenter  
Server Features  
and Virtual  
Machine

- Content Library

- Content Library

---

4 Configure vCenter Server file-  
. based backup  
8

13—Managing  
vSphere and  
vCenter Server

- vCenter Server  
Backup

---

4 . 9	Analyze basic log output from vSphere products	10—Managing and Monitoring Clusters and Resources	10—Monitoring and Managing Clusters and Resources
-------------	--	---	---

- Logging in vSphere
- Viewing the System Event Log
- System Logs Files

---

4 . 1 0	Configure vSphere Trust Authority	12—Managing vSphere Security	7—vSphere Security
------------------	-----------------------------------	------------------------------	--------------------

- Configuring and Managing vSphere Trust Authority (vTA)
- vSphere Trust Authority (vTA)

---

4 . 1 1	Configure vSphere certificates	12—Managing vSphere Security	7—vSphere Security
------------------	--------------------------------	------------------------------	--------------------

- Configuring and Managing
- ESXi Host Certificates

## vSphere Certificates

13—Managing  
vSphere and  
vCenter Server

- Verifying SSL  
Certificates for  
Legacy Hosts

---

4 . 1 1 . 1	Describe enterprise PKIs role for SSL certificates	7—vSphere Security	12—Manage vSphere Security
		■ vSphere Certificates Overview	■ Configure and Manage vSphere Certificates
4 . 1 2	Configure vSphere Lifecycle Manager/VMware Update Manager (VUM)	8—vSphere Installation	13—Managing vSphere and vCenter Server
		■ vSphere Lifecycle Manager Implementation	■ Using vSphere Lifecycle Manager

---

- About VMware Update Manager
- Update Manager Download Service (UMDS)

---

4	Securely Boot ESXi hosts	12—Managing vSphere Security	7—vSphere Security
.			
1			
3			

- Configuring UEFI Secure Boot for ESXi Hosts
- ESXi Secure Boot and TPM
- vSphere Trusted Authority (vTA)

---

4	Configure different network stacks	9—Configuring and Managing Virtual Networks	3—Network Infrastructure
.			
1			
4			

- Configuring TCP/IP Stacks
- VMkernel Networking and TCP/IP Stacks

---

4	Configure Host Profiles	8—vSphere Installation	
.			

1  
5

- Configuring ESXi Using Host Profiles

---

4 Identify boot options 8—vSphere Installation  
. 1  
6

- ESXi Kernel Options

---

4 Configure Quick Boot 13—Managing vSphere and vCenter Server  
. 1  
6  
. 1

- ESXi Quick Boot

---

## 5 Performance-tuning and Optimization

---

5 Identify resource pools use cases 4—Clusters and High Availability 10—Monitoring and Managing Clusters and Resources  
. 1

- Resource Pools



- Creating a Resource Pool
- Monitoring and Managing Resource Pool Resources

---

5 Explain shares, limits and  
 . reservations (resource  
 1 management)  
 .  
 1

4—Clusters and High Availability

10— Managing and Monitoring Clusters and Resources

- Shares, Limits, and Reservations

- Shares, Limits, and Reservations

- Creating a Resource Pool

- Monitoring and Managing Resource Pool Resources

---

5 Monitor resources of vCenter  
 . Server Appliance (VCSA) and  
 2 vSphere environment

10—Managing and Monitoring Clusters and Resources

4—Clusters and High Availability

- Cluster Concepts and Overview
  - Distributed Resource Scheduler (DRS)
- Monitoring and Managing vSphere Resources
- Monitoring and Managing vCenter Server Services

### 13—Managing vSphere and vCenter Server

- Monitoring and Managing vCenter Server

---

5 Identify and use tools for  
 . performance monitoring  
 3

### 10—Managing and Monitoring Clusters and Resources

- Monitoring and Managing vSphere Resources

5 · 4	Configure Network I/O Control	9—Configuring and Managing Virtual Networks	3—Network Infrastructure
		<ul style="list-style-type: none"> <li>■ Configuring Network I/O Control (NIOC)</li> </ul>	<ul style="list-style-type: none"> <li>■ Network I/O Control</li> </ul>
5 · 5	Configure Storage I/O Control	11—Managing Storage	2—Storage Infrastructure
		<ul style="list-style-type: none"> <li>■ Configuring and Managing SIOC</li> </ul>	<ul style="list-style-type: none"> <li>■ NIOC, SIOC, and SDRS</li> </ul>
5 · 6	Explain the performance impact of maintaining VM snapshots.	5—vCenter Server Features and Virtual Machine	14—Managing Virtual Machines
		<ul style="list-style-type: none"> <li>■ Virtual Machine Snapshots</li> </ul>	<ul style="list-style-type: none"> <li>■ Creating and Managing Virtual Machine Snapshots</li> </ul>
5	Plan for upgrading various	13—Managing	

. vSphere components  
7

vSphere and  
vCenter Server

- Using Lifecycle Manager
- Upgrading to vSphere 7.0

---

## 6 Troubleshooting and Repairing

---

### 7 Administrative and Operational Tasks

---

7 Create and manage VM snapshot  
. (consolidate, delete, etc.)  
1

14—Managing  
Virtual Machines

5—vCenter  
Server Features  
and Virtual  
Machine

- Creating and Managing Virtual Machine Snapshots
- Virtual Machine Snapshots

---

7 Create VMs using different  
. methods (OVF templates, content  
2 library, and so on)

14—Managing  
Virtual Machines

5—vCenter  
Server Features  
and Virtual  
Machine

- Managing VMs Using PowerCLI
- Deploying OVF
- Virtual Machine Cloning

and OVA  
Templates

- Deploying VMs Using Content Library
- Managing OVF Templates
- Content Library

14—Managing Virtual Machines

---

7 . 3 Manage VMs (modifying VM settings, etc.)

14—Managing Virtual Machines

5—vCenter Server Features and Virtual Machine

- Creating and Configuring Virtual Machines
- Managing Virtual Machines
- Virtual Machine Settings

---

7 . 4 Manage storage (datastores, storage policies, etc.)

11—Managing Storage

2—Storage Infrastructure

- Managing
- Datastore Types

Datastores

- Storage Policies
- Managing Storage Policies
  - Storage Multipathing and Failover
- Managing Multipathing
- Changing Path Selection Policy

---

7	Configure and modify datastores (expand/upgrade existing datastore, etc)	11—Managing Storage	2—Storage Infrastructure
.			
4			
.			
1			

- Managing Datastores
- Datastore Types

---

7	Create VM storage policies	11—Managing Storage	2—Storage Infrastructure
.			
4			
.			
2			

- Managing Storage Policies
- Storage Policies

---

7	Configure storage cluster options	11—Managing Storage	2—Storage Infrastructure
.			
4			

- Configuring and Managing Storage DRS
- SDRS
- Configuring and Managing VSAN

---

7 Create DRS affinity and anti-affinity rules for common use cases.

10—Monitoring and Managing Clusters and Resources

4—Clusters and High Availability

- Creating Affinity and Anti-Affinity Rules
- DRS Rules

---

7 Configure and perform different types of migrations (all types)

14—Managing Virtual Machines

5—vCenter Server Features and Virtual Machine

- Migrating Virtual Machines
- Virtual Machine Migration
- vMotion Details

- Storage vMotion Details

---

7	Configure role-based user management (custom permissions, on datastores, clusters, vCenter Servers, and hosts etc)	12—Managing vSphere Security	7—vSphere Security
---	--	------------------------------	--------------------

- Configuring and Managing Authentication and Authorization
- vSphere Permissions

8—vSphere Installation

- Applying Permissions to ESXi Hosts Using Host Profiles

---

7	Configure and manage the options for securing a vSphere environment (certificates, VM encryption, virtual TPM, lock-down mode, VBS, etc)	12—Managing vSphere Security	7—vSphere Security
---	--	------------------------------	--------------------

- Configuring and Managing Authentication and Authorization
- ESXi and vCenter Server Security



- Configuring and Managing ESXi Security
- Configuring and Managing vSphere Certificates
- Other Security Management

---

7 . 9	Configure and manage host profiles	8—vSphere Installation	5—vCenter Server Features and Virtual Machine
-------------	------------------------------------	------------------------	---

- Configuring ESXi Using Host Profiles
- Host Profiles

---

7 . 1 0	Utilize VUM (create baselines, applying baselines, notifications, download, remediate)	13—Managing vSphere and vCenter Server	8—vSphere Installation
		<ul style="list-style-type: none"> <li>■ Using vSphere Lifecycle Manager</li> </ul>	<ul style="list-style-type: none"> <li>■ VMware vSphere Lifecycle Manager Implementation</li> </ul>

- About VMware Update Manager

---

7 Describe vSphere Lifecycle  
 . Manager  
 1  
 1

13—Managing vSphere and vCenter Server

8—vSphere Installation

- Using vSphere Lifecycle Manager

- VMware vSphere Lifecycle Manager Implementation

14—Managing Virtual Machines

- Installing and Upgrading VMware Tools

---

7 Describe Firmware upgrades for  
 . ESXi  
 1  
 1  
 .  
 1

13—Managing vSphere and vCenter Server

8—vSphere Installation

- VMware

- Using vSphere Lifecycle Manager
- vSphere Lifecycle Manager Implementation

7	Describe ESXi updates	13—Managing vSphere and vCenter Server	8—vSphere Installation
.			
1			
1			
.			
2		<ul style="list-style-type: none"> <li>■ Using vSphere Lifecycle Manager</li> </ul>	<ul style="list-style-type: none"> <li>■ VMware vSphere Lifecycle Manager Implementation</li> </ul>

7	Describe component and driver updates for ESXi	13—Managing vSphere and vCenter Server	8—vSphere Installation
.			
1			
1			
.			
3		<ul style="list-style-type: none"> <li>■ Using vSphere Lifecycle Manager</li> </ul>	<ul style="list-style-type: none"> <li>■ VMware vSphere Lifecycle Manager Implementation</li> </ul>

7	Describe hardware compatibility check	13—Managing vSphere and vCenter Server	8—vSphere Installation
.			
1			
1			
.			
4			

- Using vSphere Lifecycle Manager
- VMware vSphere Lifecycle Manager Implementation
- 5—vCenter Server Features and Virtual Machine
- VM Hardware and Compatibility

## 14—Managing Virtual Machines

- Configuring Virtual Machine Hardware

---

7 Describe ESXi cluster image  
 . export functionality  
 1  
 1  
 .  
 5

13—Managing vSphere and vCenter Server

8—vSphere Installation

- Using vSphere
- VMware vSphere

Lifecycle  
Manager

Lifecycle  
Manager  
Implementation

4—Clusters and  
High  
Availability

- Cluster  
Concepts and  
Overview

---

7 Configure alarms  
·  
1  
2

10—Monitoring  
and Managing  
Clusters and  
Resources

- Advanced Use  
Cases for Alarms

---

## Companion Website

Register this book to get access to the Pearson IT Certification test engine and other study materials plus additional bonus content. Check this site regularly for new and updated postings written by the authors that provide further insight into the more troublesome topics on the exam. Be sure to check the box indicating 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](http://www.pearsonITcertification.com/register) and log in or create a new account.
- Step 2.** Enter the ISBN **9780135898192**
- 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 the 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 these steps, please visit [www.pearsonITcertification.com/contact](http://www.pearsonITcertification.com/contact) and select the Site Problems/Comments option. Our customer service representatives will assist you.

## How to Access the Pearson Test Prep Practice Test Software

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

- **Print book:** Look in the cardboard sleeve in the back of the book for a piece of paper with your book's unique PTP code.
- **Premium Edition:** If you purchase the Premium Edition eBook and Practice Test directly from the [www.pearsonITcertification.com](http://www.pearsonITcertification.com) website, the code will be populated on your account page after purchase. Just log in to [www.pearsonITcertification.com](http://www.pearsonITcertification.com), click Account to see details of your account, and click the Digital Purchases tab.
- **Amazon Kindle:** For those who purchase a Kindle edition from Amazon, the access code will be supplied directly from Amazon.
- **Other bookseller e-books:** Note that if you purchase an e-book version

from any other source, the practice test is not included because other vendors to date have not chosen to vend the required unique access code.

### **Note**

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

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

**Step 1.** Open this book's companion website.

**Step 2.** Click the Practice Exams button.

**Step 3.** Follow the instructions listed there both for installing the desktop app and for using the web app.

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

### **Note**

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

### **Note**

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

## Customizing Your Exams

Once you are in the exam settings screen, you can choose to take exams in one of three modes:

- **Study mode:** Enables you to fully customize your exams and review answers as you are taking the exam. This is typically the mode you would use first to assess your knowledge and identify information gaps.
- **Practice Exam mode:** Locks certain customization options, as it is presenting a realistic exam experience. Use this mode when you are preparing to test your exam readiness.
- **Flash Card mode:** Strips out the answers and presents you with only the question stem. This mode is great for late-stage preparation when you really want to challenge yourself to provide answers without the benefit of seeing multiple-choice options. This mode does not provide the detailed score reports that the other two modes do, so you should not use it if you are trying to identify knowledge gaps.

In addition to these three modes, you will be able to select the source of your questions. You can choose to take exams that cover all of the chapters or you can narrow your selection to just a single chapter or the chapters that make up specific parts in the book. All chapters are selected by default. If you want to narrow your focus to individual chapters, simply deselect all the chapters and then select only those on which you wish to focus in the Objectives area.

You can also select the exam banks on which to focus. Each exam bank comes complete with a full exam of questions that cover topics in every chapter. You can have the test engine serve up exams from all test banks or just from one individual bank by selecting the desired banks in the exam bank area. There are several other customizations you can make to your exam from the exam settings screen, such as the time of the exam, the number of questions served up, whether to randomize questions and answers, whether to show the number of correct answers for multiple-answer questions, and whether to serve up only specific types of questions. You can also create custom test banks by selecting only questions that you have marked or questions on which you have added notes.

## Updating Your Exams



If you are using the online version of the Pearson Test Prep software, you should always have access to the latest version of the software as well as the exam data. If you are using the Windows desktop version, every time you launch the software while connected to the Internet, it checks if there are any updates to your exam data and automatically downloads any changes that were made since the last time you used the software.

Sometimes, due to many factors, the exam data might not fully download when you activate your exam. If you find that figures or exhibits are missing, you might need to manually update your exams. To update a particular exam you have already activated and downloaded, simply click the Tools tab and click the Update Products button. Again, this is only an issue with the desktop Windows application. If you wish to check for updates to the Pearson Test Prep exam engine software, Windows desktop version, simply click the Tools tab and click the Update Application button. This ensures that you are running the latest version of the software engine.

# **Part I**

## **vSphere Architecture, Integration, and Requirements**

## Chapter 1

# vSphere Overview, Components, and Requirements

This chapter covers the following topics:

- [vSphere Components and Editions](#)
- [vCenter Server Topology](#)
- [Infrastructure Requirements](#)
- [Other Requirements](#)
- [VMware Cloud vs. VMware Virtualization](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.1, 1.2, 2.1, 4.1, 4.1.1, 4.1.2, and 4.4.

This chapter introduces vSphere 7.0, describes its major components, and identifies its requirements.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 1-1](#) 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 and Review Questions.”](#)

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

---

**Foundations Topics Section**

**Questions Covered in This**

	<b>Section</b>
vSphere Components and Editions	1, 2
vCenter Server Topology	3, 4
Infrastructure Requirements	5, 6
Other Requirements	7, 8
VMware Cloud vs. VMware Virtualization	9, 10

- 1.** You plan to deploy vSphere 7.0 for three ESXi hosts and want to deploy the minimum vCenter Server edition that supports vMotion. Which vCenter Server edition do you choose?
  - a.** Essentials
  - b.** Essentials Plus
  - c.** Foundation
  - d.** Standard
  
- 2.** You plan to deploy vSphere 7.0 and want to minimize virtual machine downtime by proactively detecting hardware failures and placing the host in Quarantine Mode or Maintenance Mode. Which feature do you need?
  - a.** vSphere High Availability
  - b.** Proactive HA
  - c.** Predictive DRS
  - d.** vCenter HA
  
- 3.** You are preparing to deploy and manage a vSphere environment. Which vCenter Server component provides Security Assertion Markup Language (SAML) tokens?
  - a.** vCenter Lookup Service
  - b.** VMware Directory Service
  - c.** tcServer
  - d.** STS

4. You plan to deploy another vCenter Server in your vSphere 7.0 environment and want it to use an existing vSphere Single Sign-On domain. What should you do?
  - a. During vCenter Server deployment, join an existing SSO domain.
  - b. Prior to vCenter Server deployment, deploy an external PSC.
  - c. During vCenter Server deployment, connect to an external PSC.
  - d. Configure vCenter HA.
5. You plan to deploy a vCenter Server Appliance 7.0 instance to support 350 ESXi hosts and 4500 virtual machines. What is the minimum memory you should plan for the vCenter Server Appliance instance?
  - a. 37 GB
  - b. 56 GB
  - c. 28 GB
  - d. 19 GB
6. You are interested in booting your ESXi hosts using UEFI. Which of the following is a key consideration?
  - a. After installing ESXi 7.0, you can change the boot type between BIOS and UEFI by using the direct console user interface.
  - b. ESXi boot from UEFI is deprecated in ESXi 7.0.
  - c. After installing ESXi 7.0, you can change the boot type between BIOS and UEFI by using the vSphere Client.
  - d. After you install ESXi 7.0, changing the boot type between BIOS and UEFI is not supported.
7. You are planning the backup and recovery for a new vCenter Server Appliance instance using the file-based backup feature in the vCenter Server Appliance Management Interface. Which protocol is not supported?
  - a. NFS
  - b. FTP
  - c. HTTPS

- d. SCP
8. When you are planning the procedures to manage a new vSphere 7.0 environment, which of the following is not a supported browser for the vSphere Client?
- a. Microsoft Internet Explorer 11.0.96 for Windows users
  - b. Microsoft Edge 38 for Windows users
  - c. Safari 5.0 for Mac users
  - d. Firefox 45 for Mac users
9. You need to include on-premises cloud automation software to improve the delivery of IT services and applications in your vSphere-based SDDC. Which of the following should you choose?
- a. VMware Cloud Assembly
  - b. VMware Service Broker
  - c. vCloud Director
  - d. vRealize Automation
10. You want a simple path to the hybrid cloud that leverages a common infrastructure and consistent operational model for on-premises and off-premises data centers. What should you use?
- a. vRealize Suite
  - b. VCF
  - c. vCloud Director
  - d. Cloud Automation

## Foundation Topics

### vSphere Components and Editions

VMware vSphere is a suite of products that you can use to virtualize enterprise data centers and build private clouds.

### vSphere Components

[Table 1-2](#) describes the installable VMware products that are the core components in a vSphere environment.

**Table 1-2** Installable Core vSphere Components

---

Component	Description
vCenter Server	The major management component in the vSphere environment. Its services include vCenter Server, vSphere Web Client, vSphere Auto Deploy, vSphere ESXi Dump Collector, and the components that were associated with the Platform Services Controller in prior versions: vCenter Single Sign-On, License Service, Lookup Service, and VMware Certificate Authority.
ESXi Server	The physical host (including the hypervisor) on which virtual machines run.

---

Some optional vSphere features require the deployment of additional components and specific vSphere or vCenter Server editions. [Table 1-3](#) describes two of these optional components, which require the deployment of additional virtual appliances.

**Table 1-3** Optional vSphere Components

---

Optio nal Comp onent	Description
-------------------------------	-------------

vSp here Rep licat ion	An extension to VMware vCenter Server that provides hypervisor-based virtual machine replication and recovery.
------------------------------------	--

vCe nter Hig h Ava ilabi lity	A component that provides protection for vCenter Server Appliance against host, hardware, and application failures. Provides automated active/passive failover with minimal downtime. It can also be used to significantly reduce downtime when you patch vCenter Server Appliance.
---	---

Many vSphere features, such as those described in [Table 1-4](#), require specific vSphere configuration, and some require specific licensing, but they do not require the installation or deployment of additional software or virtual appliances.

**Table 1-4** Available vSphere Features

Available vSphere Features	Description
vCenter Appliance File-Based Backup and Restore	A feature introduced in vSphere 7.0 that enables you to back up and restore the vCenter Server Appliance instances.
vMotion	A feature that provides live virtual machine migrations with negligible disruption from a source ESXi host to a target ESXi host.
<b>vSphere HA</b>	A feature provides automated failover protection for VMs against host, hardware, network, and guest OS issues. In the



event of host system failure, it performs cold migrations and restarts failed VMs on surviving hosts.

---

<b><i>Distributed Resource Scheduler (DRS)</i></b>	A feature that places and starts VMs on appropriate ESXi hosts and hot-migrates VMs using vMotion when there is contention for compute resources.
Storage vMotion	A feature that performs live migrations with negligible disruption of VMs from a source datastore to a target datastore.
Fault Tolerance (FT)	A feature that provides automated live failover protection for VMs against host, hardware, network, and guest OS issues.
Distributed Power Management (DPM)	A feature that optimizes power consumption in an ESXi cluster.
<b><i>Proactive HA</i></b>	A feature that minimizes VM downtime by proactively detecting hardware failures and placing the host in Quarantine Mode or Maintenance Mode.
Content library	A centralized repository used manage and distribute templates, ISO files, scripts, vApps, and other files associated with VMs.
Host profiles	A feature that provides a means to apply a standard configuration to a set of ESXi hosts.

---

The add-on products in [Table 1-5](#) are commonly used in a vSphere environment and are discussed in this book. These products can be sold separately from vSphere.

**Table 1-5** Add-on Products

---

<b>Product</b>	<b>Description</b>
----------------	--------------------

---

vSAN	A product that provides a SAN experience to your vSphere environment, leveraging local storage in the ESXi hosts. It tightly integrates with
------	--

A vSphere and is the leading Hyper-Converged Infrastructure (HCI)  
N solution for providing a flash-optimized, secure, and simple-to-use SAN.

---

N A product that adds software-based virtualized networking and security  
S to a vSphere environment.

---

X

---

v A suite of products that add operations (vRealize Operations Manager),  
R automation (vRealize Automation), and orchestration (vRealize  
e Orchestrator) to a vSphere environment.

---

al

iz

e

S

u

it

e

---

## Note

Although it is an add-on product, vSAN is covered in the VCP-DCV certification exam and in this book.

The vSphere Host Client is a web-based interface provided by each ESXi host. It is available immediately following the installation of a host. Its primary purpose is to provide a GUI for configuration, management, and troubleshooting purposes when vCenter Server is not available. For example, during the implementation of a new vSphere environment, you could use the vSphere Host Client to create virtual machines for running DNS, Active Directory, and vCenter Server databases prior to deploying vCenter Server. As another example, you could use the vSphere Host Client to power down, troubleshoot, reconfigure, and restart the vCenter Server virtual machine.

The HTML5-based vSphere Client is the preferred web-based GUI for managing vSphere. It is provided by services running in the vCenter Server. The flash-based vSphere Web Client used in previous vSphere versions has been deprecated and is no longer available.

## Editions and Licenses

VMware vSphere comes in many editions, and each edition is intended to address specific use cases by providing specific features. When planning for a vSphere environment, you should prepare to procure at least three line items: a vCenter Server license, a vSphere license, and support for the environment. The vCenter Server license line item should identify the desired edition and quantity (that is, the number of vCenter Server instances).

[Table 1-6](#) provides a summary of the features that are provided with each edition of vCenter Server 7.

**Table 1-6** vCenter Server Editions

<b>Feature</b>	<b>Essentials</b>	<b>Essentials Plus</b>	<b>Foundation</b>	<b>Standard</b>
Number of ESXi hosts	3 (2 CPU max)	3 (2 CPU max)	4	2000
vCenter License	Packaged with vSphere license in Essentials	Packaged with vSphere license in Essentials Plus	Sold separately from vSphere license	Sold separately from vSphere license
Basic vCenter features, like single pane of glass management, Lifecycle Manager, and VMware Converter	Supported	Supported	Supported	Supported
Common vCenter features like vMotion, vSphere HA, and vSphere Replication	Not supported	Supported	Supported	Supported
Advanced features like vCenter Server High Availability (VCHA) and vCenter Server Backup and Restore	N/A	N/A	N/A	Supported

You need to obtain a vSphere license in order to apply to license physical CPUs on your ESXi hosts. Starting with vSphere 7.0, one vSphere CPU license covers

up to 32 cores. If a CPU has more than 32 cores, you need additional CPU licenses. The number of vSphere CPU licenses consumed by an ESXi host is determined by the number of physical CPUs on the host and the number of cores in each physical CPU. For example, you can assign a vSphere license for 10 32-core CPUs to any of the following combinations of hosts:

- Five hosts with 2 CPUs and 32 cores per CPU
- Five hosts with 1 CPU with 64 cores per CPU
- Two hosts with 2 CPUs and 48 cores per CPU and two hosts with 1 CPU and 20 cores per CPU

The major editions of vSphere 7.0 are Standard and Enterprise Plus. Other editions may be licensed in different manners than the major editions. For example, the vSphere Desktop edition (for VDI environments) and VMware vSphere Remote Office Branch Office (for IT remote sites) are licensed per virtual machine.

[Table 1-7](#) lists some of the features that are provided with the major editions of vSphere 7.0.



**Table 1-7** Features in vSphere Editions

Feature	Standard	Enterprise Plus
vSphere HA, vSphere Replication, Storage vMotion, Quick Boot, vCenter Backup and Restore, vVols	Supported	Supported
Distributed Switch, Proactive HA, NIOC, SIOC, Storage DRS, DRS, DPM, VM Encryption, Cross-vCenter vMotion, Long Distance vMotion, vTrust Authority, SR-IOV, vSphere Persistent Memory	Not supported	Supported
vSphere Fault Tolerance	Supported up to 2 vCPUs	Supported up to 8 vCPUs

---

## vCenter Server Topology

This section describes the architecture for the vCenter Server.

vSphere 6.x supports multiple vCenter Server topologies and configurations, involving components and technologies such as vCenter Server Appliance, vCenter Server for Windows, embedded database (PostgreSQL), external (SQL Server or Oracle) database, external Platform Services Controller (PSC), embedded PSC, Enhanced Linked Mode, and Embedded Linked Mode. In vSphere 7.0, the vCenter Server configuration and topology are much simpler.

Beginning in vSphere 7.0, vCenter Server Appliance is required. Windows-based vCenter Servers are not supported. External PSCs are not supported. The services provided by PSC in prior vCenter Server versions are directly integrated into vCenter Server Appliance 7.0 and are no longer described as a part of the PSC in most documentation. For example, in vSphere 7.0, the *Platform Services Controller Administration* publication is replaced with the *vSphere Authentication* publication. [Table 1-8](#) describes the main services in vCenter Server Appliance and related services in the ESXi host.

**Table 1-8** Services in vCenter Server Appliance

---

### Service Description

e

---

vCenter Single Sign-On	An authentication service that utilizes a secure token exchange mechanism rather than requiring components to authenticate users per component.
------------------------	---

Sign-On

---

Security Token Service (STS)	A component that is part of vCenter Single Sign-On and provides SAML tokens to authenticate users to other vCenter components instead of requiring users to authenticate to each component. A user who authenticates to vCenter Single Sign-On is granted SAML tokens, which are then used for authentication.
------------------------------	--

Service (STS)

Administration server	A component that provides vCenter Single Sign-On administration and configuration from the vSphere Client.
vCenter Lookup Service	A service that contains the topology of the vSphere infrastructure, allowing secure communication between vSphere components.
VMware Directory Service	The directory service for the vCenter Single Sign-On (SSO) domain (vsphere.local).
vCenter Server plug-ins	Applications that add functionality to vCenter. These usually consist of server and client components.
vCenter Server database	A database that contains the status of all virtual machines, ESXi hosts, and users. It is deployed via the vCenter Server deployment wizard.
tcServer	A service that is co-installed with vCenter and is used by web services such as ICIM/Hardware status, Performance charts, WebAccess, Storage Policy Based Services, and vCenter Service status.
Licenses	A service that is used to store the available licenses and manage the

nse license assignments for the entire vSphere environment.  
Serv  
ice

---

vCe A service that is installed on an ESXi host when that host is added to  
nter vCenter's inventory. This service collects, communicates, and runs  
Serv actions initiated from the vSphere Client.  
er  
Age  
nt

---

Host An administrative agent installed with the ESXi installation.  
Age Responsible for collecting, communicating, and running actions  
nt initiated from the vSphere Host Client.

---

If you upgrade or migrate a vCenter Server deployment that uses an external PSC, you must converge the PSC into a vCenter Server Appliance instance that you specify. In domains with multiple vCenter Server instances, you must identify the SSO replication partner for each subsequent vCenter Server. If you upgrade or migrate using the GUI-based installer, the wizard prompts you to specify the replication topology. If you upgrade or migrate using the CLI-based installer, you specify the replication topology using the JSON templates. During the upgrade or migration process, the new vCenter Server Appliance 7.0 incorporates the former PSC services, enabling you to decommission the original external PSC.

## Single Sign-On (SSO) Domain

Each vCenter Server is associated with a ***vCenter Single Sign-On (SSO)*** domain, whose default name is `vsphere.local`. You can change the SSO domain name during deployment. The SSO domain is considered the local domain for authentication to vCenter Server and other VMware products, such as vRealize Operations.

During vCenter Server Appliance deployment, you must create an SSO domain or join an existing SSO domain. The domain name is used by the VMware Directory Service (vmdir) for all Lightweight Directory Access Protocol (LDAP) internal structuring. You should give your domain a unique name that is not used by OpenLDAP, Microsoft Active Directory, and other directory services.


You can add users and groups to the SSO domain. You can add an Active

Directory or LDAP identity source and allow the users and groups in that identity source to authenticate.

## Enhanced Linked Mode



You can use Enhanced Linked Mode to link multiple vCenter Server systems. With Enhanced Linked Mode, you can log in to all linked vCenter Server systems simultaneously and manage the inventories of the linked systems. This mode replicates roles, permissions, licenses, and other key data across the linked systems. To join vCenter Server systems in Enhanced Linked Mode, connect them to the same vCenter SSO domain, as illustrated in [Figure 1-1](#). Enhanced Linked Mode requires the vCenter Server Standard licensing level and is not supported with vCenter Server Foundation or vCenter Server Essentials. Up to 15 vCenter Server Appliance instances can be linked together by using Enhanced Linked Mode.



**FIGURE 1-1** Enhanced Linked Mode with Two vCenter Server Appliance 7.0 Instances

## vCenter HA

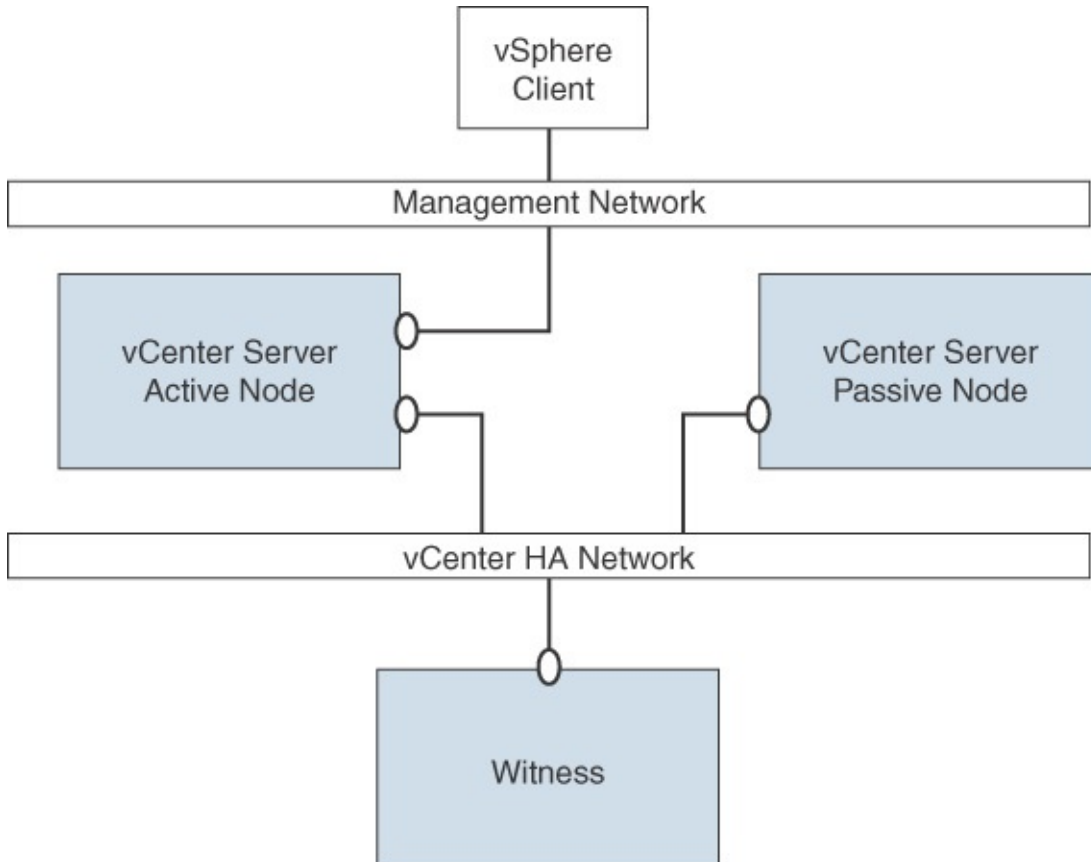
A **vCenter HA** cluster consists of three vCenter Server instances. The first instance, initially used as the Active node, is cloned twice to a Passive node and to a Witness node. Together, the three nodes provide an active/passive failover solution.

Deploying each of the nodes on a different ESXi instance protects against hardware failure. Adding the three ESXi hosts to a DRS cluster can further protect your environment.

When the vCenter HA configuration is complete, only the Active node has an active management interface (public IP address), as illustrated in [Figure 1-2](#). The



three nodes communicate over a private network called a vCenter HA network that is set up as part of the configuration. The Active node continuously replicates data to the Passive node.



**FIGURE 1-2** vCenter Server HA Nodes

All three nodes are necessary for the functioning of this feature. [Table 1-9](#) provides details for each of the nodes.

**Table 1-9** vCenter HA Node Details

<b>Node Type</b>	<b>Description</b>
Active	Is the active vCenter Server instance.  Uses a public IP address for the management interface.  Replicates data to the Passive node using the vCenter HA network.

	Communicates with the Witness node using the vCenter HA network.
Passive	<p>Is cloned from the Active node.</p> <p>Uses the vCenter HA network to constantly receive updates from the Active node.</p> <p>Automatically takes over the role of the Active node if a failure occurs.</p>
Witness	<p>Is a lightweight clone of the Active node.</p> <p>Provides a quorum to protect against a split-brain situation.</p>

## Infrastructure Requirements

This section describes some of the main infrastructure requirements that you should address prior to implementing vSphere.

### Compute and System Requirements

When preparing to implement a vSphere environment, you should prepare sufficient supported compute (CPU and memory) resources, as described in this section.

#### vCenter Server

vCenter Server Appliance 7.0 can be deployed on ESXi 6.5 hosts or later, which can be managed by vCenter Server 6.5 or later.

To prepare for deployment of vCenter Server, you should plan to address the compute specifications listed in [Table 1-10](#).

**Table 1-10** Compute Specifications for vCenter Server Appliance

Component	Number of CPUs	Memory
Tiny Environment	2	12 GB

Up to 10 hosts or 100 virtual machines

---

Small Environment	4	19 GB
-------------------	---	-------

Up to 100 hosts or 1000 virtual machines

---

Medium Environment	8	28 GB
--------------------	---	-------

Up to 400 hosts or 4000 virtual machines

---

Large Environment	16	37 GB
-------------------	----	-------

Up to 1000 hosts or 10,000 virtual machines

---

X-Large Environment	24	56 GB
---------------------	----	-------

Up to 2000 hosts or 35,000 virtual machines

---

## Note

If you want to have an ESXi host with more than 512 LUNs and 2048 paths, you should deploy a vCenter Server Appliance instance for a Large Environment or X-Large Environment component.

## ESXi

To install ESXi 7.0, ensure that the hardware system meets the following requirements:

### Key Topic

- A supported system platform, as described in the *VMware Compatibility Guide*.
- Two or more CPU cores.
- A supported 64-bit x86 processor, as described in the *VMware Compatibility Guide*.
- The CPU's NX/XD bit enabled in the BIOS.

- 4 GB or more of physical RAM. (VMware recommends 8GB or more for production environments.)
- To support 64-bit virtual machines, hardware virtualization (Intel VT-x or AMD RVI) enabled on the CPUs.
- One or more supported Ethernet controllers, Gigabit or faster, as described in the *VMware Compatibility Guide*.
- A SCSI disk or a local, non-network RAID LUN with unpartitioned space for the virtual machines.
- For Serial ATA (SATA), a disk connected through supported SAS controllers or supported on-board SATA controllers.
- A boot disk of at least 8 GB for USB or SD devices and 32 GB for other HDD, SSD, NVMe, and other device types. The boot device must not be shared between ESXi hosts.

### **Note**

SATA disks are considered remote, not local. These disks are not used as scratch partitions by default because they are considered remote. You cannot connect a SATA CD-ROM device to a virtual machine on an ESXi 7.0 host. To use the SATA CD-ROM device, you must use IDE emulation mode.

For vSphere 7.0, you should ensure that you meet the ESXi booting considerations:

- You can boot using the Unified Extensible Firmware Interface (UEFI), which enables booting from hard drives, CD-ROM drives, or USB media.
- VMware Auto Deploy supports network booting and provisioning of ESXi hosts with UEFI.
- You have boot systems from disks larger than 2 TB if the system firmware add-in card firmware supports it, according to vendor documentation.

### **Note**

Changing the host boot type between legacy BIOS and UEFI is not supported after you install ESXi 7.0.

## Storage Requirements

When preparing to implement a vSphere environment, you should ensure that you have sufficient supported storage resources, as described in this section.

### vCenter Server Appliance

As part of preparing for the deployment of vCenter Server Appliance, you should plan to address storage requirements. [Table 1-11](#) lists the storage requirements for a vCenter Server Appliance instance. It allows for Lifecycle Manager, which runs as a service in vCenter Server Appliance.

**Table 1-11** Storage Sizes for vCenter Server Appliance

Deployment Size	Default Storage Size	Large Storage Size	X-Large Storage Size
Tiny	415 GB	1490 GB	3245 GB
Small	480 GB	1535 GB	3295 GB
Medium	700 GB	1700 GB	3460 GB
Large	1065 GB	1765 GB	3525 GB
X-Large	1805 GB	1905 GB	3665 GB

## ESXi

Installing ESXi 7.0 requires a boot device that is a minimum of 8 GB. Upgrading to ESXi 7.0 requires a 4 GB minimum boot device. When booting from a local disk, SAN, or iSCSI LUN, a 32 GB disk is required to allow for the creation of the boot partition, boot banks, and a VMFS\_L ESX=OSData volume. The ESX-OSData volume replaces the legacy /scratch partition, VM-tools, and core dump location. If no local disk is found, ESXi 7.0 functions in degraded mode and places the /scratch partition on the ESXi host's RAM disk and links it to /tmp/scratch. You can reconfigure /scratch to use a separate disk or LUN. For best performance and memory optimization, do not run the ESXi host in degraded mode. Likewise, when installing ESXi 7.0 on USB and SD devices, the

installer attempts to allocate a scratch region on a local disk; otherwise, it places /scratch on the RAM disk.

## Note

You cannot roll back to an earlier version of ESXi after upgrading. If you are concerned about upgrading, create a backup of the boot device prior to upgrading; if needed, you can restore from this backup after the upgrade.

The following are recommended for ESXi 7.0 installations:

- **8 GB USB drive or SD card with a 32 GB local disk:** Boot partitions reside on USB or SD, and ESXi-OSData resides on a local disk.
- **Local disk with 32 GB minimum:** This contains boot and ESX-OSData.
- **Local disk with 142 GB or more:** This contains boot, ESX-OSData, and a VMFS datastore.

## Network Requirements

This section describes some of the key networking requirements for a successful vSphere deployment.

## Networking Concepts

In order to prepare for network virtualization in vSphere, you should understand some of the following concepts:

- **Physical network:** This is a network of physical machines that are connected so that they can send data to and receive data from each other.
- **Virtual network:** This is a network of virtual machines running on a physical machine that are connected logically to each other so that they can send data to and receive data from each other.
- **Opaque network:** This is a network created and managed by a separate entity outside vSphere. For example, logical networks that are created and managed by VMware NSX appear in vCenter Server as opaque networks of the type nsx.LogicalSwitch. You can choose an opaque network as the backing for a VM network adapter. To manage an opaque network, use the

management tools associated with the opaque network, such as VMware NSX Manager or the VMware NSX API management tools.

- **vSphere standard switch:** This type of switch works much like a physical Ethernet switch. It detects which virtual machines are logically connected to each of its virtual ports and uses that information to forward traffic to the correct virtual machines. A vSphere standard switch can be connected to physical switches by physical Ethernet adapters, also referred to as uplink adapters.
- **VMkernel TCP/IP networking layer:** This layer provides connectivity to hosts and handles the standard infrastructure traffic of vSphere vMotion, IP Storage, Fault Tolerance, and vSAN.

VMware recommends using network segmentation in vSphere environments for separating each type of VMkernel traffic and virtual machine traffic. You can implement network segments by using unique VLANs and IP subnets. Here is a set of commonly used network segments in vSphere:

- Management
- vMotion
- vSphere Replication
- vSphere High Availability Heartbeat
- Fault Tolerance
- IP Storage
- Virtual Machine (typically segregated further by application or by other factors, such as test and production)

## vCenter Server Network Requirements

Table 1-12 provides details for some of the required network connectivity involving vCenter Server. For each applicable connection, you should ensure that your network and firewall allow the described connectivity.

**Table 1-12** Required Ports for vCenter Sever

---

<b>Proto</b>	<b>Description</b>	<b>Required for</b>
--------------	--------------------	---------------------

**col/P**

**ort**

TC P 22	System port for SSHD	vCenter Server (Must be open for upgrade of the appliance.)
TC P 80	Port for direct HTTP connections; redirects requests to HTTPS port 443	vCenter Server
TC P 88	Required to be open to join Active Directory	vCenter Server
TC P/ UD P 389	LDAP port for directory services for the vCenter Server group	vCenter Server to vCenter Server
TC P 443	Default port used by vCenter Server to listen for connections from the vSphere Web Client and SDK clients	vCenter Server to vCenter Server
TC P/ UD P 514	vSphere Syslog Collector port for vCenter Server and vSphere Syslog Service port for vCenter Server Appliance	vCenter Server
TC P/ UD P 902	Default port that the vCenter Server system uses to send data to managed hosts	vCenter Server
TC P 151 4	vSphere Syslog Collector TLS port for vCenter Server	vCenter Server
TC P	Control interface RPC for Single Sign-On	vCenter Server



201  
2

---

TC P 201 4	RPC port for VMware Certificate Authority (VMCA) APIs	VMCA
---------------------	---	------

---

TC P/ UD P 202 0	Authentication framework management	vCenter Server
---------------------------------	-------------------------------------	----------------

---

TC P 548 0	vCenter Server Appliance Management Interface (VAMI)	vCenter Server
---------------------	--	----------------

---

TC P/ UD P 650 0	ESXi Dump Collector port	vCenter Server
---------------------------------	--------------------------	----------------

---

TC P 708 0, 127 21	Secure Token Service (internal ports)	vCenter Server
-----------------------------------	---------------------------------------	----------------

---

TC P 708 1	vSphere Client (internal ports)	vCenter Server
---------------------	---------------------------------	----------------

---

TC P 747 r	VMware vSphere Authentication Proxy	vCenter Server
---------------------	-------------------------------------	----------------

3,  
747  
6

TC P 808 4	vSphere Lifecycle Manager SOAP port used by vSphere Lifecycle Manager client plug-in	vSphere Lifecycle Manager
TC P 908 4	vSphere Lifecycle Manager Web Server Port used by ESXi hosts to access host patch files from vSphere Lifecycle Manager server	vSphere Lifecycle Manager
TC P 908 7	vSphere Lifecycle Manager Web SSL port used by vSphere Lifecycle Manager client plug-in for uploading host upgrade files to vSphere Lifecycle Manager server	vSphere Lifecycle Manager
TC P 944 3	vSphere Web Client HTTPS	vCenter Server

## ESXi Network Requirements

Table 1-13 provides details about some of the required network connectivity involving ESXi. For each applicable connection, you should ensure that your network and firewall allow the described connectivity.

**Table 1-13** Required Ports for ESXi

Protocol/Port	Service	Direction	Description
TCP 5988	CIM Server	Inbound	Server for Common Information Model (CIM)
TCP 5989	CIM Secure	Inbound	Secure Server for CIM

	Server		
UDP 8301, 8302	DVSSync	Inbound, outbound	Used for synchronizing states of distributed virtual ports between hosts that have VMware FT record/replay enabled
TCP 902	NFC	Inbound, outbound	ESXi uses Network File Copy (NFC) for operations such as copying and moving data between datastores
UDP 12345, 23451	vSAN Clustering	Inbound, outbound	Used by vSAN nodes for multicast to establish cluster members and distribute vSAN metadata
UDP 68	DHCP	Inbound, outbound	DHCP client for IPv4
UDP 53	DNS	Inbound	DNS client
TCP/UDP 53	DNS	Outbound	DNS client
TCP/UDP 8200, 8100, 8300	Fault Tolerance	Inbound	Traffic between hosts for vSphere Fault Tolerance (FT)
TCP/UDP 80, 8200, 8100, 8300	Fault Tolerance	Outbound	Supports vSphere Fault Tolerance (FT)
TCP 2233	vSAN Transport	Inbound	vSAN reliable datagram transport for vSAN storage I/O
TCP 22	SSH	Inbound	SSH server
TCP 902,	vSphere	Inbound	Allows user connections from vSphere Web

443	rend Web Client	Client
TCP/UDP 547	DHC Pv6	Outbound DHCP client for IPv6
UDP 9	WOL	Outbound Wake-on-LAN
TCP 3260	iSCSI	Outbound Supports software iSCSI
TCP 8000	vMotion	Outbound Supports vMotion
UDP 902	vCenter Agent	Outbound Used by the vCenter Agent

## Infrastructure Services

In addition to providing the required compute, storage, and network infrastructure, you should provide supporting infrastructure services, such as Active Directory (AD), Domain Name System (DNS), and Network Time Protocol (NTP).

### AD

In many vSphere environments, vCenter Single Sign-On (SSO) is integrated with directory services, such as Microsoft Active Directory (AD). SSO can authenticate users from internal users and groups, and it can connect to trusted external directory services such as AD. If you plan to leverage AD for an SSO identity source, you should ensure that the proper network connectivity, service account credentials, and AD services are available and ready for use.

If you plan to install vCenter Server for Windows and use AD identity sources, you should ensure that the Windows server is a member of the AD domain but is not a domain controller.

### Note

If the system you use for your vCenter Server installation belongs to a

workgroup rather than a domain, vCenter Server cannot discover all domains and systems available on the network when using some features.

## **DNS**

You might want to assign static IP addresses and resolvable fully qualified domain names (FQDNs) to your vSphere components, such as vCenter Server and ESXi hosts. Before installing these components, you should ensure that the proper IP addresses and FQDN entries are registered in your DNS server. You should configure forward and reverse DNS records.

For example, prior to deploying vCenter Server Appliance, you should assign a static IP address and host name in DNS. The IP address must have a valid (internal) domain name system (DNS) registration. During the vCenter Server installation, you must provide the fully qualified domain name (FQDN) or the static IP. VMware recommends using the FQDN. You should ensure that DNS reverse lookup returns the appropriate FQDN when queried with the IP address of the vCenter appliance. Otherwise, the installation of the Web Server component that supports the vSphere Web client fails.

When you deploy vCenter Server Appliance, the installation of the web server component that supports the vSphere Web Client fails if the installer cannot look up the FQDN for the appliance from its IP address. Reverse lookup is implemented using PTR records. If you plan to use an FQDN for the appliance system name, you must verify that the FQDN is resolvable by a DNS server.

Starting with vSphere 6.5, vCenter Server supports mixed IPv4 and IPv6 environment. If you want to set up vCenter Server Appliance to use an IPv6 address version, use the FQDN or host name of the appliance.

It is important to ensure that each vSphere Web Client instance and each ESXi host instance can successfully resolve the vCenter Server FQDN. It is also important to ensure that the ESXi host management interface has a valid DNS resolution from the vCenter Server and all vSphere Web Client instances. Finally, It is important to ensure that the vCenter Server has a valid DNS resolution from all ESXi hosts and all vSphere Web Clients.

## **NTP**

It is important to provide time synchronization between the nodes. All vCenter

Server instances must be time synchronized. ESXi hosts must be time synchronized to support features such as vSphere HA. In most environments, you should plan to use NTP servers for time synchronization. Prior to implementing vSphere, verify that the NTP servers are running and available.

Be prepared to provide the names or IP addresses for the NTP servers when installing vSphere components such as vCenter Server and ESXi. For example, during the deployment of vCenter Server Appliance, you can choose to synchronize time with NTP servers and provide a list of NTP server names or IP addresses, separated by commas. Alternatively, you can choose to allow the appliance to synchronize time with the ESXi host.

### **Note**

If a vCenter Server Appliance instance is set for NTP time synchronization, it ignores its **time\_tools-sync** Boolean parameter. Otherwise, if the parameter is TRUE, VMware Tools synchronizes the time in the appliance's guest OS with the ESXi host.

## **Other Requirements**

This section describes a few additional requirements for some of the optional components (refer to [Table 1-3](#)), available vSphere features (refer to [Table 1-4](#)), and add-on products (refer to [Table 1-5](#)).

### **Additional Requirements**

The following sections describe some of the requirements for a variety of commonly used vSphere features.

### **User Interfaces**

The vSphere Host Client and vSphere Client utilize HTML5. The flash-based vSphere Web Client is not supported in vSphere 7. For Windows users, VMware supports Microsoft Edge 38 and later, Microsoft Internet Explorer 11.0.96 and later, Mozilla Firefox 45 and later, Google Chrome 50 and later, and Safari 5.1 and later. For Mac users, VMware supports Safari 5.1 and later, Mozilla Firefox 45 and later, and Google Chrome 50 and later.

## **vCenter Server File-Based Backup and Restore**

If you plan to schedule file-based backups using the VAMI, you must prepare an FTP, FTPS, HTTP, HTTPS, or SCP server with sufficient disk space to store the backups.

## **GUI Installer**

You can use the GUI installer to interactively install vCenter Server Appliance. To do so, you must run the GUI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the instance.

## **Distributed Power Management (DPM)**

DPM requires the ability to wake a host from standby mode, which means it needs to be able to send a network command to the host to power on. For this feature, DPM requires iLO, IPMI, or a Wake-on-LAN (WoL) network adapter to be present in each participating host in the cluster. DPM must be supplied with the proper credentials to access the interface and power on the host.

## **vSphere Replication Requirements**

In order to use vSphere Replication 8.3, you must deploy a vSphere Replication Management Service (VRMS) appliance. Optionally, you can add nine additional vSphere Replication Service (VRS) appliances. You should plan for the compute, storage, and network needs of these appliances.

The VRMS appliance requires two vCPUs and 8 GB memory. Optionally, you can configure it for 4 vCPUs. Each VRS appliance requires two vCPUs and 716 MB memory. The amount of CPU and memory resources consumed by the vSphere Replication agent on each host is negligible.

Each VRMS and VRS appliance contains two virtual disks whose sizes are 13 GB and 9 GB. To thick provision these virtual disks, you must provide 22 GB storage. If you do not reserve the memory, you should provide storage for the VRMS (8 GB) and VRS (716 MB each) swap files.

Each appliance has at least one network interface and requires at least one IP address. Optionally, you can use separate network connections to allow each appliance to separate management and replication traffic.

The main storage requirement for vSphere Replication is to support the target

datastore to which the VMs will be replicated. At a minimum in the replication target datastore, you should provide enough storage to replicate each virtual disk, to support each replicated VM's swap file, and to store each VM's multiple point-in-time captures (snapshots).

## **vCenter High Availability Requirements**



The minimum software version for the nodes in a vCenter HA cluster is vCenter Server 6.5. The minimum software versions for the environment (such as a management cluster) where the vCenter HA nodes live are ESXi 6.0 and vCenter Server 6.0. Although not required, VMware recommends that you use a minimum of three ESXi hosts with DRS rules to separate the nodes onto separate hosts. You must use a vCenter Server Appliance Small or larger deployment size (not Tiny) and a vCenter Server Standard (not Foundation) license. A single vCenter Server license is adequate for a single vCenter HA cluster. vCenter HA works with VMFS, NFS, and vSAN datastores.

You must configure the appropriate virtual switch port groups prior to configuring vCenter HA. The vCenter HA network connects the Active, Passive, and Witness nodes, replicates the server state, and monitors heartbeats. The vCenter HA network must be on a different subnet than the management network, must provide less than 10 ms latency between nodes, and must not use a default gateway. The vCenter HA and management network IP addresses must be static.

You can use the Set Up vCenter HA wizard in the vSphere Client to configure vCenter HA. You have the option to perform an automatic configuration or a manual configuration. The automatic configuration requires a self-managed vCenter Server rather than a vCenter Server that resides in a management cluster that is managed by another vCenter Server. The automatic configuration automatically clones the initial (Active node) vCenter Server to create the Witness and Passive nodes. The manual configuration requires you to clone the Active node yourself but gives you more control.

When configuration is complete, the vCenter HA cluster has two networks: the management network on the first virtual NIC and the vCenter HA network on the second virtual NIC.

## **SDDC Requirements**



To build a software-defined data center (SDDC), you may plan to implement additional VMware products, such as vSAN, NSX, and vRealize Suite. Here are some of the requirements you should address.

## vSAN

When preparing to implement vSAN, verify that the ESXi hosts meet the vSAN hardware requirements. All the devices, drivers, and firmware versions in your vSAN configuration must be certified and listed in the vSAN section of the *VMware Compatibility Guide*.

[Table 1-14](#) lists the storage device requirements for vSAN hosts.

**Table 1-14** Storage Device Requirements for vSAN Hosts

Component	Requirements
Cache	One SAS or SATA solid-state disk (SSD) or PCIe flash device
Virtual machine data storage	For hybrid group configuration, at least one SAS or NL-SAS magnetic disk
	For all-flash group configuration, at least one SAS or SATA solid-state disk (SSD) or PCIe flash device
Storage controllers	One SAS or SATA host bus adapter (HBA) or a RAID controller that is in passthrough mode or RAID 0 mode

You need to prepare a network for vSAN traffic. This is the network in which you will connect a VMkernel network adapter for each ESXi host. For non-stretched vSAN clusters, the network should provide a maximum round-trip time (RTT) of 1 ms.

## NSX

When preparing to implement NSX, ensure that you address the hardware and network latency requirements.

A typical NSX Data Center for vSphere (NSX-V) implementation involves deploying NSX Manager, three NSX Controller instances, and one or more NSX Edge instances. [Table 1-15](#) lists the hardware requirements for these NSX-V Version 6.4 devices.

**Table 1-15** Hardware Requirements for NSX Appliances

Appliance	Memory	vCPUs	Disk Space
NSX Manager	16 GB	4 or 8	60 GB
NSX Controller	4 GB	4	28 GB
NSX Edge	Compact: 512 MB Large: 1 GB Quad Large: 2 GB X-Large: 8 GB	Compact: 1 Large: 2 Quad Large: 4 X-Large: 6	X-Large: 2.75 GB Other: 1 GB

You should ensure that the network latency is no higher than 150 ms RTT for NSX Manager connections with NSX Controller instances, vCenter Server, and ESXi hosts.

## vRealize Suite

vRealize Operations (vROps) is a tool that provides monitoring of and analytics for a vSphere environment. It provides smart alerts and also identifies undersized or oversized virtual machines. Many businesses use vROps to improve the operation of vSphere and SDDC. They use it for many other purposes, such as capacity planning, proactively remediating issues, reclaiming wasted resources, and compliance.

vRealize Automation is cloud automation software that speeds up the delivery of infrastructure and application resources on-premises and in the public cloud. It provides self-service and policy-based automation. Many businesses use vRealize Automation to automate processes and improve the delivery of IT services and applications.

vRealize Network Insight (vRNI) is a tool that can collect details and flows from physical and virtual network infrastructure. You can use it to help plan and monitor a software-defined network. Many businesses use vRNI for microsegmentation planning and network troubleshooting in an SDDC.

vRealize Log Insight (vRLI) is a tool that can collect and analyze logs from vSphere components, virtual machines, physical machines, and the entire

infrastructure. Many businesses use vRLI to centrally collect and analyze logs from the entire SDDC.

## **VMware Cloud vs. VMware Virtualization**

This section provides brief explanations of virtualization and cloud technologies.

### **Server Virtualization**

VMware vSphere 7.0 is the industry-leading virtualization and cloud platform. It provides virtualization (abstraction, pooling, and automation) of x86-64 based server hardware and related infrastructure, such as network switches. It provides live workload migrations, high availability, and efficient management at scale in a secured infrastructure.

### **VMware SDDC**



A software-defined data center (SDDC) is a data center that leverages logical infrastructure services that are abstracted from the underlying physical infrastructure. It allows any application to run on a logical platform that is backed by x86-64, any storage, and any network infrastructure. Pioneered by VMware, a SDDC is the ideal architecture for private, public, and hybrid clouds. It extends virtualization concepts to all data center resources and services.

The SDDC includes compute virtualization (vSphere), network virtualization (NSX), and software-defined storage (vSAN and vVols) to deliver abstraction, pooling, and automation of the compute, network, and storage infrastructure services. It includes vRealize Automation and vRealize Operations to deliver policy-based automated management of the data center, services, and applications.

### **vCloud Suite and Private Clouds**

VMware vCloud Suite is an enterprise-ready private cloud software suite that includes vSphere for data center virtualization and VMware vRealize Suite for cloud management.

### **VCF and Hybrid Clouds**

A **hybrid cloud** is a cloud that is a combination of a private cloud, public cloud, and on-premises infrastructure. It is the result of combining any cloud solution with in-house IT infrastructure.

VMware Cloud Foundation (VCF) is the industry's most advanced hybrid cloud platform. It provides a complete set of software-defined services for compute, storage, networking, security, and cloud management to run enterprise apps in private or public environments. It delivers a simple path to the hybrid cloud by leveraging a common infrastructure and consistent operational model for on-premises and off-premises data centers.

## **VMC on AWS**

**VMware Cloud (VMC)** on AWS is an integrated cloud offering jointly developed by AWS and VMware that provides a highly scalable, secure service that allows organizations to seamlessly migrate and extend their on-premises vSphere-based environments to the AWS cloud. You can use it to deliver a seamless hybrid cloud by extending your on-premises vSphere environment to the AWS cloud.

## **VMware vCloud Director**

VMware vCloud Director is a cloud service-delivery platform used by some cloud providers to operate and manage cloud-based services. Service providers can use vCloud Director to deliver secure, efficient, and elastic cloud resources to thousands of customers.

## **Cloud Automation**

VMware Cloud Assembly and VMware Service Broker are software as a service (SaaS) offerings that address similar use cases to the on-premises cases that VMware vRealize Automation addresses.

## **Exam Preparation Tasks**

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15, “Final Preparation,”](#) and the exam simulation questions on the companion website.

## Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 1-16](#) lists these key topics and the page number on which each is found.



**Table 1-16** Key Topics

Key Topic Element	Description	Page Number
<a href="#">Table 1-7</a>	vSphere editions	<a href="#">10</a>
Paragraph	Enhanced Linked Mode	<a href="#">12</a>
List	ESXi system hardware requirements	<a href="#">15</a>
Paragraph	vCenter HA requirements	<a href="#">24</a>
Paragraph	VMware SDDC	<a href="#">27</a>

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

- [vCenter Single Sign-On \(SSO\)](#)
- [VMware Cloud \(VMC\)](#)
- [hybrid cloud](#)
- [vCenter HA](#)
- [vSphere HA](#)
- [Distributed Resource Scheduler \(DRS\)](#)
- [Proactive HA](#)

## Answer Review Questions

1. You plan to implement vSphere 7.0 and use vSphere Fault Tolerance to protect virtual machines with two vCPUs. Which is the minimum vSphere edition that you need?
  - a. vSphere Essentials Plus
  - b. vSphere Foundations
  - c. vSphere Standard
  - d. vSphere Enterprise Plus
2. You are planning to deploy vSphere 7.0. Where should you place VMware Directory Service?
  - a. Nowhere as VMware Directory Service is not used
  - b. In an external PSC
  - c. Either in an external PSC or in an embedded PSC
  - d. In vCenter Server
3. You are planning to deploy ESXi in a vSphere 7.0 environment and want to minimize memory per ESXi host. What is the minimum host memory that VMware recommends for a production environment?
  - a. 4 GB
  - b. 8 GB
  - c. 16 GB
  - d. 24 GB
4. You are planning to install vCenter Server 7.0 and want to use the GUI installer. Which of the following are supported locations from which to run the installer? (Choose two.)
  - a. The vSphere Host Client on an ESXi host
  - b. The vCenter Server Appliance Management Interface
  - c. Windows
  - d. Mac
5. Which of the following is the industry's most advanced hybrid cloud

platform?

- a.** VMware Cloud Assembly
- b.** VCF
- c.** VMC on AWS
- d.** vRealize Automation

## Chapter 2

# Storage Infrastructure

This chapter covers the following topics:

- [Storage Models and Datastore Types](#)
- [vSAN Concepts](#)
- [vSphere Storage Integration](#)
- [Storage Multipathing and Failover](#)
- [Storage Policies](#)
- [Storage DRS \(SDRS\)](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.3, 1.3.1, 1.3.2, 1.3.3, 1.3.4, 1.3.5, 1.4, 1.6.5, 1.9, 1.9.1, 5.5, 7.4, 7.4.1, 7.4.2, and 7.4.3.

This chapter provides details on the storage infrastructure, both physical and virtual, involved in a vSphere 7.0 environment.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 2-1](#) 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 and Review Questions.](#)”

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

<b>Foundation Topics Section</b>	<b>Questions Covered in This Section</b>
----------------------------------	--



Storage Models and Datastore Types	1, 2
vSAN Concepts	3, 4
vSphere Storage Integration	5, 6
Storage Multipathing and Failover	7
Storage Policies	8, 9
Storage DRS (SDRS)	10

1. You need to configure a virtual machine to utilize N-Port ID Virtualization (NPIV). Which of the following are required? (Choose two.)
  - a. iSCSI
  - b. vVOLs
  - c. RDM
  - d. FCoE
  - e. vSAN
  - f. VMFS
2. You are preparing to implement vSphere with Kubernetes. Which type of virtual disk must you provide for storing logs, emptyDir volumes, and ConfigMaps?
  - a. Ephemeral
  - b. Container image
  - c. Persistent volume
  - d. Non-persistent volume
3. You are planning to implement a vSAN stretched cluster. Which of the following statements is true?
  - a. You should not enable DRS in automatic mode.
  - b. You should disable HA datastore heartbeats.
  - c. If you set PFFT to 0, you may be able to use SMP-FT.
  - d. If one of the fault domains is inaccessible, you cannot provision virtual machines.

4. You are planning to implement RAID 6 erasure coding for a virtual disk stored in a vSAN datastore. What percentage of the required capacity will be usable?
  - a. 50%
  - b. 67%
  - c. 75%
  - d. 100%
  
5. You are preparing to leverage VAAI in your vSphere environment. Which of the following primitives will not be available for your virtual machines stored in NFS datastores?
  - a. Atomic Test and Set
  - b. Full File Clone
  - c. Extended Statistics
  - d. Reserve Space
  
6. You are planning to implement vVols. Which of the following are logical I/O proxies?
  - a. Data-vVol instances
  - b. Storage providers
  - c. Storage containers
  - d. Protocol endpoints
  
7. You are explaining how vSphere interacts with storage systems. Which of the following steps may occur when VMware NMP receives an I/O request?
  - a. The PSP issues the I/O request on the appropriate physical path.
  - b. The SATP issues the I/O request on the appropriate physical path.
  - c. The PSP activates the inactive path.
  - d. The PSP calls the appropriate SATP.
  
8. In your vSphere environment where VASA is not implemented, you are planning to leverage storage policies associated with devices in your storage array. Which type of storage policies should you create?

- a. VM storage policy for host-based data services
  - b. VM storage policy for vVols
  - c. VM storage policy for tag-based placement
  - d. vSAN storage policy
9. You are configuring storage policies for use with your vSAN cluster. Which of the following is not an available option?
- a. Number of Replicas per Object
  - b. Number of Disk Stripes per Object
  - c. Primary Level of Failures to Tolerate
  - d. Secondary Level of Failures to Tolerate
10. You are testing Storage DRS (SDRS) in a scenario where the utilized space on one datastore is 82%, and it is 79% on another. You observe that SDRS does not make a migration recommendation. What might be the reason?
- a. The Space Utilization Difference threshold is set too low.
  - b. The Space Utilization Difference threshold is set too high.
  - c. The Space Utilization Difference threshold is set to 78%.
  - d. The Space Utilization Difference threshold is set to 80%.

## Foundation Topics

### Storage Models and Datastore Types

This section explains how virtual machines access storage and describes the storage models and datastore types available in vSphere.

#### How Virtual Machines Access Storage

A virtual machine communicates with its virtual disk stored on a datastore by issuing SCSI commands. The SCSI commands are encapsulated into other forms, depending on the protocol that the ESXi host uses to connect to a storage device on which the datastore resides, as illustrated in [Figure 2-1](#).

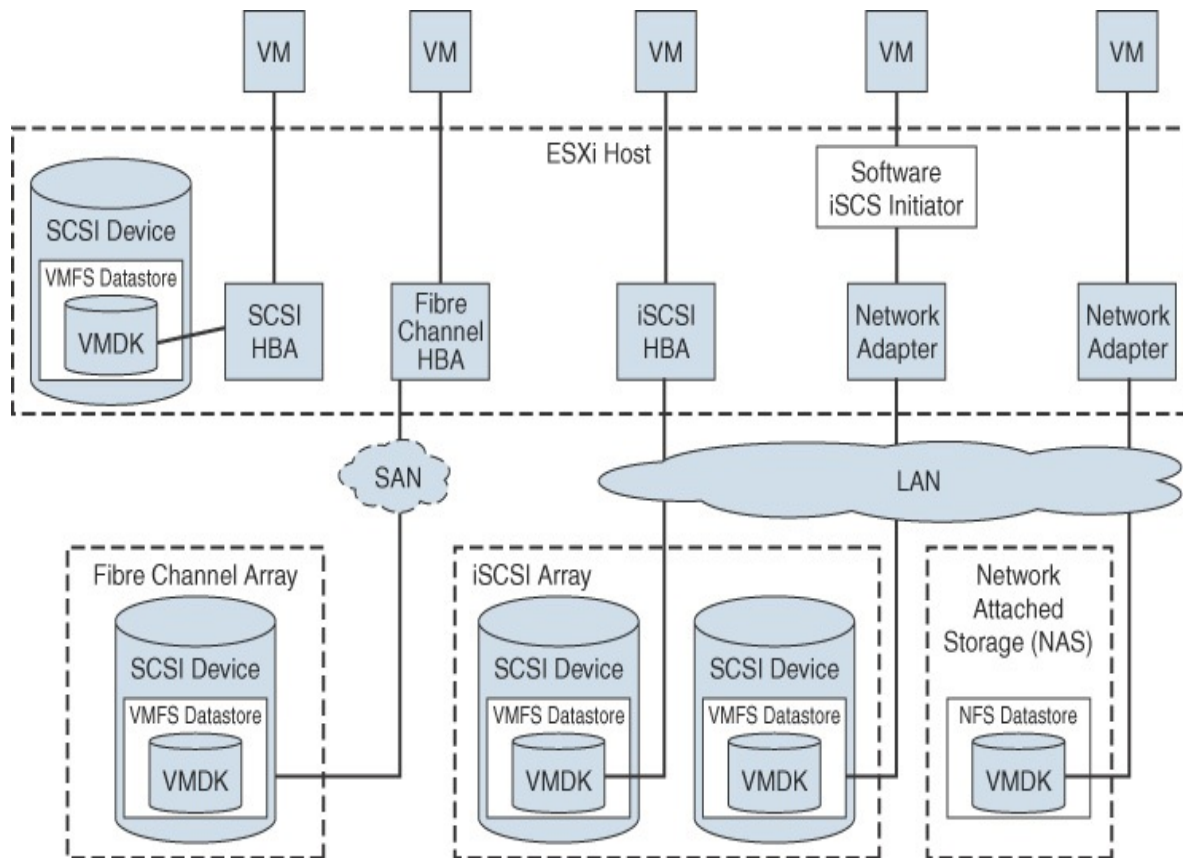


FIGURE 2-1 Virtual Machine Storage

## Storage Virtualization: The Traditional Model

Storage virtualization refers to a logical abstraction of physical storage resources and capacities from virtual machines. ESXi provides host-level storage virtualization. In vSphere environments, a traditional model is built around the storage technologies and ESXi virtualization features discussed in the following sections.

### Storage Devices (or LUNs)

In common ESXi vocabulary, the terms *storage device* and *LUN* are used interchangeably. Storage devices or LUNs are storage volumes that are presented to a host from a block storage system and are available to ESXi for formatting.

### Virtual Disk

Virtual disks are sets of files that reside on a datastore that is deployed on

physical storage. From the standpoint of the virtual machine, each virtual disk appears as if it were a SCSI drive connected to a SCSI controller. The physical storage is transparent to the virtual machine guest operating system and applications.

## **Local Storage**

Local storage can be internal hard disks located inside an ESXi host and external storage systems connected to the host directly through protocols such as SAS or SATA. Local storage does not require a storage network to communicate with the host.

## **Fibre Channel**

Fibre Channel (FC) is a storage protocol that a storage area network (SAN) uses to transfer data traffic from ESXi host servers to shared storage. It packages SCSI commands into FC frames. The ESXi host uses Fibre Channel host bus adapters (HBAs) to connect to the FC SAN, as illustrated in [Figure 2-1](#). Unless you use directly connected Fibre Channel storage, you need Fibre Channel switches to route storage traffic. If a host contains FCoE (Fibre Channel over Ethernet) adapters, you can connect to shared Fibre Channel devices by using an Ethernet network.

## **iSCSI**

Internet SCSI (iSCSI) is a SAN transport that can use Ethernet connections between ESXi hosts and storage systems. To connect to the storage systems, your hosts use hardware iSCSI adapters or software iSCSI initiators with standard network adapters.

With hardware iSCSI HBAs, the host connects to the storage through a hardware adapter that offloads the iSCSI and network processing. Hardware iSCSI adapters can be dependent and independent. With software iSCSI adapters, the host uses a software-based iSCSI initiator in the VMkernel and a standard network adapter to connect to storage. Both the iSCSI HBA and the software iSCSI initiator are illustrated in [Figure 2-1](#).

## **FCoE**

If an ESXi host contains FCoE adapters, it can connect to shared Fibre Channel devices by using an Ethernet network.

## NAS/NFS

vSphere uses NFS to store virtual machine files on remote file servers accessed over a standard TCP/IP network. ESXi 7.0 uses Network File System (NFS) Version 3 and Version 4.1 to communicate with NAS/NFS servers, as illustrated in [Figure 2-1](#). You can use NFS datastores to store and manage virtual machines in the same way that you use the VMFS datastores.

## VMFS

The datastores that you deploy on block storage devices use the native vSphere Virtual Machine File System (VMFS) format. VMFS is a special high-performance file system format that is optimized for storing virtual machines.

## Raw Device Mappings (RDMs)

### Key Topic

A **raw device mapping (RDM)** is a mapping file that contains metadata that resides in a VMFS datastore and acts as a proxy for a physical storage device (LUN), allowing a virtual machine to access the storage device directly. It gives you some of the advantages of direct access to a physical device as well as some of the management advantages of VMFS-based virtual disks. The components involved with an RDM are illustrated in [Figure 2-2](#).




FIGURE 2-2 RDM Diagram

You can envision an RDM as a symbolic link from a VMFS volume to a storage device. The mapping makes the storage device appear as a file in a VMFS volume. The virtual machine configuration references the RDM, not the storage device. RDMs support two compatibility modes:

- **Virtual compatibility mode:** The RDM acts much like a virtual disk file, enabling extra virtual disk features, such as the use of virtual machine snapshot and the use of disk modes (dependent, independent—persistent, and independent—nonpersistent).
- **Physical compatibility mode:** The RDM offers direct access to the SCSI device, supporting applications that require lower-level control.

Virtual disk files are preferred over RDMs for manageability. You should only use RDMs when necessary. Use cases for RDMs include the following.

- You plan to install in a virtual machine software that requires features inherent to the SAN, such as SAN management, storage-based snapshots, or storage-based replication. The RDM enables the virtual machine to have the required access to the storage device.
- You plan to configure a Microsoft Cluster Server (MSCS) clustering in a manner that spans physical hosts, such as virtual-to-virtual clusters and physical-to-virtual clusters. You should configure the data and quorum disks as RDMs rather than as virtual disk files.

Benefits of RDMs include the following:

- **User-friendly persistent names:** Much as with naming a VMFS datastore, you can provide a friendly name to a mapped device rather than use its device name.
- **Dynamic name resolution:** If physical changes (such as adapter hardware changes, path changes, or device relocation) occur, the RDM is updated automatically. The virtual machines do not need to be updated because they reference the RDM.
- **Distributed file locking:** VMFS distributed locking is used to make it safe for two virtual machines on different servers to access the same LUN.
- **File permissions:** Permissions are set on the mapping file to effectively apply permissions to the mapped file, much as they are applied to virtual

disks.

- **File system operations:** Most file system operations that are valid for an ordinary file can be applied to the mapping file and redirected to the mapped device.
- **Snapshots:** Virtual machine snapshots can be applied to the mapped volume except when the RDM is used in physical compatibility mode.
- **vMotion:** You can migrate the virtual machine with vMotion, as vCenter Server uses the RDM as a proxy, which enables the use of the same migration mechanism used for virtual disk files.
- **SAN management agents:** RDM enables the use of SAN management agents (SCSI target-based software) inside a virtual machine. This requires hardware-specific SCSI commands as well as physical compatibility mode for the RDM.
- **N-Port ID Virtualization (NPIV):** You can use NPIV technology, which allows a single Fibre Channel HBA port to register with the fabric by using multiple worldwide port names (WWPNs). This ability makes the HBA port appear as multiple virtual ports, each with its own ID and virtual port name. Virtual machines can claim each of these virtual ports and use them for all RDM traffic. NPIV requires the use of virtual machines with RDMs.

### Note

To support vMotion involving RDMs, be sure to maintain consistent LUN IDs for RDMs across all participating ESXi hosts.

### Note

To support vMotion for NPIV-enabled virtual machines, place the RDM files, virtual machine configuration file, and other virtual machines in the same datastore. You cannot perform Storage vMotion when NPIV is enabled.

## Software-Defined Storage Models



In addition to abstracting underlying storage capacities from VMs, as traditional storage models do, software-defined storage abstracts storage capabilities. With software-defined storage, a virtual machine becomes a unit of storage provisioning and can be managed through a flexible policy-based mechanism. Software-defined storage involves the vSphere technologies described in the following sections.

## **vSAN**

vSAN is a layer of distributed software that runs natively on each hypervisor in a cluster. It aggregates local or direct-attached capacity, creating a single storage pool shared across all hosts in the vSAN cluster.

## **vVols**

*Virtual volumes* are encapsulations of virtual machine files, virtual disks, and their derivatives that are stored natively inside a storage system. You do not provision virtual volumes directly. Instead, they are automatically created when you create, clone, or snapshot a virtual machine. Each virtual machine can be associated to one or more virtual volumes.

The *Virtual Volumes (vVols)* functionality changes the storage management paradigm from managing space inside datastores to managing abstract storage objects handled by storage arrays. With vVols, each virtual machine (rather than a datastore) is a unit of storage management. You can apply storage policies per virtual machine rather than per LUN or datastore.

## **Storage Policy Based Management**

Storage Policy Based Management (SPBM) is a framework that provides a single control panel across various data services and storage solutions, including vSAN and vVols. Using storage policies, the framework aligns application demands of virtual machines with capabilities provided by storage entities.

## **I/O Filters**

*I/O filters* are software components that can be installed on ESXi hosts and can offer additional data services to virtual machines. Depending on the implementation, the services might include replication, encryption, caching, and so on.

## Datastore Types

In vSphere 7.0, you can use the datastore types described in the following sections.

### VMFS Datastore

You can create VMFS datastores on Fibre Channel, iSCSI, FCoE, and local storage devices. ESXi 7.0 supports VMFS Versions 5 and 6 for reading and writing. ESXi 7.0 does not support VMFS Version 3. [Table 2-2](#) compares the features and functionalities of VMFS Versions 5 and 6.

**Table 2-2** Comparison of VMFS Version 5 and Version 6

<b>VMFS Features and Functionalities</b>	<b>Version 5</b>	<b>Version 6</b>
Access for ESXi hosts Version 6.5 and later	Yes	Yes
Access for ESXi hosts Version 6.0 and earlier	Yes	No
Datastores per host	512	512
512n storage devices	Yes	Yes (default)
512e storage devices	Yes (Not supported on local 512e devices.)	Yes (default)
4Kn storage devices	No	Yes
Automatic space reclamation	No	Yes
Manual space reclamation through the <b>esxcli</b> command.	Yes	Yes
Space reclamation from guest OS	Limited	Yes
GPT storage device partitioning	Yes	Yes
MBR storage device partitioning	Yes	No

For a VMFS5 datastore that has been previously upgraded from VMFS3.

Storage devices greater than 2 TB for each VMFS extent	Yes	Yes
Support for virtual machines with large-capacity virtual disks, or disks greater than 2 TB	Yes	Yes
Support of small files (1 KB)	Yes	Yes
Default use of ATS-only locking mechanisms on storage devices that support ATS	Yes	Yes
Block size	Standard 1 MB	Standard 1 MB
Default snapshots	VMFSsparse for virtual disks smaller than 2 TB  SEsparse for virtual disks larger than 2 TB	SEsparse
Virtual disk emulation type	512n	512n
vMotion	Yes	Yes
Storage vMotion across different datastore types	Yes	Yes
High Availability and Fault Tolerance	Yes	Yes
DRS and Storage DRS	Yes	Yes
RDM	Yes	Yes

When working with VMFS datastores in vSphere 7.0, consider the following:

- Datastore extents:** A spanned VMFS datastore must use only homogeneous storage devices—either 512n, 512e, or 4Kn. The spanned datastore cannot

extend over devices of different formats.

- **Block size:** The block size on a VMFS datastore defines the maximum file size and the amount of space a file occupies. VMFS Version 5 and Version 6 data-stores support a 1 MB block size.
- **Storage vMotion:** Storage vMotion supports migration across VMFS, vSAN, and vVols datastores. vCenter Server performs compatibility checks to validate Storage vMotion across different types of datastores.
- **Storage DRS:** VMFS Version 5 and Version 6 can coexist in the same datastore cluster. However, all datastores in the cluster must use homogeneous storage devices. Do not mix devices of different formats within the same datastore cluster.
- **Device Partition Formats:** Any new VMFS Version 5 or Version 6 datastore uses the GUID Partition Table (GPT) to format the storage device, which means you can create datastores larger than 2 TB. If your VMFS Version 5 datastore has been previously upgraded from VMFS Version 3, it continues to use the Master Boot Record (MBR) partition format, which is characteristic for VMFS Version 3. Conversion to GPT happens only after you expand the datastore to a size larger than 2 TB.

## NFS

You can create NFS datastores on NAS devices. ESXi 7.0 supports NFS Versions 3 and 4.1. To support both versions, ESXi 7.0 uses two different NFS clients. [Table 2-3](#) compares the capabilities of NFS Versions 3 and 4.1.

**Table 2-3** Comparison of NFS Version 3 and Version 4.1 Characteristics

NFS Characteristics	Version 3	Version 4.1
Security mechanisms	AUTH_SYS	AUTH_SYS and Kerberos (krb5 and krb5i)
Encryption algorithms with Kerberos	N/A	AES256-CTS-HMAC-SHA1-96  AES128-CTS-HMAC-SHA1-96

Multipathing	Not supported	Supported through the session trunking
Locking mechanisms	Propriety client-side locking	Server-side locking
Hardware acceleration	Supported	Supported
Thick virtual disks	Supported	Supported
IPv6	Supported	Supported for AUTH_SYS and Kerberos
ISO images presented as CD-ROMs to virtual machines	Supported	Supported
Virtual machine snapshots	Supported	Supported
Virtual machines with virtual disks greater than 2 TB	Supported	Supported

[Table 2-4](#) compares vSphere 7.0 features and related solutions supported by NFS Versions 3 and 4.1.

**Table 2-4** Comparison of NFS Version 3 and Version 4.1 Support for vSphere Features and Solutions

<b>NFS Features and Functionalities</b>	<b>Ver 3</b>	<b>Version 4.1</b>
vMotion and Storage vMotion	Y e s	Yes
High Availability (HA)	Y e s	Yes
Fault Tolerance (FT)	Y e s	Yes (Supports the new FT mechanism introduced in vSphere 6.0 that supports up to four vCPUs, not the legacy FT mechanism.)
Distributed	Y	Yes

Resource Scheduler (DRS)	e s	
Host Profiles	Y e s	Yes
Storage DRS	Y e s	No
Storage I/O Control	Y e s	No
Site Recovery Manager	Y e s	No
Virtual Volumes	Y e s	Yes
vSphere Replication	Y e s	Yes
vRealize Operations Manager	Y e s	Yes

When you upgrade ESXi from a version earlier than 6.5, existing NFS Version 4.1 datastores automatically begin supporting functionalities that were not available in the previous ESXi release, such as vVols and hardware acceleration. ESXi does not support automatic datastore conversions from NFS Version 3 to NFS Version 4.1. You can use Storage vMotion to migrate virtual machines from NFS Version 3 datastores to NFS Version 4.1 datastores. In some cases, storage vendors provide conversion methods from NFS Version 3 to Version 4.1. In some cases, you may be able to unmount an NFS Version 3 datastore from all hosts and remount it as NFS Version 4.1. The datastore can never be mounted using both protocols at the same time.

## **vVols Datastores**

You can create a vVols datastore in an environment with a compliant storage system. A virtual volume, which is created and manipulated out of band by a vSphere APIs for Storage Awareness (VASA) provider, represents a storage container in vSphere. The VASA provider maps virtual disk objects and their derivatives, such as clones, snapshots, and replicas, directly to the virtual volumes on the storage system. ESXi hosts access virtual volumes through an intermediate point in the data path called the protocol endpoint. Protocol endpoints serve as gateways for I/O between ESXi hosts and the storage system, using Fibre Channel, FCoE, iSCSI, or NFS.

## **vSAN Datastores**

You can create a vSAN datastore in a vSAN cluster. vSAN is a hyperconverged storage solution, which combines storage, compute, and virtualization into a single physical server or cluster. The following section describes the concepts, benefits, and terminology associated with vSAN.

## **Storage in vSphere with Kubernetes**

To support the different types of storage objects in Kubernetes, vSphere with Kubernetes provides three types of virtual disks: ephemeral, container image, and persistent volume.

A vSphere pod requires ephemeral storage to store Kubernetes objects, such as logs, emptyDir volumes, and ConfigMaps. The ephemeral, or transient, storage exists if the vSphere pod exists.

The vSphere pod mounts images used by its containers as image virtual disks, enabling the container to use the software contained in the images. When the vSphere pod life cycle completes, the image virtual disks are detached from the vSphere pod. You can specify a datastore to use as the container image cache, such that subsequent pods can pull it from the cache rather than from the external container registry.

Some Kubernetes workloads require persistent storage to store the data independently of the pod. Persistent volume objects in vSphere with Kubernetes are backed by the First Class Disks on a datastore. A First Class Disk (FCD), which is also called an Improved Virtual Disk, is a named virtual disk that is not associated with a VM. To provide persistent storage, you can use the Workload

Management feature in the vSphere Client to associate one or more storage policies with the appropriate namespace.

## **VMware NVMe**

NVMe storage is a low-latency, low-CPU-usage, and high-performance alternative to SCSI storage. It is designed for use with faster storage media equipped with non-volatile memory, such as flash devices. NVMe storage can be directly attached to a host using a PCIe interface or indirectly through different fabric transport (NVMe-oF).

In a NVMe storage array, a namespace represents a storage volume. An NVMe namespace is analogous to a storage device (LUN) in other storage arrays. In the vSphere Client, an NVMe namespace appears in the list of storage devices. You can use a device to create a VMFS datastore.

## **Requirements for NVMe over PCIe**

NVMe over PCIe requires the following:

- Local NVMe storage devices
- Compatible ESXi host
- Hardware NVMe over PCIe adapter

## **Requirements for NVMe over RDMA (RoCE Version 2)**

NVMe over RDMA requires the following:

- NVMe storage array with NVMe over RDMA (RoCE Version 2) transport support
- Compatible ESXi host
- Ethernet switches supporting a lossless network
- Network adapter that supports RDMA over Converged Ethernet (RoCE Version 2)
- Software NVMe over RDMA adapter
- NVMe controller



## Requirements for NVMe over Fibre Channel

NVMe over Fibre Channel requires the following:

- Fibre Channel storage array that supports NVMe
- Compatible ESXi host
- Hardware NVMe adapter (that is, a Fibre Channel HBA that supports NVMe)
- NVMe controller

## VMware High-Performance Plug-in (HPP)

VMware provides the High-Performance Plug-in (HPP) to improve the performance of NVMe devices on an ESXi host. HPP replaces NMP for high-speed devices, such as NVMe.

HPP is the default plug-in that claims NVMe-oF targets. In ESXi, the NVMe-oF targets are emulated and presented to users as SCSI targets. The HPP supports only active/active and implicit ALUA targets.

NMP is the default plug-in for local NVMe devices, but you can replace it with HPP. NMP cannot be used to claim the NVMe-oF targets. HPP should be used for NVMe-oF.

[Table 2-5](#) describes vSphere 7.0 support for HPP.

**Table 2-5** vSphere 7.0 HPP Support

<b>HPP Support</b>	<b>vSphere 7.0</b>
Storage devices	Local NVMe PCIe  Shared NVMe-oF (active/active and implicit ALUA targets only)
Multipathing	Yes
Second-level plug-ins	No
SCSI-3 persistent reservations	No
4Kn devices with software	No

emulation

---

vSAN No

---

[Table 2-6](#) describes the path selection schemes (PSS) HPP uses when selecting physical paths for I/O requests.

**Table 2-6** HPP Path Selection Schemes

<b>PSS</b>	<b>Description</b>
FIXED	A designated preferred path is used for I/O requests.
LB-RR (Load Balance—Round Robin)	After transferring a specified number of bytes or I/Os on a current path, the scheme selects the path using the round robin algorithm. You can configure the IOPS and Bytes properties to indicate the criteria for path switching. (This is the default HPP scheme.)
LB-IOPS (Load Balance—IOPS)	After transferring a specified number of I/Os on a current path (1000 by default), the scheme selects an optimal path based on the least number of outstanding bytes.
LB-BYTES (Load Balance—Bytes)	After transferring a specified amount of data on a current path (10 MB by default), the scheme selects an optimal path based on the least number of outstanding I/Os.

---

LB- Latency The scheme selects an optimal path by considering the latency evaluation time and the sampling I/Os per path.

(Load  
Balancing—  
Latency)

---

HPP best practices include the following:

- Use a vSphere version that supports HPP.
- Use HPP for NVMe local and networked devices.
- Do not use HPP with HDDs or any flash devices that cannot sustain 200,000 IOPS.
- If you use NVMe with Fibre Channel devices, follow your vendor's recommendations.
- If you use NVMe-oF, do not mix transport types to access the same namespace.
- When using NVMe-oF namespaces, make sure that active paths are presented to the host.
- Configure VMs to use VMware Paravirtual controllers.
- Set the latency sensitive threshold for virtual machines.
- If a single VM drives a significant share of a device's I/O workload, consider spreading the I/O across multiple virtual disks, attached to separate virtual controllers in the VM. Otherwise, you risk the I/O saturating a CPU core.

## vSAN Concepts

vSAN virtualizes the local physical storage resources of ESXi hosts by turning them into pools of storage that can be used by virtual machines, based on their quality of service requirements. You can configure vSAN as a hybrid or an all-flash cluster. Hybrid clusters use flash devices for the cache layer and magnetic

disks for the storage capacity layer. All-flash clusters use flash devices for both cache and capacity.

You can enable vSAN on existing host clusters as well as on new clusters. You can expand a datastore by adding to the cluster hosts with capacity devices or by adding local drives to the existing hosts in the cluster. vSAN works best when all ESXi hosts in the cluster are configured similarly, including similar or identical storage configurations. A consistent configuration enables vSAN to balance virtual machine storage components across all devices and hosts in the cluster. Hosts without any local devices can also participate and run their virtual machines on the vSAN datastore.

If a host contributes some of its local storage to a vSAN cluster, then it must contribute at least one device for cache. The drives contributed by a host form one or more disk groups. Each disk group contains a flash cache device and at least one capacity device. Each host can be configured to use multiple disk groups.

vSAN offers many features of a traditional SAN. Its main limitations are that each vSAN instance can support only one cluster. vSAN has the following benefits over traditional storage:

- vSAN does not require a dedicated storage network, such as on a FC network or SAN.
- With vSAN, you do not have to pre-allocate and preconfigure storage volumes (LUNs).
- vSAN does not behave like traditional storage volumes based on LUNs or NFS shares. You do not have to apply standard storage protocols, such as FC, and you do not need to format the storage directly.
- You can deploy, manage, and monitor vSAN by using the vSphere Client rather than other storage management tools.
- A vSphere administrator, rather than a storage administrator, can manage a vSAN environment.
- When deploying virtual machines, you can use automatically assigned storage policies with vSAN.

## **vSAN Characteristics**

vSAN is like network-distributed RAID for local disks, transforming them into

shared storage. vSAN uses copies of VM data, where one copy is local and another copy is on one of the other nodes in the cluster. The number of copies is configurable. Here are some of the features of vSAN:

- **Shared storage support:** VMware features that require shared storage (that is, HA, vMotion, DRS) are available with vSAN.
- **On-disk format:** Highly scalable snapshot and clone management are possible on a vSAN cluster.
- **All-flash and hybrid configurations:** vSAN can be used on hosts with all-flash storage or with hybrid storage (that is, a combination of SSDs and traditional HDDs).
- **Fault domains:** Fault domains can be configured to protect against rack or chassis failures, preventing all copies of VM disk data from residing on the same rack or chassis.
- **iSCSI target service:** The vSAN datastore can be visible to and usable by ESXi hosts outside the cluster and by physical bare-metal systems.
- **Stretched cluster:** vSAN supports stretching a cluster across physical geographic locations.
- **Support for Windows Server failover clusters (WSFCs):** SCSI-3 Persistent Reservations (SCSI3-PR) is supported on virtual disks, which are required for shared disks and WSFCs. Microsoft SQL 2012 or later is supported on vSAN. The following limitations apply:
  - Maximum of 6 application nodes in each vSAN cluster
  - Maximum of 64 shared disks per ESXi host
- **vSAN health service:** This service includes health checks for monitoring and troubleshooting purposes.
- **vSAN performance service:** This service includes statistics for monitoring vSAN performance metrics. This can occur at the level of the cluster, ESXi host, disk group, disk, or VM.
- **Integration with vSphere storage features:** Snapshots, linked clones, and vSphere Replication are all supported on vSAN datastores.
- **Virtual machine storage policies:** Policies can be defined for VMs on vSAN. If no policies are defined, a default vSAN policy is applied.

- **Rapid provisioning:** vSAN enables fast storage provisioning for VM creation and deployment from templates.
- **Deduplication and compression:** Block-level deduplication and compression are available space-saving mechanisms on vSAN, and they can be configured at the cluster level and applied to each disk group.
- **Data at rest encryption:** Data at rest encryption is encryption of data that is not in transit and on which no processes are being done (for example, deduplication or compression). If drives are removed, the data on those drives is encrypted.
- **SDK support:** vSAN supports an extension (written in Java) of the VMware vSphere Management SDK. It has libraries, code examples, and documentation for assistance in automating and troubleshooting vSAN deployments.

## vSAN Terminology

Be sure to know the following terminology for the Professional VMware vSphere 7.x (2V0-21.20) exam:

- **Disk group:** A *disk group* is a group of local disks on an ESXi host contributing to the vSAN datastore. It must include one cache device and one capacity device. In a hybrid cluster, a flash disk is the cache device, and magnetic disks are used for capacity devices. In all-flash clusters, flash storage is used for both cache and capacity.
- **Consumed capacity:** This is the amount of physical space used up by virtual machines at any point in time.
- **Object-based storage:** Data is stored in vSAN by way of objects, which are flexible data containers. Objects are logical volumes with data and metadata spread among nodes in the cluster. Virtual disks are objects, as are snapshots. For object creation and placement, vSAN takes the following into account:
  - Virtual disk policy and requirements are verified.
  - The number of copies (replicas) is verified; the amount of flash read cache allocated for replicas, number of stripes for replica, and location are determined.
  - Policy compliance of virtual disks is ensured.

- Mirrors and witnesses are placed on different hosts or fault domains.
- **vSAN datastores:** Like other datastores, a vSAN datastore appears in the Storage Inventory view in vSphere. A vSAN cluster provides a single datastore for all the hosts in the cluster, even for hosts that do not contribute storage to vSAN. An ESXi host can mount VMFS and NFS datastores in addition to the vSAN datastore. Storage vMotion can be used to migrate VMs between datastore types.
- **Objects and components:** vSAN includes the following objects and components:
  - **VM home namespace:** The VM home directory where all the VM files are stored
  - **VMDK:** Virtual disks for VMs
  - **VM swap object:** An object that allows memory to be swapped to disk during periods of contention and that is created at VM power on
  - **Snapshot delta VMDKs:** Change files created when a snapshot is taken of a VM
  - **Memory object:** An object created when a VM is snapshotted (and the VM's memory is retained) or suspended
- **Virtual machine compliance status:** Can be Compliant and Noncompliant, depending on whether each of the virtual machine's objects meets the requirements of the assigned storage policy. The status is available on the Virtual Disks page on the Physical Disk Placement tab.
- **Component state:** vSAN has two component states:
  - **Degraded:** vSAN detects a permanent failure of a component.
  - **Absent:** vSAN detects a temporary component failure.
- **Object state:** vSAN has two object states:
  - **Healthy:** At least one RAID 1 mirror is available, or enough segments are available for RAID 5 or 6.
  - **Unhealthy:** No full mirror is available, or not enough segments are available for RAID 5 or 6.
- **Witness:** This is a component consisting of only metadata. It is used as a tiebreaker. Witnesses consume about 2 MB of space for metadata on a

vSAN datastore when on-disk format Version 1.0 is used and 4 MB when on-disk format Version 2.0 or later is used.

- **Storage Policy Based Management (SPBM):** VM storage requirements are defined as a policy, and vSAN ensures that these policies are met when placing objects. If you do not apply a storage policy when creating or deploying VMs, the default vSAN policy is used, with Primary Level of Failures to Tolerate set to 1, a single stripe per object, and a thin provisioned disk.
- **Ruby vSphere Console (RVC):** This is a command-line interface used for managing and troubleshooting vSAN. RVC provides a cluster-wide view and is included with a vCenter Server deployment.
- **VMware PowerCLI:** vSAN cmdlets are included with PowerCLI to allow administration of vSAN.
- **vSAN Observer:** This is a web-based utility, built on top of RVC, used for performance analysis and monitoring. It can display performance statistics on the capacity tier, disk group statistics, CPU load, memory consumption, and vSAN objects in-memory and their distribution across the cluster.
- **vSAN Ready Node:** This preconfigured deployment is provided by VMware partners. It is a validated design using certified hardware.
- **User-defined vSAN cluster:** This vSAN deployment makes use of hardware selected by you.

## Note

The capacity disks contribute to the advertised datastore capacity. The flash cache devices are not included as capacity.

## What Is New in vSAN 7.0

The following new features are available in vSAN 7.0:

- **vSphere Lifecycle Manager:** vSphere Lifecycle Manager uses a desired-state model to enable simple, consistent lifecycle management for ESXi hosts, including drivers and firmware.
- **Integrated file services:** The native vSAN File Services enables you to



create and present NFS Version 4.1 and Version 3 file shares, effectively extending vSAN capabilities such as availability, security, storage efficiency, and operations management to files.

- **Native support for NVMe Hot-Plug:** The Hot-Plug plug-in provides a consistent way of servicing NVMe devices and provides operational efficiency.
- **I/O redirect based on capacity imbalance with stretched clusters:** This feature improves VM uptime by redirecting all virtual machine I/O from a capacity-strained site to the other site.
- **Skyline integration with vSphere Health and vSAN Health:** Skyline Health for vSphere and vSAN are available, enabling native, in-product health monitoring and consistent, proactive analysis.
- **Removal of EZT for shared disk:** vSAN 7.0 eliminates the prerequisite that shared virtual disks using the multi-writer flag must also use the eager zero thick format.
- **Support for vSAN memory as a metric in performance service:** vSAN memory usage is now available in Performance Charts (vSphere Client) and through the API.
- **Visibility of vSphere Replication objects:** vSphere Replication objects are visible in vSAN capacity view.
- **Support for large-capacity drives:** vSAN provides support for 32 TB physical capacity drives and up to 1 PB logical capacity when deduplication and compression are enabled.
- **Immediate repair after new witness is deployed:** vSAN immediately invokes a repair object operation after a witness has been added during a replace witness operation.
- **vSphere with Kubernetes integration:** Cloud Native Storage (CNS) is the default storage platform for vSphere with Kubernetes. This integration enables various stateful containerized workloads to be deployed on vSphere with Kubernetes Supervisor and Guest clusters on vSAN, VMFS, and NFS datastores.
- **File-based persistent volumes:** Kubernetes developers can dynamically create shared (read/write/many) persistent volumes for applications. Multiple pods can share data. The native vSAN File Services is the

foundation that enables this capability.

- **vVols support for modern applications:** You can deploy modern Kubernetes applications to external storage arrays on vSphere using the CNS support added for vVols. vSphere now enables unified management for persistent volumes across vSAN, NFS, VMFS, and vVols.
- **vSAN VCG notification service:** You can get notified through email about any changes to vSAN HCL components, such as vSAN ReadyNode, I/O controller, and drives (NVMe, SSD, HDD), and get notified through email about any changes.

### Note

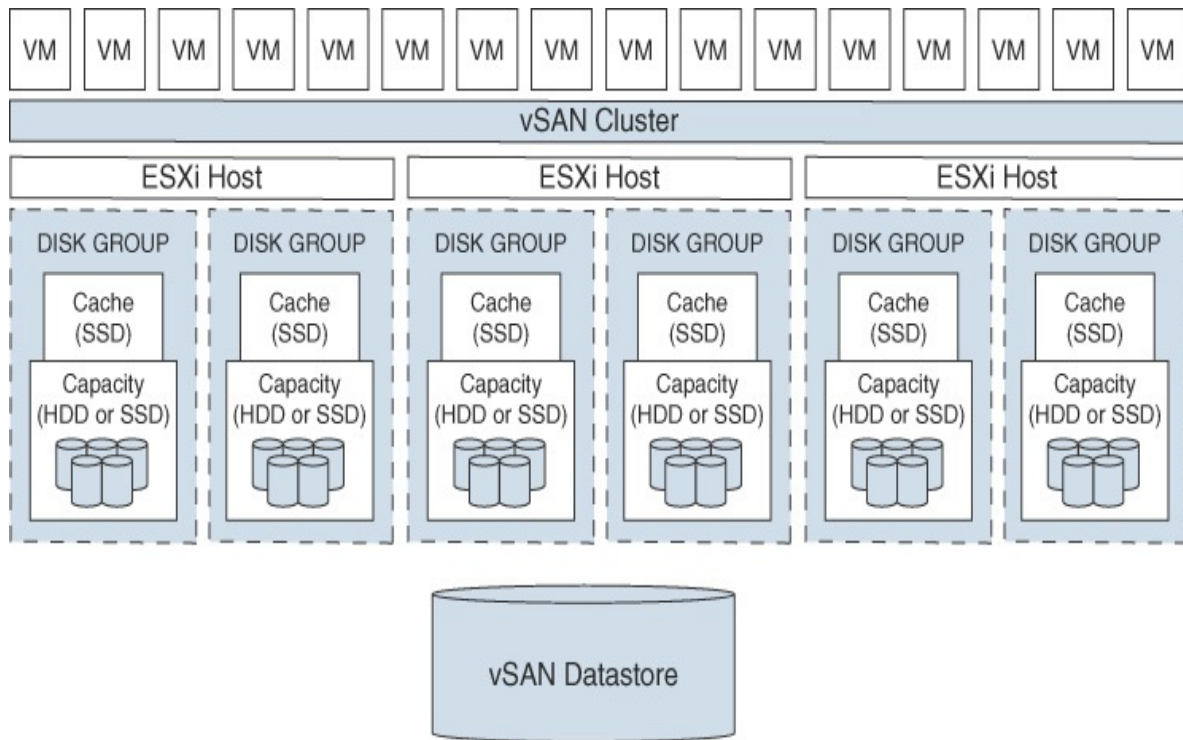
In vCenter Server 7.0.0a, vSAN File Services and vSphere Lifecycle Manager can be enabled simultaneously on the same vSAN cluster.

## vSAN Deployment Options

When deploying vSAN, you have several options for the cluster topology, as described in the following sections.

### Standard Cluster

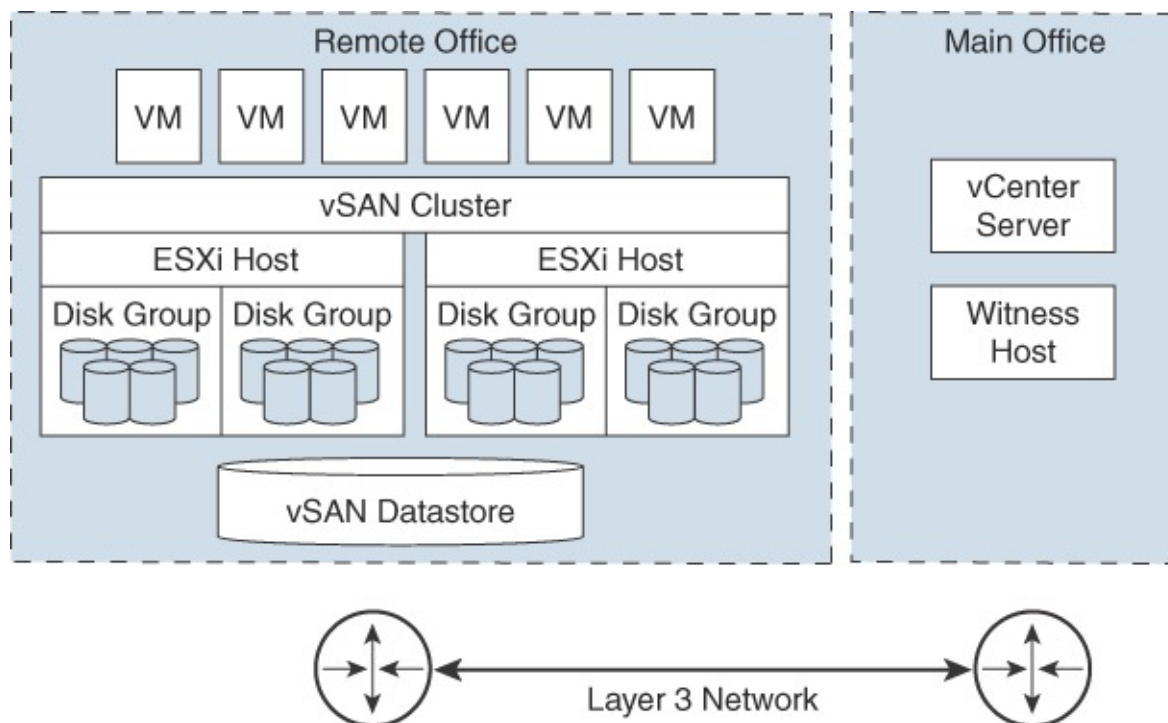
A standard vSAN cluster, as illustrated in [Figure 2-3](#), consists of a minimum of three hosts, typically residing at the same location and connected on the same Layer 2 network. 10 Gbps network connections are required for all-flash clusters and are recommended for hybrid configurations.



**FIGURE 2-3** A Standard vSAN Cluster

## Two-Host vSAN Cluster

The main use case for a two-host vSAN cluster is in a remote office/branch office environment, where workloads require high availability. A two-host vSAN cluster, as illustrated in [Figure 2-4](#), consists of two hosts at the same location, connected to the same network switch or directly connected to one another. You can configure a two-host vSAN cluster that uses a third host as a witness, and the witness can be located separately from the remote office. Usually the witness resides at the main site, along with vCenter Server. For more details on the witness host, see the following section.



**Key Topic**

**FIGURE 2-4** A Two-Node vSAN Cluster

## Stretched Cluster

**Key Topic**

You can create a stretched vSAN cluster that spans two geographic sites and continues to function if a failure or scheduled maintenance occurs at one site. Stretched clusters, which are typically deployed in metropolitan or campus environments with short distances between sites, provide a higher level of availability and inter-site load balancing.

You can use stretched clusters for planned maintenance and disaster avoidance scenarios, with both data sites active. If either site fails, vSAN uses the storage on the other site, and vSphere HA can restart virtual machines on the remaining active site.

You should designate one site as the preferred site; it then becomes the only used site in the event that network connectivity is lost between the two sites. A vSAN stretched cluster can tolerate one link failure at a time without data becoming unavailable. During a site failure or loss of network connection, vSAN

automatically switches to fully functional sites.

### Note

A link failure is a loss of network connection between two sites or between one site and the witness host.

Each stretched cluster consists of two data sites and one *witness host*. The witness host resides at a third site and contains the witness components of virtual machine objects. It contains only metadata and does not participate in storage operations. [Figure 2-5](#) shows an example of a stretched cluster, where the witness node resides at a third site, along with vCenter Server.




FIGURE 2-5 A Stretched vSAN Cluster

The witness host acts as a tiebreaker for decisions regarding availability of datastore components. The witness host typically forms a vSAN cluster with the

preferred site but forms a cluster with secondary site if the preferred site becomes isolated. When the preferred site is online again, data is resynchronized.

A witness host has the following characteristics:

- It can use low-bandwidth/high-latency links.
- It cannot run VMs.
- It can support only one vSAN stretched cluster.
- It requires a VMkernel adapter enabled for vSAN traffic with connections to all hosts in the cluster. It can have only one VMkernel adapter dedicated to vSAN but can have another for management.
- It must be a standalone host. It cannot be added to any other cluster or moved in inventory through vCenter Server.
- It can be a physical ESXi host or a VM-based ESXi host.

## Note

The witness virtual appliance is an ESXi host in a VM, packaged as an OVF or OVA, which is available in different options, depending on the size of the deployment.

Each site in a stretched cluster resides in a separate fault domain. Three default domains are used: the preferred site, the secondary site, and a witness host.

Beginning with vSAN Version 6.6, you can provide an extra level of local fault protection for objects in stretched clusters by using the following policy rules:

- **Primary Level of Failures to Tolerate (PFTT):** This defines the number of site failures that a virtual machine object can tolerate. For a stretched cluster, only a value of 0 or 1 is supported.
- **Secondary Level of Failures to Tolerate (SFTT):** This defines the number of additional host failures that the object can tolerate after the number of site failures (PFTT) is reached. For example, if PFTT = 1 and SFTT = 2, and one site is unavailable, then the cluster can tolerate two additional host failures. The default value is 0, and the maximum value is 3.

- **Data Locality:** This enables you to restrict virtual machine objects to a selected site in the stretched cluster. The default value is None, but you can change it to Preferred or Secondary. Data Locality is available only if PFTT = 0.

## Note

If you set SFTT for a stretched cluster, the Fault Tolerance Method rule applies to the SFTT. The failure tolerance method used for the PFTT is set to RAID 1.

Consider the following guidelines and best practices for stretched clusters:

- DRS must be enabled on a stretched cluster.
- You need to create two host groups, two virtual machines groups, and two VM-Host affinity rules to effectively control the placement of virtual machines between the preferred and secondary sites.
- HA must be enabled on the cluster in a manner such that it respects the VM-Host affinity rules.
- You need to disable HA datastore heartbeats.
- On-disk format Version 2.0 or later is required.
- You need to set Failures to Tolerate to 1.
- Symmetric Multiprocessing Fault Tolerance (SMP-FT) is supported when PFFT is set to 0 and Data Locality is set to Preferred or Secondary. SMP-FT is not supported if PFFT is set to 1.
- Using **esxcli** to add or remove hosts is not supported.
- If one of the three fault domains (preferred site, secondary site, or witness host) is inaccessible, new VMs can still be provisioned, but they are noncompliant until the partitioned site rejoins the cluster. This implicit forced provisioning is performed only when two of the three fault domains are available.
- If an entire site goes offline due to loss of power or network connection, you need to restart the site immediately. Bring all hosts online at approximately the same time to avoid resynchronizing a large amount of data across the

sites.

- If a host is permanently unavailable, you need to remove the host from the cluster before performing any reconfiguration tasks.
- To deploy witnesses for multiple clusters, you should not clone a virtual machine that is already configured as a witness. Instead, you can first deploy a VM from OVF, then clone the VM, and then configure each clone as a witness host for a different cluster.

The stretched cluster network must meet the following requirements:

- The management network requires connectivity across all three sites, using a Layer 2 stretched network or a Layer 3 network.
- The vSAN network requires connectivity across all three sites, using a Layer 2 stretched network between the two data sites and a Layer 3 network between the data sites and the witness host.
- The virtual machine network requires connectivity between the data sites but not the witness host. You can use a Layer 2 stretched network or Layer 3 network between the data sites. Virtual machines do not require a new IP address following failover to the other site.
- The vMotion network requires connectivity between the data sites but not the witness host. You can use a Layer 2 stretched or a Layer 3 network between data sites.

## **vSAN Limitations**



Limitations of vSAN include the following.

- No support for hosts participating in multiple vSAN clusters
- No support for vSphere DPM and storage I/O control
- No support for SE sparse disks
- No support for RDM, VMFS, diagnostic partition, and other device access features

## **vSAN Space Efficiency**



You can use space efficiency techniques in vSAN to reduce the amount of space used for storing data. These techniques include the use of any or all of the following:

- **Thin provisioning:** Consuming only the space on disk that is used (and not the total allocated virtual disk space)
- **Deduplication:** Reducing duplicated data blocks by using SHA-1 hashes for data blocks
- **Compression:** Compressing data using LZ4, which is a lightweight compression mechanism.
- **Erasur coding:** Creating a strip of data blocks with a parity block (This is similar to parity with RAID configurations, except it spans ESXi hosts in the cluster instead of disks in the host.)

## SCSI UNMAP

SCSI **UNMAP** commands, which are supported in vSAN Version 6.7 Update 1 and later, enable you to reclaim storage space that is mapped to deleted vSAN objects. vSAN supports the SCSI **UNMAP** commands issued in a guest operating system to reclaim storage space. vSAN supports offline unmaps as well as inline unmaps. On Linux, offline unmaps are performed with the **fstrim(8)** command, and inline unmaps are performed when the **mount -o discard** command is used. On Windows, NTFS performs inline unmaps by default.

## Deduplication and Compression

All-flash vSAN clusters support deduplication and compression. Deduplication removes redundant data blocks. Compression removes additional redundant data within each data block. Together, these techniques reduce the amount of space required to store data. As vSAN moves data from the cache tier to the capacity tier, it applies deduplication and then applies compression.

You can enable deduplication and compression as a cluster-wide setting, but they are applied on a disk group basis, and redundant data is reduced within each disk group.

When you enable or disable deduplication and compression, vSAN performs a rolling reformat of every disk group on every host, and this may take a long

time. You should verify that enough physical capacity is available to place your data. You should also minimize how frequently these operations are performed.

## Note

Deduplication and compression might not be effective for encrypted VMs.

The amount of storage reduction achieved through deduplication and compression depends on many factors, such as the type of data stored and the number of duplicate blocks. Larger disk groups tend to provide a higher deduplication ratio.

## RAID 5 and RAID 6 Erasure Coding

In a vSAN cluster, you can use RAID 5 or RAID 6 erasure coding to protect against data loss while increasing storage efficiency compared with RAID 1 (mirroring). You can configure RAID 5 on all-flash clusters with four or more fault domains. You can configure RAID 5 or RAID 6 on all-flash clusters with six or more fault domains.

RAID 5 or RAID 6 erasure coding requires less storage space to protect your data than RAID 1 mirroring. For example, if you protect a VM by setting Primary Level of Failures to Tolerate (PFTT) to 1, RAID 1 requires twice the virtual disk size, and RAID 5 requires 1.33 times the virtual disk size. [Table 2-7](#) compares RAID 1 with RAID 5/6 for a 100 GB virtual disk.

**Table 2-7** RAID Configuration Comparison

RAID Configuration	PF TT	Data Size	Required Capacity	Usable Capacity
RAID 1 (mirroring)	1	100 GB	200 GB	50%
RAID 5 or RAID 6 (erasure coding) with four fault domains	1	100 GB	133 GB	75%
RAID 1 (mirroring)	2	100 GB	300 GB	33%
RAID 5 or RAID 6 (erasure coding)	2	100	150 GB	67%

with six fault domains			GB	
RAID 1 (mirroring)	3	100 GB	400 GB	25%
RAID 5 or RAID 6 (erasure coding) with six fault domains	3	N/A	N/A	N/A

Before configuring RAID 5 or RAID 6 erasure coding in a vSAN cluster, you should consider the following:

- All-flash disk groups are required.
- On-disk format Version 3.0 or later is required.
- A valid license supporting RAID 5/6 is required.
- You can enable deduplication and compression on the vSAN cluster to achieve additional space savings.
- PFTT must be set to less than 3.

## vSAN Encryption

You can use data at rest encryption in a vSAN cluster, where all data is encrypted after all other processing, such as deduplication, is performed. All files are encrypted, so all virtual machines and their data are protected. Only administrators with encryption privileges can perform encryption and decryption tasks. Data at rest encryption protects data on storage devices in the event that a device is removed from the cluster.

vSAN encryption requires an external key management server (KMS), the vCenter Server system, and ESXi hosts. vCenter Server requests encryption keys from an external KMS. The KMS generates and stores the keys, and vCenter Server obtains the key IDs from the KMS and distributes them to the ESXi hosts. vCenter Server does not store the KMS keys but keeps a list of key IDs.

vSAN uses encryption keys in the following manner:

- vCenter Server requests an AES-256 key encryption key (KEK) from the KMS.
- vCenter Server stores only the ID of the KEK (not the key itself.)
- The host encrypts disk data using the industry-standard AES-256 XTS

mode.

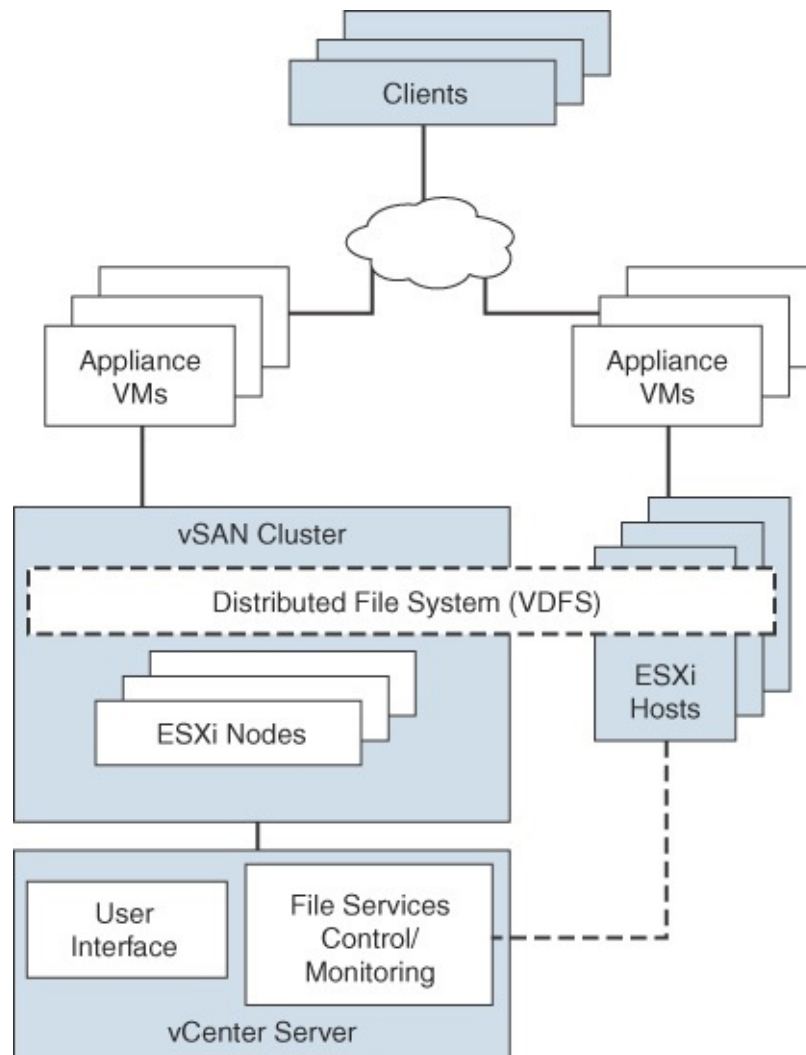
- Each disk has a unique, randomly generated data encryption key (DEK).
- A host key is used to encrypt core dumps, not data. All hosts in the same cluster use the same host key.
- When collecting support bundles, a random key is generated to re-encrypt the core dumps. You can specify a password to encrypt the random key.

## Note

Each ESXi host uses the KEK to encrypt its DEKs and stores the encrypted DEKs on disk. The host does not store the KEK on disk. If a host reboots, it requests the KEK with the corresponding ID from the KMS. The host can then decrypt its DEKs as needed.

## vSAN File Services

You can use **vSAN File Services** to provide vSAN-backed file shares that virtual machines can access as NFS Version 3 and NFS Version 4.1 file shares. It uses vSAN Distributed File System (vDFS), resilient file server endpoints, and a control plane, as illustrated in [Figure 2-6](#). File shares are integrated into the existing vSAN Storage Policy Based Management and on a per-share basis. vSAN File Services creates a single vDFS for the cluster and places a file service virtual machine (FSVM) on each host. The FSVMs manage file shares and act as NFS file servers using IP addresses from a static IP address pool.



**FIGURE 2-6** vSAN File Services Architecture

The vSAN File Services is not supported on a vSAN stretched cluster.

## **vSAN Requirements**

Prior to deploying a vSAN cluster, you should address the requirements outlined in the following sections.

## **vSAN Planning and Sizing**

When you plan the capacity for a vSAN datastore, you should consider the PFTT and the failure tolerance method, as illustrated previously in [Table 2-7](#). For RAID 1, the required data store capacity can be calculated using the following

formula:

$$\text{Capacity} = \text{Expected Consumption Size} \times (\text{PFTT} + 1)$$

For example, say that you plan to use RAID 1 for a 500 GB virtual disk that you expect to be completely filled. In this case, the required capacities are 1000 GB, 1500 GB, and 2000 GB for PFTT set to 1, 2, and 3, respectively.

Keep in mind the following guidelines for vSAN capacity sizing:

- Plan for some extra overhead, as it may be required, depending on the on-disk format version. Version 1.0 adds approximately 1 GB overhead per capacity device. Versions 2.0 and 3.0 adds up to 2% overhead per capacity device. Version 3.0 adds 6.2% overhead for deduplication and compression checksums.
- Keep at least 30% unused space to avoid vSAN rebalancing.
- Plan for spare capacity to handle potential failure or replacement of capacity devices, disk groups, and hosts.
- Reserve spare capacity to rebuild after a host failure or during maintenance. For example, with PFTT set to 1, at least four hosts should be placed in the cluster because at least three available hosts are required to rebuild components.
- Provide enough spare capacity to make it possible to dynamically change a VM storage policy, which may require vSAN to create a new RAID tree layout for the object and temporarily consume extra space.
- Plan for the space consumed by snapshots, which inherit the storage policy applied to the virtual disk.
- Plan for space consumed by the VM home namespace, which includes the virtual machine's swap file (in vSAN Version 6.7 and later).

When selecting devices to use for vSAN cache hardware (such as PCIe vs. SDD flash devices), in addition to cost, compatibility, performance, and capacity, you should consider write endurance.

When selecting storage controllers for use in a vSAN cluster, in addition to compatibility, you should favor adapters with higher queue depth to facilitate vSAN rebuilding operations. You should configure controllers for passthrough mode rather than RAID 0 mode to simplify configuration and maintenance. You

should disable caching on the controller or set it to 100% read.

When sizing hosts, consider using at least 32 GB memory for full vSAN operations based on five disk groups per host and seven capacity devices per disk group.

## **Fault Domain Planning**

If you span your vSAN cluster across multiple racks or blade server chassis, you can configure fault domains to protect against failure of a rack or chassis. A fault domain consists of one or more vSAN cluster member hosts sharing some physical characteristic, such as being in the same rack or chassis. For example, you can configure a fault domain to enable a vSAN cluster to tolerate the failure of an entire physical rack as well as the failure of a single host or other component (capacity devices, network link, or network switch) associated with the rack. When a virtual machine is configured with Primary Level of Failures to Tolerate set to 1 (PFTT=1), vSAN can tolerate a single failure of any kind and of any component in a fault domain, including the failure of an entire rack.

When you provision a new virtual machine, vSAN ensures that protection objects, such as replicas and witnesses, are placed in different fault domains. If you set a virtual machine's storage policy to  $PFTT = n$ , vSAN requires a minimum of  $2 \times n + 1$  fault domains in the cluster. A minimum of three fault domains are required to support  $PFTT = 1$ .

It is best to configure four or more fault domains in a cluster where  $PFTT=1$  is used. A cluster with three fault domains has the same restrictions as a three-host cluster, such as the inability to again protect data after a failure and the inability to use the Full Data Migration Mode.

Say that you have a vSAN cluster where you plan to place four hosts per rack. In order to tolerate an entire rack failure, you need to create a fault domain for each rack. To support  $PFTT=1$ , you need to use a minimum of 12 hosts deployed to 3 racks. To support Full Data Migration Mode and the ability to again protect after a failure, deploy a minimum of 16 hosts to 4 racks. If you want Primary Level of Failures to Tolerate set to 2, you need to configure 5 fault domains in the cluster.

When working with fault domains, you should consider the following best practices:

- At a minimum, configure three fault domains in a vSAN cluster. For best results, configure four or more fault domains.

- Each host that is not directly added to a fault domain should reside in its own single-host fault domain.
- You can add any number of hosts to a fault domain. Each fault domain must contain at least one host.
- If you use fault domains, consider creating equal-sized fault domains (with the same number of same-sized hosts).

## Hardware Requirements

You should examine the vSAN section of the *VMware Compatibility Guide* to verify that all the storage devices, drivers, and firmware versions are certified for the specific vSAN version you plan to use. [Table 2-8](#) lists some of the vSAN storage device requirements.

**Table 2-8** vSAN Storage Device Requirements

Component	Requirements
Cache	<p>One SAS or SATA SSD or PCIe flash device is required.</p> <p>For a hybrid disk group, the cache device must provide at least 10% of the anticipated storage consumed on the capacity devices in a disk group, excluding replicas.</p> <p>The flash devices used for vSAN cache must be dedicated. They cannot be used for vSphere Flash Cache or for VMFS.</p>
Capacity (virtual machine) storage	<p>For a hybrid disk group, at least one SAS or NL-SAS magnetic disk needs to be available.</p> <p>For an all-flash disk group, at least one SAS or SATA SSD or at least one PCIe flash device needs to be available.</p>
Storage controllers	<p>One SAS or SATA host bus adapter (HBA) or a RAID controller that is in passthrough mode or RAID 0 mode is required.</p> <p>If the same storage controller is backing both vSAN and non-vSAN disks, you should apply the following VMware</p>



recommendations to avoid issues:

- Do not mix the controller mode for vSAN and non-vSAN disks. If the vSAN disks are in RAID mode, the non-vSAN disks should also be in RAID mode.
- If VMFS is used on the non-vSAN disks, then use the VMFS datastore only for scratch, logging, and core dumps.
- Do not run virtual machines from a disk or RAID group that shares its controller with vSAN disks or RAID groups.
- Do not pass through non-vSAN disks to virtual machine guests as RDMs.

---

The memory requirements for vSAN depend on the number of disk groups and devices that the ESXi hypervisor must manage. According to VMware Knowledge Base (KB) article 2113954, the following formula can be used to calculate vSAN memory consumption:

$$\text{vSANFootprint} = \text{HOST\_FOOTPRINT} + \text{NumDiskGroups} \times \text{DiskGroupFootprint}$$

where:

$$\begin{aligned} \text{DiskGroupFootprint} = & \text{DISKGROUP\_FIXED\_FOOTPRINT} + \\ & \text{DISKGROUP\_SCALABLE\_FOOTPRINT} + \text{CacheSize} \times \\ & \text{CACHE\_DISK\_FOOTPRINT} + \text{NumCapacityDisks} \times \\ & \text{CAPACITY\_DISK\_FOOTPRINT} \end{aligned}$$

The ESXi Installer creates a coredump partition on the boot device, whose default size is typically adequate. If ESXi host memory is 512 GB or less, you

can boot the host from a USB, SD, or SATADOM device. When you boot a vSAN host from a USB device or SD card, the size of the boot device must be at least 4 GB. If ESXi host memory is more than 512 GB, consider the following guidelines:

- You can boot the host from a SATADOM or disk device with a size of at least 16 GB. When you use a SATADOM device, use a single-level cell (SLC) device.
- If you are using vSAN 6.5 or later, you must resize the coredump partition on ESXi hosts to boot from USB/SD devices.

Consider using at least 32 GB memory per host for full vSAN operations based on five disk groups per host and seven capacity devices per disk group. Plan for 10% CPU overhead for vSAN.

## Cluster Requirements

You should verify that a host cluster contains a minimum of three hosts that contribute capacity to the cluster. A two-host vSAN cluster consists of two data hosts and an external witness host. It is important to ensure that each host that resides in a vSAN cluster does not participate in other clusters.

## Software Requirements

For full vSAN capabilities, the participating hosts must be version 6.7 Update 3 or later. vSAN 6.7.3 and later software supports all on-disk formats.

## Network Requirements

You should ensure that the network infrastructure and configuration support vSAN as described in [Table 2-9](#).

**Table 2-9** vSAN Networking Requirements

---

Component	Requirement
-----------	-------------

---

Host bandwidth	For hybrid configuration, each host requires 1 Gbps (dedicated). For all-flash configuration, each host requires 10 Gbps (dedicated or shared).
----------------	---

Host network	Each vSAN cluster member host cluster (even those that do not contribute capacity) must have a vSAN-enabled VMkernel network adapter connected to a Layer 2 or Layer 3 network.
IP version	vSAN supports IPv4 and IPv6.
Network latency	<p>Maximum round trip time (RTT) between all the member hosts in a standard vSAN clusters is 1 ms.</p> <p>Maximum RTT between the two main sites in a stretched vSAN cluster is 5 ms.</p> <p>Maximum RTT between each main site and the witness host in a stretched cluster is 200 ms.</p>

## License Requirements

You should ensure that you have a valid vSAN license that supports your required features. If you do not need advanced or enterprise features, a standard license is sufficient. An advanced (or enterprise) license is required for advanced features such as RAID 5/6 erasure coding, deduplication, and compression. An enterprise license is required for enterprise features such as encryption and stretched clusters.

The capacity of the license must cover the total number of CPUs in the cluster.

## Other vSAN Considerations

The following sections outline some other considerations that are important in planning a vSAN deployment.

## vSAN Network Best Practices

Consider the following networking best practices concerning vSAN:

- For hybrid configurations, use dedicated network adapters (at least 1 Gbps). For best performance, use dedicated or shared 10 Gbps adapters.

- For all-flash configurations, use a dedicated or shared 10 GbE physical network adapter.
- Provision one additional physical NIC as a failover NIC.
- If you use a shared 10 GbE network adapter, place the vSAN traffic on a distributed switch with Network I/O Control (NIOC) configured.

## **Boot Devices and vSAN**

You can boot ESXi from a local VMFS on a disk that is not associated with vSAN.

You can boot a vSAN host from a USB/SD device, but you must use a high-quality 4 GB or larger USB or SD flash drive. If the ESXi host memory is larger than 512 GB, for vSAN 6.5 or later, you must resize the coredump partition on ESXi hosts to boot from USB/SD devices.

You can boot a vSAN host from a SATADOM device, but you must use a 16 GB or larger single-level cell (SLC) device.

## **Persistent Logging in a vSAN Cluster**

When you boot ESXi from a USB or SD device, log information and stack traces are lost on host reboot because the scratch partition is on a RAM drive. You should consider using persistent storage other than vSAN for logs, stack traces, and memory dumps. You could use VMFS or NFS or you could configure the ESXi Dump Collector and vSphere Syslog Collector to send system logs to vCenter Server.

## **vSAN Policies**

Storage policies are used in vSAN to define storage requirements for virtual machines. These policies determine how to provision and allocate storage objects within the datastore to guarantee the required level of service. You should assign at least one storage policy to each virtual machine in a vSAN datastore. Otherwise, vSAN assigns a default policy with Primary Level of Failures to Tolerate set to 1, a single disk stripe per object, and a thin-provisioned virtual disks.

Storage policies, including those specific to vSAN, are covered later in this chapter.

# vSphere Storage Integration

In a vSphere 7.0 environment, you have several options for integrating with supported storage solutions, including Virtual Volumes (vVols), vSphere APIs for Storage Awareness (VASA), and vSphere APIs for Array Integration (VAAI).

## VASA

Storage vendors or VMware can make use of VASA. Storage providers (VASA providers) are software components that integrate with vSphere to provide information about the physical storage capabilities. Storage providers are utilized by either ESXi hosts or vCenter to gather information about the storage configuration and status and display it to administrators in the vSphere Client. There are several types of storage providers:

- **Persistent storage providers:** These storage providers manage storage arrays and handle abstraction of the physical storage. vVols and vSAN use persistent storage providers.
- **Data storage providers:** This type of provider is used for host-based caching, compression, and encryption.
- **Built-in storage providers:** These storage providers are offered by VMware and usually do not require registration. Examples of these are vSAN and I/O filters included in ESXi installations.
- **Third-party storage providers:** If a third party is offering a storage provider, it must be registered.

The information that storage providers offer may include the following:

- Storage data services and capabilities (which are referenced when defining a storage policy)
- Storage status, including alarms and events
- Storage DRS information

Unless the storage provider is VMware, the vendor must provide the policy. There are other requirements related to implementing storage providers as well:

- Contact your storage vendor for information about deploying the storage provider and ensure that it is deployed correctly.

- Ensure that the storage provider is compatible by verifying it with the *VMware Compatibility Guide*.
- Do not install the VASA provider on the same system as vCenter.
- Upgrade storage providers to new versions to make use of new functionalities.
- Unregister and reregister a storage provider when upgrading.

Storage providers must be registered in the vSphere Client to be able to establish a connection between vCenter and the storage provider. VASA is essential when working with vVols, vSAN, vSphere APIs for I/O Filtering (VAIO), and storage VM policies.

### **Note**

If vSAN is being used, service providers are registered automatically and cannot be manually registered.

## **VAAI**

VAAI, also known as hardware acceleration or hardware offload APIs, enable ESXi hosts to be able to communicate with storage arrays. They use functions called storage primitives, which allow offloading of storage operations to the storage array itself. The goal is to reduce overhead and increase performance. This allows storage to be responsible for cloning operations and zeroing out disk files. Without VAAI hardware offloading, the VMkernel Data Mover service is used to copy data from the source datastore to the destination datastore, incurring physical network latencies and increasing overhead. The VMkernel always attempts to offload to the storage array by way of VAAI, but if the offload fails, it employs its Data Mover service.

Storage primitives were introduced in vSphere 4.1 and applied to Fibre Channel, iSCSI, and FCoE storage only. vSphere 5.0 added primitives for NAS storage and vSphere thin provisioning. The storage primitives discussed in the following sections are available in vSphere 7.0.

## **VAAI Block Primitives**

The following are the VAAI primitives for block storage:

- **Atomic Test and Set (ATS):** Replaces the use of SCSI reservations on VMFS datastores when updating metadata. With SCSI reservations, only one process can establish a lock on the LUN at a time, leading to contention and SCSI reservation errors. Metadata updates occur whenever a thin-provisioned disk grows, a VM is provisioned, or a vSphere administrator manually grows a virtual disk. With ATS, a lock is placed on a sector of the VMFS datastore when updating metadata. ATS allows larger datastores to be used without running into such contention issues. On storage arrays that do not support VAAI, SCSI reservations are still used.
- **ATS Only Flag:** Can be set on VMFS datastores that were created as VMFS Version 5 but cannot be enabled on VMFS Version 5 datastores that were upgraded from VMFS Version 3. The ATS Only Flag primitive forces ATS to be used as opposed to SCSI reservations for all metadata updates and operations. Manually enabling the ATS only flag is done via **vmkfstools**, using the following syntax:

[Click here to view code image](#)

```
vmkfstools -configATSOnly 1 [storage path]
```

- **XCOPY (Extended Copy):** Allows the VMkernel to offload cloning or Storage vMotion migrations to the storage array, avoiding use of the VMkernel Data Mover service.
- **Write Same (Zero):** Used with eager zeroed thick-provisioned virtual disks, allowing the storage device to write the zeros for the disk. This reduces overhead on the ESXi host in terms of CPU time, DMA buffers, and use of the device queue. Write same is used whenever you clone a virtual machine with eager zeroed thick-provisioned disks, whenever a thin-provisioned disk expands, or when lazy zeroed thick disks need to be zeroed out (at first write).

## VAAI NAS Primitives

The following are the VAAI primitives for NAS:

- **Full File Clone:** Works the same way as XCOPY but applies to NAS devices as opposed to block storage devices.

- **Fast File Clone/Native Snapshot Support:** Allows snapshot creation to be offloaded to the storage device for use in linked clones used in VMware Horizon View or in vCloud Director, which leverage reading from replica disks and writing to delta disks.
- **Extended Statistics:** Allows an ESXi host to have insight into space utilization on a NAS device. For example, when a NAS device is using thin provisioning without the Extended Statistics primitive, the ESXi host lacks visibility into the actual storage usage, leading you to run out of space.
- **Reserve Space:** Allows thick provisioning of virtual disks on NAS datastores. Prior to this primitive, only thin provisioning could be used on NAS storage devices.

## VAAI Thin Provisioning Primitives

If you are using thin provisioning, and VMs are deleted or migrated off a datastore, the array may not be informed that blocks are no longer in use. Multiple primitives, including the following, were added in vSphere 5.0 to add better support for thin provisioning:

- **Thin Provisioning Stun:** Prior to vSphere 5.0, if a thin-provisioned datastore reached 100% space utilization, all VMs on that datastore were paused. After the release of vSphere 5.0, only the VMs requiring extra space are paused, and other VMs are not affected.
- **Thin Provisioning Space Threshold Warning:** When a VM is migrated to a different datastore or is deleted, the SCSI **UNMAP** command is used for the ESXi host to tell the storage array that space can be reclaimed.

## Virtual Volumes (vVols)

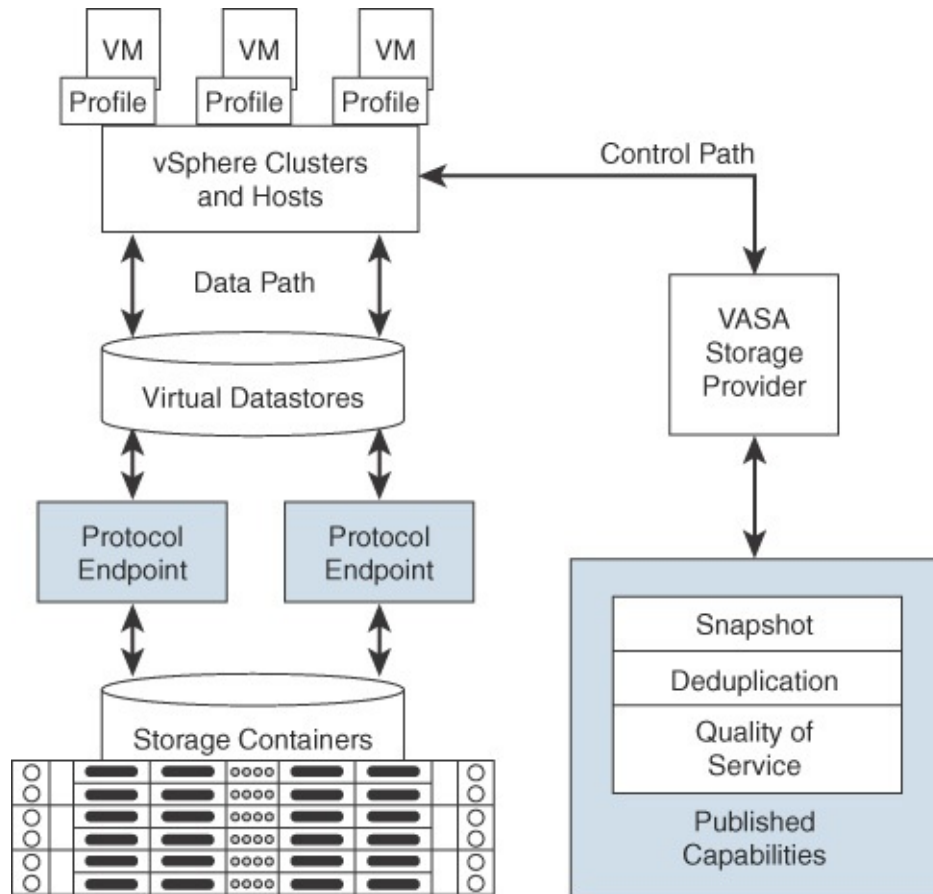
With vVols, you have a storage operational module that is similar to vSAN, and you can leverage SAN and NAS arrays. As with vSAN, with vVols you can leverage SPBM, which allows you to streamline storage operations. The VASA provider communicates with vCenter Server to report the underlying characteristics of the storage container. You can leverage these characteristics as you create and apply storage policies to virtual machines to optimize the placement and enable the underlying services (such as caching or replication).

The main use case for vVols is to simplify the operational model for virtual machines and their storage. With vVols, the operational model changes from



managing space inside datastores to managing abstract storage objects handled by storage arrays.

The major components in vVols are vVols device, protocol endpoints, storage container, VASA provider, and array. These components are illustrated in [Figure 2-7](#).



**FIGURE 2-7** vVols Architecture

These are the main characteristics of vVols:

- It has no file system.
- ESXi manages the array through VASA.
- Arrays are logically partitioned into containers, called storage containers.
- vVols objects are encapsulations of VM files, and disks are stored natively on the storage containers.
- Storage containers are pools of raw storage or aggregations of storage

capabilities that a storage device can provide to vVols.

- I/O from an ESXi host to a storage array is addressed through an access point called a protocol endpoint (PE).
- PEs are logical I/O proxies, used for communication for vVols and the virtual disk files. These endpoints are used to establish data paths on demand, by binding the ESXi hosts with the PEs.
- Bind requests must be sent from ESXi hosts or vCenter Server before a virtual volume can be used.
- Data services, such as snapshot, replication, and encryption, are offloaded to the array.
- Virtual volumes are managed through the SPBM framework. VM storage policies are required for VMs to use virtual volumes.

There are five types of virtual volumes:

- **Config-vVol:** Metadata
- **Data-vVol:** VMDKs
- **Mem-vVol:** Snapshots
- **Swap-vVol:** Swap files
- **Other-vVol:** Vendor solution specific

Limitations of vVols include the following:

- You cannot use vVols with a standalone ESXi host.
- vVols does not support RDMs.
- A vVols storage container cannot span different physical arrays.
- Host profiles that contain virtual datastores are vCenter Server specific. A profile created by one vCenter Server instance cannot be applied by another vCenter Server instance.

## Storage Multipathing and Failover

vSphere provides multipathing and failover as described in this section.

## Multipathing Overview

Multipathing is used for performance and failover. ESXi hosts can balance the storage workload across multiple paths for improved performance. In the event of a path, adapter, or storage processor failure, the ESXi host fails over to an alternate path.

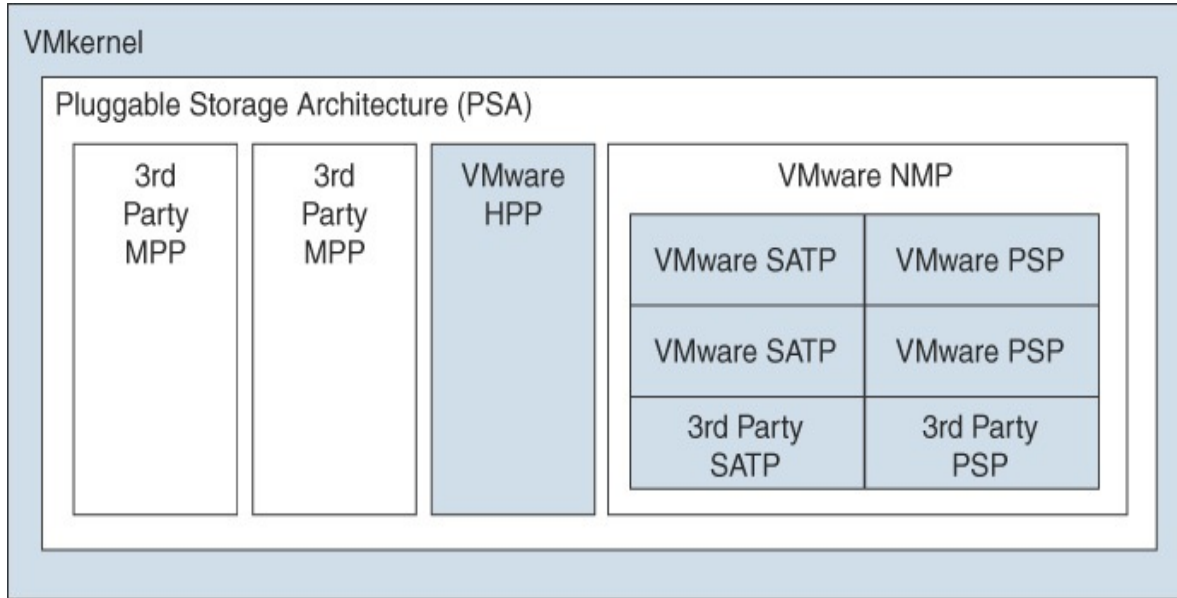
During path failover, virtual machine I/O can be delayed for a maximum of 60 seconds. Active/passive type arrays can experience longer delays than active/active arrays. vSphere supports several types of failover:

- **Fibre Channel failover:** For multipathing, hosts should have at least two HBAs in addition to redundant Fibre Channel switches (the switch fabric) and redundant storage processors. If a host has two HBAs, attached to two Fibre Channel switches, connected to two storage processors, then the datastores attached to the SAN can withstand the loss of any single storage processor, Fibre Channel switch, or HBA.
- **Host-based failover with iSCSI:** As with Fibre Channel failover, with host-based failover with iSCSI, hosts should have at least two hardware iSCSI initiators or two NIC ports used with the software iSCSI initiator. This is in addition to at least two physical switches and at least two storage processors.
- **Array-based failover with iSCSI:** On some storage systems, the storage device abstracts the physical ports from the ESXi hosts, and the ESXi hosts see only a single virtual port. The storage system uses this abstraction for load balancing and path failover. If the physical port where the ESXi host is attached should be disconnected, the ESXi host automatically attempts to reconnect to the virtual port, and the storage device redirects it to an available port.
- **Path failover and virtual machines:** When a path failover occurs, disk I/O could pause for 30 to 60 seconds. During this time, viewing storage in the vSphere client or virtual machines may appear stalled until the I/O fails over to the new path. In some cases, Windows VMs could fail if the failover is taking too long. VMware recommends increasing the disk timeout inside the guest OS registry to 60 seconds at least to prevent this.

## Pluggable Storage Architecture (PSA)

PSA was introduced in vSphere 4 as a way for storage vendors to provide their

own multipathing policies, which users can install on ESXi hosts. PSA is based on a modular framework that can make use of third-party multipathing plug-ins (MPPs) or the VMware-provided Native Multipathing Plug-in (NMP), as illustrated in [Figure 2-8](#).



**FIGURE 2-8** Pluggable Storage Architecture

VMware provides generic native multipathing modules, called VMware NMP and VMware HPP. In addition, the PSA offers a collection of VMkernel APIs that third-party developers can use. Software developers can create their own load-balancing and failover modules for a particular storage array. These third-party MPPs can be installed on the ESXi host and run in addition to the VMware native modules or as their replacement. When installed, the third-party MPPs can replace the behavior of the native modules and can take control of the path failover and the load-balancing operations for the specified storage devices.

## VMware NMP



VMware NMP supports all storage arrays listed on the VMware storage HCL and provides a default path selection algorithm based on the array type. It associates a set of physical paths with a specific storage device (LUN). NMP uses submodules, called Storage Array Type Plug-ins (SATPs) and Path Selection Plug-ins (PSPs).

NMP performs the following operations.

- Manages physical path claiming and unclaiming
- Registers and unregisters logical devices
- Maps physical paths with logical devices
- Supports path failure detection and remediation
- Processes I/O requests to logical devices:
  - Selects an optimal physical path
  - Performs actions necessary to handle path failures and I/O command retries
- Supports management tasks, such as resetting logical devices

## **Storage Array Type Plug-ins (SATPs)**

SATPs are submodules of the VMware NMP and are responsible for array-specific operations. The SATP handles path failover for a device. ESXi offers an SATP for every type of array that VMware supports. ESXi also provides default SATPs that support non-specific active/active, active/passive, ALUA, and local devices.

Each SATP performs the array-specific operations required to detect path state and to activate an inactive path. This allows the NMP module to work with multiple storage arrays without being aware of the storage device specifics.

NMP determines which SATP to use for a specific storage device and maps the SATP with the storage device's physical paths. The SATP implements the following tasks:

- Monitors the health of each physical path
- Reports changes in the state of each physical path
- Performs array-specific actions necessary for storage failover (For example, for active/passive devices, it activates passive paths.)

[Table 2-10](#) provides details on the native SATP modules.

### **Table 2-10 SATP Module Details**

---

<b>SATP Module</b>	<b>Description</b>
VMW_SATP_LOCAL	SATP for local direct-attached devices. Supports VMW_PSP_MRU and VMW_PSP_FIXED but not VMW_PSP_RR.
VMW_SATP_DEFAULT_AA	Generic SATP for active/active arrays.
VMW_SATP_DEFAULT_AP	Generic SATP for active/passive arrays.
VMW_SATP_ALUA	SATP for ALUA-compliant arrays.

### **Note**

You do not need to obtain or download any SATPs. ESXi automatically installs an appropriate SATP for any array you use. Beginning with vSphere 6.5 Update 2, VMW\_SATP\_LOCAL provides multipathing support for the local devices, except the devices in 4K native format. You are no longer required to use other SATPs to claim multiple paths to the local devices.

## **Path Selection Plug-ins (PSPs)**

VMware PSPs are submodules of NMP. PSPs handle path selection for I/O requests for associated storage devices. NMP assigns a default PSP for each logical device based on the device type. You can override the default PSP.

Each PSP enables and enforces a corresponding path selection policy. [Table 2-11](#) provides details on the native path selection policies.

**Table 2-11** VMware Path Selection Policies

<b>PSP Poli</b>	<b>Description</b>
-----------------	--------------------

<b>Module</b>	
---------------	--

V M Initially, MRU selects the first discovered working path.  
M os  
W t If the path fails, MRU selects an alternative path and does not  
- R revert to the original path when that path becomes available.  
P ec  
S en  
P tl MRU is default for most active/passive storage devices.  
- y  
M U  
R se  
U d  
( V  
M  
w  
ar  
e)

---

V Fi FIXED uses the designated preferred path if it is working. If the  
M xe preferred path fails, FIXED selects an alternative available path but  
W d reverts to the preferred path when it becomes available again.  
- ( V  
P V FIXED is the default policy for most active/active storage devices.  
S M  
P w  
- ar  
FI e)  
X  
E  
D

---

V R RR uses an automatic path selection algorithm to rotate through the  
M ou configured paths. RR sends an I/O set down the first path, sends  
W nd the next I/O set down the next path, and continues sending the next  
- R I/O set down the next path until all paths are used; then the pattern  
P ob repeats, beginning with the first path. Effectively, this allows all the  
S in I/O from a specific host to use the aggregated bandwidth of  
P ( multiple paths to a specific storage device.  
- V  
R M

- R w Both active/active and active/passive arrays use RR. With  
ar active/passive arrays, RR uses active paths. With active/active  
e) arrays, RR uses available paths.

The latency mechanism that is activated for the policy by default makes it more adaptive. To achieve better load-balancing results, the mechanism dynamically selects an optimal path by considering the I/O bandwidth and latency for the path.

RR is the default policy for many arrays.

---

## PSA Summary

To summarize, the PSA performs the following tasks:

- Loads and unloads multipathing plug-ins
- Hides virtual machine specifics from a particular plug-in
- Routes I/O requests for a specific logical device to the MPP managing that device
- Handles I/O queueing to the logical devices
- Implements logical device bandwidth sharing between virtual machines
- Handles I/O queueing to the physical storage HBAs
- Handles physical path discovery and removal
- Provides logical device and physical path I/O statistics

The following process occurs when VMware NMP receives an I/O request for one of its managed storage devices:

**Step 1.** The NMP calls the appropriate PSP.

**Step 2.** The PSP selects an appropriate physical path.

**Step 3.** The NMP issues the I/O request on the selected path.

**Step 4.** If the I/O operation is successful, the NMP reports its completion.

**Step 5.** If the I/O operation reports an error, the NMP calls the appropriate



SATP.

**Step 6.** The SATP interprets the errors and, when appropriate, activates the inactive paths.

**Step 7.** The PSP selects a new path for the I/O.

When coordinating the VMware native modules and any installed third-party MPPs, the PSA performs the following tasks:

- Loads and unloads MPPs
- Hides virtual machine specifics from MPPs
- Routes I/O requests for a specific logical device to the appropriate MPP
- Handles I/O queuing to the logical devices
- Shares logical device bandwidth between virtual machines
- Handles I/O queuing to the physical storage HBAs

## Storage Policies

Storage policies can be used to define which datastores to use when placing virtual machines disk. The following storage policies can be created:

- **VM storage policies for host-based data services:** These policies are rules for services that are offered by the ESXi hosts, such as encryption.
- **VM storage policies for vVols:** These policies allow you to set rules for VMs that apply to vVols datastores. This can include storage devices that are replicated for disaster recovery purposes or have specific performance characteristics.
- **VM storage policies for tag-based placement:** You can create custom policies for VMs and custom tags for storage devices. This is helpful for storage arrays that do not support VASA and whose storage characteristics are not visible to the vSphere client. For example, you could create a tag named Gold and use it to identify your best-performing storage.

## Storage Policy Based Management (SPBM)

You can define a required policy for a VM, such as requiring it to reside on fast storage. You can then use VASA or define storage tags manually. Then a VM can

only be placed on a storage device matching the requirements.

## Virtual Disk Types

When creating a virtual disk, you need to determine how you are going to allocate space to that virtual disk. The way space is allocated to a virtual disk is through writing zeros; this process is typically referred to as *zeroing out* the file. For example, if you wanted to create a 20 GB virtual disk and allocate all of the space up front, a VMDK file is created, and 20 GB worth of zeros are written to that file. You can determine when the zeros get written:

- **Eager zeroed thick:** The disk space for the virtual disk files is allocated and erased (zeroed out) at the time of creation. If the storage device supports VAAI, this operation can be offloaded to the storage array. Otherwise, the VMkernel writes the zeros, and this process could be slow. This method is the slowest for virtual disk creation, but it is the best for guest performance.
- **Lazy zeroed thick:** The disk space for the virtual disk files is allocated at the time of creation but not zeroed. Each block is zeroed, on demand at runtime, prior to being presented to the guest OS for the first time. This increases the time required for disk format operations and software installations in the guest OS.
- **Thin provisioned:** The disk space for the virtual disk files is not allocated or zeroed at the time of creation. The space is allocated and zeroed on demand. This method is the fastest for virtual disk creation but the worst for guest performance.

## vSAN-Specific Storage Policies

vSAN storage policies define how VM objects are placed and allocated on vSAN to meet performance and redundancy requirements. [Table 2-12](#) describes the vSAN storage policies.

**Table 2-12** vSAN Storage Policies

---

Polic	Description
-------	-------------

y

---

Pri mar y	This policy defines how many host and device failures a VM object can withstand. For $n$ failures tolerated, data is stored in $n+1$ location. (This includes parity copies with RAID 5 or 6.) If no storage policy is
-----------------	--

Level of Failures to Tolerate (PFTT)

selected at the time of provisioning a VM, this policy is assigned by default. Where fault domains are used,  $2n+1$  fault domains, each with  $n$  hosts adding to the capacity, are required. If an ESXi host isn't in a fault domain, it is considered to be in a single-host fault domain. The default setting for this policy is 1, and the maximum is 3.

---

Secondary Level of Failures to Tolerate (SF TT)

In stretched clusters, this policy defines how many additional host failures can be tolerated after a site failure's PFTT has been reached. If PFTT = 1, SF TT = 2, and one site is inaccessible, two more host failures can be tolerated. The default setting for this policy is 1, and the maximum is 3.

---

Data Locality

If PFTT=0, this option is available. The options for this policy are None, Preferred, and Secondary. This allows objects to be limited to one site or one host in stretched clusters. The default setting for this policy is None.

---

Failure Tolerance Method

This policy defines whether the data replication mechanism is optimized for performance or capacity. If RAID-1 (Mirroring)—Performance is selected, there will be more space consumed in the object placement but better performance for accessing the space. If RAID-5/6 (Erasure Coding)—Capacity is selected, there will be less disk utilization, but performance will be reduced.

d

---

Nu mb er of Dis k Stri pes per Obj ect	This policy determines the number of capacity devices where each VM object replica is striped. Setting this above 1 can improve performance but consumes more resources. The default setting for this policy is 1, and the maximum is 12.
--	---

---

Fla sh Rea d Cac he Res erv atio n	This policy defines the amount of flash capacity that is reserved for read caching of VM objects. This is defined as a percentage of the size of the VMDK. This is supported only in hybrid vSAN clusters. The default setting for this policy is 0%, and the maximum is 100%.
---	--

---

For ce Pro visi oni ng	If set to yes, this policy forces provisioning of objects, even when policies cannot be met. The default setting for this policy is no.
---------------------------------------	---

---

Obj ect Spa ce Res erv atio n	This policy defines the percentage of VMDK objects that must be thick provisioned on deployment. The options are as follows:
--	--

- Thin provisioning (default value)
- 25% reservation

- 50% reservation
- 75% reservation
- Thick provisioning

---

Disable Object Checksum

A checksum is used end-to-end in validating the integrity of the data to ensure that data copies are the same as the original. In the event of a mismatch, incorrect data is overwritten. If this policy is set to yes, a checksum is not calculated. The default setting for this policy is no.

---

IOPS Limit for Object

This policy sets a limit for IOPS of an object. If set to 0, there is no limit.

---

## Storage DRS (SDRS)

You can use vSphere Storage DRS (SDRS) to manage the storage resources of a datastore cluster. A datastore cluster is a collection of datastores with shared resources and a shared management interface. SDRS provides a number of capabilities for a datastore cluster, as discussed in the following sections.

### Initial Placement and Ongoing Balancing

SDRS provides recommendations for initial virtual machine placement and

ongoing balancing operations in a datastore cluster. Optionally, SDRS can automatically perform the recommended placements and balancing operations. Initial placements occur when the virtual machine is being created or cloned, when a virtual machine disk is being migrated to another datastore cluster, and when you add a disk to an existing virtual machine. SDRS makes initial placement recommendations (or automatically performs the placement) based on space constraints and SDRS settings (such as space and I/O thresholds).

## Space Utilization Load Balancing

You can set a threshold for space usage to avoid filling a datastore to its full capacity. When space usage on a datastore exceeds the threshold, SDRS generates recommendations or automatically performs Storage vMotion migrations to balance space usage across the datastore cluster.

## I/O Latency Load Balancing

You can set an I/O latency threshold to avoid bottlenecks. When I/O latency on a datastore exceeds the threshold, SDRS generates recommendations or automatically performs Storage vMotion migrations to balance I/O across the datastore cluster.

SDRS is invoked at the configured frequency (every eight hours by default) or when one or more datastores in a datastore cluster exceeds the user-configurable space utilization thresholds. When Storage DRS is invoked, it checks each datastore's space utilization and I/O latency values against the threshold. For I/O latency, Storage DRS uses the 90th percentile I/O latency measured over the course of a day to compare against the threshold.

## SDRS Automation Level

[Table 2-13](#) describes the available SDRS automation levels.

**Table 2-13** SDRS Automation Levels

Option	Description
No Automation (manual mode)	Placement and migration recommendations are displayed but do not run until you manually apply the recommendation.

Partially Automated	Placement recommendations run automatically, and migration recommendations are displayed but do not run until you manually apply the recommendation.
Fully Automated	Placement and migration recommendations run automatically.

## SDRS Thresholds and Behavior



You can control the behavior of SDRS by specifying thresholds. You can use the following standard thresholds to set the aggressiveness level for SDRS:

- **Space Utilization:** SDRS generates recommendations or performs migrations when the percentage of space utilization on the datastore is greater than the threshold you set in the vSphere Client.
- **I/O Latency:** SDRS generates recommendations or performs migrations when the 90th percentile I/O latency measured over a day for the datastore is greater than the threshold.
- **Space Utilization Difference:** SDRS can use this threshold to ensure that there is some minimum difference between the space utilization of the source and the destination prior to making a recommendation. For example, if the space used on datastore A is 82% and on datastore B is 79%, the difference is 3. If the threshold is 5, Storage DRS does not make migration recommendations from datastore A to datastore B.
- **I/O Load Balancing Invocation Interval:** After this interval, SDRS runs to balance I/O load.
- **I/O Imbalance Threshold:** Lowering this value makes I/O load balancing less aggressive. Storage DRS computes an I/O fairness metric between 0 and 1, with 1 being the fairest distribution. I/O load balancing runs only if the computed metric is less than  $1 - (\text{I/O Imbalance Threshold} / 100)$ .

## SDRS Recommendations

For datastore clusters, where SDRS automation is set to No Automation (manual mode), SDRS makes as many recommendations as necessary to enforce SDRS rules, balance the space, and balance the I/O resources of the datastore cluster.

Each recommendation includes the virtual machine name, the virtual disk name, the data-store cluster name, the source datastore, the destination datastore, and a reason for the recommendation.

SDRS makes mandatory recommendations when the datastore is out of space, when anti-affinity or affinity rules are being violated, or when the datastore is entering maintenance mode. SDRS makes optional recommendations when a datastore is close to running out of space or when adjustments should be made for space and I/O load balancing.

SDRS considers moving powered-on and powered-off virtual machines for space balancing. Storage DRS considers moving powered-off virtual machines with snapshots for space balancing.

## **Anti-affinity Rules**

To ensure that a set of virtual machines are stored on separate datastores, you can create anti-affinity rules for the virtual machines. Alternatively, you can use an affinity rule to place a group of virtual machines on the same datastore.

By default, all virtual disks belonging to the same virtual machine are placed on the same datastore. If you want to separate the virtual disks of a specific virtual machine on separate datastores, you can do so with an anti-affinity rule.

## **Datastore Cluster Requirements**

Datastore clusters can contain a mix of datastores having different sizes, I/O capacities, and storage array backing. However, the following types of datastores cannot coexist in a datastore cluster:

- NFS and VMFS datastores cannot be combined in the same datastore cluster.
- Replicated datastores cannot be combined with non-replicated datastores in the same SDRS-enabled datastore cluster.
- All hosts attached to the datastores in a datastore cluster must be ESXi 5.0 and later. If datastores in the datastore cluster are connected to ESX/ESXi 4.x and earlier hosts, SDRS does not run.
- Datastores shared across multiple data centers cannot be included in a datastore cluster.



- As a best practice, all datastores in a datastore cluster should have identical hardware acceleration (enabled or disabled) settings.

## NIOC, SIOC, and SDRS

In vSphere, you can use Network I/O Control (NIOC), Storage I/O Control (SIOC), and SDRS to manage I/O. These features are often confused by people in the VMware community. [Table 2-14](#) provides a brief description of each feature along with the chapter in this book where you can find more detail.

**Table 2-14** Comparing NIOC, SIOC, and SDRS

Feature	Description	Chapter
N I O C	Allows you to allocate network bandwidth to business-critical applications and to resolve situations where several types of traffic compete for common resources.	9
S I O C	Allows you to control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion by implementing shares and limits.	1 1
S D R S	Allows you to control and balance the use of storage space and I/O resources across the datastores in a datastore cluster.	1 1

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

## Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 2-15](#) lists these key topics and the page number on which each is found.



**Table 2-15** Key Topics

Key Topic Element	Description	Page Number
Paragraph	Raw device mapping	<a href="#">38</a>
<a href="#">Figure 2-4</a>	A two-node vSAN cluster	<a href="#">55</a>
Paragraph	vSAN stretched clusters	<a href="#">56</a>
List	vSAN limitations	<a href="#">59</a>
Paragraph	Native multipathing plug-in	<a href="#">75</a>
Paragraph	SDRS thresholds and behavior	<a href="#">82</a>

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

[raw device mapping \(RDM\)](#)

[I/O filter](#)

[disk group](#)

[witness host](#)

[vSAN File Services](#)

[virtual volume](#)

## Review Questions

1. You are deploying datastores in a vSphere environment and want to use the latest VMFS version that supports ESXi 6.5 and ESXi 7.0. Which version should you use?
  - a. VMFS Version 3
  - b. VMFS Version 4
  - c. VMFS Version 5
  - d. VMFS Version 6
2. You are preparing to manage and troubleshoot a vSAN environment. Which of the following is a command-line interface that provides a cluster-wide view and is included with the vCenter Server deployment?
  - a. VMware PowerCLI
  - b. vSAN Observer
  - c. Ruby vSphere Console
  - d. **esxcli**
3. You want to integrate vSphere with your storage system. Which of the following provides software components that integrate with vSphere to provide information about the physical storage capabilities?
  - a. VASA
  - b. VAAI
  - c. SATP
  - d. NMP
4. Which of the following is the default path selection policy for most active/passive storage devices?
  - a. VMW\_PSP\_MRU
  - b. VMW\_PSP\_FIXED
  - c. VMW\_PSP\_RR

**d.** VMW\_PSP\_AP

**5.** You are deploying virtual machines in a vSphere environment. Which virtual disk configuration provides the best performance for the guest OS?

**a.** Thin provisioned

**b.** Thick eager zeroed

**c.** Thick lazy zeroed

**d.** Thin eager zeroed

## Chapter 3

# Network Infrastructure

**This chapter covers the following topics:**

- [Networking Terms and Concepts](#)
- [vSphere Standard Switch \(vSS\)](#)
- [vSphere Distributed Switch \(vDS\)](#)
- [vDS Settings and Features](#)
- [Other vSphere Networking Features](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.4, 1.7, 1.7.1, 1.7.2, 1.7.3, 1.7.4, 4.2, 4.14, and 5.4.

This chapter provides details for the network infrastructure, both physical and virtual, involved in a vSphere 7 environment.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 3-1](#) 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 and Review Questions.](#)”

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

<b>Foundation Topics Section</b>	<b>Questions</b>
vSphere Standard Switch (vSS)	1, 2

vSphere Distributed Switch (vDS)	3–5
vDS Settings and Features	6–8
Other vSphere Networking Features	9, 10

1. You are configuring networking policies in your vSphere 7.0 environment. Which of the following policies are not available for a vSS?
  - a. Teaming and Failover
  - b. Security
  - c. Traffic Shaping
  - d. Port Blocking
2. You are configuring teaming policies in your vSphere 7.0 environment. Which of the following teaming options are not available for a vSS?
  - a. Route Based on Originating Virtual Port
  - b. Route Based on Physical NIC Load
  - c. Route Based on Source MAC Hash
  - d. Use Explicit Failover Order
3. You are configuring virtual networking in your vSphere 7.0 environment. Which of the following policies are available for distributed switches but not for standard switches?
  - a. Inbound traffic shaping
  - b. Outbound traffic shaping
  - c. IPv6 support
  - d. CDP
4. Using NIOC, you reserved 1.0 Gbps for virtual machine system traffic on a distributed switch with eight 10 Gbps uplinks. If you reserve 512 Mbps for a [network resource pool](#), what is the maximum reservation you can set in another network resource pool?
  - a. 0.5 Gbps

- b.** 1.0 Gbps
  - c.** 7.5 Gbps
  - d.** 79.6 Gbps
- 5.** You want to tag voice over IP data in your vSphere environment. Which step should you take?
  - a.** Use a vSS
  - b.** Implement a filtering rule
  - c.** Implement a rule with Action = Tag
  - d.** Navigate to Menu > Tags and Custom Attributes
- 6.** You want to leverage LAGs with your vDS. Which of the following is supported?
  - a.** iSCSI software port binding
  - b.** Nested ESXi
  - c.** 40 LAGs on a single distributed switch
  - d.** 40 LAGs on a single host
- 7.** You want to enable a vDS health check for virtual switch teaming. Which of the following is a requirement?
  - a.** At least two active NICs and two hosts
  - b.** At least one active NIC from each of two hosts
  - c.** At least two active NICs from at least one host
  - d.** At least one active NIC from one host
- 8.** You want to be able to use the vSphere Client to identify the non-Cisco physical switch connected to a virtual switch uplink. Which of the following should you implement?
  - a.** A vSS with LLDP
  - b.** A vSS with CDP
  - c.** A vDS with LLDP
  - d.** A vDS with CDP

9. You want to use DirectPath I/O in your vSphere 7.0 environment. Which of the following features is supported?
  - a. The virtual machine running in a vSphere cluster
  - b. Hot adding virtual devices
  - c. Snapshots
  - d. Fault tolerance
  
10. You want to create a custom TCP/IP stack for your ESXi 7.0 server. Which one of the following is an available service that you can directly enable for the stack?
  - a. vSphere Replication NFC
  - b. NFS
  - c. vSphere HA heartbeat
  - d. iSCSI

## Foundation Topics

### Networking Terms and Concepts

Computer networking is built on the TCP/IP protocol suite, which enables all types of computers, from many different vendors, to communicate with one another. This section introduces essential network terminology and concepts for a vSphere environment.

#### Traditional Networking Terminology

It is important to be familiar with the following commonly used networking terminology:

- **TCP/IP:** Transmission Control Protocol/Internet Protocol (TCP/IP) is a model used in the Internet that allows computers to communicate over long distances. It was developed by the U.S. Department of Defense (DOD) in the 1970s and is defined in RFC documents. It includes components that break data into packets, deliver the packets to the proper destination, and reassemble data from the packets.



- **RFC:** Requests for Comments (RFCs) are the rules of the Internet and are managed by the Internet Engineering Task Force (IETF). Anyone can register to attend an IETF meeting and help draft and define the standards of communication for the Internet.
- **Physical network:** A physical network consists of physical network devices and cabling that allows physical machines, such as ESXi servers, to connect and communicate with other physical machines. A key component of a physical network is the physical Ethernet switch.
- **MAC address:** Each network-connected device has a media access control (MAC) address. You can think of a MAC address as the physical network address of a device. Each server on a network has one or more network interface cards (NICs) that it uses to connect to the network. Each NIC has a unique MAC address. The MAC address is used to identify devices in an Ethernet network.
- **IP address:** In addition to having a MAC address, a device is likely assigned a unique IP address. IP addresses are used to identify devices in an IP network such as the Internet and private subnets.
- **Physical Ethernet switch:** A physical Ethernet switch manages network traffic between devices that are directly connected to its ports. A switch has multiple ports, each of which can be connected to a single device or to another switch. The switch learns the MAC address of each connected device and uses it to build a MAC address table that maps the address to its port number. An Ethernet switch (also called a Layer 2 switch) acts as a traffic cop, directing each received network packet to the proper port, based on the data in the MAC address table. Layer 3 switches use IP routing tables to route traffic between subnets.
- **Opaque network:** An opaque network is a network created and managed by a separate entity outside vSphere, such as a logical segment created by VMware NSX-T. The opaque network appears as an object in the vSphere Client, but the vCenter Server cannot manage it. You can use the NSX Manager in NSX-T to manage an opaque network.
- **EtherChannel:** An *EtherChannel* is a logical channel formed by bundling together two or more links to aggregate bandwidth and provide redundancy. Other acceptable names for EtherChannel (an IOS term) are port channel (an NX-OS term) and link aggregation group (LAG).
- **LACP:** Link Aggregation Control Protocol (LACP) is a standards-based

negotiation protocol used to dynamically build an EtherChannel. It is also known as IEEE 802.1ax (or IEEE 802.3ad). LACP is used to build EtherChannels (LAGs) dynamically. EtherChannels can also be built statically without the use of LACP.

- **IEEE 802.1ax:** This is the IEEE working group that defines port channel, EtherChannel, and link aggregation. Originally, the working IEEE group was 802.3ad, but in 2008 it was replaced by 802.1ax.
- **IEEE 802.3ad:** This is the original IEEE working group for port channel, EtherChannel, and link aggregation. Although it has been replaced with 802.1ax, it is still referred to as IEEE 802.3ad.

## Virtual NICs

Much as a physical server may have multiple NICs to connect to physical networks, a virtual machine may have multiple virtual NICs (vNICs) to connect to virtual networks. Much like a physical NIC, each vNIC has a unique MAC address. The vNIC appears as a traditional NIC to a virtual machine's guest OS. The guest OS can assign IP addresses to vNICs.

In addition to requiring network connectivity for virtual machine networking, ESXi requires network connectivity for host management activities and other purposes. To accommodate this need, you should configure one or more VMkernel virtual network adapters on each host. For example, when connecting the vCenter Server or the vSphere Host Client to an ESXi host, you provide the address (IP address or fully qualified host name) of a VMkernel virtual network adapter that is enabled for management traffic. Much as a virtual machine can use multiple virtual network adapters, each ESXi host may use multiple VMkernel network adapters.

## Virtual Switch Concepts

A virtual switch is a software construct that acts much like a physical switch to provide networking connectivity for virtual devices within an ESXi host. Each virtual machine may use vNICs to connect to virtual ports on virtual switches. To gain access to the physical network, one or more of the host's physical NICs should be connected to the virtual switch as uplinks. The virtual switch, which resides in the VMkernel, provides traffic management for all vNICs in the host and all the uplink's incoming and outgoing Ethernet frames.

A virtual switch works at Layer 2 of the OSI model. It can send and receive data,

provide VLAN tagging, and provide other networking features. A virtual switch maintains a MAC address table that contains only the MAC addresses for the vNICs that are directly attached to the virtual switch. When an Ethernet frame is received by the virtual switch, it passes the frame to the appropriate port, based on its internal MAC address table. The virtual switch updates the physical switch with its internal MAC address table using ARP.

With vSphere 4.0, VMware renamed the original virtual switch vSphere Standard Switch (vSS) and introduced the vSphere Distributed Switch (vDS). A vSS and a vDS behave similarly but are managed and controlled differently. Each vSS is configured and managed by a specific ESXi host. A vDS is configured and managed by the vCenter Server, which pushes the switch's configuration to each participating host. The vDS has many more features than the vSS. All new features are built into the new vDS, while the vSS capabilities are limited.

Each virtual switch has many virtual ports. You can configure port groups (standard port groups) on a vSS. You can configure ports (distributed ports) and port groups (distributed port groups) on a vDS.

## VLANs

A **virtual LAN (VLAN)** is a logical partition of a physical network at the data link layer (Layer 2). A VLAN is typically associated with a broadcast domain and is used to isolate the traffic from other networks. A broadcast domain is a collection of network devices that can receive traffic destined to a broadcast address. A physical switch, by default, adheres to this behavior. With VLAN technology, the switch allows a single physical network to be divided into multiple network segments. This is achieved by modifying a unique header of the Ethernet frame and adding a tag to identify the membership within a specific VLAN. A VLAN can be used to subdivide a broadcast domain to limit the number of network devices that can communicate when a broadcast packet is sent.

### Note

IEEE 802.1Q is the networking standard that supports VLANs on an Ethernet network.

ESXi provides VLAN support by allowing you to assign a VLAN ID to a virtual switch port group. The VMkernel is responsible for all tagging and untagging packets with VLAN IDs as the packets pass through the virtual switch. To properly assign VLANs, you must configure the port group to match the VLAN ID that has been configured at the physical switch. On a single virtual switch, you can configure multiple port groups and assign a unique VLAN ID to each port group. To support the multiple VLANs, the physical NIC that is used as the virtual switch's uplink should connect to a physical switch port that is configured for VLAN trunking.

On a standard port group, the valid VLAN ID range is 1 to 4095, with 4095 being unique. The vSS forwards all packets from all ports to the port group assigned VLAN 4095. One use case is to implement packet sniffing or intrusion detection services in a virtual machine. For sniffing or intrusion detection, you must also enable promiscuous mode for the port group. Another use case is to support VLAN guest tagging.

On a distributed port group, the valid VLAN ID range is 1 to 4094. You cannot assign 4095 to a distributed port group. Instead, you can directly configure VLAN trunking by using a distributed port group. On a distributed port group, you can select the VLAN Trunking option and identify the range of VLANs that are to be trunked. Setting a standard port group for VLAN 4095 is much like setting a distributed port group for VLAN trunking for VLANs 1 to 4094.

For example, say that your organization uses separate VLANs for production traffic (VLAN 101), test/development traffic (VLAN 102), QA traffic (VLAN 103), IP storage traffic (VLAN 300), vMotion traffic (VLAN 310), and management traffic (VLAN 500). You want a virtual machine running Wireshark to be able to receive all virtual machine traffic and management traffic but not IP storage or vMotion traffic. You can configure the distributed port group to use VLAN trunking range 101 to 103,500. Or, if you wanted to only exclude vMotion and IP storage traffic, you can define the range 1 to 103,500 to 4094.

## **vSphere Standard Switch (vSS)**

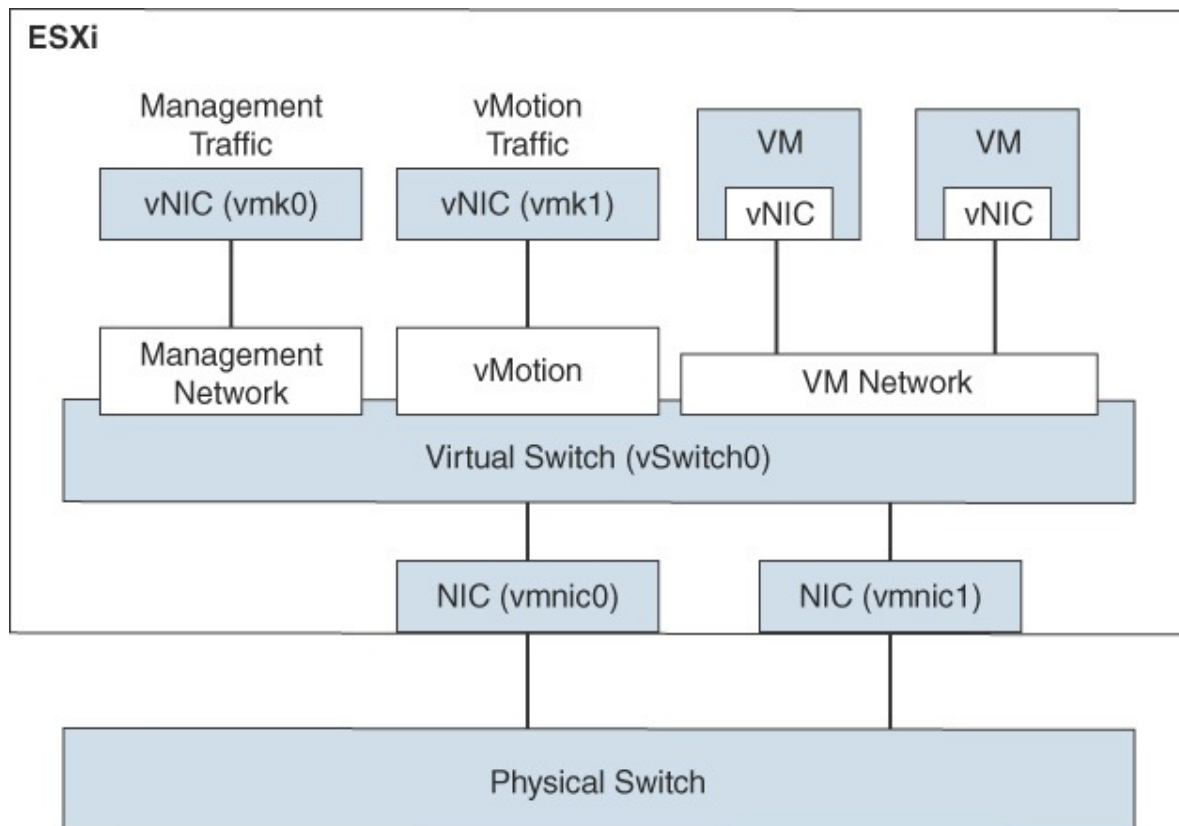
A vSphere Standard Switch (vSS) is also called a host-based switch because it is configured and maintained within a specific ESXi host. The vSS data and management planes reside within a single host.

Much like a physical Ethernet switch, a vSS maintains and uses a MAC address table. But with a vSS, the VMkernel directly registers each connected vNIC's

MAC address, along with its virtual port, in the virtual switch's MAC address table. The vSS does not observe network packets to learn the MAC address. The vSS sends all packets whose destination address does not match a connected vNIC to the physical NIC uplinks.

During a default ESXi installation, the host is configured with one virtual switch, identified as vSwitch0, with a physical NIC uplink, identified as vmnic0. The default virtual switch contains a virtual machine port group named VM Network and a port group named Management Network. VM Network is a port group for connecting virtual machines (a virtual machine port group). Management Network is a port group that contains a single VMkernel virtual network adapter (vmk0) that is used for management purposes, such as connecting to vCenter Server. A VMkernel virtual adapter is commonly called a VMkernel port. You can create additional standard switches. You can add virtual machine port groups and VMkernel virtual network adapters to existing standard switches. You can assign unique labels and policies to virtual machine port groups. These policies, which include NIC teaming, security, and traffic shaping policies, are applied to the virtual machines connected to the port group.

For example, [Figure 3-1](#) is a diagram representing a standard virtual switch with two virtual machines connected to a port group named VM Network, two VMkernel ports (one named Management Network and one named vMotion), and two physical NIC uplinks (vmnic0 and vmnic1).



**FIGURE 3-1** A Sample Virtual Switch

Each virtual machine typically has one or more vNICs to allow network communication. To connect a virtual machine to a vSS, you should connect one of its vNICs to a virtual machine port group on the vSS. To allow the virtual machines to communicate with virtual machines on other hosts, connect the port group to one or more physical NIC uplinks in the vSS. The physical NIC should be connected to a physical Ethernet switch. The inbound and outbound Ethernet frames travel through the physical NIC uplink on the vSS. Virtual machines in a port group that do not have a physical NIC uplink can only communicate with other vNICs on the same host and port group.

The vSS provides features such as VLAN tagging, NIC teaming, network security policies, and traffic shaping. The feature set provided by a vSS is smaller than the feature set provided by a vDS or a physical Ethernet switch. The vDS feature set is covered later in this chapter, along with a comparison of vSS and vDS.

A vSS does not have some of the vulnerabilities of a physical switch. For example, it is not susceptible to MAC address flooding attacks because it does

not use observed data to populate its MAC address table. The following list contains other common network vulnerabilities and brief explanations about why a vSS is safe from each vulnerability:

- **802.1q tagging attacks:** A vSS does not perform the dynamic trunking required for this type of attack.
- **Double-encapsulation attacks:** A vSS drops any double-encapsulated frames that a virtual machine attempts to send on a port configured for a specific VLAN.
- **Multicast brute-force attacks:** A vSS does not allow frames to leave their correct broadcast domain (VLAN).
- **Spanning tree attacks:** A vSS does not participate in a spanning tree protocol.

A vSS is not a managed switch, so it cannot be configured and monitored using traditional networking models and utilities. Instead, you can manage a vSS by using VMware provided tools, such as the vSphere Client, vSphere Host Client, vCLI, ESXi shell, or PowerCLI.

The main vSS switch property that you can set directly with the vSphere client is the MTU size, as explained in the next section.

## MTU

The standard size for Ethernet packets, or frames, is 1500 bytes. Using larger (jumbo) frames can provide better utilization of a fast network link. To allow jumbo frames on a vSS or vDS, you must set the virtual switch's Maximum Transmission Unit (MTU) setting to a value larger than 1500 bytes, such as 9000 bytes. To use jumbo frames, you must configure the network to support it end to end, including physical NICs and physical switches. To allow a virtual machine to use jumbo frames, you must configure the virtual machine to use the VMXNET3 virtual network adapter. (E1000 and E1000E adapters do not support jumbo frames.) You must also enable jumbo frames inside the guest operating system. To allow VMkernel services such as IP-based storage, vMotion, fault tolerance, and VSAN to use jumbo frames, you must set the MTU setting for the VMkernel network adapter to a value greater than 1500 bytes.

Failure to properly configure MTU to support jumbo frames end to end typically results in poor network performance. With a vDS, you can enable a health check



to warn of mismatched configuration between the vDS and the physical network hardware.

## vSS Network Policies

You can apply network policies to standard and distributed switches to use and configure specific features. On a vSS, you can apply policies to the switch, which automatically propagate to each port group. At the port group level, you can override policies applied to the switch and apply unique policies to a port group. On a vDS, you have more flexibility for applying policies, and you can override policies at the port level.

On a vSS, you can set the following network policies:



- Teaming and Failover
- Security
- Traffic Shaping
- VLAN

You can set the policies directly on a vSS. To override a policy at the port group level, just set a different policy on the port group.

A vDS supports additional policies. See the “vDS Network Policies” section, later in this chapter, for details.

## NIC Teaming Policies

When you connect two or more physical NICs (pNICs) to the same vSS, you should decide which NIC teaming option to utilize. The two main purposes of NIC teaming are to improve availability (by eliminating a single point of failure) and to handle large workloads (by spreading traffic across multiple uplinks). The following NIC teaming options are available on vSS and vDS:

- **Route Based on Originating Virtual Port:** This is the default option. Its load balancing is based on the originating port ID of the virtual machine. When a new virtual machine gets connected to the port group, the virtual switch uses the virtual machine’s port ID to select one of the pNICs from the active uplinks list to direct outbound traffic. Typically, this results in



nearly even virtual machine distribution among the uplinks. If an uplink fails, the switch redirects the traffic from the failed uplink to a surviving uplink.

- **Route Based on IP Hash:** With this load-balancing method, the switch reads the source and destination IP addresses of each outbound packet and uses a hash to determine on which uplink to place the packet. The selection process is resource intensive because the VMkernel examines the header of each Ethernet packet. With this method, a single VM could send packets to different destinations, using different uplinks, which can improve throughput. To support this method, the switch must support IEEE 802.3ad, and you must configure an EtherChannel on the physical switch.
- **Route Based on Source MAC Hash:** This load-balancing method is based on source MAC address. The switch processes each outbound packet header and generates a hash using the least significant bit (LSB) of the MAC address. It uses the hash to decide which active uplink should be used to place the packet.
- **Use Explicit Failover Order:** When this option is chosen, all outbound traffic uses the first uplink that appears in the active uplinks list. If the first uplink fails, the switch redirects traffic from the failed uplink to the second uplink in the list.

The vDS has additional teaming options that are addressed later in this chapter.

### Note

For each of these methods, the virtual switch does not consider the virtual machine's active workload in the load-balancing decision making. A vDS offers a load-based NIC teaming option that is addressed later in this chapter.

Table 3-2 lists some advantages and disadvantages for selecting Route Based on IP Hash.

**Table 3-2** Advantages and Disadvantages of IP Hash NIC Teaming

---

**Advantages**

**Disadvantages**

---

A more even distribution of the load compared to Route Based on Originating Virtual Port and Route Based on Source MAC Hash

Highest resource consumption compared to the other load-balancing algorithms

A potentially higher throughput for virtual machines that communicate with multiple IP addresses

Requires changes on the physical network.

Complex to troubleshoot

---

## Network Security Policies

As an administrator, you can define network security policies to protect network traffic against MAC address impersonation and unwanted port scanning. The security policies of the virtual switches are implemented in Layer 2 of the TCP/IP stack. You can use the network security policies to secure traffic against Layer 2 attacks.

It is important to understand how the virtual machine MAC addresses are defined, since they play a part in the security policies. Each vNIC in a virtual machine has an initial MAC address and an effective MAC address. The initial MAC address is assigned when the adapter is created. It can be modified from outside the guest operating system, but it cannot be changed by the guest operating system. The effective MAC address is typically assigned the same value as the initial MAC address when the adapter is created. The effective MAC address can be changed by the guest OS. The effective MAC address is used when filtering incoming packets. When the guest OS changes the effective MAC address, the adapter receives packets destined for the new effective MAC address. When sending a packet, the guest OS typically places its effective MAC address in the source MAC address field of the Ethernet frame.

A guest OS can send frames with an impersonated source MAC address, which facilitates impersonation and malicious attacks. To guard against this risk, you can leverage security policies on vSS port groups and vDS distributed port groups. There are three available options for the network security policies:

- **Promiscuous Mode:** For a vSS port group, the default value is Reject. By default, the vNIC receives only those frames that match the effective MAC address. If this option is set to Accept, the virtual switch sends all frames on

the wire to the vNIC, allowing virtual machines to receive packets that are not destined for them. This setting allows the use of tools such as tcpdump and Wireshark inside a guest operating system.

- **MAC Address Changes:** For a vSS port group, the default value is Accept. By default, ESXi accepts the effective MAC address change. If this option is set to Reject, the behavior changes such that ESXi does not honor requests to change the effective MAC address to an address that is different from the initial MAC address. Instead, it disables the virtual switch port until the effective MAC address matches the initial MAC address. The guest OS is unaware that the request was not honored.
- **Forged Transmits:** For a vSS port group, the default value is Accept. By default, ESXi does not compare source and effective MAC addresses and does not drop the packet due to a mismatch. If this option is set to Reject, ESXi compares the source and effective MAC addresses and drops the packet if the addresses do not match.

## Note

In a vDS, the default value for each of these three security options is Reject.

## Traffic Shaping Policy

A virtual machine's network bandwidth can be controlled by enabling traffic shaping. Traffic shaping is defined by average bandwidth, peak bandwidth, and burst size. Network traffic shaping can be configured on standard switches or distributed switches, but by default, traffic shaping is disabled. You can apply a traffic shaping policy to each vSS port group, each vDS distributed port, or each vDS distributed port group. On a standard switch, traffic shaping can be configured only for outbound traffic. On a vDS, it can be configured for both outbound and inbound traffic. When you enable traffic shaping for a standard switch or port group, you can configure the following options:

- **Average Bandwidth:** This is the allowed average load for a port. It is the allowed number of kilobits per second for traffic through a port, averaged over time.

- **Peak Bandwidth:** This is the allowed maximum number of kilobits per second of traffic through a port when it is sending or receiving a burst of traffic. This number tops the bandwidth that a port uses when it is using its burst bonus that is configured using the Burst Size parameter. The Peak Bandwidth setting cannot be smaller than the Average Bandwidth setting.
- **Burst Size:** This is the maximum number of kilobytes to allow in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. You can define an average bandwidth allotment for the virtual machines connected to a port group. If the virtual machines have not used the allotted bandwidth, the system may grant a burst bonus, allowing the virtual machines to send some data, limited by the burst size, at a faster rate (up to the peak bandwidth).

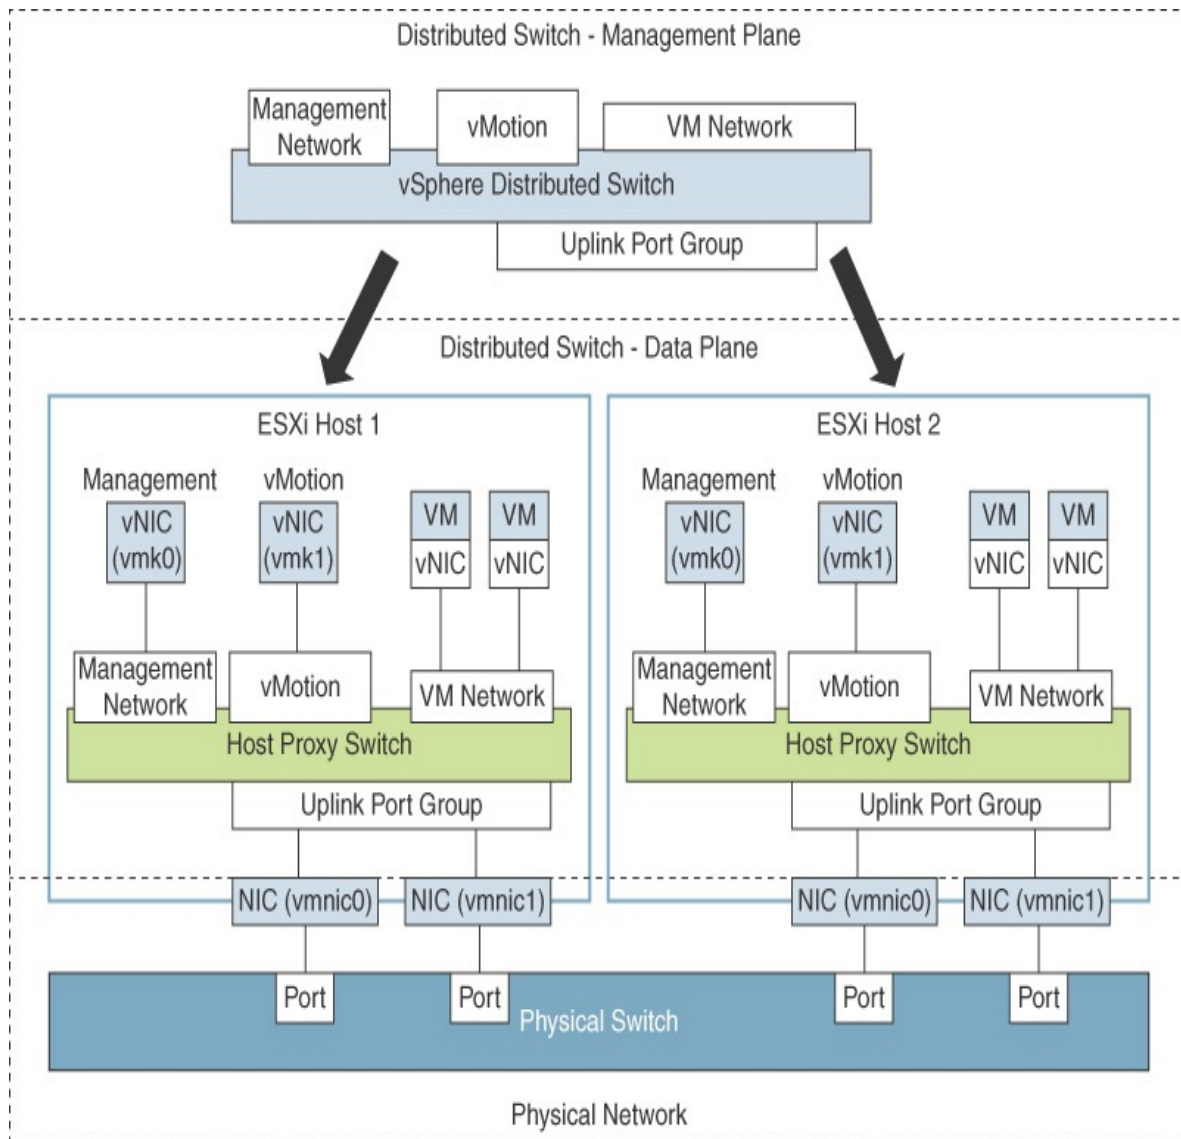
## VLAN Policies

You can apply a VLAN policy to a vSS, such that all port groups on the vSS are associated with a single VLAN. Optionally, you can override the VLAN setting per port group, such that all the virtual machines in the port group are associated with the same VLAN. A major use case for multiple virtual machine port groups is to provide a simple means to place different sets of virtual machines onto separate VLANs. In this case, you should apply a unique VLAN policy to each virtual machine port group.

VLANs are described in the “VLANs” section in this chapter.

## Distributed Virtual Switch (vDS)

A *vSphere Distributed Switch (vDS)* acts as a single virtual switch for all associated hosts in a data center. It provides centralized provisioning, monitoring, and management of virtual networks for associated hosts and virtual machines. You create and configure distributed switches on a vCenter Server system. The vCenter Server propagates the switch configuration to each connected ESXi host in the form of a host proxy switch. The ESXi host provides the data plane for the I/O traffic. The data plane implements the packet switching, filtering, tagging, and other features for the Ethernet packet. The management plane is provided by vCenter Server. [Figure 3-2](#) shows the architecture of a vDS, including its management plane and data plane.



**FIGURE 3-2** vDS Architecture

## Distributed Port Groups

The ports in a vDS, which are called *distributed ports*, are where the VMkernel virtual network adapters and virtual machine vNICs connect. A set of distributed ports is called a *distributed port group*. A vDS provides distributed port groups to simplify the configuration and management of distributed ports. You can apply unique network labels to each distributed port group. You can configure NIC teaming, VLAN, security, traffic shaping, and other policies to a distributed port group, which applies the policies to the underlying distributed ports. When you connect a vNIC to a vDS, you can choose the distributed port group rather

than a specific distributed port.

## Uplink Port Groups

An uplink port group is similar to a distributed port group, except that it is used to manage the physical connections for hosts. An uplink port group, which is defined during the creation of a distributed switch, contains one or more uplinks. An uplink (dvuplink) is where you map the hosts' physical NICs that connect the vDS to the physical network. When you change settings that involve uplink port groups (such as NIC teaming settings), the settings are automatically applied to each host and the appropriate NICs.

## vSS and vDS Comparison



You can define a custom name for each vDS, distributed port group, and uplink port group using the vSphere Client. At the switch level, you can make several more settings and take more actions with a vDS than with a vSS, such as creating private VLANs and configuring port mirroring sessions. vDS offers many advantages over vSS but requires Enterprise Plus licensing. [Table 3-3](#) provides a side-by-side comparison of the features that are available in vSS and vDS.

**Table 3-3** Comparison of vSS and vDS Features

Feature	vSS	vDS
Layer 2 switch	X	X
VLAN segmentation (802.1q tagging)	X	X
IPv6 support	X	X
NIC teaming	X	X
Outbound traffic shaping	X	X
Cisco Discovery Protocol (CDP)	X	X
Inbound traffic shaping		X
VM network port block		X
Private VLANs		X

Load-based NIC teaming	X
Data center–level management	X
Network vMotion	X
Per-port policy settings	X
Port state monitoring	X
NetFlow	X
Port mirroring	X

## vDS Network Policies

On a vDS, you have more flexibility for applying policies that you do with a vSS. You can apply policies to a distributed port group, and those policies then automatically propagate to each distributed port. At the distributed port level, you can override policies applied to the distributed port group and apply unique policies to a distributed port. Likewise, you can apply policies to uplink port groups and override the policies at the uplink port level.

At the distributed port group level, you can control which network policies can be overridden at the distributed port group level. For example, if you configure the distributed port group to allow VLAN policy override at the port level, you can set a VLAN ID directly on a distributed port. Otherwise, you cannot.

As with a vSS, you can apply the following network policies for vDS:

- Teaming and Failover
- Security
- Traffic Shaping
- VLAN

As mentioned earlier, you can apply these policies at the distributed port group and distributed port levels.

With a vDS, you can apply the following additional policies, which are addressed in upcoming sections in this chapter:

- Monitoring
- Traffic Filtering and Marking

- Resources Allocation
- Port Blocking

Compared to a vSS, a vDS provides additional teaming and failover options, which are addressed in the “Load-Based NIC Teaming” and “LACP Support” sections in this chapter. Compared to a vSS, a vDS provides additional traffic shaping options, which are addressed in the “Inbound Traffic Shaping” section in this chapter.

## **Inbound Traffic Shaping**

Distributed virtual switches can do both inbound and outbound traffic shaping, whereas standard virtual switches handle just outbound traffic shaping. Specifically, you can configure the average bandwidth, peak bandwidth, and burst size on a distributed port group and apply it to ingress or egress traffic (inbound or outbound to the distributed port).

## **Port-Blocking Policies**

When needed, you can block traffic to specific distributed ports. For example, if a virtual machine is broadcasting a lot of traffic due to a broken or hacked application and starts consuming a large portion of the network bandwidth, you may want to temporarily block traffic to its distributed port. This is particularly useful if you are a network administrator who has permission to modify the vDS but do not have permission to modify the virtual machine.

As with other policies, to allow individual port blocking, you first need to allow port blocking policy override on the distributed port group. Optionally, you can block all ports in a port group.

## **Load-Based NIC Teaming**

In addition to the NIC teaming options (Originating Virtual Port ID, Source MAC Hash, IP Hash, and Explicit Failover Order) provided by standard switches, a vDS offers load-based teaming. Load-based NIC teaming (or Route Based on Physical NIC Load) checks the current load on the uplinks and takes steps to migrate traffic from overloaded uplinks to less active uplinks. With load-based teaming, each vNIC is initially assigned to an uplink, based on the originating virtual port ID. If an uplink’s load exceeds 75% of its capacity over a 30-second interval, then the port ID of the virtual machine with the highest I/O is



moved to a different uplink.

Load-based teaming offers lower overhead than IP-based load balancing, and it does not require the physical switches to support 802.3ad (EtherChannel/link aggregation).

## Resource Allocation Policy



When you apply a resource allocation policy to a distributed port group, you can assign a network resource pool to the port group. Network resource pools leverage Network I/O Control (NIOC) to allocate resources for network traffic, as explained in the following sections.

### Network I/O Control (NIOC)

NIOC was introduced in vSphere 4.1 when VMware added support for physical NICs with speeds greater than 1 Gbps. It became common for servers to have fewer NICs with multiple types of network traffic sharing the bandwidth. VMware uses NIOC as an additional layer of packet scheduling in the hypervisor to prioritize network traffic.

When you create a vDS in vSphere 7.0, NIOC is enabled by default. Each virtual port has a network port queue where the network scheduler gets the packet and transmits it based on the shares, reservations, and limit settings for each network port. The vDS recognizes the following network traffic types: management, Fault Tolerance, NFS, vSAN, vMotion, vSphere Replication, vSphere Data Protection Backup, and virtual machine traffic. With default settings, NIOC does not directly limit or guarantee network bandwidth to any system traffic type. In other words, the reservation for each traffic type is 0, and the limit is set to unlimited. By default, NIOC provides a higher relative priority to virtual machine traffic, and this priority is applied during periods of network contention. In other words, the default virtual machine traffic shares (100) are twice as many as the shares for the other types of system traffic (50).

You can change the limit for any system traffic type to reflect the maximum bandwidth (in Mbps or Gbps) that the traffic type can consume on a single physical NIC.

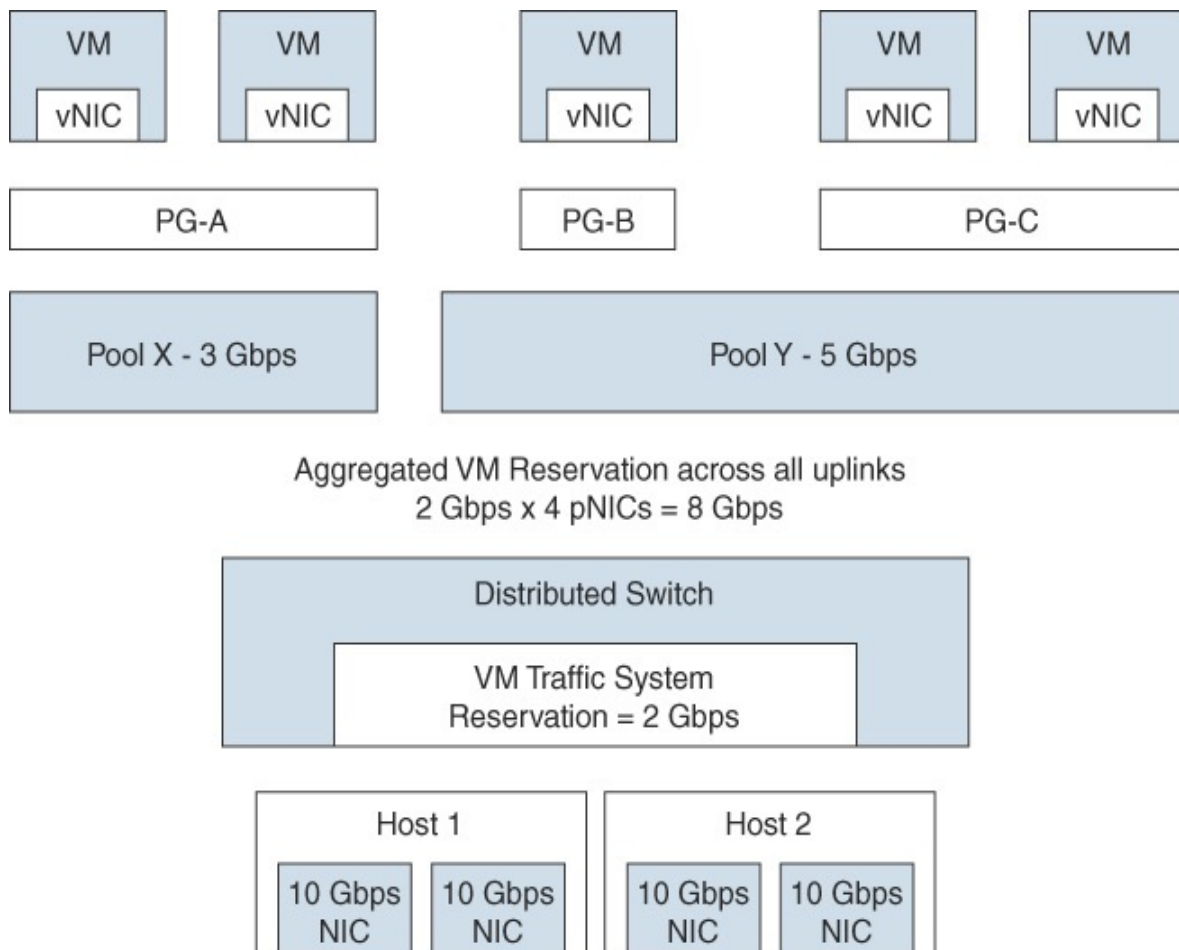
You can change the reservation for any system traffic type to establish a specific bandwidth (in Mbps) that the traffic type is guaranteed on a single physical NIC.

This bandwidth cannot exceed 75% of the bandwidth of the lowest-capacity NIC that is servicing the traffic type. If you reserve resources for virtual machine traffic, you can use network resource pools to delegate the reserved resources.

## Network Resource Pools

A network resource pool is a mechanism that enables you to apply a part of the bandwidth that is reserved for virtual machine system traffic to a set of distributed port groups. By default, no network resource pools exist.

For example, if you reserve 2 Gbps for virtual machine system traffic on a distributed switch with four uplinks, the total aggregated bandwidth available for virtual machine reservation on the switch is 8 Gbps. Each network resource pool can reserve a portion of the 8 Gbps capacity. If you create a network pool named PG-A and set it to reserve 3 Gbps, then you have 5 Gbps that you can use for other network pools, as illustrated in [Figure 3-3](#).



**FIGURE 3-3** Network Resource Pools Example

The bandwidth quota that is dedicated to a network resource pool is shared among the distributed port groups associated with the pool and is applied to the virtual machines connected to the associated distributed ports.

With NIOC Version 3, you can set shares, reservations, and limits on virtual machine vNICs, much as you can for a virtual machine's CPU and memory. To fulfill the network reservations (guarantees), NIOC uses a traffic placement engine. For virtual machines that have network reservation configured, the distributed switch attempts to place the virtual network adapter traffic on a physical adapter that can supply the required bandwidth and is in the scope of the active teaming policy. The total bandwidth reservation for virtual machines in a network resource pool cannot exceed the reservation of the pool.

The actual limit and reservation applied to a virtual network adapter may be impacted by a distributed port group's traffic shaping policy. For example, if a vNIC's limit is set to 300 Mbps, and the average bandwidth and peak bandwidth applied by traffic shaping is 200 Mbps, then the effective limit becomes 200 Mbps.

To meet the network guarantees (reservation) for a virtual machine, vSphere implements admission control at the host and cluster levels, based on bandwidth reservation and teaming policy. A distributed switch applies bandwidth admission control by using NIOC to verify the following:

- A physical adapter is available that can supply the guaranteed bandwidth to the VM network adapters.
- The virtual network adapter's reservation is less than the free quota in the network resource pool.

With NIOC Version 3, a vSphere DRS cluster places the virtual machine on a host that can fulfill the reserved bandwidth for the virtual machine, according to the active teaming policy. In the following situations, vSphere DRS migrates a virtual machine to another host to satisfy the virtual machine's bandwidth reservation:

- The reservation is changed to a value that the initial host can no longer satisfy.
- A physical adapter that carries traffic from the virtual machine is offline.

In the event of a host failure or isolation, vSphere HA powers on a virtual

machine on another host in the cluster, according to the bandwidth reservation and teaming policy. To use network admission control in a vSphere DRS or vSphere HA cluster, you should perform the following tasks:

- Allocate bandwidth for the virtual machine system traffic.
- Configure the bandwidth requirements of virtual machines connected to the distributed switch.

## NetFlow and Monitoring Policy

Cisco introduced the *NetFlow* feature in the mid-1990s, and eventually it became an open source feature controlled by the IETF. VMware added NetFlow Version 5 to the vSphere 5.0 release. NetFlow is a switch feature that collects IP network traffic as it enters or exits an interface. NetFlow data provides an overview of traffic flows, based on the network source and destination. The main purpose is to collect IP network traffic that is flowing through a vDS. The current supported version for the vDS is NetFlow Version 10, which is identified as IPFIX.

NetFlow allows administrators to forward network flows to a centralized NetFlow collector. These flows can include virtual machine-to-virtual machine flows, virtual machine-to-physical device flows, and physical device-to-virtual machine flows. The collector can be a physical device on the network or a virtual machine. The collector gathers the flows to enable network analysis and troubleshooting. You can set a sampling rate to collect network flows. For example, if you set the sampling rate to 4, NetFlow samples one packet and then drops (skips) the next four packets. If the sampling rate is zero, NetFlow samples every packet (that is, skips zero packets).

By default, NetFlow is turned off on a vDS. If you enable it, you can set the Collector IP Address, Collector Port, and other settings. To collect data from only the virtual machines on the same host, you can enable Process Internal Flows Only.

After configuring NetFlow on a vDS, you can configure monitoring policies on vDS port groups and ports.

## Traffic Filtering and Marking Policy

With a vDS, you can implement a traffic filtering and marking policy to protect the virtual network from unwanted traffic and security attacks. The traffic

filtering in a vDS uses simple filtering rules, much like the filtering in many physical switches. As an example, you could create a simple Permit rule that allows the specified traffic while blocking everything else. Or you could create a Deny rule that blocks the specified traffic while allowing everything else.

A marking policy allows you to mark traffic with a priority tag that is used during times of contention on a physical switch. In essence, it allows you to apply a tag to the Ethernet header or IP header as the Ethernet frame enters and exits your virtual switch. You can mark the traffic with a Class of Service (CoS) tag in a Layer 2 Ethernet header. Or you could mark the traffic with a Differentiated Service Code Point (DSCP) tag in a Layer 3 IP header. Higher tagged packets move to the front of the queue on a physical Ethernet switch during times of contention, and lower or untagged packets are sent using best-effort service.

With a vDS, you can apply filters on data traveling between the vNIC and the distributed port and between the uplink port and the physical NIC. These are the major steps in the process:

- Step 1.** Enable traffic filtering and marking on the distributed port group or uplink port group.
- Step 2.** Mark traffic on the distributed port group or uplink port group.
- Step 3.** Configure filters on the distributed port group or uplink port group.
- Step 4.** Define traffic rules for distributed port group or uplink port group

When marking traffic, you can create a rule in which you configure qualifiers to identify the data to be tagged and set Action to Tag. You can use system traffic qualifiers (such as NFS), MAC traffic qualifiers (to match by MAC address and VLAN ID), and IP traffic qualifiers (to match by IP version and IP address).

For example, to mark voice over IP (VoIP) traffic whose source IP is 192.168.2.0, you can create a rule that specifies the following parameters:

- **Action:** Tag
- **DSCP Value:** 26
- **Traffic Direction:** Egress
- **Traffic Qualifiers:** IP Qualifier
- **Protocol:** UDP

- **Destination Port:** 5060
- **Source Address:** IP address matches 192.168.2.0 with prefix length 24

## vDS Settings and Features

The following sections provide details on other vDS settings and features.

### Private VLANs

*Private VLANs (PVLANS)* are an extension of the VLAN standard that is not double encapsulated but that allows a VLAN to effectively be subdivided into other VLANs. This is useful for a hosting provider that has run out of VLANs or in any environment where 4094 VLANs are not enough.

A VLAN that is to be subdivided becomes known as the primary private VLAN. This primary PVLAN is then carved up into one or multiple secondary PVLANS that exist only within the primary. When a virtual machine or VMkernel port sends a packet, that packet is tagged at the distributed port group level on the vDS. Because this is not double encapsulation, packets travel with only one VLAN tag at a time. However, physical switches could be confused by seeing MAC addresses tagged with more than one VLAN tag, unless the physical switches are PVLAN aware and have their PVLAN tables configured appropriately. If the physical network is configured correctly, it identifies that the secondary PVLAN exists as part of the primary.

There are three different types of secondary PVLANS:

- **Promiscuous:** Nodes in a promiscuous secondary PVLAN can communicate with other nodes in the same PVLAN and all nodes in any secondary (community or isolated) PVLANS in the same primary PVLAN.
- **Community:** Nodes in a community secondary PVLAN can communicate with other nodes in the same PVLAN and nodes in a promiscuous secondary PVLAN in the same primary PVLAN. Nodes in a community secondary PVLAN cannot communicate with nodes in other community or isolated secondary PVLANS.
- **Isolated:** Nodes in an isolated secondary PVLAN can communicate with nodes in a promiscuous secondary PVLAN in the same primary PVLAN. Nodes in an isolated secondary PVLAN cannot communicate with other nodes in the same PVLAN or with nodes in other community or isolated

secondary PVLANs.

## **Data Center–Level Management**

The vSphere Client provides a single pane of glass for provisioning, configuring, and managing a vDS, which can be used by up to 2500 ESXi servers. At the data center level, you can perform many network management tasks, such as the following.

- Create a vDS
- Attach hosts to the vDS
- Create distributed port groups
- Assign policies to port groups
- Migrate virtual machines and VMkernel virtual network adapters to the vDS
- Monitor alerts, tasks, and events
- Monitor port state
- Manage network resources

## **Port State Monitoring**

In the vSphere Client, you can view details for each vDS port, such as the connected virtual machine, the runtime MAC address, and the link state (link up, link down, blocked, or information unavailable).

## **Port State with vMotion**

A virtual machine connected to a vDS maintains its virtual switch connection when migrating to a different ESXi host. From a virtual networking perspective, the virtual machine has not moved. It remains connected to the same distributed port, with no disruption to the policies and statistics that are tied to that distributed port.

When you perform a migration of a virtual machine connected to a vSS, the virtual machine is attached to a different port on a different virtual switch. The policies and statistics associated with the original port are not carried with the virtual machine. The virtual machine is subject to the policies associated with the destination port group and virtual switch.

## Port Mirroring



Port mirroring allows administrators to duplicate everything that is happening on one distributed port to then be visible on another distributed port. With port mirroring, you can do granular per-port network analysis by gathering the entire traffic stream coming into and going out a monitored port. To get started, you create a port mirroring session in which you identify the session type, session properties, traffic source, and destination.

For the session type, you can specify one of the following options:

- **Distributed Port Monitoring:** Mirrors packets from a set of distributed ports to other distributed port groups.
- **Remote Mirroring Source:** Mirrors packets from a set of distributed ports to specific uplinks.
- **Remote Mirroring Destination:** Mirrors packets from a set of VLANs to distributed ports.
- **Encapsulated Remote Mirroring (L3) Source:** Mirrors packets from a set of distributed ports to the IP address of a remote agent.

The session properties are dependent on the session type and include the following settings:

- **Name:** Uniquely identifies the session.
- **Status:** Enables or disables the session.
- **Description:** Describes the session.
- **Sampling Rate:** Sets the rate at which packets are sampled.
- **Normal I/O on Destination Ports:** Available only for distributed port and uplink destinations. You can disable this option to allow mirrored traffic out on destination ports but disallow mirrored traffic in on destination ports.
- **Mirrored Packet Length:** Limits the size of mirrored frames.
- **Traffic Source:** Identifies the source of the traffic.

Typically, the mirrored data is sent to a network sniffer or an intrusion detection



system.

## Port Binding and Allocation

Virtual network adapters (for virtual machines and the VMkernel) must connect to virtual switch ports to communicate on a network. Port binding determines how virtual machines are bound, or attached, to virtual switch ports. To control port assignment, you can choose to use static binding or ephemeral binding:

- **Static binding:** With static binding (which is the default), when a vNIC attaches to a port in a distributed port group, the connection is static, which means the virtual machine remains attached to the port, regardless of the power state of the virtual machine. This binding is performed and controlled by vCenter Server.
- **Ephemeral binding:** Ephemeral means there is no binding of the vNIC to a specific virtual switch port. With this setting, virtual switch ports are created and deleted on demand by the host. At any moment, the number of ports for an ephemeral distributed port group is equivalent to the number of running vNICs connected to the port group.

### Note

Having an available ephemeral port group is useful in cases where vCenter Server is down and you need to assign a virtual machine to a port group. For example, if a vCenter Server Appliance (VCSA) device is connected to a distributed port group with static binding, you may fail to reconnect the VCSA device to the network after restoring the VCSA device because vCenter is required to assign the port. In this case, you should be successful in connecting the restored VCSA device to the ephemeral distributed port group because the ESXi host will assign the port.

You can control the number of ports and the port allocation in a distributed port group. The port allocation can be fixed or elastic:

- **Elastic:** By default, ports in the port group are created and removed on demand. For example, if the port group is configured for eight ports and elastic port allocation, you can connect nine vNICs to the port group because when you connect the eighth vNIC to the port group, eight more

ports are automatically added.

- **Fixed:** The number of ports in the port group is static. Ports are not automatically created or removed. For example, if the port group is configured for eight ports and fixed port allocation, you cannot connect 9nine vNICs to the port group. When you attempt to connect the ninth vNIC to the port group, you get the error “no free port is available.”

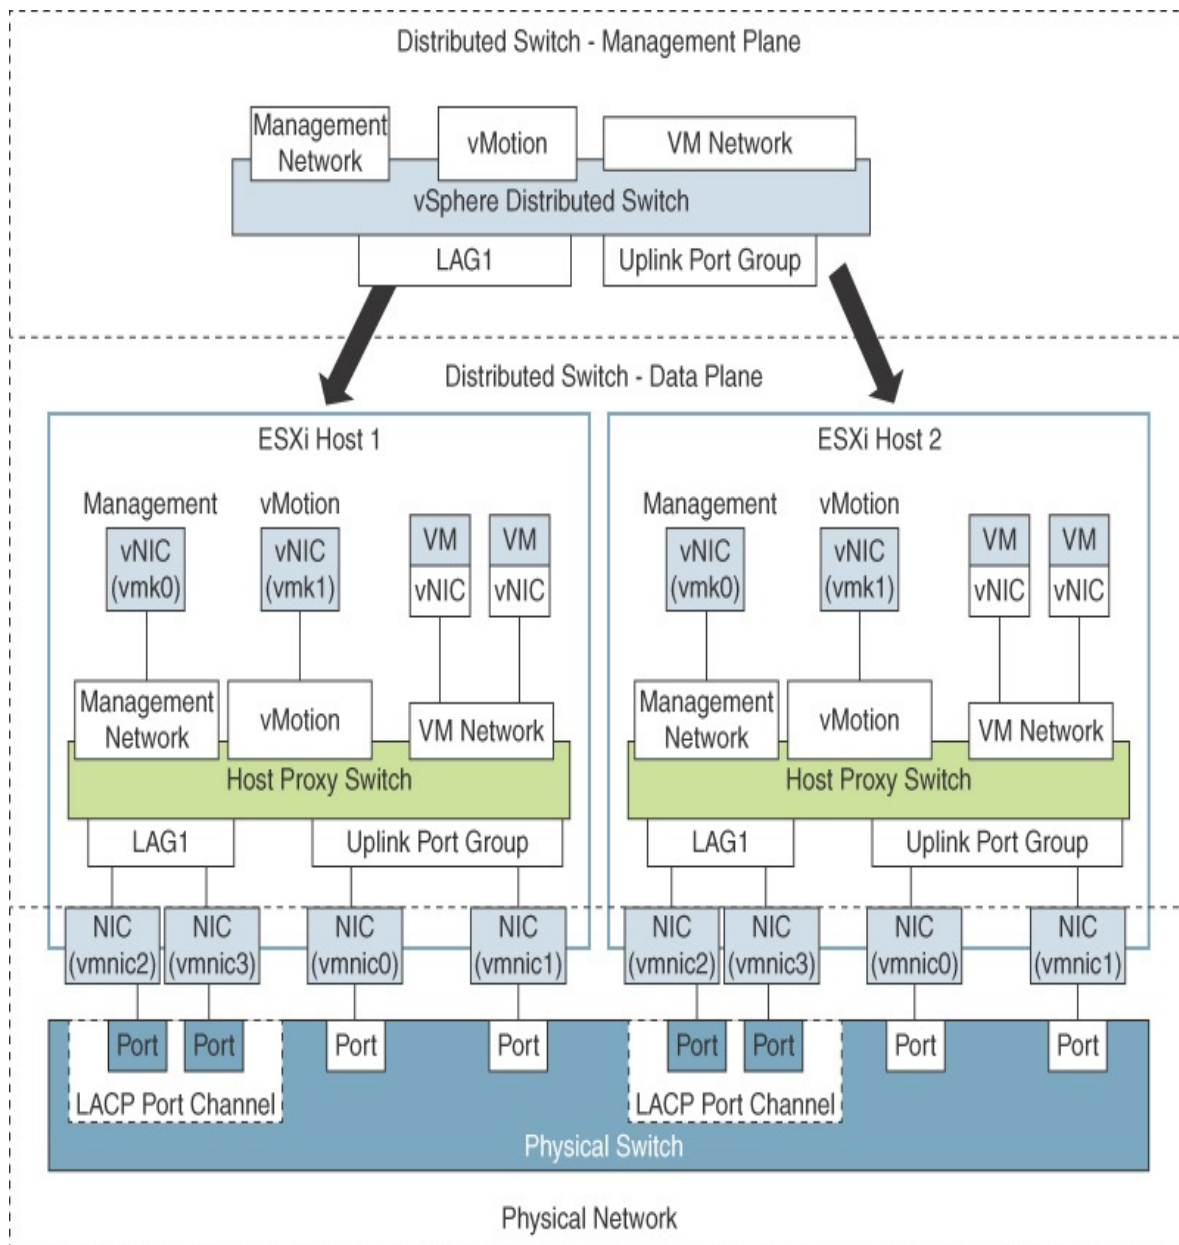
In vSphere 7.0, the default settings for a distributed port group are static binding, elastic port allocation, and eight ports.

In the past, the ephemeral setting seemed like the easiest way to go because it required the least administrative effort to address an ever-growing environment. That changed in vSphere 5.1, when static port binding became “elastic” by default.

## LACP Support

In vSphere 7.0, a vDS supports LACP. This means you can connect ESXi hosts to physical switches by using dynamic link aggregation. You can create multiple link aggregation groups (LAGs) on a distributed switch to aggregate the bandwidth of physical NICs on ESXi hosts that are connected to LACP port channels. This enables you to increase the network bandwidth, redundancy, and load balancing to the port groups. You need to configure each LAG with two or more ports and connect physical NICs to the ports. Within a LAG, the ports are teamed, such that the network traffic is load balanced between the ports using an LACP hashing algorithm.

For each LAG on a vDS, a LAG object is created on each associated host proxy switch. The same number of ports that you assign to the LAG on the vDS are assigned to the corresponding LAG on the proxy switch on each host, as illustrated in [Figure 3-4](#).



**FIGURE 3-4** LACP Configuration with a vDS

On a host proxy switch, a physical NIC can connect to just one LAG port. On a distributed switch, a LAG port can connect to multiple physical NICs from different hosts. The physical NICs that are connected to the LAG ports must connect to links that participate in an LACP port channel on the physical switch.

A vDS supports up to 64 LAGs. A host supports up to 32 LAGs. The LAG configuration for a vDS may be limited by the physical network configuration.

The following limitations apply for LACP support in on a vDS:

**Key  
Topic**

- LACP support is not compatible with software iSCSI port binding. (iSCSI multipathing is supported if iSCSI port binding is not used.)
- LACP support settings do not exist in host profiles.
- LACP is not supported with nested ESXi.
- LACP support does not work with the ESXi dump collector.
- LACP control packets (LACPDU) are not mirrored by vDS port mirroring.
- The vDS teaming and failover health check (see the following section) does not support LAG ports. (LAG port connectivity is checked by LACP.)
- Enhanced LACP support works only when a single LAG handles traffic per distributed port group.

## vDS Health Check

By default, vDS health checks are disabled, but you can enable them to check for several specific network inconsistencies. A health check identifies problems but does not attempt to fix the problems. You can use vDS health checks to regularly examine certain settings on distributed and physical switches to identify common errors in the networking configuration. The default interval between health checks is 1 minute. Health checks generate MAC addresses to use during testing. Depending on the selected options, the number of generated MAC addresses can be significant. If you disable health checking, the MAC addresses age out of your environment. For this reason, you may want to enable health checks temporarily.

[Table 3-4](#) provides a summary of the vDS health checks available.

**Table 3-4** vDS Health Checks

<b>Health Check</b>	<b>Required vDS Configuration</b>
Checks whether the VLAN trunk ranges on the distributed switch match the trunk port configuration on the connected physical switch ports.	At least two active physical NICs

Checks for matching MTU settings on the distributed switch, the physical network adapter, and the physical switch ports.	At least two active physical NICs
--	-----------------------------------

---

Checks whether the virtual switch teaming policy matches the physical switch port-channel settings.	At least two active physical NICs and two hosts
---	---

---

The potential values for VLAN Health Status, MTU Health Status, and Teaming and Failover Health Status include Normal, Warning, and Unknown. If the VLAN trunk ranges configured on a distributed switch match the trunk ranges on a physical switch, VLAN Health Status is set to Normal. If the MTU settings configured on the distributed switch match the trunk ranges on the physical switch, MTU Health Status is set to Normal. If the physical switch ports participate in an EtherChannel and the distributed port group teaming policy is IP hash, Teaming and Failover Health Status is set to Normal. Also, if the distributed port group teaming policy is set to Port Based, MAC Hash, or Explicit Failover, and physical switch ports do not participate in an EtherChannel, Teaming and Failover Health Status is set to Normal.

## Other vSphere Networking Features

The following sections provide details on other networking features supported in vSphere 7.0 that are not covered earlier in this chapter.

### Multicast Filtering Mode

vDS 6.0.0 and later supports basic multicast filtering mode and multicast snooping mode.

### Basic Multicast Filtering

In basic multicast filtering mode, a virtual switch (vSS or vDS) forwards multicast traffic for virtual machines according to the destination MAC address of the multicast group, which is based on the MAC address generated from the last 23 bits of the IPv4 address of the group. When a guest operating system joins a multicast group, it pushes the group multicast MAC address to the network through the virtual switch. The virtual switch records the port to destination multicast MAC address mapping in a local forwarding table. The

switch does not interpret IGMP messages. Instead, it sends messages directly to the local multicast router, which interprets the messages and takes the appropriate action.

The basic mode has the following restrictions:

- A virtual machine may receive packets from the wrong groups because the switch forwards packets based on the multicast group's destination MAC address, which can potentially be mapped to up to 32 IP multicast groups.
- Due to a limitation of the forwarding model, a virtual machine that is subscribed for traffic from more than 32 multicast MAC addresses may receive packets from the wrong groups.
- The switch does not filter packets based on source address, as defined in IGMP Version 3.

## **Multicast Snooping**

vDS 6.0.0 and later support multicast snooping, which forwards multicast traffic based on the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) messages from virtual machines. Multicast snooping supports IGMP Version 1, IGMP Version 2, and IGMP Version 3 for IPv4 multicast group addresses and MLD Version 1 and MLD Version 2 for IPv6 multicast group addresses. The switch dynamically detects when a virtual machine sends a packet containing IGMP or MLD membership information through a switch port and creates a record about the destination IP address of the group. The switch automatically removes the entry for a group from the lookup records whenever a virtual machine does not renew its membership on time.

When using vDS multicast snooping mode, a virtual machine can receive multicast traffic on a single switch port from up to 256 groups and 10 sources.

## **Discovery Protocol**

Switch discovery protocols help vSphere administrators identify the physical switch ports to which a vSS or vDS are connected. Cisco Discovery Protocol (CDP) support was introduced with ESX 3.x. CDP is available for standard switches and distributed switches that are connected to Cisco physical switches. Link Layer Discovery Protocol (LLDP) is supported in vSphere 5.0 and later for vDS (5.0.0 and later) but not for vSS.

CDP enables you to determine which Cisco switch port is connected to a vSS or a vDS. When CDP is enabled for a vDS, you can view the properties of the Cisco switch, such as device ID, software version, and timeout. You can set CDP Type to any of the following values:

- **Listen:** ESXi collects and displays details about the associated Cisco switch port but does not share information about the vDS with the Cisco switch.
- **Advertise:** ESXi shares information about the vDS with the Cisco switch but does not collect or display details about the associated Cisco switch port.
- **Both:** ESXi collects and displays information about the associated Cisco switch port and shares information about the vDS with the Cisco switch.

LLDP enables you to determine which physical switch port is connected to a vDS. When LLDP is enabled for a vDS, you can view the properties of the physical switch, such as chassis ID, system name and description, and device capabilities. You can set the CDP Type to Listen, Advertise, or Both, as previously explained for CDP.

In the vSphere client, to view the physical switch information, you can select a host, navigate to Configure > Networking > Physical Adapters, select an adapter, and choose either the CDP or LLDP tab in the details pane.

## TCP Segmentation Offload

TCP Segmentation Offload (TSO, which is referred to as large segment offload or large send offload [LSO] in VMXNET3's latest attributes) allows the segmentation of traffic to be offloaded from the virtual machine and from the VMkernel to the physical NIC. The network adapter separates large frames (up to 64 KB) into MTU-sized frames. One of the jobs of the TCP layer is to segment the data stream to the size of the Ethernet frame. If the frame size is 1500 bytes and an application is transmitting 8 Kbps, it is the job of the TCP layer to segment it down to 1500-byte packets. This occurs in the guest operating system. However, if a physical NIC supports TSO, the adapter can do the work instead, resulting in less overhead and better performance. To use TSO, it must be enabled along the data path, including the physical NIC, vNIC, VMkernel, and guest OS. By default, it is enabled in the VMkernel, the VMXNET2 vNIC, and the VMXNET 3 vNIC.

## DirectPath I/O



DirectPath I/O allows a virtual machine to access physical PCI functions on platforms that have an I/O memory management unit (IOMMU). You can enable DirectPath I/O passthrough for a physical NIC on an ESXi host to enable efficient resource usage and to improve performance. After enabling DirectPath I/O on the physical NIC on a host, you can assign it to a virtual machine, allowing the guest OS to use the NIC directly and bypassing the virtual switches.

## Note

Do not enable DirectPath I/O passthrough for the USB controller for an ESXi host that is configured to boot from a USB device or an SD card attached to a USB channel.

The following features are not available for DirectPath-enabled virtual machines:

- Hot addition and removal of virtual devices
- Suspend and resume
- Record and replay
- Fault tolerance
- High availability
- DRS (The virtual machine can be part of a cluster but cannot migrate across hosts.)
- Snapshots

## Single Root I/O Virtualization (SR-IOV)

**Single Root I/O Virtualization (SR-IOV)** is a feature that allows a single Peripheral Component Interconnect Express (PCIe) device to appear as multiple devices to the hypervisor (ESXi) or to a virtual machine's guest operating system. SR-IOV is useful for supporting an application in a guest OS that is sensitive to network latency.

SR-IOV-enabled devices provide virtual functions (VFs) to the hypervisor or guest operating system. VFs are lightweight PCIe functions that support data exchange with limited support for configuration changes. The virtual machine can take advantage of SR-IOV by exchanging Ethernet frames directly with the



physical network adapter, bypassing the VMkernel. This improves network performance by reducing latency and improves CPU efficiency. The downside of this approach is that the environment must meet certain requirements, and several features are not supported.

Although a virtual switch (vSS or vDS) does not handle the network traffic of an SR-IOV-enabled virtual machine, you can control the assigned VFs by using switch configuration policies at the port group level or port level. For example, the VF driver (guest OS) cannot modify the MAC address if it is not allowed in the security policy for the port group or port.

To use SR-IOV in vSphere 7.0, your environment must meet the requirements described in [Table 3-5](#).

**Table 3-5** SR-IOV Requirements

<b>Component</b>	<b>Requirements</b>
Physical host	<p>Must use an Intel or AMD processor.</p> <p>Must support IOMMU and SR-IOV.</p> <p>IOMMU and SR-IOV must be enabled in the BIOS.</p>
Physical network adapter	<p>Must be supported by the server vendor for use with the host system and SR-IOV for the specific ESXi release.</p> <p>SR-IOV must be enabled in the firmware.</p> <p>Must use MSI-X interrupts.</p>
Physical function (PF) driver in ESXi	<p>Must be certified by VMware.</p> <p>Must be installed on the ESXi host, which may require custom installation.</p>
Guest OS	<p>Must be supported by the NIC vendor for the specific ESXi release.</p>
Virtual function (VF) driver in guest OS	<p>Must be compatible with the NIC and supported on the guest OS release.</p>

Must be Microsoft WLK or WHCK certified for Windows virtual machines.

Must be installed on the operating system and may require custom installation.

---

The following features are not available for SR-IOV-enabled virtual machines, and attempts to configure these features may result in unexpected behavior:



- vSphere vMotion
- Storage vMotion
- vShield
- NetFlow
- VXLAN Virtual Wire
- vSphere High Availability
- vSphere Fault Tolerance
- vSphere DRS
- vSphere DPM
- Virtual machine suspend and resume
- Virtual machine snapshots
- MAC-based VLAN for passthrough virtual functions
- Hot addition and removal of virtual devices, memory, and vCPU
- Participation in a cluster environment
- Network statistics for a virtual machine NIC using SR-IOV passthrough

The following NICs are supported for virtual machines configured with SR-IOV:

- Products based on the Intel 82599ES 10 Gigabit Ethernet Controller Family (Niantic)

- Products based on the Intel Ethernet Controller X540 Family (Twinville)
- Products based on the Intel Ethernet Controller X710 Family (Fortville)
- Products based on the Intel Ethernet Controller XL170 Family (Fortville)
- Emulex OneConnect (BE3)

Each NIC must have SR-IOV-supported drivers and may require SR-IOV to be enabled on the firmware.

An SR-IOV NIC can operate in one of three modes in ESXi:

- **Non SR-IOV Mode:** The NIC is not used to provide VFs to virtual machines.
- **SR-IOV Only Mode:** The NIC provides VFs to virtual machines but does not back other virtual machine traffic. In the vSphere Client, the NIC appears in a separate list (External SR-IOV Adapters) in the switch topology page.
- **Mixed Mode.** The NIC services virtual machines with and without SR-IOV.

DirectPath I/O and SR-IOV offer similar performance benefits and trade-offs, but you use them to accomplish different goals. You can use SR-IOV in workloads with extremely high packet rates or very low latency requirements, where you want multiple virtual machines to share the same physical NIC (with the same physical function). With DirectPath I/O, you can map only one physical NIC to one virtual machine.

## VMkernel Networking and TCP/IP Stacks

The VMkernel networking layer provides connectivity for the hypervisor and handles system services traffic, such as management, vMotion, IP-based storage, provisioning, Fault Tolerance logging, vSphere Replication, vSphere Replication NFC, and vSAN. You can create multiple VMkernel virtual network adapters to support these services. For each VMkernel adapter, you can select which system service it supports. Each VMkernel adapter requires IP configuration and virtual switch configuration. You can choose to configure multiple system services to share networks, or you can configure a separate network for each system service. In addition to using separate virtual and physical network infrastructure to support each system service, you can use separate VMkernel TCP/IP stacks.

The VMkernel provides multiple TCP/IP stacks that you can use to isolate the system services traffic. You can use the preexisting stacks (default, vMotion, and Provisioning) and create custom stacks, as described here:

- **Default TCP/IP stack:** This stack provides networking support for management traffic and for all VMkernel traffic types. You can choose to use just this stack and ignore the other available stacks.
- **vMotion TCP/IP stack:** This stack supports the traffic for vMotion. You can use this stack to improve the isolation for the vMotion traffic. After you create a VMkernel adapter on this stack, you can no longer use the default stack for vMotion. (vMotion is automatically disabled on VMkernel adapters on the default stack.)
- **Provisioning TCP/IP stack:** This stack supports the traffic for virtual machine cold migration, cloning, and snapshot migration. It also supports the Network File Copy (NFC) traffic used for cloning virtual disks with long-distance vMotion. You can use this stack to isolate provisioning traffic and place it on a separate gateway. After you create a VMkernel adapter on this stack, you can no longer use the default stack for provisioning. (Provisioning is automatically disabled on VMkernel adapters on the default stack.)
- **Custom TCP/IP stack:** You can create custom stacks to isolate the networking of custom applications.

When you create a VMkernel virtual network adapter, you should configure the settings described in [Table 3-6](#).

**Table 3-6** VMkernel Adapter Settings

Setting	Description
IP Settings	Provide IPv4 or IPv6 configuration details, such as IPv4 address, mask, and gateway.
MTU	Set this option as described in the “MTU” section in this chapter.
TCP/IP Stack	Select a standard or custom stack, as described in this section.
Available Services	Select which of the following system services to enable for the adapter:

- Managementv
- Motion
- IP-based storage
- Provisioning
- Fault Tolerance logging
- vSphere Replication
- vSphere Replication NFC
- vSAN

---

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15, “Final Preparation,”](#) and the exam simulation questions on the companion website.

### Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 3-7](#) lists these key topics and the page number on which each is found.



**Table 3-7** Key Topics for [Chapter 3](#)

Key Topic Element	Description	Page Number
List	vSS networking policies	<a href="#">98</a>
Section	vSS and vDS comparison	<a href="#">103</a>
Section	Resource allocation policy	<a href="#">105</a>
Section	Port mirroring	<a href="#">111</a>
List	LACP with vDS limitations	<a href="#">115</a>
List	SR-IOV limitations	<a href="#">120</a>

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

- [EtherChannel](#)
- [virtual LAN \(VLAN\)](#)
- [network resource pool](#)
- [NetFlow](#)
- [Single Root I/O Virtualization \(SR-IOV\)](#)
- [vSphere Distributed Switch \(vDS\)](#)
- [private VLAN \(PVLAN\)](#)
- [port mirroring](#)

## Review Questions

1. You are configuring traffic shaping policies for your vSS. Which one

of the following is not an available setting?

- a.** Peak Bandwidth
  - b.** Minimum Bandwidth
  - c.** Average Bandwidth
  - d.** Burst Size
2. You want to implement network security policies at the lowest available level. Which approach should you choose?
- a.** Use standard switches and apply policies on the individual ports.
  - b.** Use standard switches and choose override on the port groups.
  - c.** Use distributed switches and apply policies on the individual ports.
  - d.** Use distributed switches and apply policies on the distributed port groups.
3. You created a distributed port group with default settings. Which of the following statements is true?
- a.** The port group is set for fixed allocation and 8 ports.
  - b.** The port group is set for elastic allocation and 8 ports.
  - c.** The port group is set for fixed allocation and 16 ports.
  - d.** The port group is set for elastic allocation and 16 ports.
4. You want to address an application's sensitivity to network latency by implementing a technology that presents a physical device with multiple virtual functions. Which one of the following should you implement?
- a.** DirectPath I/O
  - b.** TSO
  - c.** LACP
  - d.** SR-IOV
5. You want to control the use of NetFlow within your distributed switch. Which type of network policy should you implement?

- a.** Traffic shaping
- b.** Monitoring
- c.** Resource allocation
- d.** Filtering and marking



## Chapter 4

# Clusters and High Availability

This chapter covers the following topics:

- [Cluster Concepts and Overview](#)
- [Distributed Resources Scheduler \(DRS\)](#)
- [vSphere High Availability \(HA\)](#)
- [Other Resource Management and Availability Features](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.6, 1.6.1, 1.6.2, 1.6.3, 1.6.4, 1.6.4.1, 4.5, 4.6, 5.1, 5.1.1, 5.2, 7.5, and 7.11.5.

This chapters provides details on clusters and high availability in vSphere 7.0.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [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 and Review Questions.”](#)

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

Foundation Topics Section	Questions
Cluster Concepts and Overview	1
Distributed Resource Scheduler (DRS)	2–4
vSphere High Availability (HA)	5–7

- 1.** You are configuring EVC Mode in a vSphere cluster that uses Intel hardware. Which of the following values should you choose to set the EVC Mode to the lowest level that includes the SSE4.2 instruction set?
  - a.** Merom
  - b.** Penryn
  - c.** Nehalem
  - d.** Westmere
  
- 2.** In vSphere 7.0, you want to configure the DRS migration threshold such that it is at the minimum level at which the virtual machine happiness is considered. Which of the following values should you choose?
  - a.** Level 1
  - b.** Level 2
  - c.** Level 3
  - d.** Level 4
  - e.** Level 5
  
- 3.** Which of the following is not a good use for resource pools in DRS?
  - a.** To delegate control and management
  - b.** To impact the use of network resources
  - c.** To impact the use of CPU resources
  - d.** To impact the use of memory resources
  
- 4.** You need your resource pool to use a two-pass algorithm to allocate reservations. In the second pass, excess pool reservation is allocated proportionally to virtual machines (limited by virtual machine size). Which step should you take?
  - a.** Ensure that vSphere 6.7 or higher is used.
  - b.** Ensure that vSphere 7.0 or higher is used.
  - c.** Enable scalable shares.

- d.** Enable expandable reservations.
- 5.** You are configuring vSphere HA in a cluster. You want to configure the cluster to use a specific host as a target for failovers. Which setting should you use?
  - a.** Host Failures Cluster Tolerates
  - b.** Define Host Failover Capacity By set to Cluster Resource Percentage
  - c.** Define Host Failover Capacity By set to Slot Policy (Powered-on VMs)
  - d.** Define Host Failover Capacity By set to Dedicated Failover Hosts
  - e.** Define Host Failover Capacity By set to Disabled
- 6.** You are enabling VM Monitoring in a vSphere HA cluster. You want to set the monitoring level such that its failure interval is 60 seconds. Which of the following options should you choose?
  - a.** High
  - b.** Medium
  - c.** Low
  - d.** Normal
- 7.** You are configuring Virtual Machine Component Protection (VMCP) in a vSphere HA cluster. Which of the following statements is true?
  - a.** For PDL and APD failures, you can control the restart policy for virtual machines by setting it to Conservative or Aggressive.
  - b.** For PDL failures, you can control the restart policy for virtual machines by setting it to Conservative or Aggressive.
  - c.** For APD failures, you can control the restart policy for virtual machines by setting it to Conservative or Aggressive.
  - d.** For PDL and APD failures, you cannot control the restart policy for virtual machines.
- 8.** You want to use Predictive DRS. What is the minimum vSphere version you need?

- a. vSphere 6.0
  - b. vSphere 6.5
  - c. vSphere 6.7
  - d. vSphere 7.0
9. You are configuring vSphere Fault Tolerance (FT) in a vSphere 7.0 environment. What is the maximum number of virtual CPUs you can use with an FT-protected virtual machine?
- a. One
  - b. Two
  - c. Four
  - d. Eight
10. You are concerned about service availability for your vCenter Server. Which of the following statements is true?
- a. If a vCenter service fails, VMware Service Lifecycle Manager restarts it.
  - b. If a vCenter service fails, VMware Lifecycle Manager restarts it.
  - c. If a vCenter service fails, vCenter Server HA restarts it.
  - d. VMware Service Lifecycle Manager is a part of the PSC.

## Foundation Topics

### Cluster Concepts and Overview

A vSphere cluster is a set of ESXi hosts that are intended to work together as a unit. When you add a host to a cluster, the host's resources become part of the cluster's resources. vCenter Server manages the resources of all hosts in a cluster as one unit. In addition to creating a cluster, assigning a name, and adding ESXi objects, you can enable and configure features on a cluster, such as vSphere Distributed Resource Scheduler (DRS), VMware Enhanced vMotion Compatibility (EVC), Distributed Power Management (DPM), vSphere High Availability (HA), and vSAN.

In the vSphere Client, you can manage and monitor the resources in a cluster as a single object. You can easily monitor and manage the hosts and virtual machines in the DRS cluster.

If you enable VMware EVC on a cluster, you can ensure that migrations with vMotion do not fail due to CPU compatibility errors. If you enable vSphere DRS on a cluster, you can allow automatic resource balancing using the pooled host resources in the cluster. If you enable vSphere HA on a cluster, you can allow rapid virtual machine recovery from host hardware failures, using the cluster's available host resource capacity. If you enable DPM on a cluster, you can provide automated power management in the cluster. If you enable vSAN on a cluster, you use a logical SAN that is built on a pool of drives attached locally to the ESXi hosts in the cluster.

You can use the Quickstart workflow in the vSphere Client to create and configure a cluster. The Quickstart page provides three cards: Cluster Basics, Add Hosts, and Configure Cluster. For an existing cluster, you can use Cluster Basics to change the cluster name and enable cluster services, such as DRS and vSphere HA. You can use the Add Hosts card to add hosts to the cluster. You can use the Configure Cluster card to configure networking and other settings on the hosts in the cluster.

In addition, in vSphere 7.0 you can configure a few general settings for a cluster. For example, when you create a cluster, even if you do not enable DRS, vSphere, HA, or vSAN, you can choose to manage all hosts in the cluster with a single image. With this option, all hosts in a cluster inherit the same image, which reduces variability between hosts, improves your ability to ensure hardware compatibility, and simplifies upgrades. This feature requires hosts to already be ESXi 7.0 or above. It replaces baselines. Once it is enabled, baselines cannot be used in this cluster.

## Note

Do not confuse a vSphere cluster with a datastore cluster. In vSphere, datastore clusters and vSphere (host) clusters are separate objects. Although you can directly enable a vSphere cluster for vSAN, DRS, and vSphere HA, you cannot directly enable it for datastore clustering. You create datastore clusters separately. See [Chapter 2, "Storage Infrastructure,"](#) for details on datastore clusters.

## Enhanced vMotion Compatibility (EVC)

EVC is a cluster setting that can improve CPU compatibility between hosts for supporting vMotion. vMotion migrations are live migrations that require compatible instruction sets for source and target processors used by the virtual machine. The source and target processors must come from the same vendor class (AMD or Intel) to be vMotion compatible. The clock speed, cache size, and number of cores can differ between source and target processors. When you start a vMotion migration or a migration of a suspended virtual machine, the wizard checks the destination host for compatibility; it displays an error message if problems exist. Using EVC, you can allow vMotion between some processors that would normally be incompatible.

The CPU instruction set that is available to a virtual machine guest OS is determined when the virtual machine is powered on. This CPU feature set is based on the following items:

- The host CPU family and model
- Settings in the BIOS that might disable CPU features
- The ESX/ESXi version running on the host
- The virtual machine's compatibility setting
- The virtual machine's guest operating system

EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. If you enable the EVC cluster setting, you can configure the EVC Mode with a baseline CPU feature set. EVC ensures that hosts in a cluster use the baseline feature set when presenting an instruction set to a guest OS. EVC uses AMD-V Extended Migration technology for AMD hosts and Intel FlexMigration technology for Intel hosts to mask processor features; this allows hosts to present the feature set of an earlier generation of processor. You should configure EVC Mode to accommodate the host with the smallest feature set in the cluster.

The EVC requirements for hosts include the following.

- ESXi 6.5 or later is required.
- Hosts must be attached to a vCenter Server.
- CPUs must be from a single vendor (either Intel or AMD).

- If the AMD-V, Intel-VT, AMD NX, or Intel XD features are available in the BIOS, they need to be enabled.
- Check the *VMware Compatibility Guide* to ensure that CPUs are supported for EVC Mode.

## Note

You can apply a custom CPU compatibility mask to hide host CPU features from a virtual machine, but VMware does not recommend doing so.

You can configure the EVC settings by using the Quickstart > Configure Cluster workflow in the vSphere Client. You can also configure EVC directly in the cluster settings. The options for VMware EVC are Disable EVC, Enable EVC for AMD Hosts, and Enable EVC for Intel Hosts.

If you choose Enable EVC for Intel Hosts, you can set the EVC Mode to one of the options described in [Table 4-2](#).

**Table 4-2** EVC Modes for Intel

Level	EVC Mode	Description
L0	Intel Merom	Smallest Intel feature set for EVC mode.
L1	Intel Penryn	Includes the Intel Merom feature set and exposes additional CPU features, including SSE4.1.
L	Intel	Includes the Intel Penryn feature set and exposes additional CPU

2 tel features, including SSE4.2 and POPCOUNT.  
N  
eh  
al  
e  
m

---

L In Includes the Intel Nehalem feature set and exposes additional CPU  
3 tel features, including AES and PCLMULQDQ.  
W  
es  
t  
m  
er  
e

---

L In Includes the Intel Westmere feature set and exposes additional CPU  
4 tel features, including AVX and XSAVE.  
Sa  
nd  
y  
Br  
id  
ge

---

L In Includes the Intel Sandy Bridge feature set and exposes additional  
5 tel CPU features, including RDRAND, ENFSTRG, FSGSBASE,  
Iv SMEP, and F16C.  
y  
Br  
id  
ge

---

L In Includes the Intel Ivy Bridge feature set and exposes additional  
6 tel CPU features, including ABMX2, AVX2, MOVBE, FMA, PERMD,  
H RORX/MULX, INVPCID, and VMFUNC.  
as  
w  
ell

---

L In Includes the Intel Haswell feature set and exposes additional CPU  
7 tel features, including Transactional Synchronization Extensions,



Br  
oa  
d  
w  
ell

Supervisor Mode Access Prevention, Multi-Precision Add-Carry Instruction Extensions, PREFETCHW, and RDSEED.

---

L  
8  
S  
ky  
la  
ke

In  
tel

Includes the Intel Broadwell feature set and exposes additional CPU features, including Advanced Vector Extensions 512, Persistent Memory Support Instructions, Protection Key Rights, Save Processor Extended States with Compaction, and Save Processor Extended States Supervisor.

---

L  
9  
C  
as  
ca  
de  
L  
ak  
e

In  
tel

Includes the Intel Skylake feature set and exposes additional CPU features, including VNNI and XGETBV with ECX=1.

If you choose Enable EVC for AMD Hosts, you can set EVC Mode to one of the options described in [Table 4-3](#).

**Table 4-3** EVC Modes for AMD

Lev	EVC Mode	Description
A 0 1	AMD Opteron Generation 1	Smallest AMD feature set for EVC mode.
A 1 2	AMD Opteron Generation 2	Includes the AMD Generation 1 feature set and exposes additional CPU features, including CPMXCHG16B and RDTSCP.
A 3	AMD Opteron	Includes the AMD Generation 2 feature set and exposes additional CPU features, including SSE4A, MisAlignSSE,

	Generation 3	POPCOUNT, and ABM (LZCNT).
A 2, B 0	AMD Opteron Generation 3 (without 3DNow!)	Includes the AMD Generation 3 feature set without 3DNow support.
B 1	AMD Opteron Generation 4	Includes the AMD Generation 3 no3DNow feature set and exposes additional CPU features, including SSSE3, SSE4.1, AES, AVX, XSAVE, XOP, and FMA4.
B 2	AMD Opteron Piledriver	Includes the AMD Generation 4 feature set and exposes additional CPU features, including FMA, TBM, BMI1, and F16C.
B 3	AMD Opteron Steamroller	Includes the AMD Piledriver feature set and exposes additional CPU features, including XSAVEOPT, RDFSBASE, RDGSBASE, WRFSBASE, WRGSBAS, and FSGSBASE.
B 4	AMD Zen	Includes the AMD Steamroller feature set and exposes additional CPU features, including RDRAND, SMEP, AVX2, BMI2, MOVBE, ADX, RDSEED, SMAP, CLFLUSHOPT, XSAVES, XSAVEC, SHA, and CLZERO.
B 5	AMD Zen 2	Includes the AMD Zen feature set and exposes additional CPU features, including CLWB, UMIP, RDPID, XGETBV with ECX = 1, WBNOINVD, and GMET.

## vSAN Services

You can enable DRS, vSphere HA, and vSAN at the cluster level. The following sections provide details on DRS and vSphere HA. For details on vSAN, see [Chapter 2](#).

## Distributed Resource Scheduler (DRS)

DRS distributes compute workload in a cluster by strategically placing virtual machines during power-on operations and live migrating (vMotion) VMs when necessary. DRS provides many features and settings that enable you to control its

behavior.

You can set DRS Automation Mode for a cluster to one of the following:

- **Manual:** DRS does not automatically place or migrate virtual machines. It only makes recommendations.
- **Partially Automated:** DRS automatically places virtual machines as they power on. It makes recommendations for virtual machine migrations.
- **Fully Automated:** DRS automatically places and migrates virtual machines.

You can override Automation Mode at the virtual machine level.

## Recent DRS Enhancements

VMware added many improvements to DRS beginning in vSphere 6.5. For example, in vSphere 7.0, DRS runs once every minute rather than every 5 minutes, as in older DRS versions. The newer DRS versions tend to recommend smaller (in terms of memory) virtual machines for migration to facilitate faster vMotion migrations, whereas older versions tend to recommend large virtual machines to minimize the number of migrations. Older DRS versions use an imbalance metric that is derived from the standard deviation of load across the hosts in the cluster. Newer DRS versions focus on virtual machine happiness. Newer DRS versions are much lighter and faster than the older versions.

Newer DRS versions recognize that vMotion is an expensive operation and account for it in their recommendations. In a cluster where virtual machines are frequently powered on and the workload is volatile, it is not necessary to continuously migrate virtual machines. DRS calculates the gain duration for live migrating a virtual machine and considers the gain duration when making recommendations.

The following sections provide details on other recent DRS enhancements.

## Network-Aware DRS



In vSphere 6.5, DRS considers the utilization of host network adapters during initial placement and load balancing, but it does not balance the network load. Instead, its goal is to ensure that the target host has sufficient available network

resources. It works by eliminating hosts with saturated networks from the list of possible migration hosts. The threshold used by DRS for network saturation is 80% by default. When DRS cannot migrate VMs due to network saturation, the result may be an imbalanced cluster.

In vSphere 7.0, DRS uses a new cost modeling algorithm that is flexible and balances network bandwidth along with CPU and memory usage.

## **Virtual Machine Distribution**

Starting in vSphere 6.5, you can enable an option to distribute a more even number of virtual machines across hosts. The main use case for this is to improve availability. The primary goal of DRS—to ensure that all VMs are getting the resources they need and that the load is balanced in the cluster—remains unchanged. But with this new option enabled, DRS also tries to ensure that the number of virtual machines per host is balanced in the cluster.

## **Memory Metric for Load Balancing**

Historically, vSphere has used the Active Memory metric for load-balancing decisions. In vSphere 6.5 and 6.7, you have the option to set DRS to balance the load based on the Consumed Memory metric. In vSphere 7.0, the Granted Memory metric is used for load balancing, and no cluster option is available to change the behavior.

## **Virtual Machine Initial Placement**

Starting with vSphere 6.5, DRS uses a new initial placement algorithm that is faster, lighter, and more effective than the previous algorithm. In earlier versions, DRS takes a snapshot of the cluster state when making virtual machine placement recommendations. In the algorithm, DRS does not snapshot the cluster state, which allows for faster and more accurate recommendations. With the new algorithm, DRS powers on virtual machines much more quickly. In vSphere 6.5, the new placement feature is not supported for the following configurations:

- Clusters where DPM, Proactive HA, or HA Admission Control is enabled
- Clusters with DRS configured in Manual Mode
- Virtual machines with the Manual DRS Override setting enabled

- Virtual machines that are FT enabled
- Virtual machines that are part of a vApp

In vSphere 6.7, the new placement is available for all configurations.

## **Enhancements to the Evacuation Workflow**

Prior to vSphere 6.5, when evacuating a host entering Maintenance Mode, DRS waited to migrate templates and powered off virtual machines until after the completion of vMotion migrations, leaving those objects unavailable for use for a long time. Starting in vSphere 6.5, DRS prioritizes the migration of virtual machine templates and powered-off virtual machines over powered-on virtual machines, making those objects available for use without waiting on vMotion migrations.

Prior to vSphere 6.5, the evacuation of powered-off virtual machines was inefficient. Starting in vSphere 6.5, these evacuations occur in parallel, making use of up to 100 re-register threads per vCenter Server. This means that you may see only a small difference when evacuating up to 100 virtual machines.

Starting in vSphere 6.7, DRS is more efficient in evacuating powered-on virtual machines from a host that is entering Maintenance Mode. Instead of simultaneously initiating vMotion for all the powered-on VMs on the host, as in previous versions, DRS initiates vMotion migrations in batches of eight at a time. Each vMotion batch is issued after the previous batch completes. The vMotion batching makes the entire workflow more controlled and predictable.

## **DRS Support for NVM**

Starting in vSphere 6.7, DRS supports virtual machines running on next-generation persistent memory devices, known as non-volatile memory (NVM) devices. NVM is exposed as a datastore that is local to the host. Virtual machines can use the datastore as an NVM device exposed to the guest (Virtual Persistent Memory [vPMem]) or as a location for a virtual machine disk (Virtual Persistent Memory Disk [vPMemDisk]). DRS is aware of the NVM devices used by virtual machines and guarantees that the destination ESXi host has enough free persistent memory to accommodate placements and migrations.

## **How DRS Scores VMs**

**Key Topic**

Historically, DRS balanced the workload in a cluster based on host compute resource usage. In vSphere 7.0, DRS balances the workload based on virtual machine happiness. A virtual machine's DRS score is a measure of its happiness, which, in turn, is a measure of the resources available for consumption by the virtual machine. The higher the DRS score for a VM, the better its resource availability. DRS moves virtual machines to improve their DRS scores. DRS also calculates a DRS score for a cluster, which is a weighted sum of the DRS scores of all the virtual machines in the cluster.

In Sphere 7.0, DRS calculates the core for each virtual machine on each ESXi host in the cluster every minute. Simply put, DRS logic computes an ideal throughput (demand) and an actual throughput (goodness) for each resource (CPU, memory, and network) for each virtual machine. The virtual machine's efficiency for a particular resource is a ratio of the goodness over the demand. A virtual machine's DRS score (total efficiency) is the product of its CPU, memory, and network efficiencies.

When calculating the efficiency, DRS applies resource costs. For CPU resources, DRS includes costs for CPU cache, CPU ready, and CPU tax. For memory resources, DRS includes costs for memory burstiness, memory reclamation, and memory tax. For network resources, DRS includes a network utilization cost.

DRS compares a virtual machine's DRS score for the host on which it currently runs. DRS determines whether another host can provide a better DRS score for the virtual machine. If so, DRS calculates the cost for migrating the virtual machine to the host and factors that score into its load-balancing decision.

## DRS Rules

You can configure rules to control the behavior of DRS.

A VM–host affinity rule specifies whether the members of a selected virtual machine DRS group can run on the members of a specific host DRS group. Unlike a virtual machine–to–virtual machine (VM–VM) affinity rule, which specifies affinity (or anti-affinity) between individual virtual machines, a VM–host affinity rule specifies an affinity relationship between a group of virtual machines and a group of hosts. There are *required* rules (designated by “must”) and *preferential* rules (designated by “should”).

A VM–host affinity rule includes the following components:

- One virtual machine DRS group
- One host DRS group
- A designation of whether the rule is a requirement (“must”) or a preference (“should”) and whether it is affinity (“run on”) or anti-affinity (“not run on”)

A VM–VM affinity rule specifies whether selected individual virtual machines should run on the same host or be kept on separate hosts. This type of rule is used to create affinity or anti-affinity between individual virtual machines. When an affinity rule is created, DRS tries to keep the specified virtual machines together on the same host. You might want to do this, for example, for performance reasons.

With an anti-affinity rule, DRS tries to keep the specified virtual machines apart. You can use such a rule if you want to guarantee that certain virtual machines are always on different physical hosts. In that case, if a problem occurs with one host, not all virtual machines are at risk. You can create VM–VM affinity rules to specify whether selected individual virtual machines should run on the same host or be kept on separate hosts.

VM–VM affinity rule conflicts can occur when you use multiple VM–VM affinity and VM–VM anti-affinity rules. If two VM–VM affinity rules are in conflict, you cannot enable both of them. For example, if one rule keeps two virtual machines together and another rule keeps the same two virtual machines apart, you cannot enable both rules. Select one of the rules to apply and disable or remove the conflicting rule. When two VM–VM affinity rules conflict, the older one takes precedence, and the newer rule is disabled. DRS tries to satisfy only enabled rules and ignores disabled rules. DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.

### **Note**

A VM–VM rule does not allow the “should” qualifier. You should consider these as “must” rules.

## **DRS Migration Sensitivity**

Prior to vSphere 7.0, DRS used a migration threshold to determine when virtual machines should be migrated to balance the cluster workload. In vSphere 7.0,

DRS does not consider cluster standard deviation for load balancing. Instead, it is designed to be more virtual machine centric and workload centric rather than cluster centric. You can set the DRS Migration Sensitivity parameter to one of the following values:

#### Key Topic

- **Level 1:** DRS only makes recommendations to fix rule violations or to facilitate a host entering Maintenance Mode.
- **Level 2:** DRS expands on Level 1 by making recommendations in situations that are at or close to resource contention. It does not make recommendations just to improve virtual machine happiness or cluster load distribution.
- **Level 3:** DRS expands on Level 2 by making recommendations to improve VM happiness and cluster load distribution. This is the default level.
- **Level 4:** DRS expands on Level 3 by making recommendations for occasional bursts in the workload and reacts to sudden load changes.
- **Level 5:** DRS expands on Level 4 by making recommendations dynamic and greatly varying workloads. DRS reacts to the workload changes every time.

## Resource Pools

Resource pools are container objects in the [vSphere inventory](#) that are used to compartmentalize the CPU and memory resources of a host, a cluster, or a parent resource pool. Virtual machines run in and draw resources from resource pools. You can create multiple resource pools as direct children of a standalone host or a DRS cluster. You cannot create child resource pools on a host that has been added to a cluster or on a cluster that is not enabled for DRS.

You can use resource pools to organize VMs. You can delegate control over each resource pool to specific individuals and groups. You can monitor resources and set alarms on resource pools. If you need a container just for organization and permission purposes, consider using a folder. If you also need resource management, then consider using a resource pool. You can assign resource settings such as shares, reservations, and limits to resource pools.

## Use Cases



You can use resource pools to compartmentalize a cluster's resources and then use the resource pools to delegate control to individuals or organizations. [Table 4-4](#) provides some use cases for resource pools.

**Table 4-4** Resource Pool Use Cases

Use Case	Details
Flexible hierarchical organization	Add, remove, modify, and reorganize resource pools, as needed.
Resource isolation	Use resource pools to allocate resources to separate departments, in such a manner that changes in a pool do not unfairly impact other departments.
Access control and delegation	Use permissions to delegate activities, such as virtual machine creation and management, to other administrators.
Separation of resources from hardware	In a DRS cluster, perform resource management independently of the actual hosts.
Managing multitier applications.	Manage the resources for a group of virtual machines (in a specific resource pool), which is easier than managing resources per virtual machine.

## Shares, Limits, and Reservations

You can configure CPU and memory shares, reservations, and limits on resource pools, as described in [Table 4-5](#).

**Table 4-5** Shares, Limits, and Reservations

Opt	Description
S	Shares specify the relative importance of a virtual machine or a resource pool. If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when these two virtual machines are competing for resources.

s Shares can be thought of as priority under contention.

Shares are typically set to High, Normal, or Low, and these values specify share values with a 4:2:1 ratio. You can also select Custom and assign a specific number of shares (to express a proportional weight).

A resource pool uses its shares to compete for the parent's resources and is allocated a portion based on the ratio of the pool's shares compared with its siblings. Siblings share the parent's resources according to their relative share values, bounded by the reservation and limit.

For example, consider a scenario where a cluster has two child resource pools with normal CPU shares, another child resource pool with high CPU shares, and no other child objects. During periods of contention, each of the pools with normal shares would get access to 25% of the cluster's CPU resources, and the pool with high shares would get access to 50%.

---

R A reservation specifies the guaranteed minimum allocation for a virtual  
e machine or a resource pool. A CPU reservation is expressed in  
s megahertz, and a memory reservation is expressed in megabytes. You  
er can power on a virtual machine only if there are enough unreserved  
v resources to satisfy the reservation of the virtual machine. If the virtual  
at machine starts, then it is guaranteed that amount, even when the physical  
io server is heavily loaded.

n  
s

For example, if you configure the CPU reservation for each virtual machine as 1 GHz, you can start eight VMs in a resource pool where the CPU reservation is set for 8 GHz and expandable reservations are disabled. But you cannot start additional virtual machines in the pool.

You can use reservations to guarantee a specific amount of resources for a resource pool. The default value for a resource pool's CPU or memory reservation is 0. If you change this value, it is subtracted from the unreserved resources of the parent. The resources are considered reserved, regardless of whether virtual machines are associated with the

resource pool.

---

**E** You can enable expandable reservations to effectively allow a child  
**x** resource pool to borrow from its parent. Expandable reservations, which  
**p** are enabled by default, are considered during admission control. When  
**a** powering on a virtual machine, if the resource pool does not have  
**n** sufficient unreserved resources, the resource pool can use resources from  
**d** its parent or ancestors.

**a**  
**bl**  
**e** For example, say that in a resource pool where 8 GHz is reserved and  
**re** expandable reservations is disabled, you try to start nine virtual machines  
**s** each with 1 GHz, but the last virtual machine does not start. If you  
**er** enable expandable reservation in the resource pool, and its parent pool  
**v** (or cluster) has sufficient unreserved CPU resources, you can start the  
**at** ninth virtual machine.

---

**L** A limit specifies an upper bound for CPU or memory resources that can  
**i** be allocated to a virtual machine or a resource pool.

**m**  
**it**  
**s** You can set a limit on the amount of CPU and memory allocated to a  
resource pool. The default is unlimited. For example, if you power on  
multiple CPU-intensive virtual machines in a resource pool, where the  
CPU limit is 10 GHz, then, collectively, the virtual machines cannot use  
more than 10 GHz CPU resources, regardless of the pool's reservation  
settings, the pool's share settings, or the amount of available resources in  
the parent.

---

[Table 4-6](#) provides the CPU and memory share values for virtual machines when using the High, Normal, and Low settings. The corresponding share values for a resource pool are equivalent to those of a virtual machine with four vCPUs and 16 GB memory.

**Table 4-6** Virtual Machine Shares

---

<b>Setting</b>	<b>CPU Share Value</b>	<b>Memory Share Value</b>
High	2000 per vCPU	20 per MB

---

Normal	1000 per vCPU	10 per MB
Low	500 per vCPU	5 per MB

For example, the share values for a resource pool configured with normal CPU shares and high memory shares are 4000 (that is,  $4 \times 1000$ ) CPU shares and 327,680 (that is,  $16 \times 1024 \times 20$ ) memory shares

### Note

The relative priority represented by each share changes with the addition and removal of virtual machines in a resource pool or cluster. It also changes as you increase or decrease the shares on a specific virtual machine or resource pool.

## Enhanced Resource Pool Reservation

Starting in vSphere 6.7, DRS uses a new two-pass algorithm to allocate resource reservations to children. The old allocation model does not reserve more resources than the current demand, even when the resource pool is configured with a higher reservation. When a spike in virtual machine demand occurs after resource allocation is complete, DRS does not make the remaining pool reservation available to the virtual machine until the next allocation operation occurs. As a result, a virtual machine's performance may be temporarily impacted. In the new allocation model, each allocation operation uses two passes. In the first pass, the resource pool reservation is allocated based on virtual machine demand. In the second pass, excess pool reservation is allocated proportionally, limited by the virtual machine's configured size, which reduces the performance impact due to virtual machine spikes.

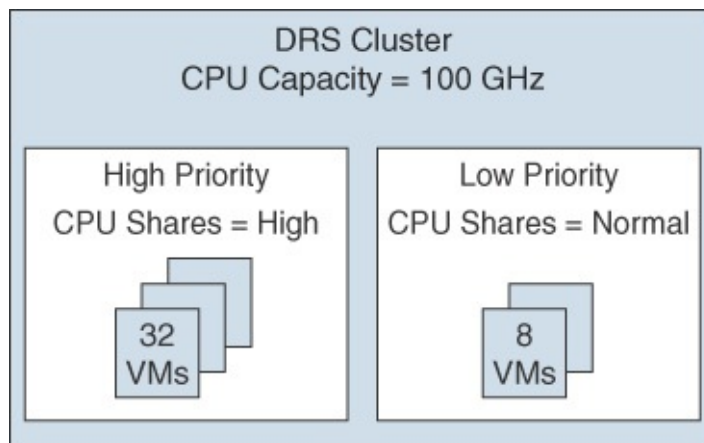
## Scalable Shares



Another new DRS feature in vSphere 7.0 is scalable shares. The main use case for scalable shares is a scenario in which you want to use shares to give high-priority resource access to a set of virtual machines in a resource pool, without concern for the relative number of objects in the pool compared to other pools. With standard shares, each pool in a cluster competes for resource allocation

with its siblings, based on the share ratio. With scalable shares, the allocation for each pool factors in the number of objects in the pool.

For example, consider a scenario in which a cluster with 100 GHz CPU capacity has a high-priority resource pool with CPU Shares set to High and a low-priority resource pool with CPU Shares set to Normal, as shown in [Figure 4-1](#). This means that the share ratio between the pools is 2:1, so the high-priority pool is effectively allocated twice the CPU resources as the low-priority pool whenever CPU contention exists in the cluster. The high-priority pool is allocated 66.7 GHz, and the low-priority pool is effectively allocated 33.3 GHz. In this cluster, 40 virtual machines of equal size are running, with 32 in the high-priority pool and 8 in the low-priority pool. The virtual machines are all demanding CPU resources, causing CPU contention in the cluster. In the high-priority pool, each virtual machine is allocated 2.1 GHz. In the low-priority pool, each virtual machine is allocated 4.2 GHz.



**FIGURE 4-1** Scalable Shares Example


If you want to change the resource allocation such that each virtual machine in the high-priority pool is effectively allocated more resources than the virtual machines in the low-priority pool, you can use scalable shares. If you enable scalable shares in the cluster, DRS effectively allocates resources to the pools based on the Shares settings and the number of virtual machines in the pool. In this example, the CPU shares for the pools provide a 2:1 ratio. Factoring this with the number of virtual machines in each pool, the allocation ratio between the high-priority pool and the low-priority pool is 2 times 32 to 1 times 8, or simply 8:1. The high-priority pool is allocated 88.9 GHz, and the low-priority pool is allocated 11.1 GHz. Each virtual machine in the high-priority pool is

allocated 2.8 GHz. Each virtual machine in the low-priority pool is allocated 1.4 GHz.

## vSphere High Availability (HA)

vSphere HA is a cluster service that provides high availability for the virtual machines running in the cluster. You can enable vSphere High Availability (HA) on a vSphere cluster to provide rapid recovery from outages and cost-effective high availability for applications running in virtual machines. vSphere HA provides application availability in the following ways:

- It protects against server failure by restarting the virtual machines on other hosts in the cluster when a host failure is detected, as illustrated in [Figure 4-2](#).



**FIGURE 4-2** vSphere HA Host Failover

- It protects against application failure by continuously monitoring a virtual machine and resetting it if a failure is detected.
- It protects against datastore accessibility failures by restarting affected virtual machines on other hosts that still have access to their datastores.
- It protects virtual machines against network isolation by restarting them if their host becomes isolated on the management or vSAN network. This protection is provided even if the network has become partitioned.

Benefits of vSphere HA over traditional failover solutions include the following:

- Minimal configuration

- Reduced hardware cost
- Increased application availability
- DRS and vMotion integration

vSphere HA can detect the following types of host issues:

- **Failure:** A host stops functioning.
- **Isolation:** A host cannot communicate with any other hosts in the cluster.
- **Partition:** A host loses network connectivity with the primary host.

When you enable vSphere HA on a cluster, the cluster elects one of the hosts to act as the primary host. The primary host communicates with vCenter Server to report cluster health. It monitors the state of all protected virtual machines and secondary hosts. It uses network and datastore heartbeating to detect failed hosts, isolation, and network partitions. vSphere HA takes appropriate actions to respond to host failures, host isolation, and network partitions. For host failures, the typical reaction is to restart the failed virtual machines on surviving hosts in the cluster. If a network partition occurs, a primary host is elected in each partition. If a specific host is isolated, vSphere HA takes the predefined host isolation action, which may be to shut down or power down the host's virtual machines. If the primary host fails, the surviving hosts elect a new primary host. You can configure vSphere to monitor and respond to virtual machine failures, such as guest OS failures, by monitoring heartbeats from VMware Tools.

### Note

Although vCenter Server is required to implement vSphere HA, the health of an HA cluster is not dependent on vCenter Server. If vCenter Server fails, vSphere HA still functions. If vCenter Server is offline when a host fails, vSphere HA can fail over the affected virtual machines.

## vSphere HA Requirements

When planning a vSphere HA cluster, you need to address the following requirements:

**Key Topic**

- The cluster must have at least two hosts, licensed for vSphere HA.
- Hosts must use static IP addresses or guarantee that IP addresses assigned by DHCP persist across host reboots.
- Each host must have at least one—and preferably two—management networks in common.
- To ensure that virtual machines can run any host in the cluster, the hosts must access the networks and datastores.
- To use VM Monitoring, you need to install VMware Tools in each virtual machine.
- IPv4 or IPv6 can be used.

**Note**

The Virtual Machine Startup and Shutdown (automatic startup) feature is disabled and unsupported for all virtual machines residing in a vSphere HA cluster.

## vSphere HA Response to Failures

You can configure how a vSphere HA cluster should respond to different types of failures, as described in [Table 4-7](#).

**Key Topic**

**Table 4-7** vSphere HA Response to Failure Settings

Option	Description
Host Failure Response > Failure Response	If Enabled, the cluster responds to host failures by restarting virtual machines. If Disabled, host monitoring is turned off, and the cluster does not respond to host failures.
Host Failure Response >	You can indicate the order in which virtual machines are restarted when the host fails (higher priority machines



Default VM Restart Priority	first).
Host Failure Response > VM Restart Priority Condition	This condition must be met before HA restarts the next priority group.
Response for Host Isolation	You can indicate the action that you want to occur if a host becomes isolated. You can choose Disabled, Shutdown and Restart VMs, or Power Off and Restart VMs.
VM Monitoring	You can indicate the sensitivity (Low, High, or Custom) with which vSphere HA responds to lost VMware Tools heartbeats.
Application Monitoring	You can indicate the sensitivity (Low, High, or Custom) with which vSphere HA responds to lost application heartbeats.

### Note

If multiple hosts fail, the virtual machines on the failed host migrate first in order of priority, and then the virtual machines from the next host.

## Heartbeats

The primary host and secondary hosts exchange network heartbeats every second. When the primary host stops receiving these heartbeats from a secondary host, it checks for ping responses or the presence of datastore heartbeats from the secondary host. If the primary host does not receive a response after checking for a secondary host's network heartbeat, ping, or datastore heartbeats, it declares that the secondary host has failed. If the primary host detects datastore heartbeats for a secondary host but no network heartbeats or ping responses, it assumes that the secondary host is isolated or in a network partition.

If any host is running but no longer observes network heartbeats, it attempts to ping the set of cluster isolation addresses. If those pings also fail, the host

declares itself to be isolated from the network.

## vSphere HA Admission Control

vSphere uses admission control when you power on a virtual machine. It checks the amount of unreserved compute resources and determines whether it can guarantee that any reservation configured for the virtual machine is configured. If so, it allows the virtual machine to power on. Otherwise, it generates an “Insufficient Resources” warning.

vSphere HA Admission Control is a setting that you can use to specify whether virtual machines can be started if they violate availability constraints. The cluster reserves resources so that failover can occur for all running virtual machines on the specified number of hosts. When you configure vSphere HA admission control, you can set options described in [Table 4-8](#).

**Table 4-8** vSphere HA Admission Control Options

Option	Description
Host Failures Cluster Tolerates	Specifies the maximum number of host failures for which the cluster guarantees failover
Define Host Failover Capacity By set to Cluster Resource Percentage	Specifies the percentage of the cluster’s compute resources to reserve as spare capacity to support failovers
Define Host Failover Capacity By set to Slot Policy (powered-on VMs)	Specifies a slot size policy that covers all powered-on VMs
Define Host Failover Capacity By set to Dedicated Failover Hosts	Specifies the designated hosts to use for failover actions
Define Host Failover Capacity By set to Disabled	Disables admission control
Performance Degradation VMs Tolerate	Specifies the percentage of performance degradation the VMs in a cluster are allowed to tolerate during a failure

If you disable vSphere HA admission control, then you enable the cluster to

allow virtual machines to power on regardless of whether they violate availability constraints. In the event of a host failover, you may discover that vSphere HA cannot start some virtual machines.

In vSphere 6.5, the default Admission Control setting is Cluster Resource Percentage, which reserves a percentage of the total available CPU and memory resources in the cluster. For simplicity, the percentage is calculated automatically by defining the number of host failures to tolerate (FTT). The percentage is dynamically changed as hosts are added to or removed from the cluster. Another new enhancement is the Performance Degradation VMs Tolerate setting, which controls the amount of performance reduction that is tolerated after a failure. A value of 0% indicates that no performance degradation is tolerated.

With the Slot Policy option, vSphere HA admission control ensures that a specified number of hosts can fail, leaving sufficient resources in the cluster to accommodate the failover of the impacted virtual machines. Using the Slot Policy option, when you perform certain operations, such as powering on a virtual machine, vSphere HA applies admission control in the following manner:

- Step 1.** HA calculates the slot size, which is a logical representation of memory and CPU resources. By default, it is sized to satisfy the requirements for any powered-on virtual machine in the cluster. For example, it is sized to accommodate the virtual machine with the greatest CPU reservation and the virtual machine with the greatest memory reservation.
- Step 2.** HA determines how many slots each host in the cluster can hold.
- Step 3.** HA determines the current failover capacity of the cluster, which is the number of hosts that can fail and still leave enough slots to satisfy all the powered-on virtual machines.
- Step 4.** HA determines whether the current failover capacity is less than the configured failover capacity (provided by the user).
- Step 5.** If the current failover capacity is less than the configured failover capacity, admission control disallows the operation.

If a cluster has a few virtual machines that have much larger reservations than the others, they will distort slot size calculation. To remediate this, you can specify an upper bound for the CPU or memory component of the slot size by using advanced options. You can also set a specific slot size (CPU size and memory size). The next section describes the advanced options that affect the

slot size.

## vSphere HA Advanced Options

You can set vSphere HA advanced options by using the vSphere Client or in the `fdm.cfg` file on the hosts. [Table 4-9](#) provides some of the advanced vSphere HA options.

**Table 4-9** Advanced vSphere HA Options

Option	Description
<code>das.isolationaddressX</code>	Provides the addresses to use to test for host isolation when no heartbeats are received from other hosts in the cluster. If this option is not specified (which is the default setting), the management network default gateway is used to test for isolation. To specify multiple addresses, you can set <code>das.isolationaddressX</code> , where <i>X</i> is a number between 0 and 9.
<code>das.ultimateisolationaddress</code>	Specifies whether to use the default gateway IP address for isolation tests.
<code>das.isolationshutdowntimeout</code>	For scenarios where the host's isolation response is to shut down, specifies the period of time that the virtual machine is permitted to shut down before the system powers it off.
<code>das.slotmemoryminimum</code>	Defines the maximum bound on the memory slot size.
<code>das.slotcpuinhibit</code>	Defines the maximum bound on the CPU slot size.

das.v  
mme  
mory  
minm  
b

Defines the default memory resource value assigned to a virtual machine whose memory reservation is not specified or is zero. This is used for the Host Failures Cluster Tolerates admission control policy.

---

das.v  
mcpu  
minm  
hz

Defines the default CPU resource value assigned to a virtual machine whose CPU reservation is not specified or is zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default of 32 MHz is used.

---

das.h  
earb  
eatds  
perho  
st

Specifies the number of heartbeat datastores required per host. The default is 2. The acceptable values are 2 to 5.

---

das.c  
onfig  
.fdm.  
isolat  
ionPo  
licyD  
elayS  
ec

Specifies the number of seconds the system delays before executing the isolation policy after determining that a host is isolated. The minimum is 30. A lower value results in a 30-second delay.

---

das.r  
espec  
tvmv  
manti  
affini  
tyrule  
s

Determines whether vSphere HA should enforce VM-VM anti-affinity rules even when DRS is not enabled.

---

## Virtual Machine Settings

To use the Host Isolation Response Shutdown and Restart VMs setting, you must install VMware Tools on the virtual machine. If a guest OS fails to shut down in 300 seconds (or a value specified by `das.isolationshutdowntimeout`), the virtual machine is powered off.

You can override the cluster's settings for Restart Priority and Isolation

Response for each virtual machine. For example, you might want to prioritize virtual machines providing infrastructure services such as DNS or DHCP.

At the cluster level, you can create dependencies between groups of virtual machines. You can create VM groups, host groups, and dependency rules between the groups. In the rules, you can specify that one VM group cannot be restarted if another specific VM group is started.

## **VM Component Protection (VMCP)**

*Virtual Machine Component Protection (VMCP)* is a vSphere HA feature that can detect datastore accessibility issues and provide remediation for affected virtual machines. When a failure occurs such that a host can no longer access the storage path for a specific datastore, vSphere HA can respond by taking actions such as creating event alarms or restarting a virtual machine on other hosts. The main requirements are that vSphere HA is enabled in the cluster and that ESX 6.0 or later is used on all hosts in the cluster.

The failures VMCP detects are permanent device loss (PDL) and all paths down (APD). PDL is an unrecoverable loss of accessibility to the storage device that cannot be fixed without powering down the virtual machines. APD is a transient accessibility loss or other issue that is recoverable.

For PDL and APD failures, you can set VMCP to either issue event alerts or to power off and restart virtual machines. For APD failures only, you can additionally control the restart policy for virtual machines by setting it to Conservative or Aggressive. With the Conservative setting, the virtual machine is powered off only if HA determines that it can be restarted on another host. With the Aggressive setting, HA powers off the virtual machine regardless of the state of other hosts.

## **Virtual Machine and Application Monitoring**

VM Monitoring restarts specific virtual machines if their VMware Tools heartbeats are not received within a specified time. Likewise, Application Monitoring can restart a virtual machine if the heartbeats from a specific application in the virtual machine are not received. If you enable these features, you can configure the monitoring settings to control the failure interval and reset period. [Table 4-10](#) lists these settings.

### **Table 4-10** VM Monitoring Settings

---

Setting	Failure Interval	Reset Period
High	30 seconds	1 hour
Medium	60 seconds	24 hours
Low	120 seconds	7 days

The Maximum per-VM resets setting can be used to configure the maximum number of times vSphere HA attempts to restart a specific failing virtual machine within the reset period.

## vSphere HA Best Practices

You should provide network path redundancy between cluster nodes. To do so, you can use NIC teaming for the virtual switch. You can also create a second management network connection, using a separate virtual switch.

When performing disruptive network maintenance operations on the network used by clustered ESXi hosts, you should suspend the Host Monitoring feature to ensure that vSphere HA does not falsely detect network isolation or host failures. You can reenable host monitoring after completing the work.

To keep vSphere HA agent traffic on the specified network, you should ensure that the VMkernel virtual network adapters used for HA heartbeats (enabled for management traffic) do not share the same subnet as VMkernel adapters used for vMotion and other purposes.

Use the `das.isolationaddressX` advanced option to add an isolation address for each management network.

## Proactive HA

**Proactive High Availability (Proactive HA)** integrates with select hardware partners to detect degraded components and evacuate VMs from affected vSphere hosts before an incident causes a service interruption. Hardware partners offer a vCenter Server plug-in to provide the health status of the system memory, local storage, power supplies, cooling fans, and network adapters. As hardware components become degraded, Proactive HA determines which hosts are at risk and places them into either Quarantine Mode or Maintenance Mode. When a host enters Maintenance Mode, DRS evacuates its virtual machines to healthy hosts, and the host is not used to run virtual machines. When a host enters Quarantine Mode, DRS leaves the current virtual machines running on the

host but avoids placing or migrating virtual machines to the host. If you prefer that Proactive HA simply make evacuation recommendations rather than automatic migrations, you can set Automation Level to Manual.

The vendor-provided health providers read sensor data in the server and provide the health state to vCenter Server. The health states are Healthy, Moderate Degradation, Severe Degradation, and Unknown.

## Other Resource Management and Availability Features

This section describes other vSphere features related to resource management and availability.

### Predictive DRS

**Predictive DRS** is a feature in vSphere 6.5 and later that leverages the predictive analytics of vRealize Operations (vROps) Manager and vSphere DRS. Together, these two products can provide workload balancing prior to the occurrence of resource utilization spikes and resource contention. Every night, vROps calculates dynamic thresholds, which are used to create forecasted metrics for the future utilization of virtual machines. vROps passes the predictive metrics to vSphere DRS to determine the best placement and balance of virtual machines before resource utilization spikes occur. Predictive DRS helps prevent resource contention on hosts that run virtual machines with predictable utilization patterns.

The following prerequisites are needed to run Predictive DRS:

- vCenter Server 6.5 or later is required.
- Predictive DRS must be configured and enabled in both vCenter Server and vROps.
- The vCenter Server and vROps clocks must be synchronized.

### Distributed Power Management (DPM)

The vSphere Distributed Power Management (DPM) feature enables a DRS cluster to reduce its power consumption by powering hosts on and off, as needed, based on cluster resource utilization. DPM monitors the cumulative



virtual machine demand for memory and CPU resources in the cluster and compares this to the available resources in the cluster. If sufficient excess capacity is found, vSphere DPM directs the host to enter Standby Mode. When DRS detects that a host is entering Standby Mode, it evacuates the virtual machines. Once the host is evacuated, DPM powers it off, and the host is in Standby Mode. When DPM determines that capacity is inadequate to meet the resource demand, DPM brings a host out of Standby Mode by powering it on. Once the host exits Standby Mode, DRS migrates virtual machines to it.

To power on a host, DPM can use one of three power management protocols: Intelligent Platform Management Interface (IPMI), Hewlett-Packard Integrated Lights-Out (iLO), or Wake-on-LAN (WoL). If a host supports multiple protocols, they are used in the following order: IPMI, iLO, WoL. If a host does not support one of these protocols, DPM cannot automatically bring a host out of Standby Mode.

DPM is very configurable. As with DRS, you can set DPM's automation to be manual or automatic.

## Note

Do not disconnect a host that is in Standby Mode or remove it from a DRS cluster without first powering it on. Otherwise, vCenter Server is not able to power the host back on.

To configure IPMI or iLO settings for a host, you can edit the host's Power Management settings. You should provide credentials for the Baseboard Management Controller (BMC) account, the IP address of the appropriate NIC, and the MAC address of the NIC.

Using WoL with DPM requires that the following prerequisites be met:

- ESXi 3.5 or later is required.
- vMotion must be configured.
- The vMotion NIC must support WoL.
- The physical switch port must be set to automatically negotiate the link speed.

Before enabling DPM, use the vSphere Client to request the host to enter Standby Mode. After the host powers down, right-click the host and attempt to power on. If this is successful, you can allow the host to participate in DPM. Otherwise, you should disable power management for the host.

You can enable DPM in a DRS cluster's settings. You can set Automation Level to Off, Manual, or Automatic. When this option is set to Off, DPM is disabled. When it is set to Manual, DPM makes recommendations only. When it is set to Automatic, DPM automatically performs host power operations as needed.

Much as with DRS, with DPM you can control the aggressiveness of DPM (that is, the DPM threshold) with a slider bar in the vSphere Client. The DRS threshold and the DPM threshold are independent of one another. You can override automation settings per host. For example, for a 16-host cluster, you might want to set DPM Automation to Automatic on only 8 of the hosts.

## **Fault Tolerance (FT)**

If you have virtual machines that require continuous availability as opposed to high availability, you can consider protecting the virtual machines with **vSphere Fault Tolerance (FT)**. FT provides continuous availability for a virtual machine (the primary VM) by ensuring that the state of a secondary VM is identical at any point in the instruction execution of the virtual machine.

If the host running the primary VM fails, an immediate and transparent failover occurs. The secondary VM becomes the primary VM host without losing network connection or in-progress transactions. With transparent failover, there is no data loss, and network connections are maintained. The failover is fully automated and occurs even if vCenter Server is unavailable. Following the failover, FT spawns a new secondary VM and reestablishes redundancy and protection, assuming that a host with sufficient resources is available in the cluster. Likewise, if the host running the secondary VM fails, a new secondary VM is deployed. vSphere Fault Tolerance can accommodate symmetric multiprocessor (SMP) virtual machines with up to eight vCPUs.

Use cases for FT include the following:

- Applications that require continuous availability, especially those with long-lasting client connections that need to be maintained during hardware failure
- Custom applications that have no other way of being clustered
- Cases in which other clustering solutions are available but are too

complicated or expensive to configure and maintain

Before implementing FT, consider the following requirements:

**Key  
Topic**

- CPUs must be vMotion compatible.
- CPUs must support hardware MMU virtualization.
- A low-latency 10 Gbps network is required for FT Logging.
- Virtual machine files other than VMDK files must be stored on shared storage.
- A vSphere Standard License is required for FT protection of virtual machines with up to two virtual CPUs.
- A vSphere Enterprise Plus License is required for FT protection of virtual machines with up to eight virtual CPUs.
- Hardware Virtualization (HV) must be enabled in the host BIOS.
- Hosts must be certified for FT.
- The virtual memory reservation should be set to match the memory size.
- vSphere HA must be enabled on the cluster.
- SSL certificate checking must be enabled in the vCenter Server settings.
- The hosts must use ESXi 6.x or later.

You should also consider the following VMware recommendations concerning vSphere FT:

- VMware recommends a minimum of two physical NICs.
- VMware recommends that the host BIOS power management settings be set to Maximum Performance or OS-Managed Performance.
- You should have at least three hosts in the cluster to accommodate a new secondary VM following a failover.

The following vSphere features are not supported for FT-protected virtual machines:

- Snapshots (An exception is that disk-only snapshots created for vStorage APIs for Data Protection [VADP] backups are supported for FT but not for legacy FT.)
- Storage vMotion
- Linked clones
- Virtual Volumes datastores
- Storage-based policy management (However, vSAN storage policies are supported.)
- I/O filters
- Disk encryption
- Trusted Platform Module (TPM)
- Virtual Based Security (VBS)–enabled VMs
- Universal Point in Time snapshots (a NextGen vSAN feature)
- Physical raw device mappings (RDMs) (However, virtual RDMs are supported for legacy FT.)
- Virtual CD-ROMs for floppy drives backed by physical devices
- USB devices, sound devices, serial ports, and parallel ports
  - N\_Port ID Virtualization (NPIV)
- Network adapter passthrough
- Hot plugging devices (Note that the hot plug feature is automatically disabled when you enable FT on a virtual machine.)
- Changing the network where a virtual NIC is connected
- Virtual Machine Communication Interface (VMCI)
- Virtual disk files larger than 2 TB
- Video devices with 3D enabled

You should apply the following best practices for FT:

- Use similar CPU frequencies in the hosts.
- Use active/standby NIC teaming settings.

- Ensure that the FT Logging network is secure (that is, FT data is not encrypted).
- Enable jumbo frames and 10 Gbps for the FT network. Optionally, configure multiple NICs for FT Logging.
- Place ISO files on shared storage.
- If vSAN is used for primary or secondary VMs, do not also connect those virtual machines to other storage types. Also, place the primary and secondary VMs in separate vSAN fault domains.
- Keep vSAN and FT Logging on separate networks.

In vSphere 6.5, FT is supported with DRS only when EVC is enabled. You can assign a DRS automation to the primary VM and let the secondary VM assume the same setting. If you enable FT for a virtual machine in a cluster where EVC is disabled, the virtual machine DRS automation level is automatically disabled. Starting in vSphere 6.7, EVC is not required for FT to support DRS.

To enable FT, you first create a VMkernel virtual network adapter on each host and connect to the FT Logging network. You should enable vMotion on a separate VMkernel adapter and network.

When you enable FT protection for a virtual machine, the following events occur:

- If the primary VM is powered on, validation tests occur. If validation is passed, then the entire state of the primary VM is copied and used to create the secondary VM on a separate host. The secondary VM is powered on. The virtual machine's FT status is Protected.
- If the primary VM is powered off, the secondary VM is created and registered to a host in the cluster but not powered on. The virtual machine FT Status setting is Not Protected, VM not Running. When you power on the primary VM, the validation checks occur, and the secondary VM is powered on. Then FT Status changes to Protected.

Legacy FT VMs can exist only on ESXi hosts running on vSphere versions earlier than 6.5. If you require legacy FT, you should configure a separate vSphere 6.0 cluster.

## **vCenter Server High Availability**

vCenter Server High Availability (vCenter HA) is described in [Chapter 1](#), “[vSphere Overview, Components, and Requirements](#).” vCenter HA implementation is covered in [Chapter 8](#), “[vSphere Installation](#).” vCenter HA management is covered in [Chapter 13](#), “[Managing vSphere and vCenter Server](#).”

## VMware Service Lifecycle Manager

If a vCenter service fails, *VMware Service Lifecycle Manager* (vmon) restarts it. VMware Service Lifecycle Manager is a service running in a vCenter server that monitors the health of services and takes preconfigured remediation action when it detects a failure. If multiple attempts to restart a service fail, the service is considered failed.

### Note

Do not confuse VMware Service Lifecycle Manager with VMware vSphere Lifecycle Manager, which provides simple, centralized lifecycle management for ESXi hosts through the use of images and baselines.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

## Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 4-11](#) lists these key topics and the page number on which each is found.



**Table 4-11** Key Topics for [Chapter 4](#)

---

Key Topic Element	Description	Page Number
-------------------	-------------	-------------

---

Section	Network-aware DRS	135
Section	How DRS scores VMs	136
List	DRS migration sensitivity	138
Section	Scalable shares	142
List	vSphere HA requirements	145
<a href="#">Table 4-7</a>	vSphere HA response to failure settings	145
List	vSphere FT requirements	154

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

- [VMware Service Lifecycle Manager](#)
- [vSphere Fault Tolerance \(FT\)](#)
- [Predictive DRS](#)
- [Proactive High Availability \(Proactive HA\)](#)
- [Virtual Machine Component Protection \(VMCP\)](#)

## Review Questions

1. You are configuring EVC. Which of the following is not a requirement?
  - a. A vSphere cluster
  - b. A DRS cluster
  - c. CPUs in the same family
  - d. CPUs with the same base instruction set

2. In vSphere 7.0, you want to configure the DRS Migration Threshold such that it is at the maximum level at which resource contention is considered, but virtual machine happiness is not. Which of the following values should you choose?
  - a. Level 1
  - b. Level 2
  - c. Level 3
  - d. Level 4
  - e. Level 5
  
3. In a vSphere cluster, which of the following statements is true if the primary host detects datastore heartbeats for a secondary host but no network heartbeats or ping responses?
  - a. The primary host declares that the secondary host is isolated.
  - b. The primary host assumes that the secondary host is isolated or in a network partition.
  - c. The primary host takes the host isolation response action.
  - d. The primary host restarts the virtual machines on the failed secondary host.
  
4. You want to configure vSphere HA. Which of the following is a requirement?
  - a. IPv4 must be used for all host management interfaces.
  - b. vMotion must be enabled on each host.
  - c. The Virtual Machine Startup and Shutdown (automatic startup) feature must be enabled on each virtual machine.
  - d. Host IP addresses must persist across reboots.
  
5. You are configuring vSphere Distributed Power Management (DPM) in your vSphere 7.0 environment. Which of the following is not a requirement for using Wake-on-LAN (WoL) in DPM?
  - a. The management NIC must support WOL.
  - b. vMotion is configured.



- c.** The vMotion NIC must support WOL.
- d.** The physical switch port must be set to auto negotiate the link speed.

## Chapter 5

# vCenter Server Features and Virtual Machines

This chapter covers the following topics:

- [vCenter Server and vSphere](#)
- [Virtual Machine File Structure](#)
- [Virtual Machine Snapshots](#)
- [Virtual Machine Settings](#)
- [Virtual Machine Migration](#)
- [Virtual Machine Cloning](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.2, 1.5, 2.3, 5.6, 7.1, 7.2, 7.3, 7.6, 7.9, and 7.11.4.

This chapter provides details on vCenter Server features that have not been covered in previous chapters. It covers virtual machine features such as file structure, migrations, and cloning. [Chapters 13, “Managing vSphere and vCenter Server,”](#) and [14, “Virtual Machine Management/Provision, Migrate, and Replication,”](#) provide details on managing vCenter Server, vSphere, and virtual machines.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 5-1](#) 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](#)

and Review Questions.”

**Table 5-1** “Do I Know This Already?” Section-to-Question Mapping

<b>Foundation Topics Section</b>	<b>Questions</b>
vCenter Server and vSphere	1, 2
Virtual Machine File Structure	3
Virtual Machine Snapshots	4
Virtual Machine Settings	5, 6
Virtual Machine Migration	7– 9
Virtual Machine Cloning	10

1. You just installed a new vCenter Server. Using the vSphere Client, which of the following objects can be the first object that you create in the inventory pane?
  - a. A cluster
  - b. A host
  - c. A virtual machine
  - d. A data center
  - e. A datastore
  - f. A virtual machine folder
2. You want to create a content library for your vCenter Server. Which type of content library cannot be modified directly?
  - a. A library backed by vSAN
  - b. A local library
  - c. A published library
  - d. A subscribed library
3. You are providing support for a virtual machine named Server01 in a vSphere 7.0 environment. Which of the following is the virtual disk data file?
  - a. Server01.vmdk

- b.** Server01-flat.vmdk
  - c.** Server01.vmx
  - d.** Server01-data.vmdk
- 4.** You have taken multiple snapshots for a virtual machine. In the vSphere Client Snapshot Manager, where is the You Are Here icon located?
  - a.** Under the parent snapshot
  - b.** Under the child snapshot
  - c.** Under the latest snapshot
  - d.** Under the associate delta file
- 5.** You are configuring a virtual machine in vSphere 7.0. Which of the following devices cannot be configured or removed?
  - a.** SIO controller
  - b.** SCSI controller
  - c.** Parallel port
  - d.** PCI device
- 6.** You are using the vSphere Client to edit a virtual machine in vSphere 7.0. Which of the following is not available on the VM Options tab?
  - a.** General Options
  - b.** Encryption Options
  - c.** Snapshot Options
  - d.** vApp Options
- 7.** You want to perform a cross-vCenter Server migration in vSphere 7.0. Which of the following statements is true?
  - a.** If separate SSO domains are used, you must use the APIs to perform the migration.
  - b.** If separate SSO domains are used, you can use the vSphere Client to perform the migration.
  - c.** If separate SSO domains are used, you cannot perform the

migration.

- d.** The vSphere and vCenter Server Enterprise licenses are required.
- 8.** You want to perform multiple simultaneous virtual machine migrations for a four-node DRS cluster with a 10 GigE vMotion network and multiple data-stores. Which of the following operations are allowed without any queuing?
  - a.** Nine simultaneous vMotion migrations
  - b.** Nine simultaneous vMotion migrations without Shared Storage
  - c.** One Storage vMotion operation and four vMotion operations
  - d.** Four simultaneous vMotion and five provisioning operations involving the same host
- 9.** You are optimizing your vSphere environment. Which of the following is not helpful for improving vMotion performance?
  - a.** Using NIOC to increase shares for vMotion traffic
  - b.** Using traffic shaping to limit the bandwidth that is available to vMotion traffic
  - c.** Using multiple-NIC vMotion
  - d.** Using jumbo frames
- 10.** You want to use instant clones in vSphere. Which of the following statements is true?
  - a.** You can use the vSphere Host Client to perform an instant clone.
  - b.** You can use the vSphere Client to perform an instant clone.
  - c.** A sample major use case for instant clones is a large-scale deployment in a VMware Horizon VDI.
  - d.** vSphere 6.5 supports instant clones.

## Foundation Topics

### vCenter Server and vSphere

Previous chapters provide details about the vSphere topology, storage infrastructure, network infrastructure, and vSphere clusters. This section provides details about other features, such as the vSphere inventory, host profiles, and content libraries.

## vSphere Managed Inventory Objects

This section describes the **vSphere inventory** and object types, which should be planned prior to implementing vSphere. It provides information on creating and configuring inventory objects during vSphere implementation.

In vSphere, the inventory is a collection of managed virtual and physical objects. Depending on the object type, you can configure each object and perform operations such as setting permissions, monitoring tasks, monitoring events, and setting alarms. You can organize many of these objects by placing them into folders, which makes managing them easier.

All inventory objects except for hosts can be renamed to represent their purposes. For example, they can be named after company departments, locations, or functions.

### Note

Many systems that rely on vCenter Server, such as VMware Horizon, also refer to vCenter objects according to their names. Take care when renaming vCenter inventory objects such as data centers, folders, and datastores if you have deployed any external systems that rely on vCenter Server.

### Note

Inventory object names cannot exceed 214 bytes (UTF-8 encoded).

## Data Centers

In the vSphere inventory, a **data center** is a container object that is an aggregation of all the different types of objects used to work in virtual infrastructure. Other than an optional folder to contain data centers, you cannot

create any object in the inventory until you create a data center.

Data centers are often used to contain all the objects in a physical data center. For example, if you use a single vCenter Server to manage vSphere assets in San Francisco and Chicago, you might want to use corresponding virtual data centers to organize each city's assets. You could create data center objects named San Francisco and Chicago and place each ESXi host, virtual machine, and other object in the appropriate data center.

Within each data center, there are four separate hierarchies:

- Virtual machines (and templates)
- Hosts (and clusters)
- Networks
- Datastores

A data center is a namespace for networks and datastores. The names for these objects must be unique within a data center. You cannot use identical datastore names within the same data center, but you can use identical datastore names within two different data centers. Virtual machines, templates, and clusters do not need to have unique names within the data center but must have unique names within their folder.

## **Folders**

In the vSphere inventory, folders are container objects that allow you to group objects of a single type. A folder can contain data centers, clusters, datastores, networks, virtual machines, templates, or hosts. For example, one folder can contain hosts and a folder containing hosts, but it cannot contain hosts and a folder containing virtual machines.

You can create data center folders directly under the root vCenter Server and use them to organize your data centers. Within each data center is one hierarchy of folders for virtual machines and templates, one for hosts and clusters, one for datastores, and one for networks.

When creating or modifying a folder, the only available setting is the folder name. You can use folders when assigning permissions and configuring alarms.

## **Clusters**

A **cluster** is a set of ESXi hosts that are intended to work together as a unit. When you add a host to a cluster, the host's resources become part of the cluster's resources. vCenter Server manages the resources of all hosts in a cluster as one unit. In addition to creating a cluster, assigning a name, and adding ESXi objects, you can enable and configure features on a cluster, such as VMware EVC, vSphere DRS, and vSphere HA.

If you enable VMware EVC on a cluster, you can ensure that migrations with vMotion do not fail due to CPU compatibility errors. If you enable vSphere DRS on a cluster, you can allow automatic resource balancing by using the pooled host resources in the cluster. If you enable vSphere HA on a cluster, you can allow rapid virtual machine recovery from host hardware failures by using the cluster's available host resource capacity.

Cluster features are covered in detail in [Chapter 4, "Clusters and High Availability."](#)

## Resource Pools

In the vSphere inventory, **resource pools** are container objects that are used to compartmentalize the CPU and memory resources of a host or cluster. Virtual machines run in resource pools, using resources provided by the resource pools. You can create multiple resource pools as direct children of a standalone host or cluster.

You can use resource pools to organize VMs. You can delegate control over each resource pool to specific individuals and groups. You can monitor resources and set alarms on resource pools. If you need a container just for organization and permission purposes, consider using a folder. If you also need resource management, then consider using a resource pool.

If DRS is enabled, you can use the vSphere Client to create resource pools in the cluster and assign resource settings, such as reservations and limits. Otherwise, you can create resource pools directly on specific ESXi hosts.

You can configure resource settings for resource pools, such as reservations, limits, and shares. See [Chapter 4](#) for more details on resource pools.

## Hosts

In the vSphere inventory, hosts are objects that represent your ESXi servers. After installing an ESXi host, you can choose to add it to the vSphere inventory,



which requires you to provide credentials for a user who is assigned the administrator role directly on the host.

The vpxa agent in the ESXi server maintains communication with vCenter Server. It is an interface between the vCenter Server and the ESXi hostd service, which drives the main operations on the host, such as powering on a virtual machine.

For maintenance and troubleshooting activities, you can disconnect a host from the vCenter Server, which does not remove it from vCenter Server but suspends related vCenter Server monitoring activities. You can connect hosts that are disconnected. If you choose to remove a host from inventory, the host and all its associated virtual machines are removed.

If the SSL certificate used by vCenter Server is replaced or changed, the vCenter Server is unable to decrypt the host passwords. You need to reconnect the certificate and resupply the host credentials.

To remove a host from the vSphere inventory, you must first enter Maintenance Mode.

## **Networks**

In the vSphere inventory, networks are objects that are used to connect a set of virtual network adapters. Each ESXi host may have multiple VMkernel virtual network adapters. Each virtual machine may have multiple virtual network adapters. Each virtual network adapter may be connected to a port group (on a standard virtual switch) or a distributed port group (on a vSphere distributed switch). All virtual machines that connect to the same port group belong to the same network in the virtual environment, even if they are on different physical servers. You can manage networks by monitoring, setting permissions, and setting alarms on port groups and distributed port groups.

[Chapter 3, “Network Infrastructure,”](#) provides details on networks.

## **Datastores**

In the vSphere inventory, datastores are objects that represent physical storage resources in the data center. A datastore is the storage location for virtual machine files. The physical storage resources can come from local SCSI disks of the ESXi host, Fibre Channel SAN disk arrays, iSCSI SAN disk arrays, or network attached storage (NAS) arrays. VMFS datastores can be backed by local

SCSI, Fibre Channel, or iSCSI. NFS datastores can be backed by NAS. vSAN datastores can be built in VSAN clusters.

[Chapter 2, “Storage Infrastructure,”](#) provides details on datastores.

## Virtual Machines

In the vSphere inventory, virtual machines are represented in a manner that reflects the current inventory view. For example, in the Hosts and Clusters view, each virtual machine is a descendant of the ESXi host on which it runs. In the Networks view, each virtual machine is a descendant of the network to which it connects.

## Templates

In the vSphere inventory, *templates* are objects that are effectively non-executable virtual machines. A template is a master copy of a virtual machine that can be used to create and provision new virtual machines. A template can have a guest operating system and application software installed. Templates are often customized during deployment to ensure that each new virtual machine has a unique name and network settings.

For more details on templates, see the “[Virtual Machine Cloning](#)” section, later in this chapter.

## vApps

A *vApp* is a container object in vSphere that provides a format for packaging and managing applications. Typically, a vApp is a set of virtual machines that runs a single application and allows you to manage the application as a single unit. You can specify a unique boot order for the virtual machines in a vApp, which allows you to gracefully start an application that spans multiple virtual machines. You can apply resource management settings to a vApp in a similar manner as you would to a resource pool.

## Host Profiles

A *host profile* is a feature that enables you to encapsulate the configuration of one host and apply it to other hosts. A host profile is especially helpful in environments where administrators manage multiple hosts and clusters with vCenter Server. The following are characteristics of host profiles:

- Host profiles are an automated and centrally managed mechanism.
- Host profiles are used for host configuration and configuration compliance.
- Host profiles can improve efficiency by reducing the use of repetitive manual tasks.
- Host profiles capture the configuration of a reference host and store the configuration as a managed object.
- Host profiles provide parameters for configuring networking, storage, security, and other host-level settings.
- Host profiles can be applied to individual hosts, a cluster, or a set of hosts and clusters.
- Host profiles make it easy to ensure that all hosts in a cluster have a consistent configuration.

You can use the following workflow to leverage a host profile to apply a consistent host configuration in your vSphere environment:



- Step 1.** Set up and configure a reference host.
- Step 2.** Create a host profile from the reference host.
- Step 3.** Attach hosts or clusters to the host profile.
- Step 4.** Check the compliance of the hosts with the host profile. If all hosts are compliant with the reference host, you do not need to take additional steps.
- Step 5.** If the hosts are not fully compliant, apply (remediate) the hosts with the host profile.

## Note

If you want a host profile to use directory services for authentication, the reference host must be configured to use a directory service.

In previous releases, vSphere requires that the reference host be available for

certain tasks, such as editing, importing, and exporting the host profile. Starting with vSphere 6.0, a dedicated reference host is no longer required for these tasks. Auto Deploy uses host profiles to configure ESXi.

## Content Libraries

A **content library** is a repository that can be used to share files such as virtual machine templates, vApps, and image files among a set of vCenter Servers. Content libraries, which were introduced in vSphere 6.0, address the fact that multiple vCenter Servers do not directly share associated files such as Open Virtualization Format (OVF) and image (ISO) files. A great use case is companies having multiple sites, each managed by a dedicated vCenter Server, where the OVF files and ISO files that are used at one site are not directly available for use at other sites. In such a case, you can create a content library at one site and publish it to serve the other sites. At the other sites, you can create subscribed libraries that automatically synchronize with the published library. For example, you can create a local content library using the main office vCenter Server, publish it, and subscribe to it from branch office vCenter Servers.

A subscribed content library can be configured to download metadata only whenever it receives notification of a change. In such a case, the subscribing library reflects the most recent changes, but it is not burdened with supplying the storage space for every published file. Instead, the administrator can choose whether to download the data for the entire item or per item.

Three types of content libraries can be used: local, published, and subscribed. A local content library is the simplest form. You can allow, modify, and delete content in a content library. A published library is a local library where content is published for subscription. A subscribed library is a library whose content you cannot change or publish. It receives its content from a published library.

Each content library is built on a single storage entity, which may be a VMFS datastore, an NFS datastore, a CIFS share, a local disk, or a vSAN datastore. In vSphere 7.0, the following maximum limitations apply:

### Key Topic

- 1000 libraries per vCenter Server
- 1000 items per library

- 16 concurrent synchronization operations per published library
- 9 virtual disk files per OVA/OVF template

After one library is set to subscribe to another library, synchronization occurs. Automatic synchronization occurs every 24 hours by default and can be modified using an API. The content library service, which is named `vmware-vdcs`, is installed as part of the vCenter Server installation and uses the same database as vCenter Server.

Simple versioning is used to keep libraries synchronized. Version numbers are assigned to the libraries and to each item in the library. These numbers are incremented whenever content is added or modified. The library does not store previous versions or provide rollback.

The following sequence occurs between a subscribed library and a published library:

- Step 1.** The library service on the subscriber connects to the library services on the publisher by using the VMware Content Subscription Protocol (VCSP) and checks for updates.
- Step 2.** The subscriber pulls the `lib.json` file from the publisher, and each library's `lib.json` files are examined to determine if discrepancies exist between the publisher and the subscriber.
- Step 3.** The library service uses VCSP to determine what data has changed and sends a request to the transfer serviced to copy the required files.
- Step 4.** The subscriber updates the versioning information in the database.

Beginning with vSphere 6.5, you can mount an ISO file directly from the content library, apply a guest OS customization specification during VM deployment, and update existing templates. The content library's performance is then improved. The Optimized HTTP Sync option stores content in a compressed format, which reduces the synchronization time. The content library leverages new features in vCenter Server 6.5, including vCenter HA and backup/restoration.

In previous versions of vSphere, content libraries supported only OVF templates. As a result, virtual machine templates and vApp templates were converted to OVF files when you uploaded them to a content library. Starting with vSphere 6.7 Update 1, content libraries support virtual machine templates. Therefore, templates in the content library can be either the OVF template type or the VM

template type. vApp templates are still converted to OVF files when you upload them to a content library. The distribution of VM templates requires that the respective vCenter Server instances be in Enhanced Linked Mode or Hybrid Linked Mode and that the respective hosts be connected through a network.

To allow a user to manage a content library and its items, you can assign the Content Library administrator role, which is a sample role, to that user as a global permission. Users who are assigned the administrator role at a vCenter Server level cannot see the libraries unless they have a read-only global permission.

## vSphere with Tanzu

By using vSphere with Tanzu, you can use a vSphere cluster as a platform for running Kubernetes workloads in dedicated resource pools. Once enabled on a vSphere cluster, vSphere with Tanzu creates a Kubernetes control plane directly in the hypervisor layer, enabling you to deploy vSphere pods and run your applications inside these clusters.

A vSphere pod, which is the equivalent of a Kubernetes pod, is a small virtual machine that runs one or more Linux containers. The storage and compute for each vSphere pod is sized precisely for its workload, with explicit resource reservations.

To use vSphere with Tanzu, you must use the VMware vSphere 7 Enterprise Plus license with an add-on for Kubernetes for all ESXi hosts that you want to use in a supervisor cluster. You must assign an NSX-T Data Center Advanced or higher license to NSX Manager.

## Virtual Machine File Structure

A virtual machine consists of several files that are stored in a datastore. The key files are the configuration file, virtual disk file, NVRAM setting file, and log file. [Table 5-2](#) provides details for virtual machine files. You can configure virtual machine settings through the vSphere Client, **esxcli**, or the vSphere Web Services SDK.

### Note

Do not directly change, move, or delete virtual machine files without guidance from a VMware Technical Support representative.

**Table 5-2** Virtual Machine Files

<b>File</b>	<b>Description</b>
vmname.vmx	Virtual machine configuration file
vmname.vmxv	Additional virtual machine configuration file
vmname.vmdk	Virtual disk characteristics (metadata) file
vmname-flat.vmdk	Virtual disk data file (commonly called a flat file)
vmname.nvram or nvram	Virtual machine BIOS or UEFI configuration file
vmname.vmsd	Virtual machine snapshot file
vmname.vmsn	Virtual machine snapshot data file
vmname.vswp	Virtual machine swap file
vmname.vms	Virtual machine suspend file
vmware.log	Current virtual machine log file
vmware-#.log	Old virtual machine log file, where # is a number starting with 1

Additional files can be created when you perform specific operations, such as creating snapshots. If you convert a virtual machine to a template, the .vmtx file replaces the virtual machine configuration file (.vmx file).

By default, when you create a virtual machine, the system creates a folder in the datastore and assigns a folder name that is similar to the virtual machine name. In cases where the default folder name is already in use, the system appends a number to the new folder to make it unique.

## Configuration File

A virtual machine's configuration file is a text file that contains all of the virtual machine's settings, including a description of the virtual hardware. For example, a portion of the contents of a VMX file for a CentOS virtual machine named server1 could include the following text:

[Click here to view code image](#)

```
displayName = "server1"
guestOS = "centos-64"
nvram = "server1.nvram"
scsi0:0.fileName = "server1.vmdk"
```

If this virtual machine is sized with two virtual CPUs and 1024 GB memory, the contents of the VMX file may also include the following text:

```
numvcpus = "2"
memSize = "1024"
```

## Virtual Disk Files

The name of the VMDK file that contains metadata for a virtual disk is included in the VMX file as shown in the previous example (`scsi0:0.fileName = "server1.vmdk"`). The VMDK metadata file is a text file that contains details about the virtual disk, such as the numbers of cylinders, heads, and sectors, as shown in the following sample content:

[Click here to view code image](#)

```
ddb.geometry.cylinders = "1305"
ddb.geometry.heads = "255"
ddb.geometry.sectors = "63"
```

The VMDK metadata file also contains the names of other files associated with the virtual disk, such as data (extent) files, as shown in the following sample content:

```
# Extent description
RW 20971520 VMFS "server1-flat.vmdk"
```

## Snapshot Files

When you take a snapshot of a virtual machine, the system creates a few files. For example, if you take a snapshot for a powered-off virtual machine named `server1` that has only one virtual disk and no previous snapshots, the following files may be created:



- **server1-000001-sesparse.vmdk**: A delta data disk that stores changes made since the creation of the snapshot
- **server1-000001.vmdk**: A VMDK metadata file for the delta disk



- **server1-Snapshot1.vmsn**: Snapshot data

The following section provides more details on virtual machine snapshots.

## Virtual Machine Snapshots

A **virtual machine snapshot** captures the state of a virtual machine and the data in the virtual machine at a specific point in time. Snapshots are useful when you want to return the state of a virtual machine to a point that was previously captured. For example, you can create a snapshot of a virtual machine just prior to installing and testing software in the virtual machine. This enables you to revert the virtual machine back to its original state when you finish testing.

You can take multiple snapshots of a virtual machine. If you take multiple snapshots without reverting the virtual machine, the snapshots are created in a linear fashion, as shown in [Figure 5-1](#). The vSphere Client represents the snapshot hierarchy of a virtual machine as a tree with the root node being the virtual machine and nodes being each snapshot. If you revert the virtual machine to a snapshot, the state of your virtual machine is associated with that snapshot, as shown in [Figure 5-2](#). If you create another snapshot, you add branches to the snapshot tree, as shown in [Figure 5-3](#).

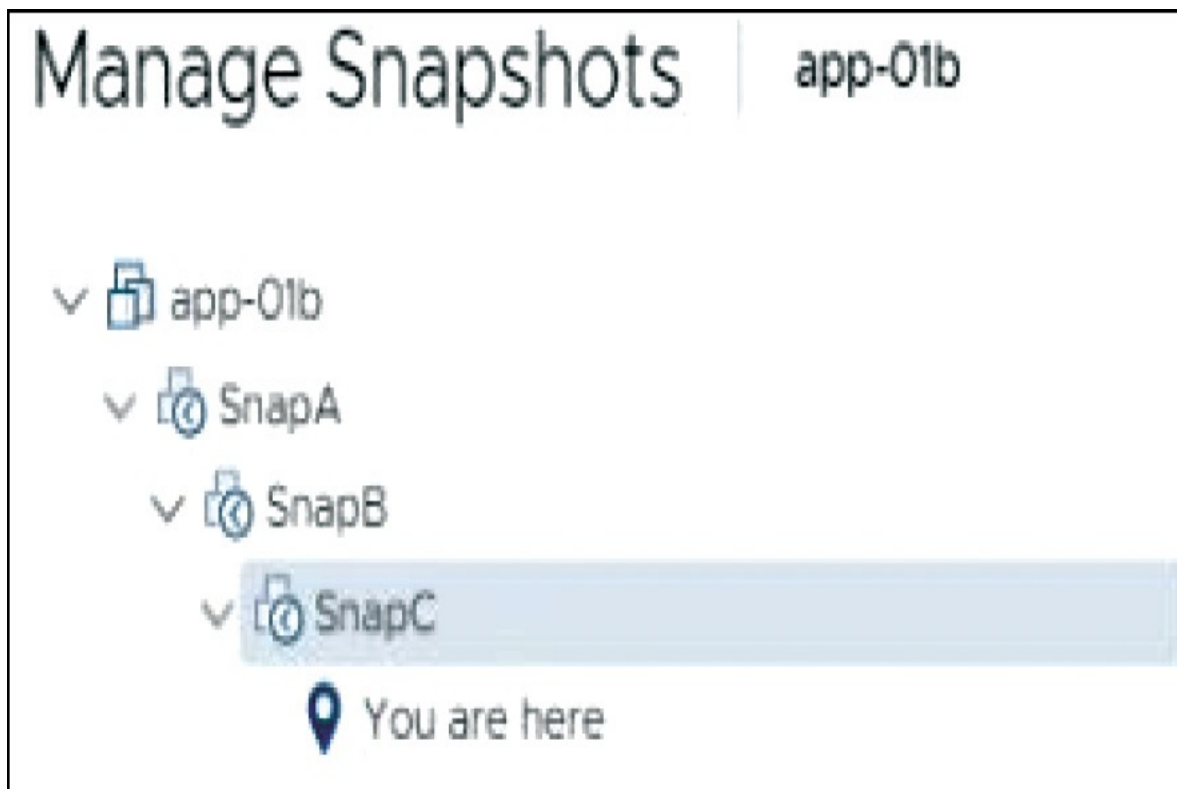
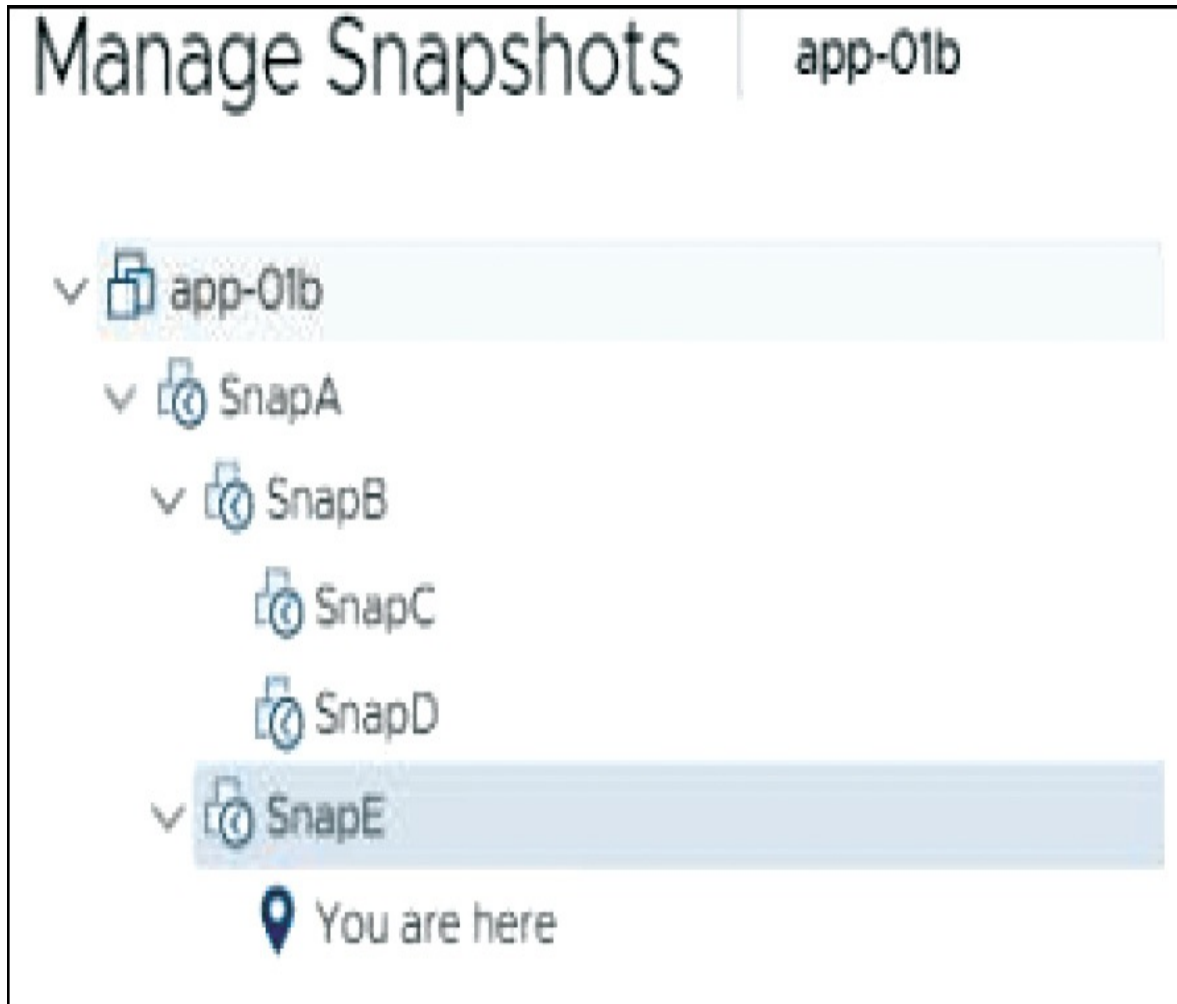


FIGURE 5-1 Linear Snapshots



FIGURE 5-2 Post-Revert Snapshot Tree



**FIGURE 5-3** New Branch in Snapshot Tree

Each branch in a snapshot tree can have up to 32 snapshots.

In the vSphere Client, you can perform several snapshot operations, including taking a snapshot, reverting to a snapshot, and deleting a snapshot. When taking a snapshot, you can choose whether to snap the memory and whether to quiesce the guest OS. In cases where no snapshot exists but delta files exist, you can choose to consolidate the disks.

## Snapshot Use Cases

The following are some common use cases for snapshots:



- **Rollback changes:** Prior to upgrading or making a configuration change to an application, you can take a virtual machine snapshot to provide a rollback option.
- **Rollback guest OS upgrade:** Prior to upgrading the guest OS, you can take a virtual machine snapshot to provide a rollback option.
- **Training and development labs:** You can take snapshots of a set of virtual machines used in a lab environment prior to allowing user access. When the user finishes experimenting, you can revert the state of the environment back to the original state for the next user.
- **Backups:** A backup utility may first trigger a virtual machine snapshot and then copy the virtual machine files without needing to deal with open files. Following the backup, the utility deletes the snapshot.
- **Troubleshooting and triage:** Taking a snapshot of a troubled virtual machine enables you to later choose to return the virtual machine to the exact state when it experienced an issue.
- **Linked clones:** Automation and virtual desktop software, such as vRealize Automation and Horizon, may leverage virtual machine snapshots to allow you to use fast provisioning (linked clone) methods by which new virtual machines share a base virtual disk. For example, to use a linked clone in a vRealize Automation blueprint, you need to identify a virtual machine snapshot.

## What a Snapshot Preserves

A snapshot preserves the following information:

- **Virtual machine settings:** The virtual machine directory, which includes the disks added or changed after you take the snapshot
- **Disk state:** The state of each virtual disk
- **Memory state:** (Optional) The contents of the virtual machine's memory
- **Power state:** Whether the virtual machine is powered on, powered off, or suspended when you take the snapshot (If you revert to a snapshot that includes the memory state, the virtual machine is returned to its preserved power state.)

## Parent Snapshots

The first virtual machine snapshot that you create is the base snapshot. Taking a snapshot creates a delta disk file for each disk attached to the virtual machine and, optionally, a memory file. The delta disk files and memory file are stored with the base VMDK file. The parent (current) snapshot is always the snapshot that appears immediately above the You Are Here icon in the Snapshot Manager. If you revert to a snapshot, that snapshot becomes the parent of the You Are Here current state. When you have multiple snapshots, each child snapshot has a parent snapshot.

### Note

The parent snapshot is not always the snapshot that you took most recently.

## Snapshot Behavior

Taking a snapshot preserves the disk state by creating a series of delta disks for each attached virtual disk or virtual raw device mapping (RDM). Taking a snapshot creates a snapshot object in the Snapshot Manager that represents the virtual machine state and settings. Each snapshot creates a delta disk for each virtual disk. When you take a snapshot, the system prevents the virtual machine from writing to the current data (VMDK) file and instead directs all writes to the delta disk. The delta disk represents the difference between the current state of the virtual disk and the state that existed at the time that you took the parent snapshot. Delta disk files can expand quickly and can become as large as the configured size of the virtual disk if the guest operating system writes to every block of the virtual disk.

When you take a snapshot, the state of the virtual machine, virtual disks, and (optionally) virtual memory is captured in a set of files, such as the delta, database, and memory files. By default, the delta disks are stored with the corresponding virtual disk files, and the memory and database files are stored in the virtual machine directory.

## Flat File

A virtual disk involves a metadata file and a data file, each with the .vmdk

extension. The metadata VMDK file contains information about the virtual disk, such as geometry and child–parent relationship information. The data VMDK file is called the flat file, and its name contains the word flat. Only the names of the metadata files appear in the vSphere Client Datastore Browser. In normal circumstances, the virtual machine’s guest OS and applications write to the flat file.

## **Delta Disk Files**

When you create a snapshot, you create a delta disk for each virtual disk. The delta (child) disk represents the difference between the current state of the virtual disk and the state that existed at the time that you took the parent snapshot. A delta disk has two VMDK files. One is a small metadata file, and the other is a data file. Delta disk data files are also called redo logs.

## **Database File**

The database file is a file with the .vmsd extension that contains snapshot details required by the Snapshot Manager. It contains details on the relationships between snapshots and child disks.

## **Memory File**

The memory file is a file with the .vmsn extension that includes the active state of the virtual machine’s memory. Capturing the memory state of the virtual machine lets you revert to a powered-on state. Memory snapshots take longer to create than nonmemory snapshots. The size of the memory impacts the amount of time required to create the snapshot.

## **Limitations**

The use of snapshots can impact a virtual machine’s performance and can be limited in some scenarios, as summarized in the following list:

- Snapshots are not supported for RDM physical mode disks or for iSCSI initiators in a guest OS.
- Snapshots of powered-on or suspended virtual machines with independent disks are not supported.
- A quiesced snapshot requires a supported guest operating system and active

VMware Tools services.

- Snapshots are not supported with PCI vSphere DirectPath I/O devices.
- Snapshots are not supported for virtual machines configured for bus sharing.
- Although snapshots may be a useful step for a backup utility, a snapshot is not a backup by itself. A snapshot does not provide a redundant copy of data. If the base flat file is lost or corrupted, you cannot restore the virtual machine by reverting to a snapshot.
- Snapshots can negatively affect the performance of a virtual machine. The performance degradation is impacted by factors such as the age of the snapshot, the depth of the snapshot tree, and the amount of data in the delta files.
- Snapshot operations can take much longer to finish when they involve virtual disks larger than 2 TB.
- Deleting a large snapshot that is part of the current path (as indicated by You Are Here in the Snapshot Manager) can negatively impact the performance and the health of the virtual machine. To minimize risk, you can shut down the virtual machine prior to deleting the snapshot.

## Virtual Machine Settings

This section provides information on virtual machine settings in vSphere 7.0.

### VM Hardware/Compatibility

You can configure a virtual machine's compatibility setting to control which ESXi host versions can be used to run the virtual machine. In the vSphere Client, you can set the Compatible With option for a virtual machine to a compatible ESXi version, such as ESXi 7.0 and later or ESXi 6.7 Update 2 and later. The compatibility setting determines which ESXi host versions the virtual machine can run on and the hardware features available to the virtual machine. At the host, cluster, or data center level, you can set the Default VM Compatibility setting. (See [Chapter 14](#) for details.)

Virtual hardware devices perform the same function for the virtual machines as physical hardware devices do for traditional servers. Each virtual machine has CPU, memory, and disk resources. All modern operating systems provide support for virtual memory, allowing software to use more memory than is

present in the server hardware. Similarly, ESXi can provide to its virtual machines VM memory totaling more than the capacity of the host's physical memory.

You can add virtual hardware devices to a virtual machine by editing the virtual machine's settings in the vSphere Client. Not all devices are configurable. For example, the PCI and SIO virtual hardware devices are part of the virtual motherboard but cannot be configured or removed. You can enable the Memory Hotplug or CPU Hotplug settings in order to add memory or CPU resources to a running virtual machine. Memory Hotplug is supported on all 64-bit operating systems, but some guest operating systems may not be able to make use of the added memory without restarting. The ESXi license and other factors for the host where the virtual machine runs may impact the available devices for the virtual machine. For a list of hardware devices and their functions, see [Table 5-3](#).

**Table 5-3** Virtual Machine Hardware Devices

Device	Description
CPU	At least one vCPU but not more than the number of logical CPUs in the host.  You can set advanced CPU features, such as the CPU identification mask and hyperthreaded core sharing.
Chip set	The virtual motherboard consists of VMware-proprietary virtual devices, based on the following chips: <ul style="list-style-type: none"><li data-bbox="321 1444 1170 1480">■ Intel 440BX AGPset 82443BX host bridge/controller</li><li data-bbox="321 1570 1089 1606">■ Intel 82371AB (PIIX4) PCI ISA IDE Xcelerator</li><li data-bbox="321 1696 1409 1774">■ National Semiconductor PC87338 ACPI 1.0 and PC98/99 Compliant SuperI/O</li><li data-bbox="321 1864 1328 1900">■ Intel 82093AA I/O advanced programmable interrupt controller</li></ul>



DVD /CD- ROM M drive	<p>One by default.</p> <p>You can configure the virtual DVD/CD-ROM device to connect to client devices, host devices, or datastore ISO files.</p> <p>You can add and remove virtual DVD/CD-ROM devices.</p>
Hard disk	A virtual disk is backed by a set of files, as discussed earlier in this chapter.
IDE 0, IDE 1	Two virtual Integrated Drive Electronics (IDE) interfaces are present by default.
Keyb oard	The virtual keyboard is mapped to the user keyboard when you connect to the virtual machine console.
Mem ory	The size of the virtual memory becomes the size of memory that the guest OS perceives to be physical memory.
Netw ork adapt er	You can configure the number of virtual network adapters (NICs) and the adapter type used by each virtual machine.
Paral lel port	You can add, remove, and configure virtual parallel ports.
PCI contr oller	One PCI controller, which is located on the virtual motherboard, is presented to the virtual machine. It cannot be removed or modified.
PCI devic e	If you configured devices to be reserved for PCI passthrough on the host, you can add up to 16 PCI vSphere DirectPath devices to a virtual machine.
Point ing devic	The virtual pointing device is mapped to the user's pointing device when you connect to the virtual machine console.

e

---

Serial port	You can configure a virtual machine with up to 32 virtual serial ports.
-------------	---

You can add, remove, or configure virtual serial ports.

---

SATA	Provides access to virtual disks and DVD/CD-ROM devices.
------	--

controller	The SATA virtual controller appears to the guest OS as an AHCI SATA controller.
------------	---

---

SCSI	Provides access to virtual disks.
------	-----------------------------------

controller	The SCSI virtual controller appears to the guest OS as different types of controllers, including LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual.
------------	---

---

SIO	Provides serial and parallel ports and floppy devices and performs system management activities.
-----	--

One SIO controller is available to the virtual machine, but it cannot be configured or removed.

---

USB	The virtual USB controller is the software virtualization of the USB host controller function in the virtual machine.
-----	---

---

USB device	You can add multiple virtual USB devices to a virtual machine that you map to USB devices connected to an ESXi host or a client computer.
------------	---

---

VMCI	The Virtual Machine Communication Interface (VMCI) device provides a high-speed communication channel between a virtual machine and the hypervisor.
------	---

You cannot add or remove VMCI devices.

---

NVMe	NVM Express (NVMe) is a logical device interface specification for accessing non-volatile storage media attached through a PCI Express (PCIe) bus in real and virtual hardware.
------	---

---

controller	
------------	--

NVD IMM contr oller	Provides access to the non-volatile memory resources of the host.
NVD IMM devic e	You can add up to 64 virtual non-volatile dual in-line memory module (NVDIMM) devices to a virtual machine.
TPM devic e	You can add a virtual Trusted Platform Module (TPM) 2.0 device to a virtual machine to allow the guest OS to store sensitive information, perform cryptographic tasks, or attest the integrity of the guest platform.

## Virtual Disk Provisioning

You can configure the provisioning type for a virtual disk to thin provisioned, lazy zeroed thick provisioned, or eager zeroed thick provisioned, as described in the section “Virtual Disk Type” in [Chapter 2](#):

- With thin provisioning, storage blocks are not allocated during disk creation, which allows fast provisioning but requires allocation and zeroing during runtime.
- With thick eager zeroed, storage blocks are allocated and zeroed during provisioning, which allows fast runtime.
- With thick lazy zeroed provisioning, storage blocks are pre-allocated but not pre-zeroed.

Your choice for the provisioning type depends on each virtual machine’s use case. For example, if you want to minimize the virtual machine startup time and minimize its risk, you may choose thick provision lazy zeroed.

## VMware Tools

**VMware Tools** is a set of software modules and services, including services that can communicate with the VMkernel. This communication allows integration with vSphere for activities such as customizing the guest OS, running scripts in the guest OS, and synchronizing time. If you use guest operating systems without VMware Tools, many VMware features are not available. VMware

Tools enhances the performance of the guest OS by enabling the latest drivers for virtual devices, enabling memory functions (such as ballooning), and more. It includes drivers such as SVGA, Paravirtual SCSI, VMXNet NIC, mouse, audio, guest introspection, and memory control drivers. Prior to upgrading the hardware for a virtual machine, you should upgrade VMware Tools.

VMware Tools includes the VMware user process named **vmtoolsd**, which enables copy and paste and mouse control and automatically sets screen resolution for some non-Windows guests. It enhances the performance of the virtual machine's guest operating system and improves management of the virtual machine. It includes device drivers and other software that is essential for the VM. With VMware Tools, you have more control over the virtual machine interface.

## Virtual Machine Options

To edit a virtual machine setting, you can navigate to and manipulate settings on the VM Options tab. Many of these options have dependencies with the ESXi hosts, data centers, clusters, or resource pools on which the virtual machine resides. [Table 5-4](#) describes the available virtual machine options.

**Table 5-4** Virtual Machine Options

Category	Description
General Options	Settings include virtual machine name, configuration file location, and the working directory location.
Encryption Options	Settings allow you to enable or disable virtual machine encryption or vMotion encryption.
Power Management	Settings allow you to choose how to respond when the guest OS is placed on standby. The choices are to suspend the virtual machine or put the guest OS into standby mode.
VMware Tools	Settings allow you to choose how to respond to specific power operations. For example, you can choose whether to power off the virtual machine or shut down the guest when the red power-off button is clicked.
Virtualization	For virtual machines running the modern Windows OS versions, you can enable VBS to add an extra level of protection.

Based  
Security  
(VBS)

---

Boot Options	Settings include firmware, boot delay, and failed boot recovery parameters.
Advanced Options	Settings include logging, debugging, swap file location, and configuration parameters.
Fibre Channel NPIV	Settings allow the virtual machine to use N_Port ID Virtualization (NPIV), including whether to generate new worldwide names (WWNs).
vApp Options	Settings allow you to control vApp functionality for the virtual machine, such as enable/disable and IP allocation policy. vApp settings that are made directly to a virtual machine override settings made on the vApp.

---

## Virtual Machine Advanced Settings

As indicated in [Table 5-4](#), you can use the vSphere Client to edit the advanced settings for a virtual machine on its VM Options tab. You can set a virtual machine's advanced settings to enable or disable logging. You can enable or disable hardware acceleration. You can set debugging and statistics to Run Normally, Record Debugging Information, Record Statistics, or Record Statistics and Debugging. For applications that are highly sensitive to latency, you can set Latency Sensitivity to High.

Under Advanced Settings, you can select Configuration Parameters, where you can directly manipulate the virtual machine's low-level settings, as illustrated in [Figure 5-4](#).

# Configuration Parameters



pciBridge0.present	TRUE
svga.present	TRUE
pciBridge4.present	TRUE
pciBridge4.virtualDev	pcieRootPort
pciBridge4.functions	8
pciBridge5.present	TRUE
pciBridge5.virtualDev	pcieRootPort
pciBridge5.functions	8
pciBridge6.present	TRUE
pciBridge6.virtualDev	pcieRootPort
pciBridge6.functions	8
pciBridge7.present	TRUE
pciBridge7.virtualDev	pcieRootPort
pciBridge7.functions	8

CANCEL OK

**FIGURE 5-4** Sample Virtual Machine Configuration Parameters

## Virtual Machine Migration

This section provides information such as concepts, prerequisites, and data flow details for each migration type. [Chapter 14](#) provides information on performing the migrations.

### Virtual Machine Migration

You can migrate virtual machines from one compute resource or storage location to another while the virtual machine is stopped (cold) or running (hot). For example, if you want to balance the workload, you can migrate some virtual machines from busy ESXi hosts or datastores (or both) to other hosts and datastores. As another example, if you want to perform maintenance (such as an upgrade), you can migrate all virtual machines from an ESXi host or datastore, perform the maintenance, and optionally migrate virtual machines back to the original location.

Moving a virtual machine between the inventory folder or resource pools in the same data center is not considered a migration. Cloning and copying virtual machines are also not forms of migration.

Each migration type involves a unique set of requirements, such as minimum privileges required to perform the operation.

#### Note

To migrate virtual machines with disks larger than 2 TB, the source and destination ESXi hosts must be Version 6.0 and later.

### Cold Migrations

Moving a powered-off or suspended virtual machine to a new host, new datastore, or both is considered a cold migration. The required privilege is Resource.Migrate Powered Off Virtual Machine.

### Hot Migrations

Moving a powered-on virtual machine to a new host, new datastore, or both is considered a hot migration. During the migration, vCenter Server must take steps to ensure that active connections and services of the virtual machine are not interrupted.

## **Cross-Host Migrations**

Moving a virtual machine, whether hot or cold, to a new host is considered a cross-host migration. In vSphere Client wizards that involve cross-host migrations, you can choose a destination host. Alternatively, when available and properly configured, you can choose a DRS cluster, resource pool, or vApp as the destination.

The cross-host migration wizards include a Compatibility panel to identify any compatibility issues or warnings. If the panel displays the message “Compatibility Checks Succeeded,” you can proceed with no concern. If the panel displays an error, the migration is disabled for the associated hosts. If it displays a warning message, the migration is not disabled, and you can proceed, bearing in mind the warning. For hot migrations, the compatibility check accommodates vMotion CPU compatibility checking.

For a virtual machine using an NVDIMM device and PMem storage, the destination host or cluster must have available PMem resources to pass the compatibility check. For a cold migration involving a virtual machine that does not have an NVDIMM device but uses PMem storage, you can choose a target host or cluster without available PMem resources. The hard disks use the storage policy and data-store selected for the virtual machine’s configuration files.

## **Cross-Datastore Migrations**

Moving a virtual machine, whether hot or cold, to a new datastore is considered a cross-datastore migration.

## **Cross-vCenter Server Migrations**

Moving a virtual machine, whether hot or cold, to a new vCenter Server is considered a cross-vCenter Server migration. To perform a cross-vCenter Server migration, you must meet the following requirements:





- The associated vCenter Servers and ESXi hosts must be 6.0 or later.
- The cross-vCenter Server and long-distance vMotion features require an Enterprise Plus license.
- The vCenter Server instances must be time-synchronized with each other for correct vCenter Single Sign-On token verification.
- For migration of compute resources only, both vCenter Server instances must be connected to the shared virtual machine storage.
- When using the vSphere Client, both vCenter Server instances must be in Enhanced Linked Mode, and they must be in the same vCenter Single Sign-On domain.

## Note

If the vCenter Server instances exist in separate vCenter Single Sign-On domains, you can use vSphere APIs/SDK to migrate virtual machines.

## Virtual Machine Migration Limitations



vCenter Server limits the number of simultaneous virtual machine migration and provisioning operations that occur per host, network, and datastore. Each of the network, datastore, and host limits must be satisfied for the operation to proceed. vCenter Server uses a costing method by which each migration and provisioning operation is assigned a cost per resource. Operations whose costs cause resources to exceed their limits are queued until other operations complete.

Limits depend on the resource type, ESXi version, migration type, and other factors, such as network type. ESXi Versions 5.0 to 7.0 have consistent limits:

- **Network limits:** Network limits apply only to vMotion migrations. Each vMotion migration has a network resource cost of 1. The network limit depends on the network bandwidth for the VMkernel adapter enabled for vMotion migration. For 1 GigE the limit is 4, and for 10 GigE it is 8.
- **Datastore limits:** Datastore limits apply to vMotion and Storage vMotion migrations. Each vMotion migration has a resource cost of 1 against the

shared datastore. Each Storage vMotion migration has a resource cost of 16 against both the source and destination datastores. The datastore limit per datastore is 128.

- **Host limits:** Host limits apply to vMotion, Storage vMotion, and cold migrations. They also apply to virtual machine provisioning operations, including new deployments, and cloning. Provisioning and vMotion operations have a host cost of 1. Storage vMotion operations have a host cost of 4. The host limit per host is 8.

For costing purposes, a hot migration that is both a cross-host and cross-datastore migration (vMotion migration without shared storage) is considered to be a combination of a vMotion and Storage vMotion migration and applies the associated network, host, and datastore costs. vMotion migration without shared storage is equivalent to Storage vMotion migration with a network cost of 1.

Consider the following examples for a four-node DRS cluster with a 10 GigE vMotion network:

- If you perform nine simultaneous vMotion migrations, the ninth migration is queued due to the network limit, even if different hosts are involved.
- If you perform nine simultaneous hot cross-host and cross-datastore migrations involving the same datastore, the ninth migration is queued due to the datastore limit, even if the migrations are split as to whether the datastore is the source or the target.
- You can simultaneously perform one Storage vMotion and four vMotion operations involving a specific host.

## TCP/IP Stacks

You can use the vMotion TCP/IP stack to isolate vMotion traffic and assign it to a dedicated default gateway, routing table, and DNS configuration. To use the vMotion TCP/IP stack, select vMotion from the TCP/IP Stack drop-down menu when configuring the associated VMkernel virtual network adapter. When you assign a VMkernel virtual network adapter to the vMotion stack, you cannot use the adapter for purposes other than vMotion. Likewise, you can use the provisioning TCP/IP stack to isolate traffic for cold migration, cloning, and snapshots. To use the provisioning TCP/IP stack, select Provisioning from the TCP/IP Stack drop-down menu when configuring the associated VMkernel virtual network adapter. When you assign a VMkernel virtual network adapter to

the provisioning stack, you cannot use the adapter for purposes other than provisioning.

## vMotion Details

This section provides details on the vMotion feature in vSphere.

### vMotion Overview

A hot cross-host migration is called a **vMotion** migration. A hot migration across hosts and datastores is often called a vMotion migration without shared storage. A hot cross-vCenter Server migration is often called a cross-vCenter Server vMotion migration. Although the term *vMotion migration* may be used to describe any hot cross-host migration, this section provides details on just the traditional vMotion migration, in which shared storage is used and cross-datastore migration does not occur.

During a vMotion migration, the entire state of the virtual machine is moved to the new host. The state includes the current memory content and all the information that defines and identifies the virtual machine. The memory content includes the components of the operating system, applications, and transaction data that are in the memory. The state includes all the data that maps to the virtual machine hardware elements, such as BIOS, devices, CPU, MAC addresses for the Ethernet cards, chipset states, and registers. The associated virtual disk remains in the original location on storage that is shared between the source and destination hosts. After the virtual machine state is migrated to the destination host, the virtual machine continues execution on the destination host.

### vMotion Requirements



As explained in the section “Enhanced vMotion Compatibility (EVC)” in [Chapter 4](#), vMotion requires that the destination host’s processors be compatible with the source host’s processors to support live migration. Specifically, the destination processors must come from the same family and provide the same instruction set as the source processors. You can enable EVC in the cluster to broaden the vMotion compatibility.

Before using vMotion, you must address its host configuration requirements. Each host must meet the licensing, shared storage, and networking requirements

for vMotion.

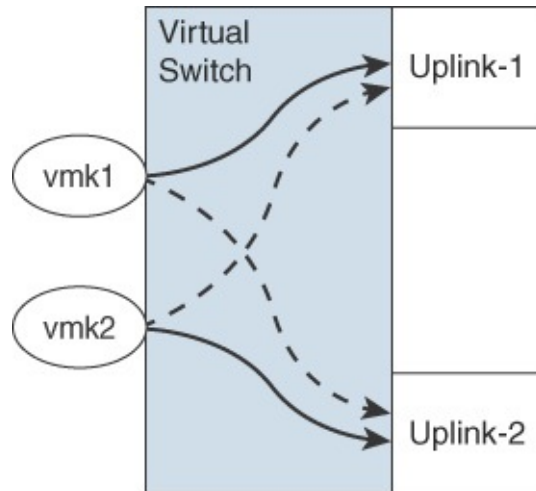
For standard vMotion migration, you must configure the source and destination hosts with shared storage to enable the migrated virtual machines to remain in the same datastore throughout the migration. Shared storage may be implemented with Fibre Channel, iSCSI, or NAS storage. The datastore may be VMFS or NFS. You can also leverage a vSAN datastore to meet the shared storage requirement for vMotion migrations between cluster members.

### **Note**

Hot migrations that are cross-host and cross-datastore migrations, which are often called vMotion migrations without shared storage, do not require shared storage.

For vMotion migration, you must configure each host with a VMkernel virtual network interface connected to a virtual switch with an uplink that uses at least one physical network interface card (NIC). VMware recommends that the network connection be made to a secured network. The vMotion network must provide at least 250 Mbps of dedicated bandwidth per concurrent vMotion session. For long-distance migrations, the maximum supported network round-trip time for vMotion migrations is 150 milliseconds. For faster vMotion migrations, consider using 10 Gbps NICs instead of 1 Gbps NICs.

To improve vMotion migration times even further, consider implementing multi-NIC vMotion. With multi-NIC vMotion, multiple paths are used simultaneously to carry the vMotion workload. To configure multi-NIC vMotion, you can enable vMotion traffic for two VMkernel virtual network adapters that are configured to use separate paths. For example, you can apply the following steps to enable multi-NIC vMotion, as illustrated in [Figure 5-5](#):



**FIGURE 5-5** Multi-NIC vMotion

- Step 1.** On a virtual switch, attach two uplink adapters connected to the vMotion network.
- Step 2.** Connect two VMkernel adapters enabled for vMotion.
- Step 3.** For the first VMkernel adapter, set the first uplink path to Active and the second uplink path to Standby.
- Step 4.** For the second VMkernel adapter, set the first uplink path to Standby and the second uplink path to Active.

For more vMotion performance improvements, you can use Network I/O Control (NIOC) to guarantee network bandwidth to vMotion traffic. You can also use jumbo frames. To avoid network saturation, you can use traffic shaping to limit the average and peak bandwidth available to vMotion traffic.

By default, you cannot use vMotion to migrate a virtual machine that is attached to a standard switch with no physical uplinks. To change this behavior, you can set the `config.migrate.test.CompatibleNetworks.VMOnVirtualIntranet` advanced settings of vCenter Server to False.

### Note

During a vMotion migration without shared storage, the virtual disk data is transferred over the vMotion network.

## vMotion Migration and Data Flow Details

During a vMotion migration, the state of the running virtual machines is copied to the destination host over the designated vMotion network, the virtual machine is stopped on the source ESXi host, and the VM is resumed on the target ESXi host. The process involves the following phases:

- **Compatibility check:** Intended to ensure that requirements are met and that the destination host can run the virtual machine.
- **Pre-copy:** Briefly stuns the source memory and starts memory trackers. Copies memory page from source to target. Tracks which source pages are modified after the pre-copy so these pages (dirty pages) can be re-sent later.
- **Iterations of Pre-copy:** If dirty pages exist, repeats the pre-copy of just the dirty pages and scans for new dirtied pages. Continue iteration until no dirty pages exist or until vMotion determines that the final page copy can be completed in less than 500 ms.
- **Switchover:** Quiesces and suspends the virtual machine execution on the source host, transfers checkpoint data, and starts the execution of the virtual machine using the checkpoint data on the target host.

The stun time (the time at which the virtual machine is not running anywhere) is typically between 100 ms and 200 ms.

## Encrypted vMotion

When migrating encrypted virtual machines, vSphere vMotion always uses encryption. For non-encrypted virtual machines, you can select one of the following encrypted vMotion options:

- **Disabled:** Do not use encryption.
- **Opportunistic:** Use encryption if the source and destination hosts support it.
- **Required:** If the source or destination host does not support encrypted vMotion, do not allow the migration.

### Note

Only ESXi Versions 6.5 and later use encrypted vSphere vMotion. To use vMotion to migrate encrypted virtual machines across vCenter Server instances, you must use the vSphere API.

## Storage vMotion Details

This section provides details on the Storage vMotion feature in vSphere.

### Storage vMotion Overview

**Storage vMotion** migration is a hot cross-datastore migration. Storage vMotion enables you to migrate a virtual machine and its disk files from one datastore to another while the virtual machine is running. Use cases for Storage vMotion include preparing for datastore maintenance (such as upgrading the underlying storage array), optimizing performance (redistribution of storage load), and transforming the virtual disk provisioning type. When you use Storage vMotion on a virtual machine, you can migrate all the virtual machine files to a single location, migrate individual virtual disks, or separate virtual disks from other virtual machine files.

#### Note

Migration with Storage vMotion changes virtual machine files on the destination datastore to match the inventory name of the virtual machine. The migration renames all virtual disk, configuration, snapshot, and NVRAM files. If the new names exceed the maximum filename length, the migration fails.

### Storage vMotion Requirements and Limitations

The following are the major requirements and limitations for Storage vMotion in vSphere 7.0:

- Virtual disks in nonpersistent mode are not supported for Storage vMotion. For virtual compatibility mode RDMs, you can migrate just the mapping file or include the migration of the data to a virtual disk file. For physical compatibility mode RDMs, you can only migrate the mapping file.

- Storage vMotion migration is not supported during VMware Tools installation.
- You cannot use Storage vMotion to migrate virtual disks larger than 2 TB from a VMFS Version 5 datastore to a VMFS Version 3 datastore.
- The source host that is running must have a license that includes Storage vMotion.
- ESXi 4.0 and later hosts do not require vMotion configuration to perform Storage vMotion migrations.
- The host on which the virtual machine is running must have access to both the source and target datastores.

## Storage vMotion Data Flow Details

The following major steps are automatically performed by Storage vMotion in vSphere 7.0:

- Step 1.** The virtual machine's home directory is copied to the destination datastore.
- Step 2.** A hidden (shadow) virtual machine starts using the copied files. The underlying processes (worlds) are visible to the **esxtop** utility. The virtual machine continues to run in preexisting worlds.
- Step 3.** An initial copy of the source virtual disk is made to the destination data-store, and change block tracking (CBT) is leveraged to track blocks that are changed after they are copied.
- Step 4.** Step 4 is repeated until the number of changed blocks is small enough to support a fast switchover.
- Step 5.** The system invokes a fast suspend and resume operation that transfers the running virtual machine to the idling hidden virtual machine. The virtual machine now runs in the new worlds. The preexisting worlds that were associated with the virtual machine are removed.

## Virtual Machine Cloning

vSphere provides many cloning options, as described in this section.

### Clones



When you clone a virtual machine, vCenter Server creates a virtual machine that is a copy of the original virtual machine. The virtual disk files, configuration file, and other files are copied from the original virtual machine to the new virtual machine. The new virtual machine is commonly referred to as a *clone*. The new virtual machine files are named and stored based on parameters you provide during the deployment. You can choose to make some configuration changes and customizations during the cloning process. The contents of some of the files, such as the configuration file, are modified. At the end of the operation, you can manage both the original virtual machine and the new virtual machine as inventory objects in vCenter Server.

## Cold Clones

A cold clone occurs when the source virtual machine is powered down prior to starting the clone operation. In this case, vCenter Server does not have to worry about interrupting the execution of the source virtual machine.

## Hot Clones

A hot clone occurs when the source virtual machine is running during a clone operation. In this case, the vCenter Server must avoid disrupting the execution of the source virtual machine. To do so, it takes a virtual machine snapshot prior to copying data and removes the snapshot at the end of the operation.

## Linked Clones

A linked clone is a virtual machine that is cloned in such a manner that it shares its virtual disk files with the original virtual machine (parent). The shared files are static. Much like a virtual machine that has a snapshot, a linked clone writes its virtual disk changes to separate data files. Compared to a full clone, a linked clone operation is faster and conserves disk space. You cannot use the vSphere Client to directly create linked clones. You can use PowerCLI (via the - **LinkedClone** parameter with the **New-VM** command) or other VMware products to create linked clones. For example, in VMware Horizon you can create desktop pools based on linked clones, and in vCloud Director you can use fast provisioning.

## Rapid Provisioning with Templates

As stated previously, templates are objects in the vSphere inventory that are

effectively non-executable virtual machines.

You can convert a virtual machine to a template and vice versa. But the main use case for templates is for rapid deployment of new similar virtual machines from a single template. In such a case, you are effectively cloning the template again and again, allowing the template to remain unchanged and ready for future use. To update a template, such as to install the most recent guest OS updates, you can temporarily convert the template to a virtual machine, apply the updates, and convert back to a template.

When you deploy a virtual machine from a template, vCenter Server creates a virtual machine that is a copy of the original template. The virtual disk files, configuration file, and other files are copied from the template to the new virtual machine. The new virtual machine files are named and stored based on parameters you provide during the deployment. You can choose to make some configuration changes and customizations during the cloning process. The contents of some of the files, such as the configuration file, are modified. At the end of the operation, you can manage both the original template and the new virtual machine as inventory objects in vCenter Server.

Deploying a virtual machine from a template is a lot like cloning a virtual machine. In either case, you create a new virtual machine by copying a source object. For template deployments, the source object is a template. For virtual machine cloning, the source object is a virtual machine.

## Instant Clones



Starting with vSphere 6.7, you can use the instant clone technology to hot clone a running virtual machine in a manner that is like a combination of vMotion and linked clone technology. The result of an instant clone operation is a new virtual machine (destination virtual machine) that is identical to the source virtual machine. The processor state, virtual device state, memory state, and disk state of the destination virtual machine match those of the source virtual machine. To avoid network conflicts, you can customize the MAC addresses of the virtual NICs, but the guest customization feature is not supported for instant clones. You cannot use the vSphere Client to perform an instant clone operation.

A common use case for instant clones is just-in-time deployment in a VMware Horizon virtual desktop infrastructure (VDI). Instant clones enable you to

perform large-scale deployments by creating virtual machines from a controlled point in time. For example, VMware Horizon uses Instant Clone to improve the provisioning process for virtual desktops. Compared to View Composer, which uses linked clones, instant clones eliminate some steps (such as reconfiguration and checkpoints) and replace other steps to greatly reduce the provisioning time. Other use cases are large deployments of identical virtual servers in the cloud and situations where you want to reduce boot storms and provisioning times.

During an instant clone (vmFork) operation, the system quiesces and stuns the source virtual machine, creates and transfers a checkpoint, customizes the destination MAC address and UUID, and forks the memory and disk. The destination virtual machine shares the parent virtual machine's disk and memory for reads. For writes, the destination machine uses copy on write (COW) to direct disk and memory changes to delta files and private memory space.

The requirements for instant clones may depend on the software applications that use the API to perform the cloning operations. For example, VMware Horizon 7.1 requires static port binding, ESXi 6.0 Update 1 or later, and a distributed virtual switch.

Instant cloned virtual machines are fully independent vCenter Server inventory objects. You can manage instant clone destination virtual machines as you would regular virtual machines, without any restrictions.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

## Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 5-5](#) lists these key topics and the page number on which each is found.



**Table 5-5** Key Topics for [Chapter 5](#)

---

<b>Key Topic Element</b>	<b>Description</b>	<b>Page Number</b>
Procedure	Host profile workflow	171
List	Content library limitations	172
List	Files in a virtual machine snapshot	175
List	Use cases for snapshots	177
List	Requirements for cross-vCenter Server migrations	187
Section	Virtual machine migration limitations	187
Section	vMotion requirements	189
Section	Instant clones	195

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

- [vSphere inventory](#)
- [data center](#)
- [cluster](#)
- [resource pool](#)
- [template](#)
- [vApp](#)
- [host profile](#)
- [content library](#)
- [virtual machine snapshot](#)
- [VMware Tools](#)
- [vMotion](#)

## Review Questions

1. Which of the following is not a valid use case for virtual machine snapshots?
  - a. Rolling back guest OS changes
  - b. Recovering from the accidental deletion of a flat file
  - c. Troubleshooting
  - d. Linking a clone in a vRA blueprint
2. You are troubleshooting a virtual machine by using the vSphere Client. Which of the following is not a valid debugging and statistics advanced setting?
  - a. Record Trivial
  - b. Record Debugging
  - c. Run Normal
  - d. Record Statistics
3. You want to migrate a virtual machine with a 2.5 TB virtual disk. What is the minimum ESXi version that supports this?
  - a. 6.0
  - b. 6.5
  - c. 6.7
  - d. 7.0
4. You want to hot migrate a virtual machine from one ESXi host and VMFS datastore to another ESXi host and VMFS datastore. Which of the following statements is true?
  - a. This operation is not supported in vSphere 7.0.
  - b. The virtual disk data is transferred over the management network.
  - c. The virtual disk data is transferred over the vMotion network.
  - d. You must perform the operation in two separate steps: vMotion

and Storage vMotion.

5. You are supporting thousands of virtual machines in a vSphere environment. Which of the following features is associated with vmFork?
- a. Instant clone
  - b. Linked clone
  - c. Snapshot
  - d. Cross-vCenter vMotion

## Chapter 6

# VMware Product Integration

This chapter covers the following topics:

- [vSphere Add-ons](#)
- [vRealize Suite](#)
- [Desktop and Application Virtualization](#)
- [Replication and Disaster Recovery](#)
- [Private Public and Hybrid Clouds](#)
- [Networking and Security](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 2.2, 2.3, 2.4, and 2.5.

This chapter provides information on vSphere 7.0 integration with other VMware products, including vRealize Suite, Site Recovery Manager, Horizon, and NSX.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 6-1](#) 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 and Review Questions.](#)”

**Table 6-1** “Do I Know This Already?” Section-to-Question Mapping

<b>Foundation Topics Section</b>	<b>Questions</b>
----------------------------------	------------------

vSphere Add-ons	1
vRealize Suite	2-3
Desktop and Application Virtualization	4-6
Replication and Disaster Recovery	7
Private Public and Hybrid Clouds	8-9
Networking and Security	10

1. You want to streamline the development of modern applications using a familiar single stack for containers and virtual machines. Which of the following products should you use?
  - a. VMware Horizon
  - b. VMware App Volumes
  - c. VMware AppStack
  - d. vSphere with Tanzu
2. You want to provide continuous performance optimization and intelligent remediation in your vSphere software-defined data center. Which one of the following products should you use?
  - a. vRLI
  - b. vROps
  - c. vRA
  - d. vRNI
3. You want to decrease time and effort spent on root cause analysis in your data center. Which of the following products should you use?
  - a. vRLI
  - b. vROps
  - c. vRA
  - d. vRNI
4. You want to deliver VDI using stateless virtual desktops and just-in-time delivery of user profile data and applications. Which products should you choose? (Choose two.)



- a. VMware Horizon
  - b. VCF
  - c. vSphere Replication
  - d. HCX
  - e. App Volumes
5. You want to use App Volumes in your Horizon VDI environment. Which other environments can benefit from App Volumes? (Choose two.)
- a. VCF
  - b. VMware on AWS
  - c. Azure VMware Solution
  - d. Citrix XenApp
  - e. RDSH
6. You want to provide replication for your vSphere virtual machines to a remote site. Which of the following includes the required software?
- a. Site Recovery Manager
  - b. vRealize Suite
  - c. vSphere Foundations
  - d. vSphere Standard
7. Which of the following are use cases for VMware Site Recovery Manager? (Choose two.)
- a. Replicate data
  - b. Planned migrations
  - c. Disaster recovery
  - d. VDI
  - e. Data center automation
8. You want to use a platform that provides Cloud Builder and SDDC Manager. Which product should you choose?

- a. HCX
  - b. vRealize Automation
  - c. VCF
  - d. vRNI
9. You want to implement a workload mobility platform that simplifies application migration, workload rebalancing, and business continuity across hybrid clouds. Which product should you implement?
- a. HCX
  - b. vRealize Automation
  - c. VCF
  - d. vRNI
10. You want to adopt zero-trust security and automated network deployment. Which of the following products should you use?
- a. VMware NSX
  - b. HCX
  - c. VCF
  - d. AppDefense

## Foundation Topics

### vSphere Add-ons

This section addresses the following products that are directly related to vSphere but are not covered in other chapters of this book:

- **vSphere with Tanzu:** A vSphere edition that natively provides support for containers in the hypervisor
- **vCenter Converter:** A product that facilitates the conversion of physical and other servers into virtual machines running in vSphere
- **vSphere Replication:** A virtual machine replication feature that is included with specific vSphere editions

- **VMware Skyline:** A proactive support offering for many VMware products, including vSphere

## vSphere with Tanzu

By using vSphere with Tanzu, you can implement vSphere as a platform for natively running Kubernetes workloads. When enabled on a vSphere cluster, vSphere with Tanzu enables you to run Kubernetes workloads directly in the ESXi hypervisor and create Kubernetes clusters using dedicated resource pools. This works by creating a Kubernetes control plane directly in the hypervisor. A vSphere cluster that is enabled with vSphere for Kubernetes is called a *supervisor cluster*. The supervisor cluster runs on top of ESXi for compute, NSX-T Data Center for networking, and vSAN (or another shared storage solution) for storage.

With a Kubernetes control plane, you can create namespaces on the supervisor cluster, run containers in vSphere pods, and manage Kubernetes clusters using Tanzu Kubernetes Grid Service. A vSphere pod is a specialized virtual machine for running containers. In vSphere, you can manage and monitor vSphere pods and Tanzu Kubernetes clusters running in different namespaces.

## vSphere with Tanzu Use Cases



- vSphere with Tanzu is commonly used to provide a familiar single stack for containers and virtual machines and to streamline development of modern applications.

## vSphere with Tanzu Integration

To use vSphere with Tanzu, the main step is to install vSphere 7 with Kubernetes instead of installing vSphere 7. To run Kubernetes workloads on a vSphere cluster, you must enable the cluster with vSphere with Tanzu. In the vSphere Client, select Workload Management, select a vCenter Server, select Enable, select a cluster, and complete the wizard. In the wizard, you need to configure the cluster, network, and storage settings.

## vCenter Converter

VMware **vCenter Converter** (also called Converter Standalone) is a free solution that automates the process of converting existing Windows and Linux machines into virtual machines running in a vSphere environment. The source machines can be physical servers or virtual machines in non-ESXi environments. You can use Converter to convert virtual machines running in VMware Workstation, VMware Fusion, Hyper-V, and Amazon EC2 Windows to virtual machines running in vSphere.

With Converter, you can hot clone Windows servers without disrupting users of the source Windows Server. With hot cloning, you can minimize downtime when converting existing Windows and Linux servers to virtual machines running in vSphere.

Converter offers a centralized management console that allows users to queue and monitor multiple simultaneous remote and local conversions.

## **vCenter Converter Use Cases**

The following list identifies common use cases for vCenter Converter.

- Creation of virtual machine templates based on existing servers
- Physical-to-virtual server conversions
- Virtual machine migrations from non-vSphere environments

## **vCenter Converter Integration**

Converter is a standalone product that you can install on a Windows system. You can install it locally on a Windows desktop or on a server instance. To enable remote creation and management of tasks, you can install Converter's server and worker components on a Windows server in your data center and install the client component on multiple desktops. The server component installs an agent component on each Windows source machine prior to hot cloning.

In the Converter user interface, you can specify a vCenter Server as the destination for a conversion operation. You must provide credentials with sufficient privileges to create the virtual machine in the vSphere environment.

## **VMware vSphere Replication**

VMware vSphere Replication is included in multiple editions of vSphere. See the section "[Replication and Disaster Recovery](#)," later in this chapter, for details.

## **VMware SkyLine**

VMware [Skyline](#) is a proactive support technology, developed by VMware Global Services, that is available to customers with an active Production Support or Premier Services agreement. Skyline helps you avoid problems before they occur and reduces the amount of time spent on support requests.

The Skyline architecture includes a standalone on-premises virtual appliance (Skyline Collector) for secure, automatic data collection. It also includes a self-service web portal (Skyline Advisor) for accessing your VMware inventory, proactive findings, recommendations, and risks. You can segment data by factors, such as region and lines of business. You can use VMware Cloud Services Console to control user access and permissions. With a Premier Services agreement, you can access executive summary reports and view more powerful recommendations.

You can use Skyline Advisor to access Skyline Log Assist, which automatically (with your permission) uploads support log bundles to VMware Technical Support and eliminates manual procedures for log gathering and uploading. If you approve the request in Skyline Advisor, the requested logs are automatically uploaded to VMware Support. Likewise, you can choose to proactively push log files to VMware Support by using Log Assist within Skyline Advisor.

## **VMware SkyLine Use Cases**

Skyline is commonly used to avoid issues and streamline resolution in a vSphere environment.

## **VMware SkyLine Integration**

See [Chapter 10](#), “[Managing and Monitoring Clusters and Resources](#),” for instructions on configuring vCenter Server integration with Skyline Health.

## **vRealize Suite**

This section covers the vRealize Suite, which is a set of products that provides a layer for operations, automation, and analysis for software-defined data centers and hybrid clouds.

## **vRealize Operations Manager (vROps)**

vRealize Operations Manager (vROps) provides operations management for private, hybrid, and multi-cloud environments in a single pane of glass. It offers full-stack visibility from physical, virtual, and cloud infrastructure to virtual machines, containers, and applications. vROps provides continuous performance optimization, efficient capacity planning, cost management, and integrated compliance. It offers self-driving operations and intelligent remediation. It is available as an on-premises offering and as an SaaS offering.

vROps provides intelligent alerts, trending, and forecasting functionality. It uses current, historical, and predicted data for capacity analysis. For example, it calculates and provides a Days Remaining metric for many managed objects, such as vSphere clusters and datastores. This metric represents the number of days until the resource is predicted to have insufficient capacity. vROps applies customizable policies for everything it does. For example, you can use policies to adjust headroom buffers and provisioning lead times that impact capacity analysis.

vROps provides many native dashboards for management and troubleshooting. In addition, it offers many management packs to extend its operations beyond just vSphere. For example, you can install VMware-provided and third-party management packs that support operations for vSAN, NSX, third-party network gear, third-party storage system, and third-party software applications. Each management pack has its own unique requirements, which typically include installing the pack and connecting it to a management endpoint for the managed entity.

## **vROps Use Cases**

The following list identifies common uses cases for vROps:



- Continuous performance optimization
- Integrated compliance
- Next-generation operations platform
- Capacity and cost management and planning
- Intelligent remediation

## **vROps Integration**

In the vROps user interface, you can use the Solutions page to add a vCenter Server adapter instance (cloud account). You configure the instance by providing the address and user credentials for connecting to vCenter Server. At a minimum, the user account must have **Read** privileges assigned at the data center or vCenter Server level. To collect virtual machine guest OS metrics, the credential must have Performance > Modify Intervals permission enabled in the target. Additional requirements exist to allow vROps to perform automated actions in vSphere.

## **vRealize Log Insight (vRLI)**

*vRealize Log Insight (vRLI)* is a software product that provides intelligent log management for infrastructure and applications for any environment. It is a highly scalable log management solution that provides intuitive dashboards, analytics, and third-party extensibility. It collects and automatically identifies structure in all types of machine-generated log data, such as application logs, network traces, configuration files, messages, performance data, and system state dumps. It builds a high-performance index for performing analytics. It monitors and manages machine data at scale.

vRLI is especially useful in a large environment with multiple vCenter Server instances and complex infrastructure. You can configure vRLI to collect and analyze data from the vCenter Servers, ESXi hosts, guest OS, network infrastructure, storage infrastructure, and more. vRLI provides a single pane of glass you can use to analyze data from the entire stack when troubleshooting.

## **vRLI Use Cases**

vRLI is commonly used to decrease time and effort spent on root cause analysis and to aid in centralized log management and analysis.

## **vRLI Integration**

To collect alarms, events, and tasks data from a vSphere environment, you must connect vRealize Log Insight to one or more vCenter Server systems. vRLI can collect events, tasks, alerts, and other structured data directly from the vCenter Server. It can also collect unstructured data from ESXi hosts and the vCenter Server via syslog.

When connecting the vCenter Server to vRLI, you must provide a service account with appropriate privileges. To collect structured data from the vCenter Server, the service account must have the System.View privilege. To collect syslog data from ESXi hosts, the account must have the following privileges:

- Host.Configuration.Change settings
- Host.Configuration.Network configuration
- Host.Configuration.Advanced Settings
- Host.Configuration.Security Profile and Firewall

See the “[vRealize Log Insight \(vRLI\)](#)” section in [Chapter 10](#) for instructions on configuring vRLI to integrate with vCenter Server.

## **vRealize Automation (vRA)**

**vRealize Automation (vRA)** is an automation platform for private and multi-cloud environments. It delivers self-service automation, DevOps for infrastructure, and network automation that helps you increase your business’s agility, productivity, and efficiency. With vRA, your internal IT teams, DevOps engineers, developers, and others get the infrastructure, applications, and resources they need as a service with a public cloud–like experience. Customers benefit from increased speed, flexibility, reliability, and scalability, and you maintain security and control.

vRA 8.x includes vRA Cloud Assembly, vRA Service Broker, and vRA Code Stream.

You can use vRA Cloud Assembly to iteratively develop and deploy blueprints for your vSphere environment and other clouds. You can use vRA Service Broker to create, manage, and use self-service catalog items. You can use vRA Code Stream to create pipelines that automate your entire DevOps lifecycle, including automation of software testing and release.

With vRA Cloud Assembly, you can build blueprints that automatically provision virtual machines based on existing virtual machines and templates in your vSphere environment. In a blueprint, you define the provisioning method, such as full clone or linked clone. You can configure the blueprint to provision multiple virtual machines, applications, and networks. For example, you can develop a blueprint to deploy a multitier application involving multiple virtual machines, networks, and software components. You can publish the blueprints



and use vRA Service Broker to make them available as a service in the self-service catalog. You can configure vRA to allow consumers to provision the multitier application and its networks on demand and to destroy it when it is no longer needed.

## vRA Use Cases

The following list identifies common use cases for vRA.



- **Self-service private and hybrid clouds:** vRA provides a self-service catalog for delivering IaaS in your on-premises vSphere environment, private clouds built on VMware Cloud Foundation, and VMware Cloud Foundation on AWS.
- **Multi-cloud automation with governance:** An organization that uses vRA can extend self-service automation to multiple public clouds, including Amazon Web Services, Microsoft Azure, and Google Cloud Platform.
- **DevOps:** vRA provides automation for continuous integration/continuous development (CI/CD) pipelines.
- **Kubernetes automation:** vRA can be used to automate Kubernetes cluster and namespace provisioning management and support.

## vRA Integration

To begin using vRA, you need to deploy a vRA instance to a management vSphere cluster. A vRA 8.x deployment typically involves three vRA virtual appliances and three VMware Identity Manager (vIDM) appliances. To facilitate the deployment of these appliances, you can deploy and use the vRealize Lifecycle Manager (LCM) appliance.

To provide vSphere automation using vRA 8.x, you need to add at least one vCenter cloud account. (In vRA 7.x, you create a vSphere endpoint.) The vCenter Server that you use for a cloud account manages the user workload vSphere clusters. The cloud account provides the credentials that vRA uses to connect to vCenter Server. [Table 6-2](#) lists the required permissions for the vCenter cloud account.

**Table 6-2** Required Permissions for the vCenter Cloud Account

---

<b>Object</b>	<b>Permissions</b>
Datastore	<ul style="list-style-type: none"><li>▪ Allocate space</li><li>▪ Browse datastore</li><li>▪ Low level file operations</li></ul>
Datastore cluster	<ul style="list-style-type: none"><li>▪ Configure a datastore cluster</li></ul>
Folder	<ul style="list-style-type: none"><li>▪ Create folder</li><li>▪ Delete folder</li></ul>
Global	<ul style="list-style-type: none"><li>▪ Manage custom attributes</li><li>▪ Set custom attribute</li></ul>
Network	<ul style="list-style-type: none"><li>▪ Assign network</li></ul>
Permissions	

- Modify permission

---

## Resource

- Assign VM to resource pool
- Migrate powered-off virtual machine
- Migrate powered-on virtual machine

---

## Content library

- Add library item
- Create local library
- Create subscribed library
- Delete library item
- Delete local library
- Delete subscribed library
- Download files
- Evict library item
- Evict subscribed library

- Probe subscription information
- Read storage
- Sync library item
- Sync subscribed library
- Type introspection
- Update configuration settings
- Update files
- Update library
- Update library item
- Update local library
- Update subscribed library
- View configuration settings

---

Tags

- Assign or unassign vSphere tag
  - Create a vSphere tag
  - Create a vSphere tag category
  - Delete vSphere tag
- 

- Delete vSphere tag category
  - Edit vSphere tag
  - Edit vSphere tag category
  - Modify UsedBy field for category
  - Modify UsedBy field for tag
- 

## vApp

- Import
  - vApp application configuration.
- 

## Virtual machine inventory

- Create from existing

- Create new
- Move
- Remove

---

### Virtual machine interaction

- Configure CD media
- Console interaction
- Device connection
- Power off
- Power on
- Reset
- Suspend
- Tools install

---

### Virtual machine configuration

- Add existing disk

- Add new disk
- Add or remove
- Remove disk
- Advanced
- Change CPU count
- Change resource
- Extend virtual disk
- Disk change tracking
- Memory

- 
- Modify device settings
  - Rename
  - Set annotation
  - Settings

- Swapfile placement

---

## Virtual machine provisioning

- Customize
- Clone template
- Clone virtual machine
- Deploy template
- Read customization specs

---

## Virtual machine state

- Create snapshot
- Remove snapshot
- Revert to snapshot

---

## **vRealize Orchestrator (vRO)**

vRO is a modern workflow automation platform that simplifies complex infrastructure processes. It is a key component of vRA for providing custom workflows within on-demand services and providing anything as a service (XaaS). It can be used independently to run prebuilt workflows and to create custom workflows. It automates management and operational tasks of VMware and third-party systems, such as ticketing systems, change management systems,



and IT asset management systems.

In a vSphere environment, you may frequently perform some operational tasks. For example, say that you frequently receive requests to support the update procedure for a complex application involving multiple virtual machines. For each update, you are required to take the following actions:

- Shut down the virtual machines, one by one, in a specific order, ensuring that each shutdown operation completes prior to beginning the next shutdown.
- Create a snapshot of each virtual machine.
- Power on the virtual machines, one by one, in a specific order, ensuring that the guest OS and application services for each one are running prior to beginning the next power on.
- Inform the application team that the application is ready for update.
- Following a successful update, delete the snapshots.

With vRO, you can build workflows to automate all or portions of such an operation. For example, vRO provides out-of-the-box workflows for virtual machine power and snapshot operations. You can build a custom workflow that leverages the existing workflows as nested workflows. In the custom workflow, you can add data input, conditional paths, looping, and monitoring.

## vRO Use Cases



vRO is commonly used to orchestrate common vSphere operations tasks, orchestrate common data center infrastructure and application tasks, and provide XaaS for a vRA environment.

## vRO Integration

You can configure vRO 8.x to use vRA authentication or vSphere authentication. To use vSphere authentication, in the vRO Control Center, set Configure Authentication Provider > Authentication Mode to vSphere and configure it to use the credentials of the local administrator account of the vCenter Single Sign-On (SSO) domain ([administrator@vsphere.local](mailto:administrator@vsphere.local) by default). With vRO 8.x, you

must use vCenter Server 6.0 or later. To add a vCenter Server instance, you run the provided Add a vCenter Server Instance workflow.

## **vRealize Network Insight (vRNI)**

vRNI helps you build an optimized, highly available and secure network infrastructure for hybrid clouds and multi-clouds. It accelerates microsegmentation planning and implementation. It provides auditing for changes to the security posture and helps you ensure compliance. It facilitates troubleshooting across network infrastructure (virtual and physical) and security infrastructure.

vRNI provides network visibility and analytics to accelerate microsegmentation security, minimize risk during application migration, optimize network performance, and manage NSX, SD-WAN VeloCloud, and Kubernetes deployments.

## **vRNI Use Cases**

The following list identifies common use cases for vRNI.

- **Plan application security and migration:** vRNI accelerates microsegmentation deployment for private clouds and public clouds.
- **Optimize and troubleshoot virtual and physical networks:** vRNI enables you to reduce the mean time to resolution for application connectivity issues, eliminate network bottlenecks, and audit network and security changes.
- **Manage and scale NSX:** vRNI covers multiple NSX Manager instances and increases your availability by proactively detecting configuration issues.

## **vRNI Integration**

You can add VMware managers, such as vCenter Server, VMware NSX Manager, and VMware NSX-T Manager, to vRNI for data collection. To add a vCenter Server to vRNI as a data source, you need to have the following privileges applied and propagated at the root level:

- System.Anonymous
- System.Read

- System.View
- Global.Settings

To support IPFIX, you also need the Modify and Port Configuration Operation privilege on the distributed switches and Modify and Policy Operation on the distributed port groups.

To identify VM-to-VM paths, you must install VMware Tools in the virtual machines.

## Desktop and Application Virtualization

This section addresses VMware products for desktop and application virtualization.

### VMware Horizon

**VMware Horizon** is a platform for securely delivering virtual desktops and applications in private clouds and hybrid clouds. It enables provisioning and management of desktop pools that have thousands of virtual desktops each. It streamlines the management of images, applications, profiles, and policies for desktops and their users. It integrates with VMware Workspace ONE Access, which establishes and verifies end-user identity with multifactor authentication and serves as the basis for conditional access and network microsegmentation policies for Horizon virtual desktops and applications.

Horizon includes instant clones which work with VMware Dynamic Environment Manager and VMware App Volumes to dynamically provide just-in-time (JIT) delivery of user profile data and applications to stateless desktops.

Horizon provisions large pools of virtual desktops from a small set of base virtual desktops by integrating with vCenter Server. Horizon makes the provisioning requests, which are carried out by vCenter Server in the appropriate vSphere clusters. vSphere provides the environment, including the compute, storage, and network resources for running the virtual desktops. With vSphere DRS and vSphere HA, it provides load balancing and high availability.

### VMware Horizon Use Cases



The following list identifies common use cases for VMware Horizon.

- Remote users
- Kiosk and task users
- Call center
- Bring-your-own-device (BYOD) deployments
- Graphics-intensive applications

## VMware Horizon Integration

To get started with Horizon, in the vSphere environment, you need to prepare vSphere clusters to be used as resources for virtual desktop provisioning. You must add vCenter Server instances using the Horizon console. When adding the vCenter Server instance, you need to provide the vCenter Server address and appropriate user credentials. You can use the administrator account in the SSO domain ([administrator@vsphere.local](mailto:administrator@vsphere.local) by default) or, preferably, an account that is assigned the minimum privileges. [Table 6-3](#) provides the required privileges when you are not using instant clones. The use of instant clones requires additional privileges, such as all virtual machine configuration and inventory privileges.

**Table 6-3** Required vCenter Server Privileges for Horizon (without instant clones)

Privilege Group	Privileges to Enable
Folder	<ul style="list-style-type: none"><li>■ Create Folder</li><li>■ Delete Folder</li></ul>
Datastore	<ul style="list-style-type: none"><li>■ Allocate space</li></ul>
Virtual Machine	In Configuration:

- Add or remove device
- Advanced
- Modify device settings

---

In Interaction:

- Power off
- Power on
- Reset
- Suspend
- Perform wipe or shrink operations

In Inventory:

- Create new
- Create from existing
- Remove

In Provisioning:

- Customize
- Deploy template
- Read customization specifications
- Clone template
- Clone virtual machine

Resource	Assign virtual machine to resource pool
Global	Act as vCenter Server
Host (for Storage Accelerator)	Advanced settings (in Configuration)
Profile Driven Storage (for vSAN or Virtual Volumes)	All privileges

## App Volumes

VMware **App Volumes** is a set of application and user management solutions for VMware Horizon, Citrix Virtual Apps and Desktops, and Remote Desktop Services Host (RDSH) virtual environments. It streamlines your ability to deliver, update, assign, and manage applications and users across virtual desktop infrastructure (VDI) and published application environments. With App Volumes, you install an application once, using a provisioning computer, collect the application components in AppStacks, and centrally control the mapping of AppStacks to desktops.

AppStacks and companion writable volumes are stored in virtual disk files and attached to virtual machines to deliver applications. Updates to applications involve updating or replacing AppStacks or their mappings to desktops.

In RDSH environments, applications are installed on servers and delivered via Remote Desktop. Using App Volumes with RDSH simplifies the installation and management of the application on the server. Instead of attaching AppStacks to desktops, you attach AppStacks to RDSH servers and allow RDSH to deliver the application to the user.

## App Volumes Use Cases



The following list identifies common use cases for App Volumes.

- Application virtualization in VMware Horizon VDI environments
- Application virtualization in Citrix XenDesktop and XenApp environments
- Virtualization for RDSH-delivered applications

## App Volumes Integration

In the App Volumes management console, you can add and register a vCenter Server as a machine manager. To register vCenter Server and to allow App Volumes Manager to function, you must allow the privileges listed in [Table 6-4](#).

**Table 6-4** Required vCenter Server Privileges for App Volumes Manager

Privilege Group	Privileges to Enable
-----------------	----------------------

## Datastore

- Allocate space
- Browse datastore
- Low level file operations
- Remove file
- Update virtual machine files

---

Global	Cancel task
Host (Local Operations)	Reconfigure virtual machine
Sessions	View and stop sessions
Tasks	Create task
Virtual machine	In Configuration:

---

- Add existing disk
- Add new disk

- 
- Add or remove device
  - Query unowned files



- Change resource
- Remove disk
- Settings
- Advanced

#### In Inventory:

- Create new
- Move
- Register
- Remove
- Unregister

#### In Provisioning:

- Promote disks

---

## Replication and Disaster Recovery

This section addresses VMware products for replication and disaster recovery.

### vSphere Replication

**vSphere Replication** is an extension to VMware vCenter Server that provides hypervisor-based virtual machine replication and recovery. It provides virtual machine replication between the following source and destination combinations:

- Data center to data center
- Cluster to cluster within a data center
- Multiple source sites to a shared target site

vSphere Replication provides a number of benefits over storage-based replication:

- Lower cost per virtual machine
- Flexibility in storage vendor selection at the source and target sites
- Lower overall cost per replication

vSphere Replication is compatible with most vSphere features, including vMotion, Storage vMotion, vSphere HA, DRS, Storage DRS, vSAN, and DPM. It is not compatible with vSphere Fault Tolerance.

### vSphere Replication Use Cases

vSphere Replication is commonly used for disaster recovery and for data center migrations.

### vSphere Replication Integration

vSphere Replication does not require separate licensing. Instead, it is included as a feature of the following vSphere license editions, with no limit on the number

of replicated virtual machines:

- vSphere Essentials Plus
- vSphere Standard
- vSphere Enterprise
- vSphere Enterprise Plus

A minimum vSphere Replication deployment involves a single virtual appliance per site that provides vSphere Replication Management Service (VRMS) and vSphere Replication Service (VRS). It requires that specific network ports—including TCP 80, 443, and 902—be open for ESXi hosts. Likewise, it requires TCP ports 80, 443, 10443, and 7444 to be open for the vCenter Server.

vSphere Replication provides hypervisor-based replication, in which a vSphere Replication agent in the hypervisor collects and transmits changed blocks to the target VRS appliance. The method is similar to the Change Block Tracking (CBT) feature in ESXi, but it is a mechanism that is unique to vSphere Replication and prevents compatibility issues with technologies leveraging CBT.

It is possible to connect up to nine virtual appliances running just VRS per vCenter Server instance. A single VRMS appliance and nine VRS appliances can work in unison to provide replication for a single vCenter Server environment for up to 2000 replicated virtual machines.

Prior to installing vSphere Replication, you need to set the vCenter Server advanced setting `VirtualCenter.FQDN` to the fully qualified domain name of the vCenter Server. To install vSphere Replication, you use the standard vSphere OVF deployment wizard in the vSphere Client. Then you use the vCenter Server Appliance Management Interface (VAMI) to register the appliance with vCenter Single Sign-On. Successful registration produces a Site Recovery option on the vSphere Client home page. To replicate between sites, you deploy vSphere Replication to both sites and configure a vSphere Replication connection between the sites, using the Site Recovery page in the vSphere Client.

Configuring vSphere Replication involves the following steps:

**Step 1.** In the vSphere Client, navigate to **Home > Site Recovery > Open Site Recovery**.

**Step 2.** Select a vSphere connected pair and click **View Details**.

- Step 3.** Click the **Replications** tab.
- Step 4.** Select **Outgoing** or **Incoming** and click the **Create New Replication** icon.
- Step 5.** Complete the wizard to configure the replication settings, such as those for the target, seed, recovery point objective (RPO), point in time instances, and quiescing options.

To recover a virtual machine, you follow these steps:

- Step 1.** In the vSphere Client, navigate to **Home > Site Recovery > Open Site Recovery**.
- Step 2.** Select a vSphere connected pair and click **View Details**.
- Step 3.** Click the **Replications** tab.
- Step 4.** Select a replication instance in the Incoming section.
- Step 5.** Click the **Recover** icon.
- Step 6.** Choose **Synchronize Recent Changes** (to perform a final synchronization) or **Use Latest Available Data** (to continue without performing a final synchronization).
- Step 7.** Optionally, select **Power on the Virtual Machine After Recovery**.
- Step 8.** Complete the wizard by selecting the target folder, compute resource, and other options.

You can leverage alarms in vCenter Server to get alerts about issues in vSphere Replication, such as issues with the connection, a VRS instance, or a specific replication. For example, you configure an alarm to be triggered whenever a configured replication exceeds the configured RPO.

## Site Recovery Manager (SRM)

VMware *Site Recovery Manager (SRM)* is a business continuity solution that you can use to orchestrate planned migrations, test recoveries, and disaster recoveries. For data replication, SRM integrates with vSphere Replication and supported storage-based replication products. In SRM you can build recovery plans that include recovery steps, virtual machine priority groups, dependencies, IP address changes, and resource mappings. You can run a single plan in one of three modes:

- **Planned migration:** In planned migration mode, SRM automatically shuts down the source virtual machines prior to migration, performs a final data synchronization, and stops if errors occur.
- **Disaster recovery:** In disaster recovery mode, SRM attempts to shut down and synchronize the source virtual machines but continues with recovery.
- **Test:** In test mode, SRM leaves the source machines running while it brings up another instance of each virtual machine, using snapshots at the recovery site in an isolated network. During a test recovery, the source machines continue to be replicated and protected. After a test recovery, you should run a cleanup to shut down and remove the target site snapshots.

SRM is tightly integrated with vSphere Replication in vSphere 7.0. To use SRM, you begin by navigating to Home > Site Recovery in the vSphere Client.

## SRM Use Cases

SRM is commonly used for disaster recovery and data center migrations.

## SRM Integration

Prior to installing SRM, you should implement a supported replication technology, such as EMC RecoverPoint or vSphere Replication. You need to deploy SRM to both the source and target sites. You can install a Windows-based version of SRM in a supported Windows server, or you can deploy the SRM virtual appliance. In most cases, you should deploy the SRM appliance, which includes an embedded vPostgreSQL database that supports a full-scale SRM environment.

At each site, you need to deploy an SRM server and register it with a vCenter Server. SRM requires a separate vCenter Server at the source site and at the target site.

SRM uses Transport Layer Security (TLS) and solution user authentication for secured connections with vCenter Server. It assigns a private key and a certificate to the solution user and registers it with the vCenter Single Sign-On service. When you pair SRM instances across vCenter Servers that do not use Enhanced Linked Mode, Site Recovery Manager creates an additional solution user at the remote site.

# Private, Public, and Hybrid Clouds

This section addresses VMware products for private, public, and hybrid clouds.

## VMware Cloud Foundation (VCF)

VCF is a hybrid cloud platform built on full-stack hyperconverged infrastructure (HCI) technology. It provides a single easy-to-deploy architecture that enables consistent, secure infrastructure and operations across private and public clouds.

VCF provides the following features:

- Automated bring-up of the software stack, including vSphere, vCenter Server, vSAN, NSX-T, and vRealize Suite
- Simplified provisioning in workload domains built on vSphere, vSAN, and NSX-T
- Application-focused management, leveraging vSphere with Tanzu to support virtual machines and containers in the same platform
- With Automated Lifecycle Management (LCM), simplified updates for all components in the stack
- Multi-instance management, which allows multiple VCF instances to be managed together

The following are the main components in a private cloud powered by VCF 4.0:

- Cloud Builder
- SDDC Manager
- vSphere
- vSAN
- NSX-T Data Center
- vRealize Suite

Cloud Builder is the VCF component that automates the deployment of the entire software-defined stack. SDDC Manager is the VCF component that automates the entire system lifecycle and simplifies day-to-day management and operations.

The standard model for VCF uses separate virtual infrastructure domains for running management and user workloads. VCF also supports a consolidated model, in which the management and user workloads run in the same virtual infrastructure domain.

## VCF Use Cases

VCF is commonly used for private clouds, hybrid clouds, modern applications, and VDI.

## VCF Integration

To get started with VCF, you should prepare ESXi hosts for the implementation of the management domain, address network and environment prerequisites, fill in the deployment parameter workbook, deploy the VMware Cloud Builder appliance, and use Cloud Builder to deploy the management domain, including vCenter Server.

## VMware Hybrid Cloud Extension (HCX)

VMware HCX is a workload mobility platform that simplifies application migration, workload rebalancing, and business continuity across on-premises data centers, private clouds, and hybrid clouds. HCX enables you to migrate thousands of virtual machines, migrate from non-vSphere platforms, upgrade vSphere versions, balance workload between on-premises and cloud, and implement replication to protect against disaster.

VMware HCX enables you to schedule and migrate thousands of vSphere virtual machines within and across data centers without requiring a reboot. HCX offers a number of services with each VMware HCX license (Advanced or Enterprise), as listed in [Table 6-5](#).

**Table 6-5** VMware HCX Services

Service	Licenses	Description
Interconnect	Advanced	Creates secured connections between HCX instances, supporting migration, replication, disaster recovery, and management operations.

a  
n  
c  
e  
d  
Deployed as a virtual appliance.

---

WAN Optimization      A      Optimizes the performance of the connection  
d      provided by HCX Interconnect through a  
v      combination of deduplication, compression, and  
a      line conditioning techniques.  
n  
c      Deployed as a virtual appliance.  
e  
d

---

Network Extension      A      Extends (that is, provides Layer 2 adjacency)  
d      the virtual machine networks between source  
v      and remote HCX-enabled environments.  
a  
n      Deployed as a virtual appliance.  
c  
e  
d

---

Bulk Migration      A      Migrates a set of virtual machines using  
d      VMware vSphere Replication in parallel  
v      between HCX-enabled sites.  
a  
n  
c  
e  
d

---

vMotion Migration      A      Migrates a single virtual machine between  
d      HCX-enabled sites with no service interruption,  
v      using vMotion.  
a  
n  
c  
e  
d

---



Disaster Recovery      A  
d  
v  
a  
n  
c  
e  
d      Protects virtual machines from disaster by using replication and recovery.

---

Mobility Groups      E  
n  
t  
e  
r  
p  
r  
i  
s  
e      Allows you to group virtual machines by application, network, or other aspects for migration and monitoring.

---

OS Assisted Migration      E  
n  
t  
e  
r  
P  
r  
i  
s  
e      Leverages HCX Sentinel software in the guest OS to migrate Windows and Linux virtual machines to a vSphere-enabled data center.  
  
Uses a gateway appliance at the source and a receiver appliance at the destination.

---

Replication Assisted vMotion (RAV)      E  
n  
t  
e  
r  
P  
r  
i  
s  
e      Migrates a set of virtual machines in parallel, using VMware vSphere Replication and vMotion between HCX-enabled sites with no service interruption.

---

Site Recovery Manager (SRM) Integration      E  
n  
t  
e  
r  
P      Integrates HCX functionality with the VMware SRM for protection and orchestrated recovery operations.

ri  
s  
e

---

Traffic Engineering: Application Path Resiliency and TCP Flow Conditioning	E n te r P ri s e	Optimizes network traffic for HCX Interconnect and Network Extension services.  The Application Path Resiliency service creates multiple tunnel flows for both Interconnect and Network Extension traffic.  The TCP Flow Conditioning service adjusts and optimizes the segment size to reduce fragmentation and reduce the overall packet rate.
Mobility Optimized Networking (MON)	E n te r p ri s e	Integrates HCX Network Extension with NSX Dynamic Routing to enable optimal networking between migrated virtual machines and other virtual machines.  Works with new or existing network extensions to NSX-T 3.0 Data Center.

---

## HCX Use Cases

HCX is commonly used for cloud adoption and migration, workload rebalancing, and business continuity.

## HCX Integration

To integrate HCX into an on-premises vSphere environment, you need to implement HCX components that connect the environment to another environment, such as a hosted private cloud or hybrid cloud. The following are the key components, and their services are described in [Table 6-5](#):

- HCX Connector and HCX Cloud Installation
- HCX-IX Interconnect Appliance

- HCX WAN Optimization Appliance
- HCX Network Extension Virtual Appliance

VMware HCX is used in VMware on AWS, Azure VMware Solution, and other hybrid cloud solutions.

## **VMware on AWS**

VMware Cloud (VMC) on Amazon Web Services (AWS) is an integrated cloud offering jointly developed by AWS and VMware. You can migrate and extend your on-premises VMware vSphere-based environments to the AWS Cloud running on Amazon Elastic Compute Cloud (Amazon EC2). With VMC on AWS, you can deploy a software-defined data center (SDDC) on demand. You configure the network and security to suit your needs and then begin deploying virtual machines. VMware provides support for VMC on AWS. You can open the VMware Cloud Services console to get support.

VMC on AWS provides workloads with access to more than 175 AWS services, including database, AI/ML, and security services. It provides simplicity for hybrid cloud operations by enabling you to use the same VCF technologies (vSphere, vSAN, NSX, vCenter Server) across the on-premises environment and the AWS cloud. It does not require custom, on-premises hardware. It improves availability and accelerates cloud migration by enabling workloads to be directly migrated between an on-premises environment and AWS. To migrate virtual machines from an on-premises vSphere environment to VMC on AWS, you can perform a live migration via vMotion or use VMware HCX.

## **Azure VMware Solution**

Azure VMware Solution combines VMware's SDDC software with the Microsoft Azure global cloud service ecosystem to provide a hosted private cloud. Azure VMware Solution is managed to meet performance, availability, security, and compliance requirements. Currently, you cannot use on-premises vCenter Server to manage the hosted private cloud. Instead, you use vCenter Server and NSX Manager in a hosted private cloud.

You can set up VMware HCX for an Azure VMware Solution private cloud. HCX enables migration of VMware workloads to the cloud and other connected sites. If you meet standard cross-vCenter vMotion requirements, you can migrate on-premises virtual machines to the hosted private cloud. To configure HCX,

you deploy an HCX virtual appliance in your on-premises vSphere environment and connect it to HCX in your hosted private cloud.

You can use Azure Migrate to migrate on-premises vSphere virtual machines to Azure.

## Networking and Security

This section addresses VMware products for networking and security.

### AppDefense

VMware AppDefense provides data center endpoint security that protects applications running in a virtualized environment. It learns good behavior for each of your virtual machines and applications so it can detect and respond to deviations from that good behavior. It can respond with actions on the virtual machine, such as blocking networking, suspending, quarantining, snapshotting, powering off, or killing a suspicious process in the guest OS.

AppDefense is a SaaS product that works with an on-premises virtual appliance and a vCenter Server plug-in. The on-premises virtual appliance connects to vCenter Server, NSX Manager, and other optional components. You can deploy an AppDefense module to each protected host. Likewise, you can deploy an AppDefense agent to guest operating systems. AppDefense can tie into provisioning systems such as vRealize Automation or Puppet to define appropriate and allowed behaviors.

### AppDefense Use Cases



The following list provides common use cases for AppDefense

- Accelerate security operations
- Protect critical workloads
- Increase the security posture of every workload
- Enhance the effectiveness of microsegmentation

### App Defense Integration

AppDefense Appliance must be able to connect over the Internet with AppDefense Manager using TCP port 443.

To get started, you should subscribe to the AppDefense service, log on to the AppDefense Manager website, and download the AppDefense Appliance OVA file. Use the vSphere Client to deploy the appliance (by selecting Deploy OVF Template) into your management environment. Back in the AppDefense Manager website, choose Provision Appliance and identify the appliance name. The provisioning process generates an UUID, an API key, and a URL. Use a web browser to access the AppDefense Appliance GUI and register the appliance with the protected vCenter Server. In the registration wizard, provide the UUID and API key. You should deploy AppDefense Appliance for each protected vCenter Server.

When you deploy AppDefense Appliance, it registers with vCenter Server and collects inventory. You can then install the AppDefense host and guest modules. To install the host modules, download the vSphere Information Bundle (VIB) from AppDefense Manager and use the **esxcli software vib install** command.

To deploy the guest module to protected virtual machines, you must install VMware Tools, ensure that the virtual machine hardware version is 13 or later, and install the AppDefense host module to the host running the virtual machine. Use AppDefense Manager to select the virtual machine and then select Enable Guest Integrity. Finally, install the guest module in the guest OS. In Windows, use Control Panel > Programs and Features to modify the VMware Tools program and enable AppDefense.

## NSX

VMware NSX Data Center (NSX) is a network virtualization and security platform that enables a software-defined approach to networking that extends across data centers, clouds, and application frameworks. NSX enables you to provision and manage networks independently of the underlying hardware, much as you do with virtual machines. You can reproduce a complex network in seconds and create multiple networks with diverse requirements.

NSX provides a new operational model for software-defined networking and extends it to the virtual cloud network. It provides a complete set of logical networking, security capabilities, and services, such as logical switching, routing, firewalling, load balancing, virtual private networking, quality of service (QoS), and monitoring.

VMware NSX-T Data Center (NSX-T) is the recommended product for practically all new virtualized networking use cases. Although it was originally developed for non-vSphere environments, it now supports vSphere. Most NSX customers are migrating or starting to consider migrating to NSX-T. The NSX-T platform provides the following components:

- NSX-T managers
- NSX-T edge nodes
- NSX-T distributed routers
- NSX-T service routers
- NSX-T segments (logical switches)

Guest Module uses a data plane, a control plane, and a management plane.

## NSX Use Cases



The following list identifies the common uses cases for NSX.

- Adopt zero-trust security
- Multi-cloud networking
- Automated network deployment
- Network and security for cloud-native applications (containers)

## NSX Integration

To prepare for an NSX-T installation, you need to meet the requirements for deploying its components, such as the NSX-T managers and edge nodes. Typically, a three-node NSX-T Manager cluster is deployed to a management vSphere cluster, and the NSX-T edges are deployed in a shared edge and compute cluster.

After deploying the required virtual appliances from OVF, you log in to NSX-T Manager and add a vCenter Server as a compute manager. When adding the vCenter Server compute manager, you should use the administrator account of the Single Sign-On domain ([administrator@vsphere.local](mailto:administrator@vsphere.local) by default) or use a

custom account configured with the appropriate privileges. Next, you deploy NSX-T edges to vSphere clusters managed by the vCenter Server and create the transport zones and transport nodes.

With NSX-T, you implement NSX-T Virtual Distributed Switches (N-VDS), which are logical switches that are decoupled from the vCenter Server to provide cross-platform support. They function much like a vSphere Distributed Switch (vDS), in that they provide uplinks to host physical NICs, multiple teaming policies, VLAN support, and more, but they can reside in a non-vSphere environment.

With vSphere 7.0 and NSX-T 3.0, you can run NSX-T directly on a vDS 7.0. This provides simpler integration in vCenter Server as well as some other benefits. When creating transport zone nodes on ESXi hosts, you can choose between N-VDS and VDS as the host switch type.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

## Review All Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 6-6](#) lists these key topics and the page number on which each is found.



**Table 6-6** Key Topics for [Chapter 6](#)

Key Topic Element	Description	Page Number
List	Use cases for vSphere with Tanzu	<a href="#">204</a>
List	Use cases for vROps	<a href="#">207</a>
List	Use cases for vRA	<a href="#">210</a>
List	Use cases for vRO	<a href="#">214</a>

List	Use cases for Horizon	216
List	Use cases for App Volumes	218
List	Use cases for AppDefense	227
List	Use cases for NSX	229

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

- [vCenter Converter](#)
- [Skyline](#)
- [vRealize Log Insight \(vRLI\)](#)
- [vRealize Automation \(vRA\)](#)
- [VMware Horizon](#)
- [App Volumes](#)
- [vSphere Replication](#)
- [Site Recovery Manager \(SRM\)](#)

## Review Questions

1. You want to build custom workflows to support XaaS. Which product should you use?
  - a. vRLI
  - b. vRO
  - c. vROps
  - d. App Volumes
2. You need to provide virtual desktops and applications to remote users



and call centers. Which product should you implement?

- a. VCF
  - b. vRealize Suite
  - c. AppDefense
  - d. Horizon
3. You want to configure vSphere Replication using the vSphere Client. Which of the following is the correct navigation path?
- a. Home > vCenter Server > vSphere Replication
  - b. Home > Site Recovery > Open Site Recovery
  - c. Home > Host and Clusters > Replications
  - d. Home > Administration > Replication
4. Which of the following products provides connection, WAN optimization, and bulk migrations?
- a. vRealize Suite
  - b. vSphere Replication
  - c. SRM
  - d. HCX
5. You want to provide data center endpoint security to protect applications running in a vSphere environment. Which one of the following products should you use?
- a. AppDefense
  - b. HCX
  - c. VCF
  - d. vRealize Operations

## Chapter 7

# vSphere Security

This chapter covers the following subjects:

- [vSphere Certificates](#)
- [vSphere Permissions](#)
- [ESXi and vCenter Server Security](#)
- [vSphere Network Security](#)
- [Virtual Machine Security](#)
- [Available Add-on Security](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.10, 1.11, 4.10, 4.11, 4.11.1, 4.13, 7.7, and 7.8.

This chapter covers topics related to hardening a vSphere environment.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 7-1](#) 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 and Review Questions.](#)”

**Table 7-1** “Do I Know This Already?” Foundation Topics Section-to-Question Mapping

Foundations Topics Section	Questions Covered in This Section
vSphere Certificates	1, 2

vSphere Permissions	3, 4
ESXi and vCenter Server Security	5, 6
vSphere Network Security	7
Virtual Machine Security	8, 9
Available Add-on Security	10

1. You are preparing to import certificates for your vSphere environment. Which of the following is not a requirement?
  - a. x509 Version 3
  - b. PKCS8 and PKCS1 PEM format
  - c. Digital signature and key encipherment keys
  - d. 1024- to 16,384-bit keys
2. You are making plans for ESXi host certificates. Which of the following is not a valid certificate mode?
  - a. VMware Endpoint Certificate Store Mode
  - b. VMware Certificate Authority Mode
  - c. Custom Certificate Authority Mode
  - d. Thumbprint Mode
3. You are preparing to apply permissions in vCenter Server. Which of the following is a system role?
  - a. Read-only
  - b. Virtual machine user
  - c. Datastore consumer
  - d. Content library administrator
4. You are configuring permissions in vCenter Server. Which privilege is required for a user to use Storage vMotion to migrate a virtual machine?
  - a. Resource.Migrate Powered On Virtual Machine
  - b. Resource.Migrate Powered Off Virtual Machine

- c. Resource.Assign Virtual Machine to Resource Pool on the Cluster
  - d. Resource.Assign Virtual Machine to Resource Pool on the VM Folder
5. You are hardening your ESXi hosts. Which of the following is true concerning normal lockdown mode?
- a. All users with administrator privileges on the host can access the DCUI.
  - b. All users in the Exception Users list can access the DCUI.
  - c. No one can access the DCUI.
  - d. Users identified in the host's DCUI.Access advanced option can access the DCUI.
6. You are creating user accounts in the vCenter SSO domain. With default settings, which of the following is a valid password?
- a. VMware1!
  - b. VMworld!
  - c. VMwareR0cks
  - d. VMwarerocks!!
7. You are configuring IPsec on your ESXi hosts. Which of the following commands can you use to list the available security associations on an ESXi host?
- a. **esxcli network ipsec sa list**
  - b. **esxcli network ip ipsec sa list**
  - c. **esxcli network ip ipsec list**
  - d. **esxcli network ip sa list**
8. You want to migrate virtual machines across vCenter instances. Which of the following statements is true concerning vMotion migration across vCenter Server instances?
- a. For encrypted vMotion migration, you can use the vSphere Client.

- b. For encrypted vMotion migration, you must use the vSphere APIs.
  - c. vMotion migration of encrypted virtual machines is not supported.
  - d. Encrypted vMotion migration of non-encrypted virtual machines is not supported.
9. You are hardening virtual machines in your vSphere 7 environment. Which of the following options can be set to TRUE to disable an unexposed feature?
- a. tools.guestlib.enableHostInfo
  - b. tools.setInfo.sizeLimit
  - c. vmx.log.keepOld
  - d. isolation.tools.ghi.launchmenu.change
10. You want to use microsegmentation to protect the applications and data in your vSphere environment. What should you implement?
- a. VMware AppDefense
  - b. VMware NSX
  - c. VMware vRealize Automation
  - d. VMware vRealize Log Insight

## Foundation Topics

### vSphere Certificates

This section describes the use of certificates in a vSphere environment.

#### vSphere Certificates Overview

In vSphere 7.0, you can use the default approach to provision vCenter Server components and ESXi hosts with VMware Certificate Authority (VMCA) certificates. You can also use custom certificates stored in the **VMware Endpoint Certificate Store (VECS)**. vCenter Server supports custom certificates that are generated and signed from your own enterprise public key infrastructure (PKI).

vCenter Server also supports custom certificates that are generated and signed by trusted third-party certificate authorities (CAs), such as VeriSign or GoDaddy. vSphere uses certificates to do the following:

- Encrypt communications nodes, such as a vCenter Server and ESXi hosts
- Authenticate vSphere services
- Perform internal actions such as signing tokens

Table 7-2 describes the core identity services in vSphere.

**Table 7-2** Core Identity Services in vSphere

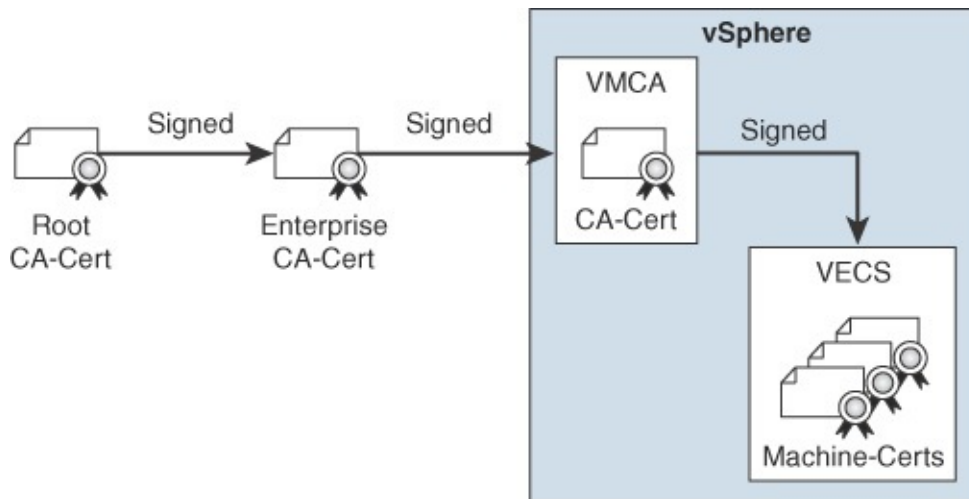
Service	Description
<i>VMware Directory Service (vmdir)</i>	Serves as an identity source that handles SAML certificate management for authentication with vCenter Single Sign-On.
<i>VMware Certificate Authority (VMCA)</i>	Issues certificates for VMware solution users, machine certificates for machines on which services are running, and ESXi host certificates. VMCA can be used as is, or it can be used as an intermediary certificate authority.
VMware Authentication Framework Daemon (VMAFD)	Includes VMware Endpoint Certificate Store (VECS) and several internal authentication services.

VECS is a local (client-side) repository for certificates, private keys, and other certificate information that can be stored in a keystore. You can choose not to use VMCA as your certificate authority and certificate signer, but you must use VECS to store all vCenter certificates and keys. ESXi certificates are stored locally on each host and not in VECS. The stores included in VECS are described in [Table 8-9](#) in [Chapter 8, “vSphere Installation.”](#)

VMCA, which runs in every vCenter Server Appliance, is vSphere’s internal certificate authority. It provides all the required certificates for vCenter Server and ESXi. VMCA’s default configuration provides the lowest operational overhead for certificate management and immediately secures the solution

without any other modification. vSphere provides a mechanism to renew expired certificates and to replace specific certificates with your own certificates.

You can replace the VMCA root certificate with a certificate that is signed by an enterprise CA or a third-party CA, in which case VMCA signs the custom root certificate each time it provisions certificates, making VMCA an intermediate CA (see [Figure 7-1](#)).



**FIGURE 7-1** VMCA as an Intermediate CA

You can replace the existing VMCA-signed certificates with custom certificates. If you do so, you bypass VMCA and become responsible for all certificate provisioning and monitoring.

VMware recommends that you replace only the SSL certificate that provides encryption between nodes. VMware does not recommend replacing either solution user certificates or STS certificates. VMware does not recommend using a subordinate CA in place of VMCA. If you fail to follow these recommendations, you might encounter significant complexity and an increase in your operational risk. [Table 7-3](#) summarizes VMware’s recommended modes for managing certificates.

**Table 7-3** Recommended Modes for Managing Certificates

<b>Mode</b>	<b>Description</b>	<b>Advantages</b>
VMCA Default Certificate	VMCA provides all the certificates for vCenter Server and ESXi hosts.	Lowest-overhead option. VMCA manages the

s

certificate life cycle for vCenter Server and ESXi hosts.

---

VMCA Default Certificates with External SSL Certificates (hybrid mode)	You replace the vCenter Server SSL certificates and allow VMCA to manage certificates for solution users and ESXi hosts. Optionally, for high-security-conscious deployments, you can replace the ESXi host SSL certificates as well.	VMCA manages internal certificates, but you get the benefit of using your corporate-approved trusted SSL certificates.
--	---	--

---

## Certificate Requirements

The following requirements apply to all imported certificates:

- The key size is 2048 bits to 16,384 bits.
- VMware supports PKCS8 and PKCS1 (RSA key) PEM formats. When you add keys to VECS, they are converted to PKCS8.
- x509 Version 3 is required.
- SubjectAltName must contain DNS Name=machine\_FQDN.
- CRT format is required.
- The digital signature and key encipherment keys are available.
- Enhanced Key Usage can either be empty or contain Server Authentication.

VMCA does not support the following certificates:

- Certificates with wildcards
- The algorithms md2WithRSAEncryption 1.2.840.113549.1.1.2, md5WithRSAEncryption 1.2.840.113549.1.1.4, and sha1WithRSAEncryption 1.2.840.113549.1.1.5
- The algorithm RSASSA-PSS with OID 1.2.840.113549.1.1.10

If you do not generate certificate signing requests (CSRs) using Certificate



Manager, you need to ensure that a CSR includes the fields listed in [Table 7-4](#).

**Table 7-4** Required Fields for a CSR

String	X.500 Attribute Type
CN	commonName
L	localityName
ST	stateOrProvinceName
O	organizationName
OU	organizationalUnitName
C	countryName
STREET	streetAddress
DC	domainComponent
UID	userid

If you use VMCA as an intermediate CA, you can use the vSphere Certificate Manager to create a CSR or you can create a CSR manually. When you create a CSR manually, in addition to the previously stated requirements, you should consider the requirements in [Table 7-5](#), which are based on the specific certificate types.

**Table 7-5** Requirements for Certificates When VMCA Is an Intermediate CA

Certificate Type	Additional Requirements
Root certificate	<p>Set CA extension to true and include CertSign. For example, use the following in the CSR:</p> <p>basicConstraints = critical, CA:true</p> <p>keyUsage = critical, digitalSignature,keyCertSign</p>
Machine SSL certificate	No additional requirements

Solution user certificate      Use a different Name value for each solution user, which may appear as CN under Subject, depending on the tool

---

**Note**

Do not use CRL distribution points, authority information access, or certificate template information in any custom certificates.

VMCA provisions an environment with certificates, including machine SSL certificates for secure connections, solution user certificates for service authentication with vCenter Single Sign-On, and ESXi host certificates (see [Table 7-6](#)).

**Table 7-6** Certificates in vSphere

<b>Certificate</b>	<b>Provisioned</b>	<b>Details</b>
ESXi certificate	VM CA (default)	Stored locally on an ESXi host in the /etc/vmware/ssl directory when the host is first added to vCenter Server and when it reconnects.
Machine SSL certificate	VM CA (default)	<p>Stored in VECS.</p> <p>Used to create SSL sockets for SSL client connections, for server verification, and for secure communication such as HTTPS and LDAPS.</p> <p>Used by the reverse proxy service, the vCenter Server service (vpxd), and the VMware Directory service (vmdir).</p> <p>Uses X.509 Version 3 certificates to encrypt session information.</p>
Solution user	VM	Stored in VECS.

certificate	CA (default)	Used by solution users to authenticate to vCenter Single Sign-On through SAML token exchange.
vCenter Single Sign-On SSL signing certificate	During installation	Used throughout vSphere for authentication, where a SAML token represents the user's identity and contains group membership information.  You can manage this certificate from the command line. Changing this certificate in the file system leads to unpredictable behavior.
VMware Directory Service (vmdir) SSL certificate	During installation	Starting with vSphere 6.5, the machine SSL certificate is used as the vmdir certificate.
vSphere Virtual Machine Encryption Certificates	Depends	Used for virtual machine encryption, which relies on an external key management server (KMS).  Depending on how the solution authenticates to the KMS, it might generate certificates and store them in VECS.

A solution user presents a certificate to vCenter Single Sign-On when it first authenticates, after a reboot, and after a timeout has elapsed. The timeout (Holder-of-Key Timeout) can be set from the vSphere Client and defaults to 2,592,000 seconds (30 days).

The following solution user certificate stores are included in VECS:

- **Machine:** Used by the license server and the logging service.
- **vpxd:** Used by the vCenter service (vpxd) to authenticate to vCenter Single Sign-On.
- **vpxd-extension:** Used by the Auto Deploy service, inventory service, and other services that are not part of other solution users.

- **vsphere-webclient:** Used by the vSphere Client and some additional services, such as the performance chart service.
- **wcp:** Used by vSphere with Kubernetes.

### Note

Do not confuse the machine solution user certificate with the machine SSL certificate. The machine solution user certificate is used for SAML token exchange. The machine SSL certificate is used for secure SSL connections for a machine.

## ESXi Host Certificates

In vSphere 6.0 and later, vCenter Server supports a number of certificate modes for ESXi hosts, as described in [Table 7-7](#).

**Table 7-7** Certificate Modes for ESXi Hosts

Mode	Description
VMware Certificate Authority Mode (default)	Use this mode when VMCA provisions an ESXi host as either the top-level CA or an intermediate CA. In this mode, you can refresh and renew certificates from the vSphere Client.
Custom Certificate Authority Mode (CA)	Use this mode with custom certificates signed by a third-party or an enterprise CA. In this mode, you cannot refresh and renew certificates from the vSphere Client.
Thumbprint Mode	Use this legacy (vSphere 5.5) mode only for troubleshooting. In this mode, vCenter Server checks the certificate format, not the certificate's validity. For example, expired certificates are accepted. Some vCenter 6.x and later services might not work correctly in this mode.

## Note

If you apply custom certificates to hosts but do not change the certificate mode to Custom Certificate Authority, VMCA might replace custom certificates when you select Renew in the vSphere Client.

You can use the vSphere Client to view expiration data for certificates, whether they are signed by VMCA or a third party. The vCenter Server raises yellow alarms for hosts where certificates expire shortly (that is, less than 8 months) and red alarms where certificates are in the Expiration Imminent state (that is, expire in less than 2 months).

ESXi hosts that boot from installation media have autogenerated certificates. When a host is added to the vCenter Server system, it is provisioned with a certificate that is signed by VMCA as the root CA.

## vSphere Permissions

This section describes the permissions model in vSphere.

### Authentication and Authorization

vCenter Single Sign-On (SSO) is responsible for authenticating vCenter Server users. The user accounts may be defined directly in the SSO domain or in a supported identity source. vCenter Server uses permissions and roles to provide authorization, and this controls what an authenticated user can do. It allows you to assign a permission to an object in the vCenter Server inventory by specifying which privileges a specific user or group has on that object.

The default, the SSO domain name, is `vsphere.local`, but you can change it during domain creation. Initially, only the SSO domain administrator is authorized to log in to vCenter Server. By default, the SSO domain administrator is [administrator@vsphere.local](mailto:administrator@vsphere.local). You can create additional users in the SSO domain. You can add supported identity sources to SSO, including Active Directory over LDAP, a native Active Directory (Integrated Windows Authentication) domain, or an OpenLDAP directory service.

Starting in vSphere 7.0, vCenter Server supports federated authentication, where you configure a connection to an external identity provider to replace vCenter Server as the identity provider. Currently, vCenter Server supports only Active

Directory Federation Services (AD FS) as an external identity provider.

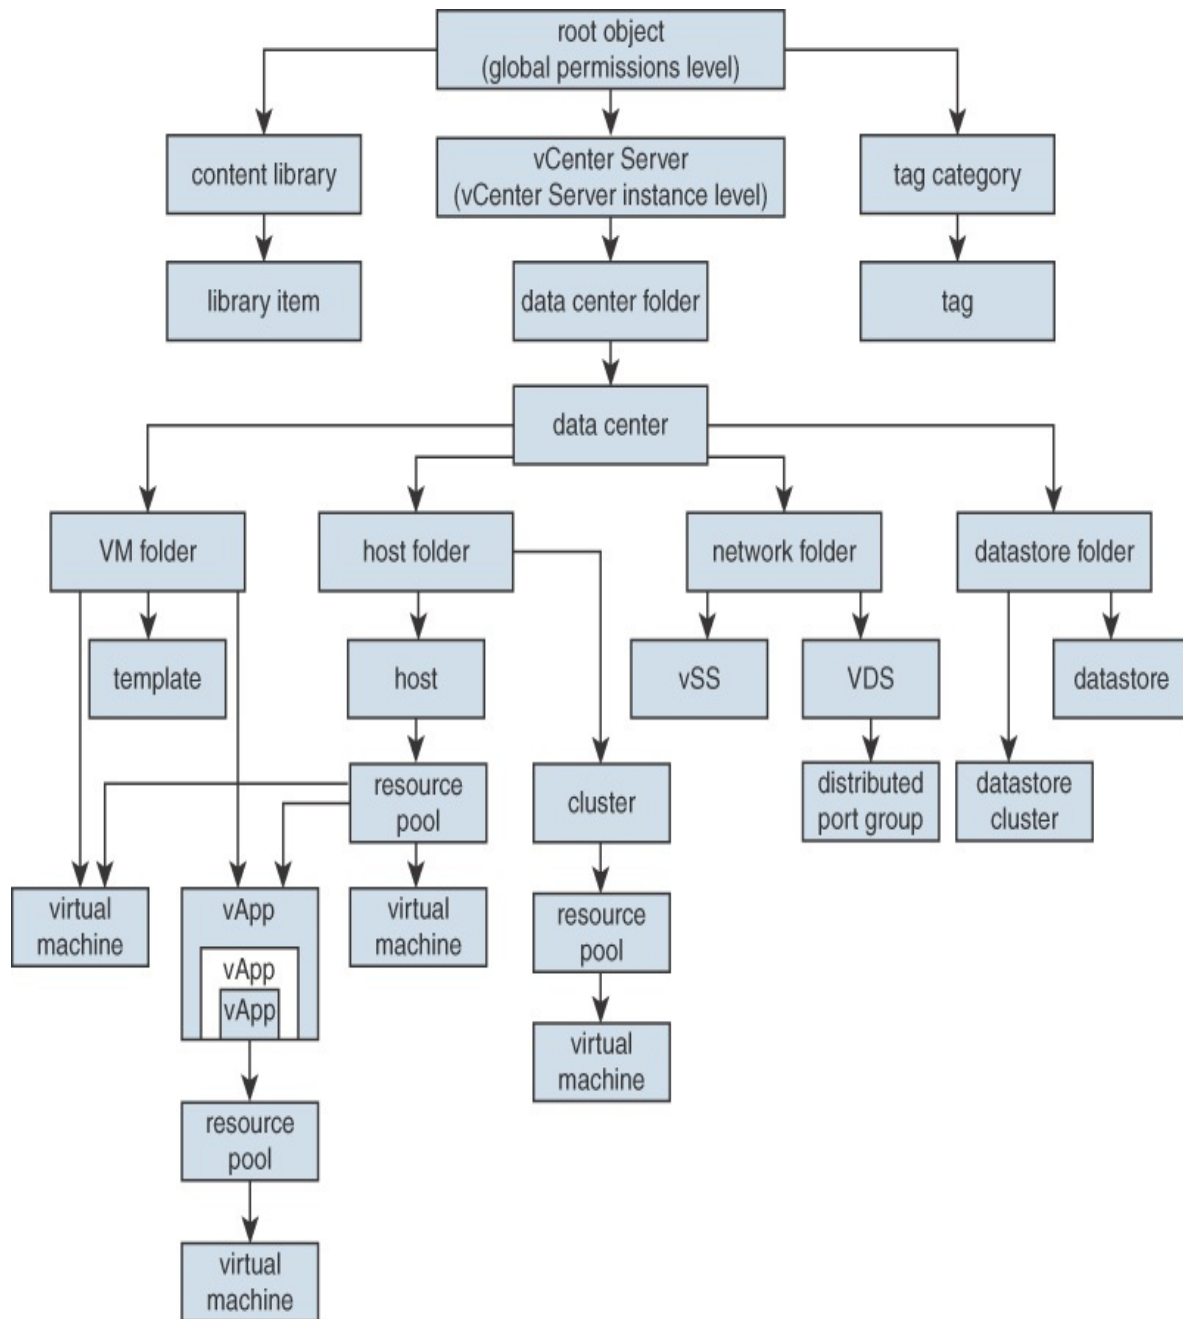
The permissions model for vCenter Server systems relies on assigning permissions to objects in the object hierarchy. A permission involves the assignment of a user (or group) and a role to an inventory object.

When you add a new identity source to SSO, all users can be authenticated, but they effectively have the no access role to the vCenter Server inventory.

## Inventory Hierarchy and Objects

You can assign permissions to objects at different levels of the inventory hierarchy, such as ESXi hosts, clusters, virtual machines, folders, resource pools, datastores, and networks. You can also assign permissions to a global root object to apply the permissions to all objects in all solutions. You can apply permissions to container objects and can optionally allow the permissions to propagate to descendant objects. Most objects inherit permissions from their parents via a single path, but virtual machines inherit permissions from virtual machine folders, hosts, resource pools, and so on, as you can see in [Figure 7-2](#). If an object inherits permissions from two parent objects, then its inherited permissions are determined by the union of the permissions. [Figure 7-2](#) is a diagram from the *VMware vSphere 7.0 Security Guide* that shows the vSphere inventory hierarchy.





**FIGURE 7-2** vSphere Inventory Hierarchy

An object might have multiple permissions but only one permission for each user or group. In other words, you cannot assign to an object two permissions that specify the same group. If multiple permissions are applied to a specific object using multiple groups and if a specific user belongs to more than one of these groups, then the effective permission for that user on that object is the

union of the privileges in applicable roles.

Privileged users can define permissions on managed objects including the following:

- Clusters
- Data centers
- Datastores
- Datastore clusters
- Folders
- Hosts
- Networks (except vSphere Distributed Switches)
- Distributed port groups
- Resource pools
- Templates
- Virtual machines
- vSphere vApps

## Privileges and Roles

Privileges are the lowest-level access controls, and they can be used to define the actions that a user can take on an object in the vSphere inventory. [Table 7-8](#) lists some of the available privilege categories and a few sample privileges in each category.

**Table 7-8** Sample Privileges

Category	Sample Privileges
Virtual machine configuration	Virtual Machine.Configuration.Add Existing Disk
	Virtual Machine.Configuration.Add New Disk
	Virtual Machine.Configuration.Change CPU Count
Datastore	Datastore.Allocate Space



	Datastore.Browse Datastore
	Datastore.Remove File
Virtual machine snapshot	Virtual Machine.Snapshot Management.Create Snapshot
	Virtual Machine.Snapshot Management.Rename Snapshot
	Virtual Machine.Snapshot Management.Revert to Snapshot

A role is composed of a set of privileges. Out of the box, the vCenter Server provides many roles. You cannot modify the vCenter Server system roles, which are described in [Table 7-9](#). You can modify the sample roles, but VMware recommends that you not modify these roles directly but instead clone the roles and modify the clones to suit your case.

### Note

Changes to roles take effect immediately, even for users who are currently logged in to vCenter Server. One exception is with searches, where a change is not realized until the next time the user logs in to vCenter Server.

**Table 7-9** System Roles in vCenter Server 7.0

System Role	Description
Read-only	Allows the user to view the state of an object and details about the object. For example, users with this role can view virtual machine attributes but cannot open the VM console.
Administrator	Includes all privileges of the read-only role and allows the user to view and perform all actions on the object. If you have the administrator role in an object, you can assign privileges to individual users and groups. If you have the administrator role in vCenter Server, you can assign privileges to users and groups in the default SSO identity source. By

default, the [administrator@vsphere.local](mailto:administrator@vsphere.local) user has the administrator role in both vCenter Single Sign-On and vCenter Server.

---

No  
acc  
ess

Prevents users from viewing or interacting with the object. New users and groups are effectively assigned this role by default.

---

No  
cry  
pto  
gra  
phy  
ad  
mi  
nist  
rat  
or

Includes all privileges of the administrator role, except for cryptographic operations privileges. This role allows administrators to designate users who can perform all administrative tasks except encrypting or decrypting virtual machines or accessing encrypted data.

---

Tru  
ste  
d  
infr  
astr  
uct  
ure  
ad  
mi  
nist  
rat  
or  
rol  
e

Allows users to perform VMware vSphere Trust Authority operations on some objects. Membership in the TrustedAdmins group is required for full vSphere Trust Authority capabilities.

---

These are the sample roles in vCenter Server 7.0:

**Key  
Topic**

- Resource pool administrator (sample)
- Virtual machine user (sample)
- VMware consolidated backup user (sample)

- Datastore consumer (sample)
- Network administrator (sample)
- Virtual machine power user (sample)
- Content library administrator (sample)
- Content library registry administrator (sample)

To get familiar with the privileges in a sample role, you can edit the role and explore the privileges that are included in the role. For example, if you edit the virtual machine console user role, you see that it only includes some privileges in the Virtual Machine > Interaction category and no other privileges. Specifically, it includes only these privileges:

- Answer Question
- Configure CD media
- Configure floppy media
- Connect devices
- Console interaction
- Install VMware tools
- Power off
- Power on
- Reset
- Suspend

### **Note**

If you create a role, it does not inherit privileges from any of the system roles.

## **Permissions**

The permissions model for vCenter Server systems relies on assigning permissions to objects in the object hierarchy. A permission is the assignment of

a user (or group) and a role to an inventory object. A permission is set on an object in the vCenter object inventory. Each permission associates the object with a group (or user) and a role, as illustrated in Figure 7-3. For example, you can select a virtual machine object, add one permission that gives the read-only role to Group 1, and add a second permission that gives the administrator role to User 2.

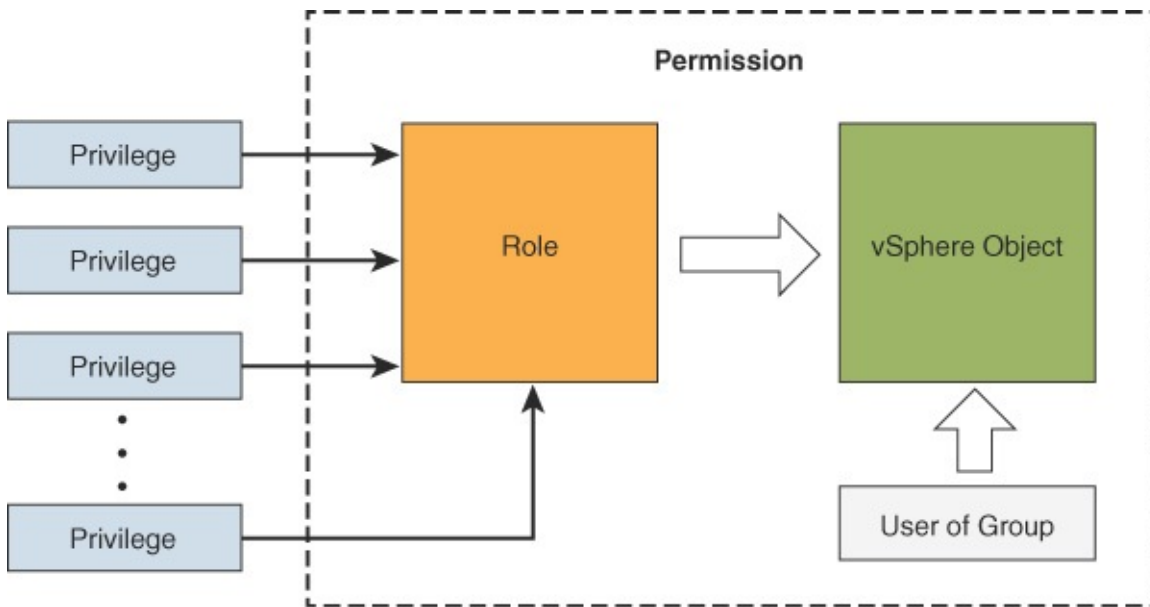


FIGURE 7-3 vSphere Permission Diagram

## Global Permissions

Most entities that appear in the vCenter Server inventory are managed objects whose access can be controlled using permissions. You cannot modify permissions on entities that derive permissions from the root vCenter Server system, such as the following:

- Custom fields
- Licenses
- Roles
- Statistics intervals
- Sessions

The global root object is used to assign permissions across solutions. The

vCenter Server is an example of a solution, and it is attached as a child to the global root object in the hierarchy. The content library and tag category objects are also attached as children to the global root object. Global permissions are permissions that are applied to the global root object and span solutions. For example, a global permission can be applied to both vCenter Server and vRealize Orchestrator. Each solution has its own root object in the hierarchy, whose parent is the global root object. You can give a group of users read permissions to all objects in both object hierarchies.

## **Best Practices for Roles and Permissions**

VMware recommends the following best practices when configuring roles and permissions in a vCenter Server environment:

- Where possible, assign roles to groups rather than to individual users.
- Grant permissions to users (groups) only on the objects where they are required. Use the minimum permissions to meet the required functionality.
- If you assign a restrictive role to a group, check to ensure that the group does not contain the administrator user or other users who require administrative privileges.
- Use folders to group objects. For example, to grant modify permission on one set of hosts and view permission on another set of hosts, place each set of hosts in a folder.
- Use caution when adding a permission to the root vCenter Server objects. Users with privileges at the root level have access to global data on vCenter Server, such as roles, custom attributes, and vCenter Server settings.
- Consider enabling propagation when you assign permissions to an object. Propagation ensures that new objects in the object hierarchy inherit permissions. For example, you can assign a permission to a virtual machine folder and enable propagation to ensure that the permission applies to all VMs in the folder.
- Use the no access role to mask or hide specific areas of the hierarchy. The no access role restricts access for the users or groups with that role.

### **Note**

Changes to licenses propagate to all linked vCenter Server systems in the same vCenter Single Sign-On domain.

## Required Privileges for Common Tasks

Many tasks require permissions on multiple objects in the inventory. Consider the following:

- To perform any operation that consumes storage space, such as taking a snapshot, you must have the Datastore.Allocate Space privilege on the target datastore in addition to having the directly required privileges on the major object.
- Moving an object in the inventory hierarchy requires appropriate privileges on the object itself, the source parent object (such as a folder or cluster), and the destination parent object.
- Deploying a virtual machine directly to a host or cluster requires the Resource.Assign Virtual Machine to Resource Pool privilege because each host or cluster has its own implicit resource pool.

[Table 7-10](#) shows the required privileges for a few common tasks.

**Table 7-10** Required Permissions for Common Tasks

Task	Required Privileges
Create a virtual machine	On the destination folder or in the data center: <ul style="list-style-type: none"><li>■ Virtual Machine.Inventory.Create New</li><li>■ Virtual Machine.Configuration.Add New Disk</li><li>■ Virtual Machine.Configuration.Add Existing Disk</li></ul>

- Virtual Machine.Configuration.Raw Device

On the destination host or cluster or in the resource pool:

- Resource.Assign Virtual Machine to Resource Pool

On the destination datastore or in the datastore folder:

- Datastore.Allocate Space

On the network:

- Network.Assign Network

---

Deploy a virtual machine from a template

On the destination folder or in the data center:

- Virtual Machine.Inventory.Create from Existing

- Virtual Machine.Configuration.Add New Disk

On a template or in a template folder:

- Virtual Machine.Provisioning.Deploy Template

On the destination host or cluster or in the resource pool:

- Resource.Assign Virtual Machine to Resource Pool

On the destination datastore or in a datastore folder:

- Datastore.Allocate Space

On the network that the virtual machine will be assigned to:

- Network.Assign Network



Take a virtual machine snapshot

On the virtual machine or in a virtual machine folder:

- Virtual Machine.Snapshot Management.Create Snapshot

On the destination datastore or in a datastore folder:

- Datastore.Allocate Space

---

Move a virtual machine into a resource pool

On the virtual machine or in a virtual machine folder:

- Resource.Assign Virtual Machine to Resource Pool
- Virtual Machine.Inventory.Move

In the destination resource pool:

- Resource.Assign Virtual Machine to Resource Pool

---

Install a guest operating system on a virtual machine

On the virtual machine or in a virtual machine folder:

- Virtual Machine.Interaction.Answer Question
- Virtual Machine.Interaction.Console Interaction
- Virtual Machine.Interaction.Device Connection
- Virtual Machine.Interaction.Power Off
- Virtual Machine.Interaction.Power On
- Virtual Machine.Interaction.Reset
- Virtual Machine.Interaction.Configure CD Media
- Virtual Machine.Interaction.Configure Floppy Media
- Virtual Machine.Interaction.Tools Install

On a datastore containing the installation media ISO image:

- Datastore.Browse Datastore

On the datastore to which you upload the installation media ISO image:

- Datastore.Browse Datastore
- Datastore.Low Level File Operations

---

Migrate a virtual machine with vMotion

On the virtual machine or in a virtual machine folder:

- Resource.Migrate Powered on Virtual Machine
- Resource.Assign Virtual Machine to Resource Pool

On the destination host or cluster or in a resource pool:

- Resource.Assign Virtual Machine to Resource Pool

---

Cold migrate (relocate) a virtual machine

On the virtual machine or in a virtual machine folder:

- Resource.Migrate Powered Off Virtual Machine
- Resource.Assign Virtual Machine to Resource Pool

On the destination host or cluster or in a resource pool:

- Resource.Assign Virtual Machine to Resource Pool

On the destination datastore:

- Datastore.Allocate Space

---

Migrate a virtual machine with Storage vMotion

On the virtual machine or in a virtual machine folder:

- Resource.Migrate Powered On Virtual Machine

On the destination datastore:

- Datastore.Allocate Space

---

Move a host into a cluster

On the host:

- Host.Inventory.Add Host to Cluster

On the destination cluster:

- Host.Inventory.Add Host to Cluster
- Host.Inventory.Modify. cluster

---

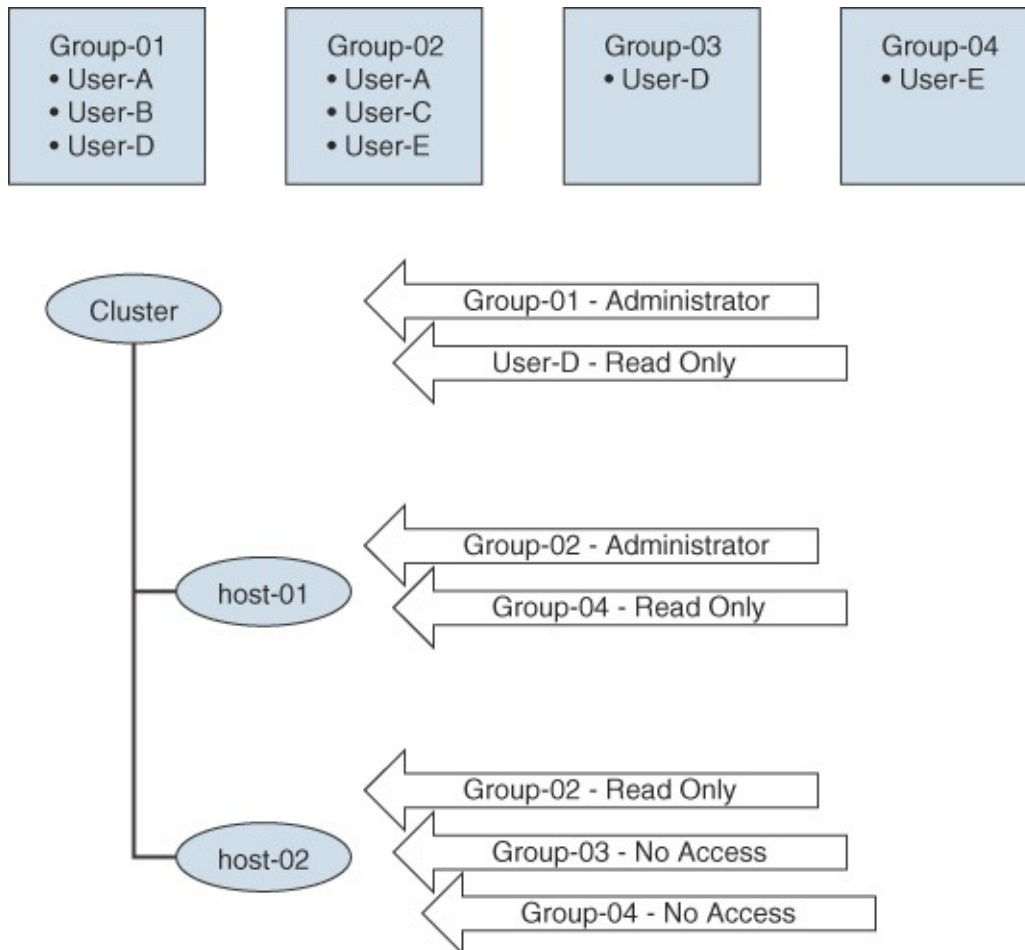
## How Permissions Are Applied by vCenter Server

As you assign each permission, you can choose whether to allow the permission to propagate to child objects. This setting is made for each permission and cannot be universally applied. The default setting is to allow propagation to child objects. The propagation is applied to the vSphere inventory hierarchy as shown in [Figure 7-2](#).

### Key Topic

If conflicting permissions are applied to an object and to its ancestors, the permissions that are assigned at a lower-level object in the inventory hierarchy override permissions assigned at a higher-level object. If multiple permissions are assigned to the same object in different groups that contain a specific user, then that user's effective permission is the union of the associated privileges. Permissions assigned to a user override permissions assigned to groups containing the user when the permissions are applied to the same object. The no

access permission is given precedence over all other roles. For example, consider the following scenario, which is illustrated in [Figure 7-4](#):



**FIGURE 7-4** vSphere Permission Example

- One cluster exists in the inventory, and it contains host-01 and host-02.
- The user account User-A is a member of groups Group-01 and Group-02.
- The user account User-B is a member of group Group-01.
- The user account User-C is a member of group Group-02.
- The user account User-D is a member of groups Group-01 and Group-03.
- The user account User-E is a member of groups Group-02 and Group-04.

Propagate to Child Objects is enabled for each of the following permissions:

- A permission assigns Group-01 the administrator role on the cluster.
- A permission assigns Group-02 the administrator role on host-01.
- A permission assigns Group-02 the read-only role on host-02.
- A permission assigns User-D the read-only role on the cluster.
- A permission assigns Group-03 the no access role on host-02.
- A permission assigns Group-04 the read-only role on host-01.
- A permission assigns Group-04 the no access role on host-02.

In this scenario, the following effective permissions apply:

- User-A:
  - Can perform all tasks on the cluster object
  - Can perform all tasks on the host-01 object
  - Can only view the host-02 object
- User-B:
  - Can perform all tasks on the cluster object
  - Can perform all tasks on the host-01 object
  - Can perform all tasks on the host-02 object
- User-C:
  - Cannot view or perform any task on the cluster object
  - Can perform all tasks on the host-01 object
  - Can only view the host-02 object
- User-D
  - Can only view the cluster object
  - Can only view the host-01 object
  - Cannot view or perform any task on the host-02 object
- User-E

- Cannot view or perform any task on the cluster object
- Can perform all tasks on the host-01 object
- Cannot view or perform any task on the host-02 object

## ESXi and vCenter Server Security

ESXi has many built-in security features, such as CPU isolation, memory isolation, and device isolation. An ESXi host is protected with a firewall that is intended to permit only required network traffic. Starting with vSphere 6.0, ESXi hosts participate in the certificate infrastructure and, by default, are provisioned with certificates that are signed by VMCA.

Optionally, you can further harden ESXi by configuring features such as lockdown mode, certificate replacement, and smart card authentication for enhanced security. You should consider limiting direct access to ESXi hosts, using security profiles, using host profiles, and managing certificates. In addition, you can take other security measures, such as using multiple networks to segregate ESXi network functions and implementing UEFI Secure Boot for ESXi hosts

### Built-in Security Features



The following list identifies some vSphere built-in security features.

- ESXi Shell and SSH are disabled by default.
- By default, ESXi runs only services that are essential to managing its functions.
- By default, all ports that are not required for management access to the host are closed.
- By default, weak ciphers are disabled and communications from clients are secured by SSL. Default certificates created on ESXi use PKCS#1 SHA-256 with RSA encryption as the signature algorithm.
- A Tomcat web service is used internally by ESXi to support access by web clients. ESXi is not vulnerable to the Tomcat security issues reported in other use cases because the service has been modified to run only functions



that a web client requires for administration and monitoring.

- VMware monitors all security alerts that can affect ESXi security and issues security patches when needed.
- Secure services such as SSH and SFTP are available and should be used instead of insecure counterparts, such as Telnet and FTP.
- ESXi provides the option of using UEFI Secure Boot.
- When a TPM 2.0 chip is available in the hardware and configured in the system BIOS, ESXi works with Secure Boot to enhance security and trust assurance rooted in hardware.

## Security Profiles

You can customize many of the essential security settings for a host through the Security Profile panel in the vSphere Client. You can use security profiles to customize services and configure the ESXi firewall. [Table 7-11](#) describes the services that are available to you to view and manage using the vSphere Client for a default vSphere installation, along with the default state for each of them. You can use the vSphere Client to start, stop, and restart individual services.

**Table 7-11** ESXi Security Profile Services

Service	Default State	Description
Direct Console User Interface (DCUI)	Running	Allows you to interact with an ESXi host from the local console host using text-based menus
ESXi Shell	Stopped	Is available from the DCUI or from SSH
SSH	Stopped	Allows remote connections through Secure Shell
Load-Based Teaming Daemon	Running	Enables load-based teaming
attestd	Stopped	Enables the vSphere Trust Authority Attestation Service
kmxd	Stop	Enables the vSphere Trust Authority Key

	ped	Provider Service
Active Directory Service	Stopped	Is started on hosts after you configure ESXi for Active Directory
NTP Daemon	Stopped	Enables the Network Time Protocol daemon
PC/SC Smart Card Daemon	Stopped	Is started on hosts after you enable the host for smart card authentication
CIM Server	Running	Can be used by Common Information Model (CIM) applications
SNMP Server	Stopped	Enables the SNMP daemon
Syslog Server	Stopped	Enables the syslog daemon
VMware vCenter Agent (vpxa)	Running	Connects the host to vCenter Server
X.Org Server	Stopped	Internally used for virtual machine 3D graphics

Table 7-12 lists the firewall ports that are installed by default in ESXi 7.0. On a specific host, the list of actual services and firewall ports can be impacted by the currently installed VMware Installation Bundles (VIBs).

**Table 7-12** Incoming and Outgoing Firewall Ports

Firewall Service	Incoming Port(s)	Outgoing Port(s)
CIM Server	5988 (TCP)	
CIM Secure Server	5989 (TCP)	
CIM SLP	427 (TCP,UDP)	427 (TCP,UDP)
DHCPv6	546 (TCP,UDP)	547 (TCP,UDP)
DVSSync	8301, 8302 (UDP)	8301, 8302 (UDP)
HBR		44046, 31031 (TCP)
NFC	902 (TCP)	902 (TCP)
WOL		9 (UDP)

vSAN Clustering	12345, 23451 (UDP)	12345, 23451 (UDP)
DCHP Client	68 (UDP)	68 (UDP)
DNS Client	53 (UDP)	53 (TCP,UDP)
Fault Tolerance	8100, 8200, 8300 (TCP,UDP)	80, 8100, 8200, 8300 (TCP,UDP)
NSX Distributed Logical Router Service	6999 (UDP)	6999 (UDP)
Software iSCSI Client		3260 (TCP)
rabbitmqproxy		5671 (TCP)
vSAN Transport	2233 (TCP)	2233 (TCP)
SNMP Server	161 (UDP)	
SSH Server	22 (TCP)	
vMotion	8000 (TCP)	8000 (TCP)
VMware vCenter Agent		902 (UDP)
vSphere Web Access	80 (TCP)	
vsanvp	8080 (TCP)	8080 (TCP)
RFB Protocol	5900–5964 (TCP)	
vSphere Life Cycle Manager	80, 9000 (TCP)	80, 9000 (TCP)
I/O Filter	9080 (TCP)	

The RFB protocol (TCP 5900–5964) and OpenWSMAN daemon (TCP 8889) are firewall ports for services that are not visible in the vSphere Client by default.

## ESXi Password Hardening

One step in hardening an ESXi host is to harden the password required to use its predefined local administrator account, which is called root. By default, the ESXi host enforces passwords for its local user accounts, which may be used to access the host via the Direct Console User Interface (DCUI), the ESXi Shell, Secure Shell (SSH), or the vSphere Client. You can modify the ESXi password

requirements by setting the Security.PasswordQualityControl advanced option for the host. For example, you can set Security.PasswordQualityControl to configure the ESXi host to accept passphrases, which it does not accept by default.

## **Joining an ESXi Host to a Directory Service**

You can join an ESXi host to a directory service, such as Active Directory, and configure permissions to allow the associated users to connect directly to the ESXi host using DCUI, ESXi Shell, SSH, or the vSphere Host Client. The main reason for doing this is to reduce the number of local ESXi user accounts that you must create and manage. Another reason is to provide users with the means to directly access ESXi with an existing user account whose password is already hardened.

## **vSphere Authentication Proxy**

You can add ESXi hosts to an Active Directory domain by using the vSphere Authentication Proxy instead of adding the hosts explicitly to the Active Directory domain. To do this, you can add the host's IP address to the vSphere Authentication Proxy access control list. By default, the vSphere Authentication Proxy authorizes the host, based on its IP address. You can enable client authentication to have vSphere Authentication Proxy check the host's certificate. If you are using Auto Deploy, you can configure a reference host to point to the Authentication Proxy, set up a rule that applies the reference host's profile to other hosts provisioned by Auto Deploy, let Auto Deploy store the host's IP address in the access control list, and join the host to the AD domain.

## **ESXi Host Access**

You can implement lockdown mode to force operations to be performed through vCenter Server. You can choose to use strict lockdown mode, which disables the DCUI service, or normal lockdown mode, which allows DCUI access for some users. In normal lockdown mode, user accounts that are in the Exception Users list and have administrator privileges on the host can access the DCUI. A common use case for this is to provide access to service accounts, such as backup agents. Also, in normal lockdown mode, users identified in the host's DCUI.Access advanced option can access the DCUI. If the ESXi Shell or SSH is enabled and the host is placed in lockdown mode, accounts in the Exception Users list who have administrator privileges can use these services. For all other

users, ESXi Shell or SSH access is disabled. The main use case for this is to provide user access in the event of catastrophic failure. Starting with vSphere 6.0, ESXi or SSH sessions for users who do not have administrator privileges are terminated.

## **Control MOB Access**

The vCenter Server Managed Object Browser (MOB) provides a means to explore the vCenter Server object model. Its primary use is for debugging. It provides the ability to make some configuration changes, so it may be considered a vulnerability for malicious attacks. The MOB is disabled by default. You should enable it only for debugging or for tasks that require it, such as extracting an old certificate from a system. To enable MOB, you can use the vSphere Client to set the host's advanced system setting `Config.HostAgent.plugins.solo.enableMob`. You should not use the **vim-cmd** tool in the ESXi Shell for this purpose.

## **ESXi Secure Boot and TPM**

UEFI Secure Boot is a mechanism which ensures that only trusted code is loaded by the EFI firmware prior to OS handoff. When Secure Boot is enabled, the UEFI firmware validates the digitally signed kernel of an OS against a digital certificate stored in the UEFI firmware. Starting with vSphere 6.5, ESXi supports Secure Boot if it is enabled in the hardware. ESXi version 6.5 and later supports UEFI Secure Boot at each level of the boot stack.

ESXi is composed of digitally signed packages called vSphere Installation Bundles (VIBs). These packages are never broken open. At boot time, the ESXi file system maps to the content of those packages. By leveraging the same digital certificate in the host UEFI firmware that is used to validate the signed ESXi kernel, the kernel validates each VIB by using the Secure Boot verifier against the firmware-based certificate, ensuring a cryptographically "clean" boot.

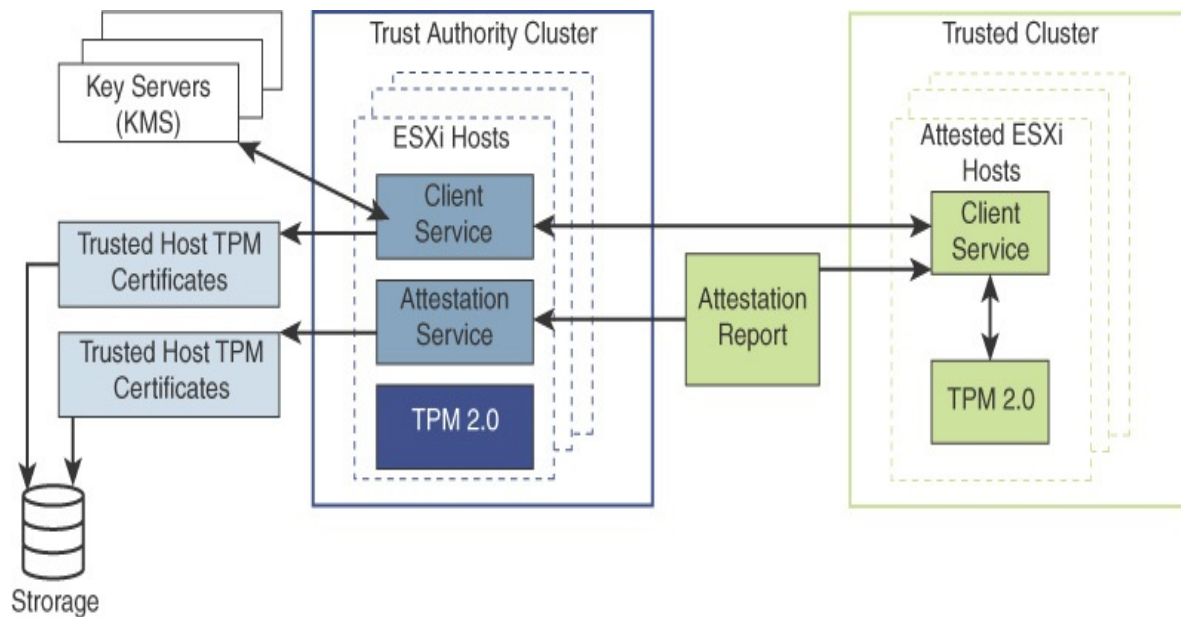
When Secure Boot is enabled, ESXi prevents the installation of unsigned code on ESXi. To install unsigned code such as beta drivers, you must disable Secure Boot. When Secure Boot is enabled, the Secure Boot verifier runs, detects the unsigned VIB, and crashes the system, which produces the purple screen of death (PSOD) event, which identifies the VIB that must be removed. To remediate, you need to boot the ESXi host with Secure Boot disabled, remove the VIB, and reboot with Secure Boot enabled.

ESXi can use **Trusted Platform Module (TPM)** chips, which are secure cryptoprocessors that enhance host security by providing a trust assurance rooted in hardware as opposed to software. TPM is an industry standard for secure cryptoprocessors. TPM chips are found in most of today's computers, including laptops, desktops, and servers. vSphere 7.0 supports TPM version 2.0. A TPM 2.0 chip attests to an ESXi host's identity. Host attestation is the process of authenticating and attesting to the state of the host's software at a given point in time. UEFI Secure Boot, which ensures that only signed software is loaded at boot time, is a requirement for successful attestation.

## vSphere Trust Authority (vTA)



In an environment where TPM attestation is used, you can implement the vSphere Trust Authority (vTA), which uses its own management cluster to serve as a hardware root of trust. Ideally the vTA trusted hosts cluster is small and separate from all other clusters and has very few administrators. In vSphere 6.7, you can leverage TPM and vCenter Server to identify hosts that failed attestation, but you cannot automatically prevent secured/encrypted workloads from migrating to those hosts. Also, you cannot encrypt vCenter Server. In vSphere 7.0 with vTA, you can enable the trusted hosts cluster to handle the attestation of other hosts and to take over the distribution of the encryption keys from the key management servers (KMSs). This removes vCenter Server from the critical path for key distribution and enables you to encrypt vCenter Server. A trusted infrastructure consists of at least one vSphere Trust Authority cluster, at least one Trusted cluster, and at least one external KMIP-compliant KMS, as illustrated in [Figure 7-5](#), which is a diagram from the *VMware vSphere 7.0 Security Guide*.



**FIGURE 7-5** vSphere Trusted Authority Architecture

## vCenter Server Security

You should follow VMware guidelines, summarized in this section, to ensure security of the vCenter Server environment.

### User Access

The user accounts defined in the local operating system (localos) of the Linux-based vCenter Server Appliance have no permissions defined in the vCenter Server environment. The localos user accounts—such as root, sshd, and vdtc—are not members of any SSO domain (vsphere.local) group to which permissions are applied. No one should attempt to use these accounts to log in to the vSphere Client. You should not use these accounts when configuring permissions or group memberships. Do not allow users to log in directly to the localos of the vCenter Server appliance. Log in only locally when required.

By default, the only accessible user account in the SSO domain is administrator, which has full control of the environment. If you use the default SSO domain name, the user account is [administrator@vsphere.local](mailto:administrator@vsphere.local). Ideally, you should integrate vSphere with a supported enterprise directory service, such as Active Directory, to allow users seamless access without requiring additional user accounts. Alternatively, you can create other user accounts in the SSO domain for your users. You should ensure that each user can access the environment with

a unique account that is assigned the minimally required privileges.

## Note

Do not confuse the administrator (root) of the localos with the SSO administrator ([administrator@vsphere.local](mailto:administrator@vsphere.local) by default). By default, no localos user account has full administrator privileges in vCenter Server.

For users who require the administrator role, you should assign the role to the appropriate user accounts or group accounts to avoid using the SSO administrator account.

The vCenter Server connects to each ESXi host with the vpxuser account defined on the host. By default, vCenter Server changes the vpxuser password automatically every 30 days on each connected ESXi host. To change this behavior, you can change the value of the vCenter Server advanced setting `VimPasswordExpirationInDays`.

## vCenter SSO Password Policy

The password for the SSO administrator account and other SSO domain user accounts is controlled by the SSO password policy. By default, this password must meet the following requirements:

- At least eight characters
- At least one lowercase character
- At least one numeric character
- At least one special character

In addition, the password cannot use more than 20 characters and cannot contain non-ASCII characters. SSO administrators can change the default password policy.

## Restricting Administrative Privileges

You should use permissions to assign the administrator role to just the specific users and group that truly require it. You should create and use custom roles with only the required privileges when creating permissions. In other words, you



should apply the principle of least privileges when configuring permissions in vCenter Server.

By default, a user with the administrator role can interact with files and applications in a virtual machine's guest operating system. If administrators do not require this interaction, consider applying a role without the guest operations privilege.

## **Restricting vCenter Server Access**

You should minimize users who can log directly in to the vCenter Server localos, as they could intentionally or unintentionally cause harm by altering settings and modifying processes. Allow only users with legitimate purposes to log in to the system and ensure that the login events are audited.

You should secure the network where vCenter Server is connected by applying the information in the “vSphere Network Security” section, later in this chapter.

## **Controlling Datastore Browser Access**

Assign the Datastore.Browser datastore privilege only to users and user groups who truly require it.

## **vCenter Server and Client Certificates**

You should ensure that vSphere Client users and other client applications heed certificate verification warnings to prevent man-in-the-middle (MITM) attacks.

You should remove any expired or revoked certificates from the vCenter Server to avoid MITM attacks.

## **Time Synchronization**

You should ensure that all systems, such as vCenter Server, ESXi, and supporting services, use the same relative time source. The time source must be in sync with an acceptable time standard, such as Coordinated Universal Time (UTC). Time synchronization is critical for many vSphere features, such as vSphere HA. It is also critical for securing vSphere.

Time synchronization is essential for certificate validation. Time synchronization simplifies troubleshooting and auditing. Incorrect time settings make it difficult to analyze and correlate log files related to detecting attacks and conducting

security audits.

## vSphere Network Security

You can use firewalls, segmentation, VLANs, and other measures to secure the networks used by your virtual machines and the vSphere environment. It is important to put vCenter Server on the management network only. In addition, you should avoid putting the vCenter Server system on other networks, such as your production network or storage network, or on any other network that has access to the Internet. vCenter Server does not need access to the network where vMotion operates.

### Firewalls

You can use traditional (physical) firewalls, virtual machine–based firewalls, and hypervisor-based firewalls (such as NSX Distributed Firewall) to protect traffic to and from the vCenter Server, ESXi hosts, virtual machines, and other vSphere components. Ideally, you could use firewalls to allow only the required traffic between specific vSphere components, virtual machines, and network segments.

### Segmentation and Isolation

You should keep different virtual machine zones within a host on different network segments to reduce the risk of data leakage and threats. Such threats include Address Resolution Protocol (ARP) spoofing, in which an attacker manipulates the ARP table to remap MAC and IP addresses and gains access to network traffic to and from a host. Attackers use ARP spoofing to generate MITM attacks, perform denial of service (DoS) attacks, and hijack the systems. You can implement segmentation by using one of two approaches:

- **Use separate physical network adapters for virtual machine zones:** This may be the most secure method.
- **Set up virtual local-area networks (VLANs) for virtual machine zones:** This may be the most cost-effective method.

You should isolate the vSphere management network, which provides access to the management interface on each component. In most cases, you should place the vSphere management port group in a dedicated VLAN and ensure that the network segment is not routed except to other management-related networks.

Likewise, you should isolate IP-based storage traffic and vMotion traffic.

## Internet Protocol Security

You can configure Internet Protocol Security (IPsec) on ESXi hosts to enable authentication and encryption of incoming and outgoing packets. You can configure security associations to control *how* the system encrypts the traffic. For each association, you configure a name, a source, a destination, and encryption parameters. You can configure security policies to determine *when* the system should encrypt traffic. Security policies include information such as source, destination, protocol, direction, mode, and a security association.

To list the available security associations, you can use this command in ESXi:

```
esxcli network ip ipsec sa list
```

To add a security association, you can use the **esxcli network ip ipsec sa add** command with one or more of the options listed in [Table 7-13](#).

**Table 7-13** IPsec Options

Option	Description
<b>--sa-source</b> = <i>source address</i>	Required. Specify the source address.
<b>--sa-destination</b> = <i>destination address</i>	Required. Specify the destination address.
<b>--sa-mode</b> = <i>mode</i>	Required. Specify the mode: either transport or tunnel.
<b>--sa-spi</b> = <i>security parameter index</i>	Required. Specify the security parameter index, in hexadecimal.
<b>--encryption-algorithm</b> = <i>encryption algorithm</i>	Required. Specify the algorithm as one of the following parameters: <ul style="list-style-type: none"><li>▪ <b>3des-cbc</b></li><li>▪ <b>aes128-cbc</b></li></ul>

- **null**

(no encryption)

<b>--integrity- algorithm=authentication algorithm</b>	Required. Specify the authentication algorithm: either <b>hmac-sha1</b> or <b>hmac-sha2-256</b> .
<b>--integrity- key=authentication key</b>	Required. Specify the authentication key. You can enter keys as ASCII text or in hexadecimal.
<b>--sa-name=name</b>	Required. Provide a name for the security association.

## General Networking Security Recommendations

The following are additional general networking security recommendations:

### Key Topic

- If spanning tree is enabled, ensure that physical switch ports are configured with PortFast.
- Ensure that NetFlow traffic for a vSphere Distributed Switch (vDS) is sent only to authorized collector IP addresses.
- Ensure that only authorized administrators have access to virtual networking components by using the role-based access controls.
- Ensure that port groups are not configured to the value of the native VLAN.
- Ensure that port groups are not configured to VLAN values reserved by upstream physical switches.
- Ensure that port groups are not configured to VLAN 4095 except in the case of Virtual Guest Tagging (VGT).
- On distributed virtual switches, restrict port-level configuration overrides. The port-level override option is disabled by default.
- Ensure that vDS port mirror traffic is sent only to authorized collector ports

or VLANs.

## Network Security Policies

You should connect virtual machines to standard virtual switch port groups or distributed virtual switch port groups that are configured with an appropriate security policy. The network security policy provides three options, which may be set to Reject or Accept, as described in [Table 7-14](#).

**Table 7-14** Network Security Policies

Option	Setting	Description
Permitted	Accept	The virtual switch forwards all frames to the virtual network adapter.
Used	Reject	The virtual switch forwards only the frames that are addressed to the virtual network adapter.
Modified	Accept	If the guest operating system changes the effective MAC address of the virtual adapter to a value that differs from the MAC address assigned to the adapter in the VMX file, the virtual switch allows the inbound frame to pass.
Selected	Reject	If the guest operating system changes the effective MAC address of the virtual adapter to a value that differs from the MAC address assigned to the adapter in the VMX file, the virtual switch drops all inbound frames to the adapter. If the guest OS changes the MAC address back to its original value, the virtual switch stops dropping the frames and allows inbound traffic to the adapter.

F A The virtual switch does not filter outbound frames. It permits all  
or c outbound frames, regardless of the source MAC address.  
g c  
e e  
d p  
T t

---

ra R The virtual switch drops any outbound frame from a virtual  
ns ej machine virtual adapter that uses a source MAC address that differs  
m e from the MAC address assigned to the virtual adapter in the VMX  
it ct file.  
s

---

On a distributed virtual switch, you can override the security policy per virtual port.

## Virtual Machine Security

To harden a virtual machine, you can follow best practices, configure UEFI, implement security policies, protect against denial of service attacks, and implement encryption, as described in this section.

### Virtual Machine Hardening Best Practices

To harden virtual machines, you should apply the following best practices:



- **General protection:** In most respects, you should treat virtual machine as you would a physical server when it comes to applying security measures. For example, be sure to install guest operating systems patches, protect with antivirus software, and disable unused serial ports.
- **Templates:** Carefully harden the first virtual machine deployment of each guest O/S and verify hardening completeness. Convert the virtual machine into a template and use the template to deploy virtual machines, as needed.
- **Virtual machine console:** Minimize the use of this console. Use it only when required. Use remote tools such as SSH and Remote Desktop to access virtual machines. Consider limiting the number of console connections to just one.

- **Virtual machine resource usage:** Prevent virtual machines from taking over resources on the ESXi host to minimize the risk of denial of service to other virtual machines. Configure each virtual machine with enough virtual hardware but not more virtual hardware resources than needed. For example, configure each virtual machine with enough virtual memory to handle its workload and meet application vendor recommendations, but do not provide much more memory than you expect it will need. Consider setting reservations or shares to ensure that critical virtual machines have access to enough CPU and memory.
- **Unnecessary services:** Disable or uninstall any function for the guest operating system that is not required in order to reduce the number of components that can be attacked and to reduce its resource demand. For example, turn off screen savers, disable unneeded guest operating system services, and disconnect the CD/DVD drive.
- **Unnecessary hardware devices:** To minimize potential attack channels, disable any hardware devices that are not required, such as floppy drives, serial ports, parallel ports, USB controllers, and CD-ROM drives.

## Configuring UEFI Boot



Starting with vSphere 6.5, if the operating system supports UEFI Secure Boot, you can configure your VM to use UEFI boot. Prerequisites are UEFI firmware virtual hardware version 13 or later, VMware Tools version 10.1 or later, and an operating system that supports UEFI Secure Boot. For Linux virtual machines, the VMware host guest file system is not supported in Secure Boot mode and should be removed from VMware Tools before you enable Secure Boot. If you turn on Secure Boot for a virtual machine, you can load only signed drivers into that virtual machine.

In a guest operating system that supports UEFI Secure Boot, each piece of boot software is signed, including the bootloader, the operating system kernel, and operating system drivers. The default configuration includes several code-signing certificates, including a Microsoft certificate for booting Windows, a Microsoft certificate for third-party code, and a VMware certificate for booting ESXi inside a virtual machine. The virtual machine default configuration includes one certificate: the Microsoft Key Exchange Key (KEK) certificate.

If you turn on Secure Boot for a virtual machine, you can only load signed drivers in the guest OS.

## Disabling Unexposed Features

Some virtual machine settings are useful for other platforms (such as VMware Workstation and VMware Fusion) but can be disabled in a vSphere environment. To reduce potential risk, consider setting the following advanced virtual machine options to TRUE:

- `isolation.tools.unity.push.update.disable`
- `isolation.tools.ghi.launchmenu.change`
- `isolation.tools.memSchedFakeSampleStats.disable`
- `isolation.tools.getCreds.disable`
- `isolation.tools.ghi.autologon.disable`
- `isolation.bios.bbs.disable`
- `isolation.tools.hgfsServerSet.disable`

## Other Common Settings

You should consider the following settings, which are commonly set to address specific potential security threats:

- **Disk shrinking:** Because disk shrinking, which involves reclaiming unused disk space from a virtual machine, can take considerable time to complete and its invocation can result in a temporary denial of service, you should disable disk shrinking by using the following lines in the VMX file:

[Click here to view code image](#)

```
isolation.tools.diskwiper.disable = "TRUE"
```

```
isolation.tools.diskShrink.disable = "TRUE"
```

- **Copying and pasting:** Copy and paste operations are disabled by default in new virtual machines. In most cases, you should retain this default to ensure that one user of the virtual machine console cannot paste data that was originally copied from a previous user. Ensure that the following lines remain in the VMX files:



[Click here to view code image](#)

```
isolation.tools.copy.disable = "TRUE"
```

```
isolation.tools.paste.disable = "TRUE"
```

- **Connecting devices:** By default, the ability to connect and disconnect devices is disabled. One reason is to prevent one user from accessing a sensitive CD-ROM device that was left in the drive. Another reason is to prevent users from disconnecting the network adapter, which could lead to denial of service. Ensure that the following lines remain in the VMX file:

[Click here to view code image](#)

```
isolation.device.connectable.disable = "TRUE"
```

```
isolation.device.edit.disable = "TRUE"
```

- **Logging:** Uncontrolled virtual machine logging could lead to denial of service if the associated datastore runs out of disk space. VMware recommends keeping 10 log files. To set this on a virtual machine, set the following in the VMX file:

```
vmx.log.keepOld = "10"
```

Alternatively, to limit the number of log files for virtual machines on an ESXi host, add the previous line to the host's `/etc/vmware/config` file. A more aggressive measure is to disable virtual machine logging with the following statement in the VMX file:

```
logging = "FALSE"
```

- **VMX file size:** By default, the size of each VMX file is 1 MB because uncontrolled file sizes can lead to denial of service if the datastore runs out of disk space. Occasionally, *setinfo* messages that define virtual machine characteristics or identifiers are sent as name/value pairs from the virtual machine to the VMX file. If needed, you can increase the size of the VMX file limit by using the following statement in the VMX file and replacing the numeric value with a larger value:

[Click here to view code image](#)

```
tools.setInfo.sizeLimit = "1048576"
```

If `tools.setInfo.sizeLimit` is not set in the virtual machine's advanced options, the default size applies. In most cases, you can keep the default setting as a security measure.

- **Performance counters:** VMware Tools provides performance counters on CPU and memory from the ESXi host into the virtual machine for use by PerfMon. This feature is disabled by default because an adversary could potentially make use of this information to attack the host. Ensure that the following line remains in the VMX files, as it blocks some (but not all) performance metrics:

[Click here to view code image](#)

```
tools.guestlib.enableHostInfo = "FALSE"
```

## Virtual Machine Risk Profiles

VMware's *vSphere 6.0 Hardening Guide* provides guidelines for addressing vulnerabilities based on risk profiles. When you can apply this hardening guide to your environment, the first step is to apply the appropriate risk profile, based on the sensitivity of your environment and data. The hardening guide offers three risk profiles:

- **Risk Profile 1:** Intended to be implemented in just the most secure environments, such as top-secret government environments.
- **Risk Profile 2:** Intended to be implemented in sensitive environments to protect sensitive data such as data that must adhere to strict compliance rules.
- **Risk Profile 3:** Intended to be implemented in all production environments.

For vSphere 6.7, the *vSphere Hardening Guide* is replaced with the *vSphere 6.7 Update 1 Security Configuration Guide*. The risk profiles are removed, primarily because the only remaining Risk Profile 1 setting is `ESXi.enable-strict-lockdown-mode`. Instead of identifying risk profiles, the new guide simply lists the current 50 guideline IDs alphabetically and includes a vulnerability discussion for each guideline. Because no *vSphere 7.0 Security Configuration Guide* is available as of July 2020, you should get familiar with the risk profile data described in this section.

## Protecting Virtual Machines Against Denial-of-Service Attacks

As mentioned earlier, the size limit of the virtual machine configuration file (VMX file) is 1 MB by default, but you can change it by using the `tools.setInfo.sizeLimit` parameter to avoid filling the datastore and causing denial of service.

Virtual Machine Communication Interface (VMCI) is a high-speed communication mechanism for virtual machine-to-ESXi host communication. In some VMware products, including ESXi 4.x, VMCI also provides high-speed communication between virtual machines on the same ESXi host. In ESXi 5.1, the guest-to-guest VMCI is removed. In a VMX file, the `vmci0.unrestricted` parameter is used to control VMCI isolation for virtual machines running on ESX/ESXi 4.x and ESXi 5.0, but it has no effect on virtual machines running on ESXi 5.1 and later. Any DoS concerns related to VMCI in previous vSphere versions do not apply to vSphere 7.0.

Non-administrative users in the guest operating system can shrink virtual disks to reclaim the disk's unused space. However, if you shrink a virtual disk repeatedly, the disk can become unavailable, leading to denial of service. To prevent this, you could disable the ability to shrink virtual disks by following these steps:

- Step 1.** Shut down the virtual machine.
- Step 2.** Modify the advanced settings in the virtual machine options.
- Step 3.** Set `isolation.tools.diskWiper.disable` and `isolation.tools.diskShrink.disable` to TRUE.

## Controlling VM Device Connections

As mentioned earlier in this chapter, the ability to connect and disconnect devices is disabled by default for new virtual machines. In most cases, you should not change this behavior. You should verify that the following settings exist in your VMX files, especially if the virtual machines were deployed from a non-hardened template or were originally built on older ESXi hosts:

[Click here to view code image](#)

```
isolation.device.connectable.disable = "TRUE"  
  
isolation.device.edit.disable = "TRUE"
```

If these parameters are set to FALSE, then in a guest operating system, any user or process, with or without root or administrator privileges, could use VMware Tools to change device connectivity and settings. The user or process could then connect or disconnect devices, such as network adapters and CD-ROM drives and modify device settings. This functionality could allow the user or process to connect a CD-ROM with sensitive data or disconnect a network adapter, which could lead to denial of service for other users.

## Virtual Machine Encryption

Starting with vSphere 6.5, you can protect your virtual machines, virtual disks, and other virtual machines files by using virtual machine encryption. In vSphere 6.5 and 6.7, you must set up a trusted connection between vCenter Server and a KMS. The KMS generates and stores keys. It passes the keys to vCenter Server for distribution. Starting in vSphere 7.0, you can remove the need for vCenter to request keys from the KMS by configuring vTA and making encryption keys conditional to cluster attestation.

You can use the vSphere Client or the vSphere API to add key provider instances to a vCenter Server system. vCenter Server uses Key Management Interoperability Protocol (KMIP) to allow flexibility in choosing a KMS. If you use multiple key provider instances, all instances must be from the same vendor and must replicate keys. If you use different KMS vendors in different environments, you can add a key provider for each KMS and specify a default key provider. The first key provider that you add becomes the default key provider, but you can change it.

Only vCenter Servers (not the ESXi hosts) have the credentials for logging in to the KMS. vCenter Server obtains keys from the KMS and pushes them to the hosts. Two types of keys are used for virtual machine encryption:

- **Data encryption keys (DEKs):** DEKs are internal keys generated by the ESXi host and used to encrypt virtual machines and disks. DEKs are XTS-AES-256 keys.
- **Key encryption key (KEKs):** KEKs are the keys that vCenter Server requests from the KMS. KEKs are AES-256 keys. vCenter Server stores only the ID of each KEK and not the key itself. These keys are used to encrypt the DEKs as they are written to an encrypted virtual machine's VMX file.

You can encrypt an existing virtual machine or virtual disk by changing its storage policy. Encryption works with any guest OS because encryption occurs at the hypervisor level. Encryption keys and configuration are not contained in the virtual machine's guest OS. Encryption works with any supported storage type, including VMware vSAN.

You can encrypt virtual disks only for encrypted virtual machines. You cannot encrypt the virtual disk of an unencrypted VM. You can encrypt virtual machine files (NVRAM, VSWP, and VMSN files), virtual disk files, and coredump files. Log files, virtual machine configuration files, and virtual disk descriptor files are not encrypted. For each virtual machine, you can use the vSphere Client to encrypt and decrypt virtual disks independently.

Coredumps are always encrypted on ESXi hosts where encryption mode is enabled. Coredumps on the vCenter Server system are not encrypted. To perform cryptographic operations, you must be assigned the cryptographic operations privilege.

ESXi uses KEKs to encrypt the internal keys and stores the encrypted internal keys on disk. ESXi does not store the KEK on disk. If a host reboots, vCenter Server requests the KEK with the corresponding ID from the KMS and makes it available to ESXi, which decrypts the internal keys as needed. In addition to VMDK files, most virtual machine files that contain guest data are encrypted, such as the NVRAM, VSWP, and VMSN files. The key that vCenter Server retrieves from the KMS unlocks an encrypted bundle in the VMX file that contains internal keys and other secrets.

VM encryption uses vSphere APIs for I/O Filtering (VAIO), which is typically called IOFilter. IOFilter is an ESXi framework that allows for the interception of virtual machine I/O in the virtual SCSI (vSCSI) emulation layer, which is just below the virtual machine and above the file system. It enables VMware and third-party developers to develop services using virtual machine I/O, such as encryption, caching, and replication. It is implemented entirely in user space, which cleanly isolates it from the core architecture and core functionality of the hypervisor. In the event of any failure, only the virtual machine in question is impacted. Multiple filters can be enabled for a particular virtual machine or a virtual disk, which are typically chained in a manner such that I/O is processed serially by each of these filters before the I/O is passed down to VMFS or completed within one of the filters.

The default administrator system role includes all cryptographic operations

privileges. A new default role, the no cryptography administrator, supports all administrator privileges except for the cryptographic operations privileges. You can create a custom role that contains granular cryptographic operations privileges such as by setting Cryptographic Operations to Encrypt (which allows a user to encrypt a virtual machine or virtual disk) or setting Cryptographic Operations to Add Disk (which allows a user to add a disk to an encrypted virtual machine).

The vSphere Client can be used to encrypt and decrypt virtual machines. To re-encrypt a virtual machine, you must use the API. You can use the API to perform a deep re-encryption (replacing the DEK and KEK) or a shallow re-encryption (replacing just the KEK) of a virtual machine. A deep re-encryption requires that the virtual machine be powered off and be free from snapshots. A shallow re-encryption is permitted on a virtual machine with one snapshot (not multiple snapshots). The **crypto-util** command-line utility can be used to decrypt coredumps, check for file encryption, and perform management tasks on the ESXi host.

When a user performs an encryption task, such as creating an encrypted virtual machine, the following events occur:

- The vCenter Server requests a new key from the default KMS to use as the KEK.
- The vCenter Server stores the key ID and passes the key to the ESXi host. If the host is part of a cluster, vCenter Server sends the KEK to each host in the cluster.
- The key itself is not stored on the vCenter Server system. Only the key ID is known.
- The ESXi host generates internal keys (DEKs) for the virtual machine and its disks. It uses the KEKs to encrypt internal keys and keeps the internal keys in memory only (never on disk). Only encrypted data is stored on disk.
- The ESXi host uses the encrypted internal keys to encrypt the virtual machine.
- Any hosts that can access the encrypted key file and the KEK can perform operations on the encrypted virtual machine or disk.

## Encrypted vSphere vMotion

Encrypted vSphere vMotion provides confidentiality, integrity, and authenticity of the data that is transferred with vSphere vMotion. Starting with vSphere 6.5, vSphere vMotion uses encryption when migrating encrypted virtual machines. You cannot turn off encrypted vSphere vMotion for encrypted virtual machines. For virtual machines that are not encrypted, you can set Encrypted vMotion to one of the following states:

- **Disabled:** Do not use encrypted vSphere vMotion.
- **Opportunistic:** Use encrypted vSphere vMotion if the source and target hosts support it. (This is the default.)
- **Required:** If the source or destination host does not support encrypted vSphere vMotion, do not allow migration with vSphere vMotion.

**Key  
Topic**

The following rules concerning encrypted vMotion across vCenter Server instances apply:

- You must use the vSphere APIs.
- Encrypted vMotion of unencrypted and encrypted virtual machines is supported.
- The source and destination vCenter Server instances must share the KMS cluster that was used to encrypt the virtual machine.
- The name of the shared KMS cluster must be the same on each vCenter Server instance.
- You must have the Cryptographic Operations.Migrate privilege on the virtual machine.
- You must have the Cryptographic Operations.EncryptNew privilege on the destination vCenter Server.
- If the destination ESXi host is not in safe mode, then you also need the Cryptographic Operations.RegisterHost privilege on the destination vCenter Server.
- You cannot change the virtual machine storage policy or perform a key change.

## Note

Only ESXi Versions 6.5 and later use encrypted vSphere vMotion.

When using vTA, the following requirements apply:

- The destination host must be configured with vTA and must be attested.
- Encryption cannot change on migration.
- You can migrate a standard encrypted virtual machine onto a trusted host.
- You cannot migrate a vTA encrypted virtual machine onto a non-trusted host.

## virtual Trusted Platform Module (vTPM)

A *virtual Trusted Platform Module (vTPM)* is a software-based representation of a physical TPM 2.0 chip. A vTPM uses software to perform the same functions that a TPM performs in hardware. A vTPM uses the NVRAM file, which is encrypted using virtual machine encryption, as its secure storage. A hardware TPM includes a preloaded key called the endorsement key (EK), which includes a private key and a public key. For vTPM, the EK is provided either by VMCA or by a third-party CA.

You can add a vTPM to either a new virtual machine or an existing virtual machine, which enables the guest operating system to create and store keys that are private. The keys are not exposed to the guest operating system itself even if the guest operating system is compromised. The keys can be used only by the guest operating system for encryption or signing. With an attached vTPM, a third party can remotely attest to (validate) the identity of the firmware and the guest operating system.

When you configure a vTPM, VM encryption automatically encrypts the virtual machine files but not the disks. The backup of a VM with a vTPM must include all virtual machine data, including the NVRAM file. In order to successfully restore the VM, the backup must include the NVRAM file, and you must ensure that the encryption keys are available.

To use a vTPM, the following requirements must be met:



- Virtual machine hardware Version 14 using EFI firmware
- vCenter Server 6.7 or greater
- Virtual machine encryption (for home files) and KMS
- Windows Server 2016 (64 bit) or Windows 10 (64 bit)

You can add a vTPM as you create a virtual machine by selecting **Customize Hardware > Add New Device > Trusted Platform Module**. Likewise, you can add a vTPM to an existing (powered-down) virtual machine. In the vSphere Client, you can identify which virtual machines are enabled with vTPM by using **Show/Hide Column** for a selected object, such as a host or cluster.

## **virtual Intel Software Guard Extension (vSGX)**

Intel *Software Guard Extension (SGX)* is a processor-specific technology for application developers to protect code and data from disclosure or modification. It allows user-level code to define enclaves, which are private regions of memory. It prevents code running outside the enclave from accessing content in the enclave.

If Intel SGX technology is available on your hardware, your virtual machines can use Virtual Intel SGX (vSGX). To enable vSGX for a virtual machine, the following requirements must be met:

- Virtual machine hardware Version 17 and EFI firmware
- vCenter Server 7.0 and ESXi 7.0
- Linux, Windows Server 2016 (64 bit) or later, or Windows 10 (64 bit) guest OS
- Intel Coffee Lake CPU or later

When vSGX is enabled on a virtual machine, the following features are not supported for that machine:

- vMotion/DRS migration
- Virtual machine suspend and resume
- Memory snapshots (Virtual machine snapshots are supported without snapshotting the memory.)

- Fault Tolerance
- Guest Integrity (GI) (the platform foundation for VMware AppDefense 10)

## Available Add-on Security

You can further secure your environment by procuring and implementing additional features that are not provided natively in vSphere, including the VMware products vRealize Operations Manager, NSX, and AppDefense.

## Compliance Using vRealize Operations Manager

You can implement vRealize Operations Manager (vROps) to provide a single-pane-of-glass monitoring solution for your virtual infrastructure, applications, storage, and network devices. vROps provides an open and extensible platform supported by third-party management packs. It monitors performance and availability metrics, performs predictive analysis of the data, and enables proactive remediation of emerging issues. In addition, you can use vROps to monitor objects in your vSphere environment—such as vCenter Servers, hosts, virtual machines, distributed port groups, and datastores—to ensure compliance with the appropriate standards. You can use vROps to define and analyze compliance standards.

You can customize vROps policies in accordance with the *vSphere Security Configuration Guide*, to enable vSphere alerts for ESXi hosts, vCenter Server, and virtual machines that are in violation of the guide. In addition, hardening guides for regulatory standards are delivered as management packs (PAK files) that you can upload, license, and install. For example, you can install management packs for the following regulatory standards:

- Health Insurance Portability and Accountability Act (HIPAA)
- Payment Card Industry Data Security Standard (PCI DSS)
- CIS security standards
- Defense Information Systems Agency (DISA) security standards
- Federal Information Security Management Act (FISMA) security standards
- International Organization for Standardization (ISO) security standards

vROps collects compliance data from vSphere objects, generates compliance

alerts, and creates reports based on the compliance results.

## VMware NSX

You can implement VMware NSX Data Center for vSphere (NSX) to add a distributed logical firewall, microsegmentation, and additional security measures to your vSphere environment.

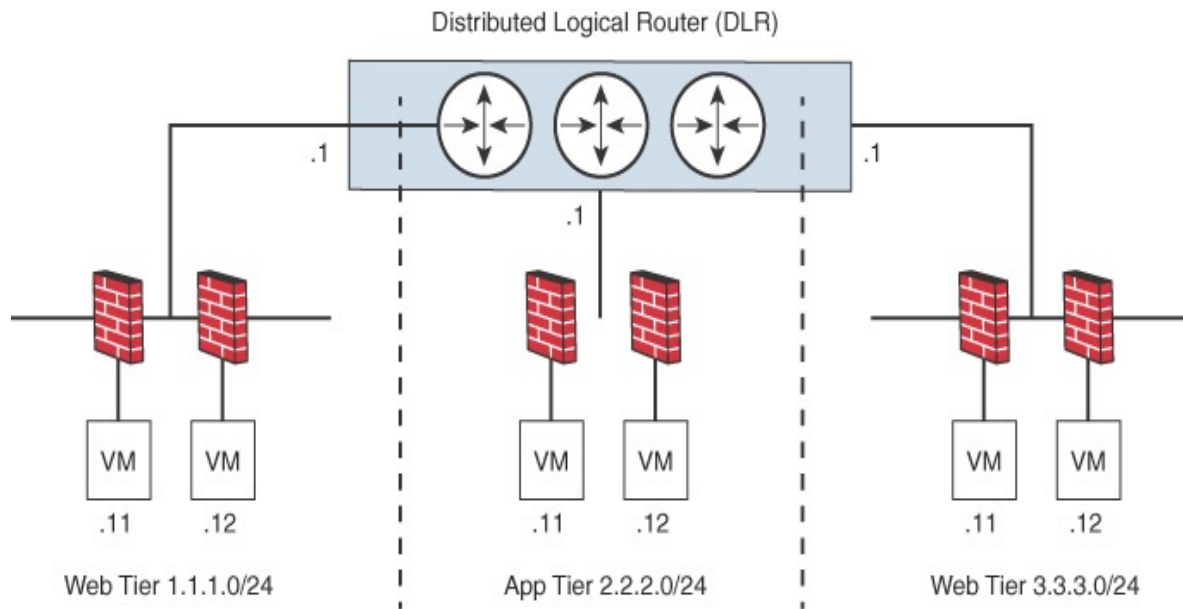
NSX provides Distributed Firewall (DFW), which runs in the VMkernel as a VIB package on all NSX-prepared ESXi hosts. DFW offers near-line-rate performance, virtualization, identity awareness, automated policy creation, advanced service insertion, and other network security features. DFW enhances your physical security by removing unnecessary hairpinning from the physical firewalls and reduces the amount of traffic on the network. It enables **microsegmentation**, which effectively enables you to place a firewall on each VM network connection.

Microsegmentation decreases the level of risk and increases the security posture of a vSphere environment. Microsegmentation involves the following capabilities:

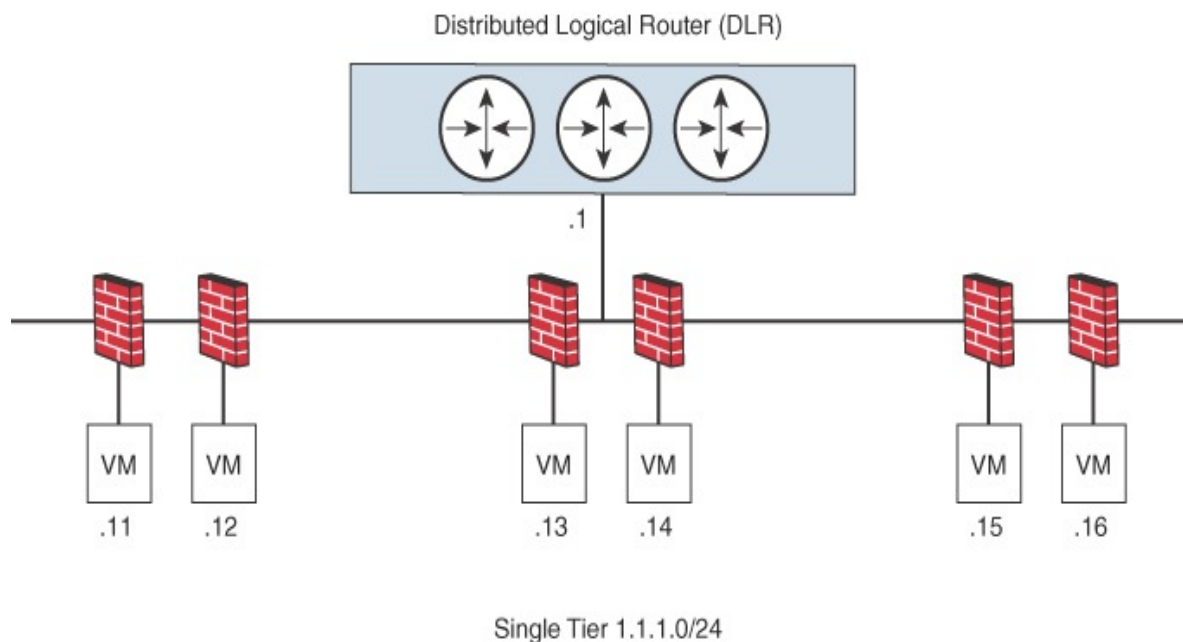
- Distributed stateful firewalling
- Topology-agnostic segmentation
- Centralized policy control
- Granular control
- Network-based isolation

With NSX, isolation can be achieved by leveraging VXLAN technology and virtual networks (that is, logical switches). Isolation can also be achieved by using traditional networking methods, such as implementing ACLs, firewall rules, and routing policies. For example, in a brownfield environment, you could choose to keep existing VLAN segmentation to isolate VMkernel traffic and VM zones while using NSX DFW to implement application segmentation.

With NSX, you can implement virtual machine-to-virtual machine protection, which is commonly referred to as east-west protection, in more than one manner. For example, you could implement multiple Layer 2 segments with Layer 3 isolation (see [Figure 7-6](#)) or implement a single Layer 2 segment and use DFW rules for isolation (see [Figure 7-7](#)).



**FIGURE 7-6** Example with Multiple Layer 2 Segments



**FIGURE 7-7** Example with a Single Layer 2 Segment

NSX provides other security features, such as Service Composer, which you can use to configure security groups and security policies. A security policy is a collection of firewall rules, endpoint services, and network introspection services. Security groups may be populated statically or dynamically based on

containers (such as folders and clusters), security tags, Active Directory groups, and regular expressions. You can map a security policy to a security group.

NSX includes other security features, such as SpoofGuard, the Edge firewall, and a virtual private network (VPN).

## AppDefense

You can secure your vSphere environment further by using VMware [AppDefense](#), which is a data center endpoint security product that protects applications running in vSphere. AppDefense understands an application's intended state and behavior, and it monitors for changes to that intended state that indicate a threat. When a threat is detected, AppDefense automatically responds based on the implemented policies. You can use AppDefense to define “good behavior” and to trigger automated custom actions when other behavior is detected. For vSphere 7.0, AppDefense is available only as a separate product. For vSphere 6.7, AppDefense is included in the vSphere Platinum edition.

Key features of AppDefense include the following:

- It understands the intended state of each application and runs inside the hypervisor, where it has an authoritative understanding of how data center endpoints are meant to behave. This means it is the first to know when changes are made.
- It is hypervisor based, and so it runs in an isolated, protected environment, reducing the likelihood that it will be compromised.
- When a threat is detected, it takes the action that you preconfigure, leveraging vSphere and NSX, such as:
  - Block VM network communication
  - Snapshot a VM
  - Suspend or shut down a VM

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

## Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 7-15](#) lists these key topics and the page number on which each is found.



**Table 7-15** Key Topics for [Chapter 7](#)

Key Topic Element	Description	Page Number
<a href="#">Figure 7-2</a>	vSphere inventory hierarchy	<a href="#">243</a>
List	vCenter Server 7.0 sample roles	<a href="#">246</a>
Paragraph	How permissions are applied by vCenter Server	<a href="#">251</a>
List	Built-in security features	<a href="#">254</a>
Paragraph	vSphere Trust Authority (vTA)	<a href="#">258</a>
List	General networking security recommendations	<a href="#">263</a>
List	Virtual machine hardening best practices	<a href="#">265</a>
Paragraph	Configure UEFI boot	<a href="#">266</a>
List	Rules for encrypted vMotion across vCenter Server instances	<a href="#">272</a>

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

Trusted Platform Module (TPM)  
VMware Directory Service (vmdir)  
VMware Certificate Authority (VMCA)  
VMware Endpoint Certificate Store (VECS)  
Virtual Trusted Platform Module (vTPM)  
Intel Software Guard Extension (SGX)  
microsegmentation  
AppDefense

## Review Questions

1. You are preparing to implement certificates in your vSphere environment. Which of the following does VCMA support in custom certificates when it is used as a subordinate CA?
  - a. CRL distribution points
  - b. Authority information access
  - c. CRT format
  - d. Certificate template information
2. On which of the following items can you set permissions in vCenter Server?
  - a. Licenses
  - b. Datastores
  - c. Roles
  - d. Sessions
3. You are examining the default security profile in your vSphere environment. Which of the following services are stopped by default?
  - a. DCUI
  - b. Load-Based Teaming Daemon
  - c. CIM Server
  - d. SNMP Server
4. You are hardening a vCenter Server and see that it contains some

expired certificates. What is the main purpose of removing expired and revoked certificates from vCenter Server?

- a.** To avoid DoS attacks
  - b.** To avoid MITM attacks
  - c.** To avoid automatic virtual machine shutdown due to expired certificates
  - d.** To avoid ARP spoofing
- 5.** You want to enable UEFI boot for your virtual machines. Which of the following is a requirement?
- a.** Virtual hardware Version 11 or later
  - b.** VMware Tools Version 11 or later
  - c.** Virtual hardware Version 12 or later
  - d.** VMware Tools Version 10.1 or later



# **Part II**

# **vSphere Installation/Configuration**

## Chapter 8

# vSphere Installation

This chapter covers the following topics:

- [Installing ESXi hosts](#)
- [Deploying vCenter Server Components](#)
- [Configuring Single Sign-On \(SSO\)](#)
- [Initial vSphere Configuration](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.1, 1.2, 1.8, 4.1.1, 4.1.2, 4.3, 4.3.1, 4.3.2, 4.3.3, 4.4, 4.6, 4.12, 4.15, 4.16, 7.9, 7.10, 7.11, 7.11.1, and 7.11.2.

This chapter discusses the procedures involved in installing and configuring a vSphere 7.0 environment.

### “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 8-1](#) 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 and Review Questions.](#)”

**Table 8-1** “Do I Know This Already?” Foundation Topics Section-to-Question Mapping

Foundations Topics Section	Questions Covered in This Section
Installing ESXi hosts	1, 2

Deploying vCenter Server Components 3, 4

---

Configuring Single Sign-On (SSO) 5–7

---

Initial vSphere Configuration 8–10

---

1. You are preparing to deploy vSphere 7.0. Which of the following is a prerequisite for installing ESXi interactively?
  - a. Download the ESXi installer ISO.
  - b. Download the ESXi installer OVF.
  - c. Download the GUI installer for Windows or Mac.
  - d. Download the ESXi MSI.
2. You are preparing to do a scripted installation of ESXi 7.0. Where is the default installation script located?
  - a. `/etc/vmware/weasel/ks.py`
  - b. `/etc/vmware/weasel/ks.cfg`
  - c. `/etc/vmware//ks.cfg`
  - d. `/etc/vmware/ks.py`
3. You are preparing to install vCenter Server 7.0 using a deployment command. To perform a pre-deployment check, which command should you use?
  - a. **`vcsa-deploy-precheck path-to-JSON-file`**
  - b. **`vcsa-deploy install --precheck path-to-JSON-file`**
  - c. **`vcsa-deploy install --verify-only path-to-JSON-file`**
  - d. **`vcsa-deploy precheck path-to-JSON-file`**
4. You are installing vSphere 7.0 and want to document the location of certificates. Where are the ESXi certificates stored?
  - a. Locally on the ESXi hosts
  - b. In the VECS
  - c. In the vCenter Server database
  - d. In the VMCA

5. You are adding an OpenLDAP authentication source for your recently deployed vCenter Server. Which of the following is not a requirement?
  - a. All users must have the object of class inetOrgPerson.
  - b. All groups must have the object of class groupOfUniqueNames.
  - c. All groups must have the group membership attribute uniqueMember.
  - d. All users must be members of the OpenLDAP group.
6. You are deploying a new vSphere environment and need to control which users can manage certificates. Which vCenter Server Single Sign-On domain group membership should you manipulate?
  - a. DCAdmins
  - b. SolutionUsers
  - c. CAAdmins
  - d. SystemConfiguration\_Administrators
7. You are adding an Active Directory (integrated Windows authentication) identity source and want to ensure that future machine name changes do not cause issues. Which of the following allows you to do this?
  - a. Select Use Service Principal Name (SPN) and provide a UPN.
  - b. Select Use Machine Account and provide a UPN.
  - c. Select Use Service Principal Name (SPN) and provide a base DN for users.
  - d. Select Use Machine Account and provide a base DN for users.
8. You are deploying vCenter Server in a secured network that has no Internet access. What do you need to install in order to download updates?
  - a. Update Manager Download Service
  - b. Update Manager Proxy Service
  - c. Lifecycle Manager Download Service
  - d. Lifecycle Manager Proxy Service

9. You are implementing vCenter HA. How will you connect the nodes to the vCenter HA network?
- a. Connect NIC 1 on the Active and Passive nodes and NIC 0 to the vCenter HA network. Do not connect the Witness node to the vCenter HA network.
  - b. Connect NIC 1 on the Active and Passive nodes and NIC 0 on the Witness node to the vCenter HA network.
  - c. Connect NIC 0 on the Active, Passive, and Witness nodes to the vCenter HA network.
  - d. Connect NIC 1 on the Active, Passive, and Witness nodes to the vCenter HA network.
10. You are installing new ESXi hosts and want to configure boot options. Which of the following **kernelopt** options has been deprecated in ESXi 7.0?
- a. **autoCreateDumpFile**
  - b. **autoPartitionCreateUSBCoreDumpPartition**
  - c. **skipPartitioningSsds**
  - d. **autoPartitionOnlyOnceAndSkipSsd**

## Foundation Topics

### Installing ESXi Hosts

To begin a vSphere deployment, you should install and configure at least one ESXi host using the information in this section. Optionally, you can apply the information here to install and configure additional ESXi hosts. In many cases, administrators choose to deploy the first ESXi host and then deploy vCenter Server and use vCenter Server along with other tools, such as host profiles, to facilitate the deployment and configuration of the remaining ESXi hosts.

You have several choices for installing ESXi, such as using the interactive wizard, using scripts, and using Auto Deploy. These choices are covered in this section. Using host profiles to configure ESXi hosts after installation is discussed later in this chapter.

## Installing ESXi Interactively

You can use the following procedure—which is very useful in small environments with fewer than five ESXi hosts—to install ESXi interactively:



- Step 1.** Verify that all the target machine hardware is supported and meets minimum requirements.
- Step 2.** Gather and record the information that will be required during the installation (see [Table 8-2](#)).

**Table 8-2** Information Required for ESXi Installation

Information	Required or Optional	Details
Keyboard layout	Required	Default: US English
VLAN ID	Optional	Range: 0–4094  Default: None
IP address	Optional	Default: DHCP
Subnet mask	Optional	Default: Based on the configured IP address
Gateway	Optional	Default: Based on the configured IP address and subnet mask
Primary DNS	Optional	Default: Based on the configured IP address and subnet mask
Secondary DNS	Optional	Default: None
Host name	Required for static IP settings	Default: None
Install location	Required	At least 5 GB if you install on a single disk  Default: None

Migrate existing ESXi settings; preserve VMFS datastore	Required if you are installing ESXi on a drive with an existing ESXi installation	Default: None
Root password	Required	Must contain at least 8 to 40 characters and meet other requirements
		Default: None

- Step 3.** Verify that the server hardware clock is set to **UTC**. This setting is in the system BIOS.
- Step 4.** Download the ESXi installer ISO and prepare the hardware system to boot from it.
- Step 5.** Start the machine so that it boots from the ESXi installer.
- Step 6.** On the Select a Disk page, select the drive on which to install ESXi and press **Enter**.
- Step 7.** Select the keyboard type and language for the host.
- Step 8.** Enter a password to be used by the root account.
- Step 9.** When prompted to do so, remove the bootable media and press **Enter** to reboot the host.

The default behavior is to configure the ESXi management network using DHCP. You can override the default behavior and use static IP settings for the management network after the installation is completed. If your host is not yet assigned an IP address or if you want to change it, you can use the following procedure to select the appropriate network adapter, configure the VLAN, and configure the IP configuration for the host's management network interface:

- Step 1.** Log on to the Direct Console User Interface (DCUI), which appears on the host's monitor.
- Step 2.** If needed, use the DCUI to change the network adapter used for management:
  - a.** Select **Configure Management Network** and press **Enter**.

- b. Select **Network Adapters** and press **Enter**.
- c. Select a network adapter and press **Enter**.

**Step 3.** If needed, use the DCUI to change the VLAN used for management:

- a. Select **Configure Management Network** and press **Enter**.
- b. Select **VLAN** and press **Enter**.
- c. Enter the appropriate VLAN ID for the network connection.

**Step 4.** If needed, use the DCUI to change the IP configuration used for management:

- a. Select **Configure Management Network** and press **Enter**.
- b. Select **IPv4 Configuration** and press **Enter**.
- c. Select **Set Static IPv4 Address and Network Configuration**.
- d. Enter the IP address, subnet mask, and default gateway and press **Enter**.

## Note

It is important for the IP address of the management network to remain consistent. If you do not select a static IP address, make sure that you create a DHCP reservation for the MAC address of your first physical NIC, which ESXi calls VMNIC0.

You can use the DCUI to configure DNS by following this procedure:

**Step 1.** Select **Configure Management Network** and press **Enter**.

**Step 2.** Select **DNS Configuration** and press **Enter**.

**Step 3.** Select **Use the Following DNS server Addresses and Hostname**.

**Step 4.** Enter the primary server, an alternative server (optional), and the host name.

After ESXi is installed and the management network is configured, you can manage the host and make other configuration changes by using the vSphere Host Client.



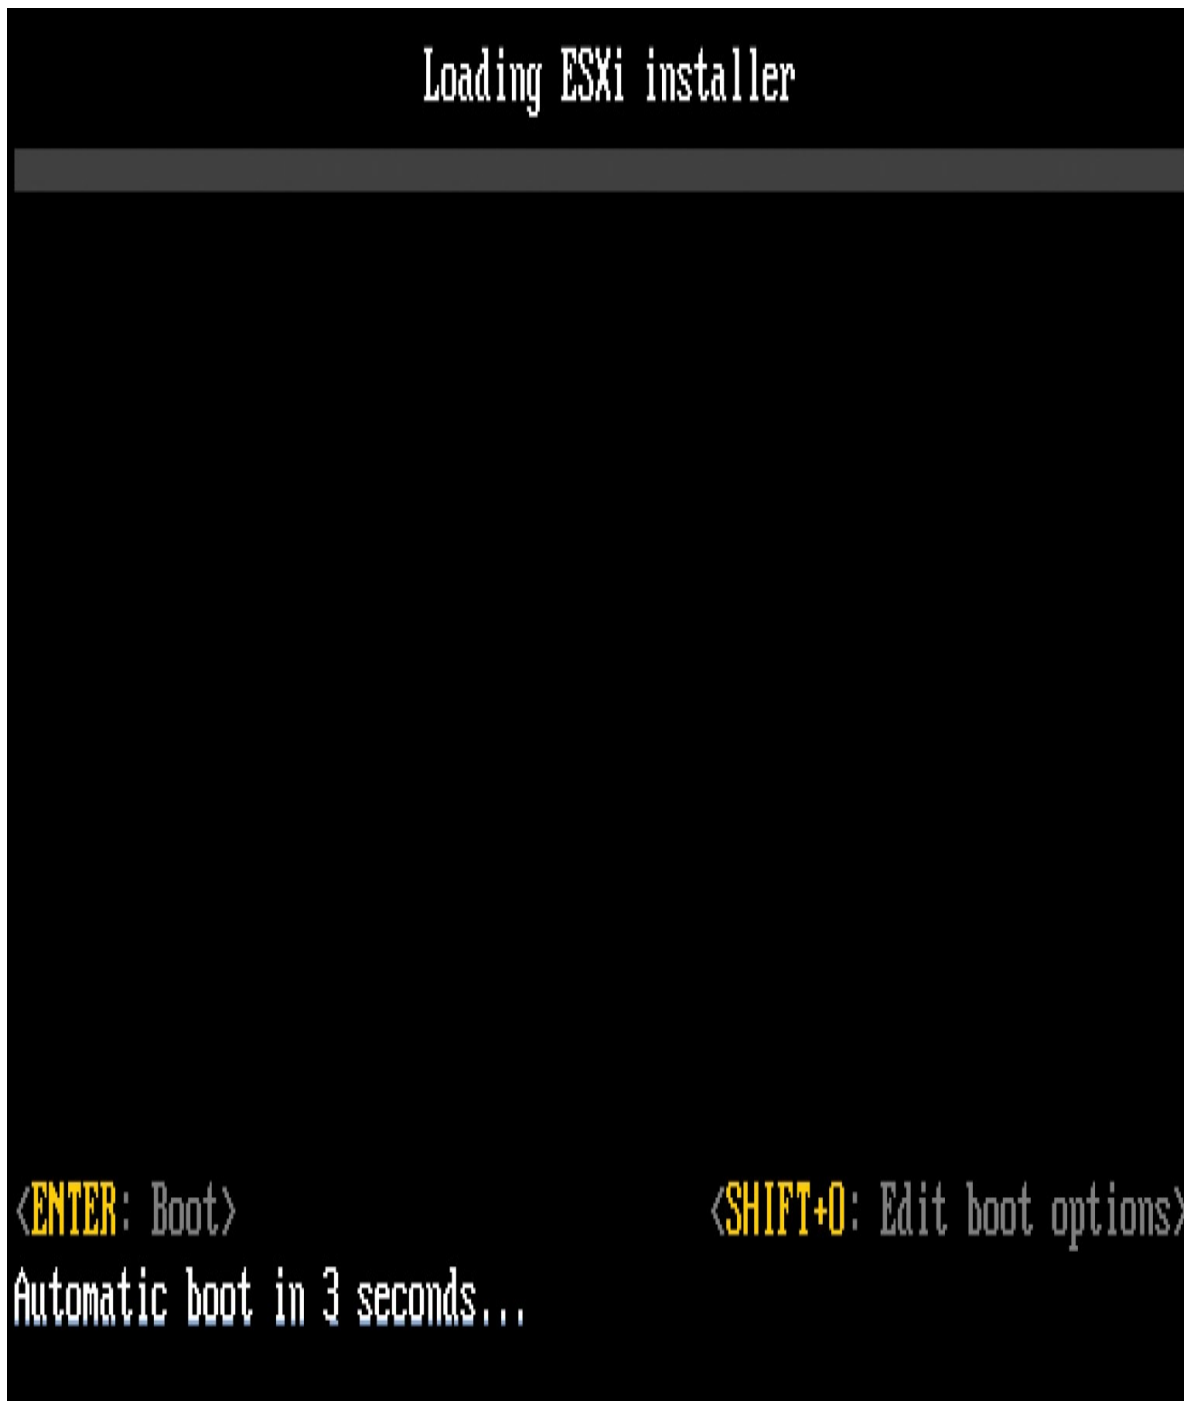
## Scripted ESXi Installation

Installation scripts provide an efficient way to deploy multiple hosts and/or to deploy hosts remotely.

You can use an installation script that includes the settings for installing ESXi. The script can be applied to all of the hosts that need to have the same configuration. Only supported commands can be used in the installation script. This can be modified for settings that need to be unique for each host. The installation script can be in any of the following locations:

- FTP server
- HTTP/HTTPS server
- NFS server
- USB flash drive
- CD-ROM drive

To start the installation script, you can enter boot options at the ESXi installer boot command line. At boot time, you can press Shift+O in the boot loader (see [Figure 8-1](#)) to specify boot options and access the kickstart file. If you are installing using PXE boot, options can be passed through the **kernelopts** line of the boot.cfg file. The location of the installation script is defined by setting the **ks=filepath** option, where *filepath* is the location of the kickstart file. If **ks=filepath** is not included in the script, the text installer runs.



**FIGURE 8-1** ESXi Installer

Follow these steps to start the installation script:

**Step 1.** Start the host.

**Step 2.** When the Loading ESXi installer window appears (see [Figure 8-1](#)),

press **Shift+O** to define the options.

**Step 3.** At the **runweasel** command prompt, enter **ks=** along with the path to the installation script and the command-line options. For example, you could enter the following options to boot the host from a script named **ks-script-01** residing on the server **192.168.1010.10** and set the host's IP address to **192.168.100.101**:

[Click here to view code image](#)

```
ks=http://192.168.100.10/kickstart/ks-script-01.cfg
nameserver=192.168.1.100 ip=192.168.100.101
netmask=255.255.255.0 gateway=192.168.100.101
```

To successfully perform a scripted installation, you may need to enter boot options to access the script file. [Table 8-3](#) shows some of these options.

**Table 8-3** Boot Options for ESXi Scripted Installation

<b>Boot Option</b>	<b>Description</b>
<b>BOOTIF=</b> <i>hwtype-</i> <i>MAC</i> <i>address</i>	Is similar to the <b>netdevice</b> option except in the PXELINUX format.
<b>gateway=</b> <i>ip</i> <i>address</i>	Serves as a gateway for downloading the installation script.
<b>ip=</b> <i>ip</i> <i>address</i>	Indicates the IP address used for downloading the installation script.
<b>ks=</b> <i>cdrom</i> <i>:/path</i>	Performs a scripted installation with the script at <i>path</i> , which resides on the CD in the CD-ROM drive.
<b>ks=</b> <i>file://</i> <i>path</i>	Performs a scripted installation with the script at <i>path</i> .
<b>ks=</b> <i>proto</i> <i>col://serv</i> <i>erpath</i>	Performs a scripted installation with a script located on the network at the given URL. <i>protocol</i> can be <b>http</b> , <b>https</b> , <b>ftp</b> , or <b>nfs</b> (for example, <b>ks=nfs://host/porturl-path</b> ).
<b>ks=</b> <i>usb</i>	Performs a scripted installation, accessing the script from an

attached USB drive.

---

<b>ks=usb:// path</b>	Performs a scripted installation with the script file at the specified path, which resides on a USB drive.
---------------------------	--

---

<b>ksdevice= device</b>	Tries to use a network adapter device when looking for an installation script and installation media. For <i>device</i> , specify a MAC address, such as, 00:50:56:C0:00:01, or a vmnic## name.
-----------------------------	---

or

**netdevice=  
device**

---

<b>nameser ver=ip address</b>	Specifies the domain name server to be used for downloading the installation script and installation media.
---------------------------------------	---

---

<b>netmask= subnet mask</b>	Specifies the subnet mask used for downloading the installation script.
-------------------------------------	---

---

<b>vlanid=vl anid</b>	Specifies the VLAN used for downloading the installation script.
---------------------------	--

---

There is a default installation script included with the ESXi installer, and it can be used to install ESXi onto the first disk that is detected. The default ks.cfg installation script is in the initial RAM disk at /etc/vmware/weasel/ks.cfg. The location of the default ks.cfg file can be defined with the

**ks=file://etc/vmware/weasel/ks.cfg** boot option. When using the ks.cfg script for the installation, the default root password is [myp@ssw0rd](#). The installation script on the installation media can't be modified. After the ESXi host has been installed, the vSphere Host Client or the vSphere Web Client logged in to the vCenter Server that manages the ESXi host can be used to change any of the default settings.

[Example 8-1](#) shows the contents of the default script.

### Example 8-1 Default Installation Script

[Click here to view code image](#)

```
#  
# Sample scripted installation file
```

```

#
# Accept the VMware End User License Agreement
vmaccepteula
# Set the root password for the DCUI and Tech Support Mode
rootpw myp@ssw0rd
# Install on the first local disk available on machine
install --firstdisk --overwritevmfs
# Set the network to DHCP on the first network adapter
network --bootproto=dhcp --device=vmnic0
# A sample post-install script
%post --interpreter=python --ignorefailure=true
import time
stampFile = open('/finished.stamp',

```

You can see that this default script sets the root password to [myp@ssw0rd](#), installs on the first disk, overwrites any existing VMFS datastore, and sets the network interface to use DHCP. When creating your own script, you can specify many options, a few of which are shown in [Table 8-4](#).

**Table 8-4** Sample Options for ESXi Installation Script

Comma nd	Option	Description
<b>clear</b>	--	Removes partitions on all drives except those specified.
<b>part</b> (option al)	<b>igno</b> <b>redr</b> <b>ives</b> =	
	--	Allows overwriting of VMFS partitions on the specified drives.
	<b>over</b> <b>writ</b> <b>evm</b> <b>fs</b>	
<b>dryru</b> <b>n</b>	N/A	Parses and checks the installation script but does not perform the installation.
<b>install</b>	-- <b>disk</b> =	Specifies the disk to partition. Acceptable values can use various forms, as in these examples:

Path:

**--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0**

MPX name:

**--disk=mpx.vmhba1:C0:T0:L0**

VML name:

**--disk=vml.000000034211234**

vmkLUN UID:

**--disk=vmkLUN\_UID**

---

**--  
ignores  
disk** Excludes solid-state disks from eligibility for partitioning.

---

**--  
overwrite  
evsa  
n** Used to install ESXi on a disk—either SSD or HDD (magnetic)—that is already in a vSAN disk group.

---

**network  
boot  
proto=  
o=** Specifies whether the IP address should be set statically or via DHCP.

---

**--ip=** Sets an IP address for the machine to be installed, in the form *xxx.xxx.xxx.xxx*. Required with the **bootproto=static** option and ignored otherwise.

---

**--** Designates the primary name server as an IP address. Used

**nam** with the **bootproto=static** option. You can omit this option  
**eser** if you do not intend to use DNS.  
**ver**

---

## Using Auto Deploy

vSphere Auto Deploy makes it possible to install ESXi on hundreds of physical hosts. Large environments can be managed efficiently by experienced administrators utilizing Auto Deploy. Hosts use network booting to boot from a central Auto Deploy server. Hosts can be configured with a host profile created from a reference host, if desired. This host profile could be created to prompt for input. After the hosts boot and are configured, they are then managed by vCenter Server, like other ESXi hosts. Auto Deploy can also be configured for stateless caching or stateful installations.

With *stateless caching*, which is the default, Auto Deploy does not store ESXi configuration or state data within the host. Instead, Auto Deploy uses image profiles and host profiles to maintain the host configuration. During subsequent boots, the host must connect to the Auto Deploy server to retrieve its configuration. If a network boot fails, the ESXi host can use a local cache to boot from the last known ESXi image.

With stateful installations, Auto Deploy is used to boot the host, but the installation and configuration are written to a local disk. On subsequent boots, the host boots from the local disk where this host configuration is stored.

Auto Deploy can be configured and managed using a graphical user interface (GUI) in vSphere 6.5 and later. The PowerCLI method is still available, but the GUI provides an easier-to-use option. For the Auto Deploy GUI to be visible in the vSphere Web Client, both the Image Builder and Auto Deploy services must be running when you're logging in to vCenter Server. The Image Builder feature in the GUI enables you to download ESXi images from the VMware public repository or to upload ZIP files containing ESXi images or drivers. You can customize the images by adding or removing components and optionally export images to ISO or ZIP files for use elsewhere. You can compare two images to see how their contents differ.

You can use the Deployed Hosts tab to view hosts that are provisioned with Auto Deploy and to perform tests and remediations.

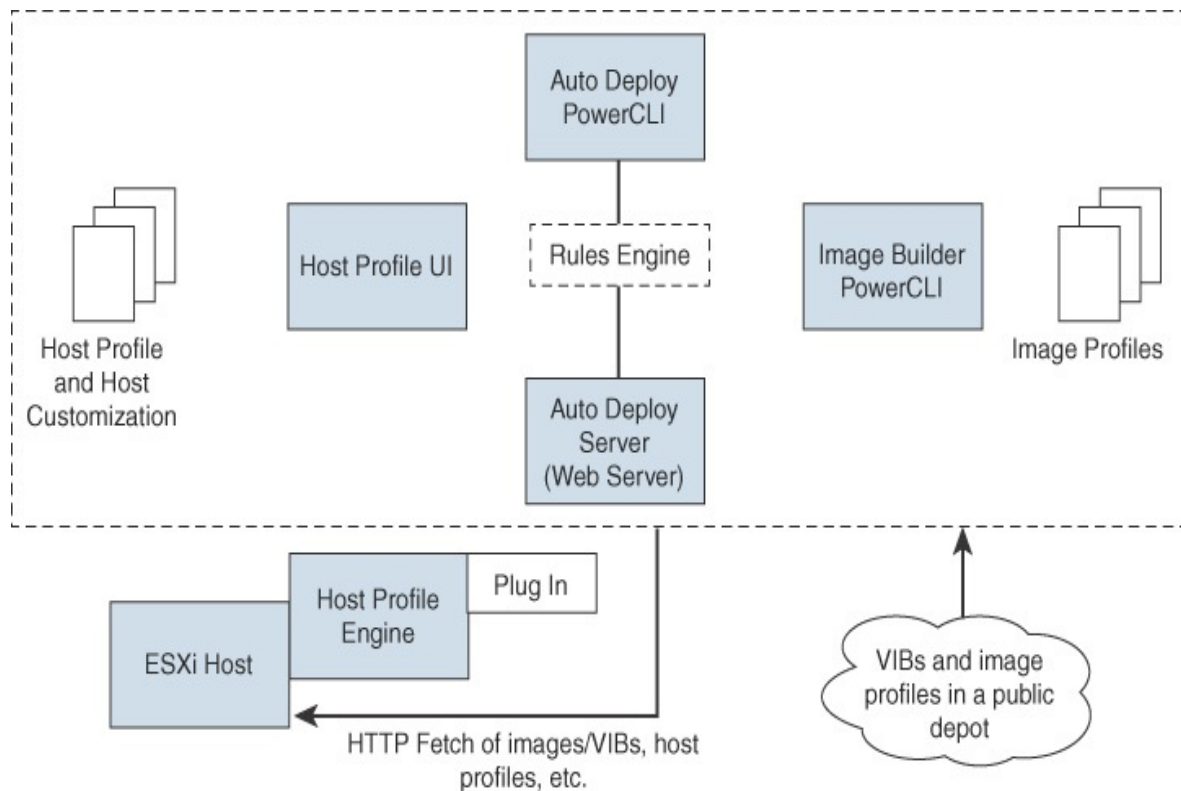
The architecture for Auto Deploy includes many components, as described in [Table 8-5](#) and illustrated in [Figure 8-2](#).

**Table 8-5** Auto Deploy Components

<b>Component</b>	<b>Description/Purpose</b>
Auto Deploy server	Uses a rules engine, a set of images, a set of host profiles, and required infrastructure to manage ESXi deployments.
Rules engine	Assigns image profiles and host profiles to each host.
Host profile	Defines host-specific configurations, such as networking, NTP, and host permissions. You can use host customization in conjunction with host profiles to provide details that are unique to each host, such as IP address.
Auto Deploy PowerCLI	Serves as a command-line engine for driving Auto Deploy.
Image Builder PowerCLI	Serves as a command-line engine for building images.
vCenter Server	Manages the vSphere inventory and provides host profiles.
DHCP server	Provides IP configuration to the host and redirects the host to the PXE server.



PXE server	Boots the host and directs it to the TFTP server.
TFTP server	Provides the appropriate boot image.
Software depot	Holds a collection of VIBs either online (accessible via HTTP) or offline (accessible via a USB drive or CD/DVD).
Image profile	Holds a collection of VIBs used to install the ESXi server and saved as ZIP files or ISO images. You can obtain image profiles from VMware and VMware partners, and you can create custom image profiles by using ESXi Image Builder.
vSphere Installation Bundle (VIB)	<p>Packages a collection of files (such as drivers) into an archive similar to a ZIP file. Each VIB is released with an acceptance level that cannot be changed. The host acceptance level assigned to each host determines which VIBs can be installed to the host. These are the acceptance levels, from highest to lowest:</p> <ul style="list-style-type: none"> <li>■ VMwareCertified</li> <li>■ VMwareAccepted</li> <li>■ Partner Supported</li> <li>■ CommunitySupported</li> </ul>



**FIGURE 8-2** Auto Deploy Architecture

You control the behavior of the vSphere Auto Deploy server by using rules. The rules engine checks the rule set for matching host patterns to decide which items (image profile, host profile, vCenter Server location, or script object) to use to provision each host. Rules can assign image profiles and host profiles to a set of hosts. A rule can identify target hosts by boot MAC address, Basic Input/Output System (BIOS), universally unique identifier (UUID), System Management BIOS (SMBIOS) information, vendor, model, or fixed DHCP IP address. You can create rules by using the vSphere Web Client or vSphere Auto Deploy cmdlets with the PowerCLI. For example, to create a new deployment rule named Rule-01 that places all hosts in a folder named Auto-deployed Hosts, you can use the following PowerCLI command:

[Click here to view code image](#)

```
New-DeployRule -Name Rule-01 -Item "Auto-deployed Hosts" -allhosts
```

To modify the rule so that it applies only to a set of hosts in a specific IP range, you can use the Set-DeployRule cmdletL

[Click here to view code image](#)

```
Set-DeployRule -DeployRule Rule-01 -Pattern  
"ipv4=192.168.100.101-192.168.100.112"
```

Table 8-6 describes some of the common Auto Deploy PowerCLI cmdlets.

**Table 8-6** Common Auto Deploy PowerCLI cmdlets

cmdlet	Description
<b>Get-DeployCommand</b>	Returns a list of Auto Deploy cmdlets.
<b>New-DeployRule</b>	Creates a new rule with the specified items and patterns.
<b>Set-DeployRule</b>	Updates an existing rule with the specified items and patterns. You cannot update a rule that is part of a rule set.
<b>Get-DeployRule</b>	Retrieves the rules with the specified names.
<b>Copy-DeployRule</b>	Clones and updates an existing rule.
<b>Add-DeployRule</b>	Adds one or more rules to the working rule set and, by default, also to the active rule set. Use the <b>NoActivate</b> parameter to add a rule only to the working rule set.
<b>Remove-DeployRule</b>	Removes one or more rules from the working rule set and from the active rule set. Run this command with the <b>-Delete</b> parameter to completely delete the rule.
<b>Set-</b>	Explicitly sets the list of rules in the working rule set.

## DeployRuleSet

---

**Get-DeployRuleSet** Retrieves the current working rule set or the current active rule set.

---

**Switch-ActiveDeployRuleSet** Activates a rule set so that any new requests are evaluated through the rule set.

---

**Get-VMHostMatchingRules** Retrieves rules that match a pattern. For example, you can retrieve all rules that apply to a host or hosts. Use this cmdlet primarily for debugging.

---

**Test-DeployRuleSetCompliance** Checks whether the items associated with a specified host are in compliance with the active rule set.

---

**Repair-DeployRuleSetCompliance** Given the output of **Test-DeployRuleSetCompliance**, this cmdlet updates the image profile, host profile, and location for each host in the vCenter Server inventory. The cmdlet might apply image profiles, apply host profiles, or move hosts to the prespecified folders or clusters on the vCenter Server system.

---

**Apply-ESXiImage** Associates the specified image profile with the specified host.

---

## Profile

---

**Get-VMHostImageProfile** Retrieves the image profile in use by a specified host. This cmdlet differs from the **Get-EsxImageProfile** cmdlet in the Image Builder PowerCLI.

---

**Repair-DeployImageCache** Deploys a new image cache. Use this cmdlet only if the Auto Deploy image cache is accidentally deleted.

---

**Get-VMHostAttributes** Retrieves the attributes for a host that are used when the Auto Deploy server evaluates the rules.

---

**Get-DeployMachineIdentity** Returns a string value that Auto Deploy uses to logically link an ESXi Host in vCenter to a physical machine.

---

**Set-DeployMachineIdentity** Logically links a host object in the vCenter Server database to a physical machine. Use this cmdlet to add hosts without specifying rules.

---

**Get-DeployOption** Retrieves the Auto Deploy global configuration options. This cmdlet currently supports the **vlan-id** option, which specifies the default VLAN ID for the ESXi Management Network of a host provisioned with Auto Deploy. Auto Deploy uses the value only if the host boots without a host profile.

---

**Set-  
Deplo  
yOpti  
on** Sets the value of a global configuration option. Currently supports the **vlan-id** option for setting the default VLAN ID for the ESXi Management Network.

---

The first time a host boots using Auto Deploy, the following sequence of events occurs:

**Key  
Topic**

1. The host starts a PXE boot sequence. The DHCP server assigns an IP address and redirects the host to the TFTP server.
2. The host downloads and executes the iPXE file (configured by the DHCP server) and applies the associated configuration file.
3. The host makes an HTTP boot request to the vSphere Auto Deploy server. The HTTP request includes hardware and network information.
4. The vSphere Auto Deploy server queries the rules engine and streams data (the ESXi image) from the image profile and the host profile.
5. The host boots using the image profile. If the vSphere Auto Deploy server provided a host profile, the host profile is applied to the host.
6. vSphere Auto Deploy adds the host to the proper inventory location and cluster in the vCenter Server system.
7. If the host is part of a DRS cluster, virtual machines from other hosts might be migrated to the host.

## Note

If a host profile requires a user to specify certain information, such as a static IP address, the host is placed in maintenance mode when the host is added to the vCenter Server system. You must reapply the host profile and update the host customization to have the host exit maintenance mode.

## Deploying vCenter Server Components

VMware vCenter 7.0 is only available as a prebuilt Linux-based virtual appliance. It is no longer available to be installed on Windows systems. This

section provides information for deploying vCenter Server components.

## **vCenter Server Database**

As of vCenter 7.0, only one database can be used: an included version of the VMware-specific distribution of PostgreSQL.

### **Note**

The vCenter installation program allows migration from a Windows-based vCenter Server to the vCenter Server Appliance, including migration from Oracle or Microsoft SQL to the embedded PostgreSQL database.

## **Platform Services Controller (PSC)**

In vSphere 6.x, the Platform Services Controller (PSC) component could be deployed externally to vCenter Server. Beginning with vSphere 7.0, the services that the PSC provided in vSphere 6.x, such as SSO, VMCA and License Service, are consolidated into vCenter Server. These services cannot be deployed as a separate virtual appliance. Because the services are now part of vCenter, they are no longer listed as part of the PSC. For example, in vSphere 7.0, the *vSphere Authentication* publication replaces the *Platform Services Controller Administration* publication.

## **vCenter Server Appliance**

As of vSphere 7.0, the vCenter Server is only available as a prebuilt virtual appliance that consists of the following:

- Photon OS 3.0
- vSphere Authentication Services
- PostgreSQL
- The VMware vSphere Lifecycle Manager extension for the vSphere Client
- VMware vCenter Lifecycle Manager

vCenter Version 7.0 deploys with VM hardware Version 10, supporting up to 64 virtual CPUs per VM. The deployment wizard helps you to determine the

resource allocation for the virtual appliance, based on the size of your environment.

To prepare for a vCenter Server Appliance deployment, you should download the vCenter Server Appliance installer ISO file and mount it to a virtual machine or physical machine from which you want to perform the deployment. To use the vCenter Server Appliance GUI (or CLI) installer, you can use a machine that is running a supported version of a Windows, Linux, or Mac operating system, as shown in [Table 8-7](#).

**Table 8-7** Requirements for GUI/CLI Installers

Operating System	Supported Versions	Minimum Hardware Configuration
Windows	<ul style="list-style-type: none"> <li>■ Windows 8, 8.1, 10</li> <li>■ Windows 2012 x86-64 bit</li> <li>■ Windows 2012 R2 x86-64 bit</li> <li>■ Windows 2016 x86-64 bit</li> </ul>	4 GB RAM, 2 CPUs having 4 cores with 2.3 GHz, 32 GB hard disk, 1 NIC
Linux	<ul style="list-style-type: none"> <li>■ SUSE 15</li> <li>■ Ubuntu 16.04 and 18.04</li> </ul>	4 GB RAM, 1 CPU having 2 cores with 2.3 GHz, 16 GB hard disk, 1 NIC  With the CLI installer, 64-bit OS
Mac		



- macOS 10.13, 10.14, 10.15
- macOS High Sierra, Mojave, Catalina

---

You can use the GUI or CLI installers to do the following:

- Deploy or upgrade vCenter Server Appliance.
- Converge prior vCenter installations with external PSCs to the current vCenter Server version.
- Migrate a Windows-based vCenter Server to the Linux-based vCenter Server Appliance.
- Restore a vCenter Server Appliance from a file-based backup.

## Deploying VCSA Using the GUI Installer

You can use the GUI installer to deploy a vCenter Server Appliance. To perform a GUI-based deployment, you download the vCenter Server Appliance installer on a network client machine, run the deployment wizard from the client machine, and provide the required information.

### Note

vCenter 7 incorporates all services that were part of the PSC to a single vCenter VM.

Using the GUI installer involves two stages. In the first stage, you navigate through the installation wizard, choose the deployment type, provide the appliance settings, and deploy the OVA. Alternatively, you could use the vSphere Web Client or the vSphere Host Client to deploy the OVA.

In the second stage of using the GUI installer, you use a wizard to configure the appliance time synchronization, configure vCenter Single Sign-On (SSO), and

start the services in the newly deployed appliance. Alternatively, you can use a web browser to access the appliance's vCenter Server Appliance Management Interface (VAMI) at <https://FQDN:5480>. If you use the alternative approach in the first stage, then you must use it in the second stage. To use the GUI installer to deploy the VCSA, you can follow these steps:



**Step 1.** In the vCenter Server Appliance installer, navigate to the appropriate subdirectory in the **vcsa-ui-installer** directory and run the installer executable file:

For Windows OS, use the **win32** subdirectory and the **installer.exe** file.

For Linux OS, use the **lin64** subdirectory and the **installer** file.

For Mac OS, use the **mac** subdirectory and the **Installer.app** file.

**Step 2.** On the Home page, click **Install**.

**Step 3.** On the next page, click **Next**.

**Step 4.** Read and accept the license agreement and click **Next**.

**Step 5.** Connect to the target server where you want to deploy the appliance and make one of two choices:

- Provide the FQDN (or IP address) and credentials for the target ESXi host and provide the appropriate certificate.
- Provide the FQDN (or IP address) and credentials for the target vCenter Server (that is managing the hosts on which this new vCenter Server will be deployed), provide the appropriate certificate, and specify the appropriate location in the vSphere inventory.

**Step 6.** On the next page, set the appliance's name and root user password, following these rules:

- The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.
- The password must contain only lowercase ASCII characters without spaces, must be at least eight characters, and must contain a number, uppercase and lowercase letters, and a special character, such as an

exclamation point (!), hash symbol (#), at sign (@), or parentheses (()).

- Step 7.** Select the deployment size: Tiny, Small, Medium, Large, or X-Large (refer to [Chapter 1, “vSphere Overview, Components, and Requirements”](#)).
- Step 8.** Select the storage size for the appliance (as explained in [Chapter 1](#)).
- Step 9.** Select an available datastore and select the disk provisioning type.
- Step 10.** On the Configure Network settings page, fill in the network settings, such as virtual switch port group, IP configuration, and communication ports.
- Step 11.** On the Ready to Complete page, click **Finish**.
- Step 12.** Wait for the OVA to deploy and then click **Continue** to proceed with Stage 2.
- Step 13.** On the Introduction page, click **Next**.
- Step 14.** Choose a time configuration option.:
- Synchronize time with the ESXi host
  - Synchronize time with NTP servers
- Step 15.** Optionally, enable SSH connections into the appliance.
- Step 16.** Create a new SSO domain or join an existing domain:
- **Create a new SSO domain:** Enter the domain (such as vsphere.local), set the SSO administrator account ([administrator@vsphere.local](#) by default) password, provide a SSO site name, and confirm the password.
  - **Join an existing SSO domain:** Enter the PSC FQDN containing the SSO server, provide the HTTPS port that PSC will use, provide the target SSO domain name (such as vsphere.local), and enter the SSO administrator account password.
- Step 17.** Optionally, choose the option to join the VMware Customer Experience Improvement Program (CEIP).

**Step 18.** On the Ready to Complete page, click **Finish** and then click **OK**.

## Deploying VCSA Using the CLI

You can use the CLI installer to perform a silent VCSA deployment. The CLI deployment process includes downloading the installer, preparing a JSON configuration file with the deployment information, and running the deployment command. The VCSA installer contains JSON templates for all deployment types. This enables you to deploy an appliance with minimum effort by copying the appropriate JSON template, changing a few values, and using it with the CLI installer. The steps are as follows:

**Step 1.** In the vCenter Server Appliance installer, navigate to one of the following directories:

- In Windows, **/vcsa-cli-installer/win32**
- In Linux, **/vcsa-cli-installer/lin64**
- In macOS, **/vcsa-cli-installer/mac**

**Step 2.** Copy the templates from the **install** subfolder to your desktop.

**Step 3.** Use a text editor to modify the JSON template for your use case. Modify the default parameter values with your appropriate values and add additional parameters, as necessary. For example, to use an IPv4 DHCP assignment, in the **network** subsection of the template, change the value of the **mode** parameter to **dhcp** and remove the default configuration parameters that are used for a static assignment, as shown here:

```
"network": {  
  "ip_family": "ipv4",  
  "mode": "dhcp"  
},
```

**Step 4.** Save the file in UTF-8 format.

[Table 8-8](#) shows some of the available JSON templates.

**Table 8-8** JSON Templates

Template	Description
PSC_repli	Contains the minimum configuration parameters that are

cation_on _ESXi.json	required for deployment of a PSC appliance joining an existing vCenter SSO domain on an ESXi host.
PSC_repli cation_on _VC.json	Contains the minimum configuration parameters that are required for deployment of a PSC appliance joining an existing vCenter SSO domain on a vCenter Server instance.
vCSA_on _ESXi.jso n	Contains the minimum configuration parameters that are required for deployment of a vCenter Server Appliance with an external PSC on an ESXi host.
vCSA_on _VC.json	Contains the minimum configuration parameters that are required for deployment of a vCenter Server Appliance with an external PSC on a vCenter Server instance.

## Note

When using the CLI installer, you must strictly use only ASCII characters for the command-line and JSON configuration file values, including usernames and passwords.

Prior to running the deployment command, you can run a pre-deployment check by using this command:

[Click here to view code image](#)

```
vcsa-deploy install --verify-only path-to-JSON-file
```

When ready, you can run the deployment command:

[Click here to view code image](#)

```
vcsa-deploy install --accept-eula --acknowledge-ceip  
optional_arguments path-to-JSON-file
```

## vCenter Server Post-Installation

After installing vCenter Server, you should be able to access the vSphere Client at [https://vcenter\\_server\\_ip\\_address\\_or\\_fqdn/ui](https://vcenter_server_ip_address_or_fqdn/ui) and the vSphere Web Client at [https://vcenter\\_server\\_ip\\_address\\_or\\_fqdn/vsphere-client](https://vcenter_server_ip_address_or_fqdn/vsphere-client).

The **VMware Enhanced Authentication Plug-in** provides integrated Windows authentication and Windows-based smart card functionality. In the vSphere 6.5 release, the VMware Enhanced Authentication Plug-in replaced the Client Integration Plug-in.

To install the VMware Enhance Authentication Plug-in, you can use the following procedure:

- Step 1.** In a web browser, open the vSphere Web Client.
- Step 2.** At the bottom of the vSphere Web Client login page, click **Download Enhanced Authentication Plug-in**.
- Step 3.** Save the plug-in to your computer and run the executable.
- Step 4.** Follow the wizard for both the VMware Enhanced Authentication Plug-in and the VMware Plug-in Service, which are run in succession.
- Step 5.** When the installations are complete, refresh the browser.
- Step 6.** In the External Protocol Request dialog box, click **Launch Application**.

If you install the plug-in from an Internet Explorer browser, you must first disable Protected Mode and enable pop-up windows on your web browser.

## Configuring and Managing VMware Certificate Authority (VMCA)

**VMware Certificate Authority (VMCA)**, which runs in the vCenter Server Appliance, is responsible for issuing certificates for VMware solution users, certificates for machines running required services, and certificates for ESXi hosts. **VMware Endpoint Certificate Service (VECS)** is a local repository for certificates and private keys. VECS is a mandatory component that is used when VMCA is not signing certificates. VECS includes a set of keystores, including machine SSL certificates, trusted roots, certificate revocation lists (CRLs), and solution users (including machine, vpxd, vpx-extension, and vSphere-webclient). VECS does not store ESXi certificates. ESXi certificates are stored locally on the ESXi hosts in the /etc/vmware/ssl directory. [Table 8-9](#) describes the stores included in VECS.

**Table 8-9** VECS Stores

---

Store	Description
-------	-------------

Machine SSL store (MACHINE_SSL_CERT)	Used by the reverse proxy service on each ESXi host and by the vmdir service.
Trusted root store (TRUSTED_ROOTS)	Contains all trusted root certificates.
Solution user stores:	VECS includes one store for each solution user.

- | Machine
- | vpxd
- | vpxd-extension
- | vsphere-webclient

vSphere Certificate Manager utility backup store (BACKUP_STORE)	Used by VMCA to support certificate reversion.
Other stores	Other stores might be added by solutions. For example, the Virtual Volumes solution adds an SMS store.

With VMCA, you can deal with certificates in three different manners:

- You can let VMCA operate in a default manner, so that it uses a self-signed root certificate, issues certificates to the vSphere components, and serves as the certificate authority (CA) to vSphere.
- You can configure VMCA to operate as a subordinate CA on behalf of the enterprise's CA and to use a subordinate CA signing certificate.
- You can bypass VMCA and use only third-party certificates, which you need to store in the VECS, except for ESXi host certificates. When

necessary, you can use `vecs-cli` commands to explicitly manage certificates and keys.

## Note

The VMCA in vSphere 7.x does not support the use of CRLs, and it does not have the concept of certificate revocation. If you suspect that one certificate was compromised, you should remove it and consider replacing all certificates.

When you use VMCA in the default manner, so that it acts as the CA for vSphere, no real configuration is required other than to configure web browsers to trust VMCA. The VMCA can handle all certificate management in vSphere environments, where historically the administrator has elected not to replace certificates. During an upgrade to vSphere 6.0, all self-signed certificates are replaced with certificates signed by VMCA.

Using VMCA in a subordinate CA manner requires you to replace the VMCA root certificate with a certificate signed by a third-party CA, making the VMCA certificate an intermediate certificate of the CA. To use VMCA in the subordinate CA manner, follow this procedure:



- Step 1** Launch vSphere 6.0 Certificate Manager.
- Step 2** Select **Option 2**, which is to replace the VMCA root certificate with a custom signing certificate and replace all certificates.
- Step 3** When prompted to do so, provide the password for the SSO domain administrator account.
- Step 4** Select **Option 1** to generate a certificate signing request (CSR) and key. When prompted to do so, specify the directory where you want to save the CSR and key.
- Step 5** Provide the CSR (`root_signing_cert.csr`) to your CA to generate the subordinate signing certificate.
- Step 6** Use a text editor to copy content from intermediate CA certificates and the root CA certificate into a single file (`root_signing_chain.cer`).



- Step 7** In Certificate Manager, select **Option 1** to continue on to importing custom certificates.
- Step 8** Import the root signing certificate (root\_signing\_chain.cer) and root signing key (root\_signing\_cert.key).
- Step 9** When prompted to do so, provide a value for each item, such as country, name, and organization.

When you complete these steps, the VMCA root certificate is replaced with a custom signing certificate.

For more details on this procedure, see <https://kb.vmware.com/s/article/2112016>.

## Configuring Single Sign-On (SSO)

In addition to deploying one or more VCSAs and creating a vCenter Single Sign-On (SSO) domain in a new vSphere environment, you need to configure SSO. Configuring SSO includes adding and editing SSO identity sources, configuring SSO users, and configuring SSO policies.

### SSO and Identity Sources Overview

To access vCenter Server, users must log in using SSO domain user accounts or user accounts from identity sources registered in SSO. The acceptable identity sources are Active Directory (with integrated Windows authentication), Active Directory as a LDAP server, Open LDAP, and local OS.

The local OS identity source is available immediately following the installation of SSO. Local operating system users are local to the operating system where the vCenter Single Sign-On server is running. The local operating system identity source exists only in basic vCenter Single Sign-On server deployments and is not available in deployments with multiple vCenter Single Sign-On instances. Only one local operating system identity source is allowed.

The SSO domain is also available immediately as an identity source. This domain was called vsphere.local in vSphere 5.5, but in vSphere 6.x and higher, you can assign the SSO domain name during installation.

### Adding, Editing, and Removing SSO Identity Sources

You can use the *vSphere Client* to add SSO identity sources by following these steps:

- Step 1.** Log in with the vSphere Client to the vCenter Server using [administrator@vsphere.local](mailto:administrator@vsphere.local) or another member of the SSO administrators group.
- Step 2.** Navigate to **Home > Administration > Single Sign On > Configuration**.
- Step 3.** Click **Identity Sources** and then click **Add Identity Source**.
- Step 4.** Select one of the following available identity sources and enter the appropriate settings.
- Active Directory (integrated Windows authentication)
  - Active Directory over LDAP
  - OpenLDAP
  - Local operating system of the SSO server

**Step 5.** Click **Add**.

To remove an SSO identity source, follow these steps:

- Step 1.** In the vSphere Web Client, navigate to **Administrator > Single Sign-On > Configuration > Identity Sources** tab and select the identity source.
- Step 2.** Click the **Delete Identity Source** icon.
- Step 3.** When to do so prompted, click **Yes** to confirm.

You can configure a default domain for SSO. The default SSO domain allows users to authenticate without identifying a domain name. Users from other identity sources must identify the domain name during authentication. To configure a default domain using the vSphere Client, follow these steps:

- Step 1.** Navigate to **Home > Administration > Single Sign On > Configuration**.
- Step 2.** Click **Identity Sources** and then click **Add Identity Source**.
- Step 3.** Select an identity source and click **Set as Default**.

## **Adding an Active Directory Identity Source**

To permit Active Directory authentication in vSphere, you need to add one or

more Active Directory domains as identity sources in SSO. In scenarios where the SSO server is a member of an Active Directory domain, that domain may be added as an Active Directory (integrated Windows authentication) identity source. You can add other Active Directory domains to SSO as Active Directory LDAP server identity sources.

Before you can add an integrated Active Directory identity source, you need to ensure that the server where SSO is installed is in the domain. You add the vCenter Server Appliance to the domain by using the following procedure:

- Step 1.** Log on to the vSphere Web Client using the SSO domain administrator account (such as [administrator@vsphere.local](mailto:administrator@vsphere.local)).
- Step 2.** From the left pane, select **Administration > System Configuration > Nodes**.
- Step 3.** Select the appropriate node and click the **Manage** tab.
- Step 4.** Select **Active Directory** and click **Join**.
- Step 5.** Enter the Active Directory details, such as the domain, organizational unit, username, and password.
- Step 6.** Click **OK**.
- Step 7.** Right-click the node and select **Reboot**.

After the appliance reboots, you can add an Active Directory (integrated Windows authentication) identity source.

When adding an Active Directory (integrated Windows authentication) identity source, you need to provide information for the following parameters:

- **Domain Name:** Enter the FDQN of the domain.
- **Use Machine Account:** Select this option to use the local machine account as the Server Principal Name (SPN). Do not use this option if you plan to rename the machine.
- **Use Service Principal Name (SPN):** Select this option, instead of Use Machine Account, if you prefer to specify a unique SPN instead of using the machine name as the SPN. If you choose this option, you must also provide the SPN, UPN, and password, as follows:
  - **Service Principal Name (SPN):** If you selected the Use Service

Principal Name option, you need to provide a unique name that includes the domain name, such as [STS/domain.com](#).

- **User Principal Name (UPN):** If you selected the Use Service Principal Name option, you need to provide a username that can authenticate the Active Directory domain.
- **Password:** If you selected the Use Service Principal Name option, you need to provide a password that is associated with the UPN.

## Note

The user account must have read access to the organizational units containing users and group. If a user's account does not have sufficient permission or is locked or disabled, then authentications and searches in the Active Directory domain fail.

When adding an Active Directory over LDAP identity source, you need to provide information for the following parameters:

### Key Topic

- **Name:** Specify the logical name for the identify source.
- **Base DN for users:** Specify the base distinguished name for users.
- **Base DN for groups:** Specify the base distinguished name for groups.
- **Domain Name:** Specify the FDQN of the domain.
- **Domain Alias:** Specify the Domain's NetBIOS name.
- **Username:** Specify a username in the domain that has at least read access to the specified user and group base DNs.
- **Password:** Specify the password that is associated with the username.
- **Connect to:** Specify which domain controller to connect.
- **Primary Server URL:** Specify the primary domain controller's URL, in the form `ldap://hostname:port` or `ldaps://hostname:port`.
- **Secondary Server URL:** Specify the secondary domain controller's URL,

in the form of `ldap://hostname:port` or `ldaps://hostname:port`.

- **SSL certificate:** When using LDAPS in the URL parameters, specify the certificate.

You can add additional user accounts from other identity sources to the SSO administrators group. To add additional user accounts from other identity sources to the administrators group in the SSO domain, you can follow these steps:

- Step 1.** Log in to vSphere Web Client with the SSO domain administrator account.
- Step 2.** Navigate to **Home > Administration > Single Sign-On > Users and Groups**.
- Step 3.** Navigate to **Group > Administrators > Group Members** and select the **Add Member** icon.
- Step 4.** Select the additional identity source from the **Domain** drop-down menu.
- Step 5.** Select the account you would like to add.
- Step 6.** Click **OK**.

## **Adding an LDAP Authentication Source**

To use OpenLDAP for authentication, one or more LDAP authentication sources have to be added in vCenter. In order to use OpenLDAP for authentication, the following requirements must be met.

- The OpenLDAP schema must be RFC 4519 compliant.
- All users must have the object class `inetOrgPerson`.
- All groups must have the object class `groupOfUniqueNames`.
- All groups must have the group membership attribute `uniqueMember`.
- All users and group objects must have `entryUUID` configured.

When configuring OpenLDAP, you need to provide information for the following parameters:

- **Name:** Specify the logical name for the identify source.

- **Base DN for users:** Specify the base distinguished name for users.
- **Base DN for groups:** Specify the base distinguished name for groups.
- **Domain Name:** Specify the FDQN of the domain.
- **Domain Alias:** Specify the domain name in capital letters if no alias is defined.
- **Username:** Specify a username in the domain that has at least read access to the specified user and group base DN's.
- **Password:** Specify the password that is associated with the username.
- **Primary Server URL:** Specify the primary server's URL, in the form of `ldap://hostname:port` or `ldaps://hostname:port`.
- **Secondary Server URL:** Specify the secondary server's URL, in the form of `ldap://hostname:port` or `ldaps://hostname:port`.
- **SSL certificate:** When using LDAPS in the URL parameters, specify the certificate.

## Enabling and Disabling Single Sign-On (SSO) Users

To manage SSO users, you can use the vSphere Client. For example, to add an SSO user, follow these steps:

- Step 1.** Log on to the vCenter Server connected using [administrator@vsphere.local](mailto:administrator@vsphere.local) or another user in the SSO administrators group.
- Step 2.** Navigate to **Home > Administration > Single Sign-On > Users and Groups**.
- Step 3.** Select the **Users** tab and click **Add User**.
- Step 4.** Provide the username and password. Optionally provide values for the other fields.
- Step 5.** Click **OK**.

In a similar manner, you can create an SSO group by selecting the Users tab in step 3 and providing details in step 4. You can also use the Groups tab to select a group and use the Add Member icon (in the details section) to add users to the group. When adding a user to a group, use the Domain drop-down to select the

SSO domain or another identity source and select a user account from the provided list.

To disable or enable an SSO user account, select the user account in Users and Groups, click the ellipsis icon, and click Disable or Enable.

The SSO domain—which is called vsphere.local in vSphere 5.5 but may be named differently in vSphere 6.0—provides several predefined groups. You can add users from Active Directory domains or other identity sources to these predefined groups. Some SSO privileges are determined solely by the membership of these groups. For example, a user who is a member of the CAAdmins group can manage VMCA, and a user who is a member of the LicenseService.Administrators group can manage licenses.

The SSO domain contains many predefined groups, including the following:

- **Users:** This group contains all users in the SSO domain.
- **DCAdmins:** Members of this group can perform domain controller administrator actions on VMware Directory Service.
- **SolutionUsers:** Each solution user authenticates individually to vCenter Single Sign-On with a certificate. By default, VMCA provisions solution users with certificates. Do not add members to this group explicitly.
- **CAAdmins:** Members have administrator privileges for VMCA. Adding members to these groups is not usually recommended, but a user must be a member of this group to perform most certificate management operations, such as using the **certool** command.
- **SystemConfiguration.BashShellAdministrators:** Members can enable and disable access to the BASH Shell.
- **SystemConfiguration.Administrators:** Members can view and manage the system configuration and perform tasks such as restarting services.
- **LicenseService.Administrators:** Members have full write access to all licensing-related data and can add, remove, assign, and un-assign serial keys for all product assets registered in licensing service.
- **Administrators:** Members can perform SSO administration tasks for VMware Directory Service (vmdir).

## Configuring SSO Policies

SSO provides policies that enforce security rules in the environment. You can configure SSO password policies, SSO lockout policies, and SSO token policies. To configure these policies, you can use the vSphere Client to select Administration > Single Sign-On > Configuration and then select Password Policy, Lockout Policy, or Token Policy and click Edit. For each set of policies, you need to set the appropriate password policy parameters, as described in [Table 8-10](#).

**Table 8-10** SSO Policies and Parameters

SSO Policy Parameter	Policy Setting	Details
Password Policy	Description	Password policy description.
	Maximum lifetime	Maximum number of days a password can exist before the user must change it.
	Restrict reuse	Number of the user's previous passwords that cannot be selected.
	Maximum length	Maximum number of characters that are allowed in the password.
	Minimum length	Minimum number of characters that are allowed in the password, which must be no fewer than the combined minimum of alphabetic, numeric, and special character requirements.
	Character requirements	Minimum number of different character types that are required in the password. The types include special, alphabetic, uppercase, lowercase, and numeric.
	Identical adjacent characters	The number of identical adjacent characters that are supported in a password. The value must be greater than 0.
Lockout	Description	Description of the lockout policy.



t Poli cy	Max number of failed login attempts	Maximum number of failed login attempts that are allowed before the account is locked.
	Time interval between failures	Time period in which failed login attempts must occur to trigger a lockout.
	Unlock time	The amount of time the account stays locked. The value 0 specifies that an administrator must explicitly unlock the account.
Tok en Poli cy	Clock tolerance	Time difference, in milliseconds, that SSO tolerates between a client clock and a domain controller clock. If the time difference is greater than the specified value, SSO declares the token to be invalid.
	Maximu m token renewal count	Maximum number of times a token may be renewed before a new security token is required.
	Maximu m token delegatio n count	Maximum number of times a single holder-of-key token can be delegated.
	Maximu m bearer token lifetime	The lifetime value of a bearer token before the token must be reissued.
	Maximu m holder- of-key token lifetime	The lifetime value of a holder-of-key token before the token is marked invalid.

## Configuring Identity Federation

When added to an identity provider, the vSphere Client redirects a user's login attempt to that provider's login page. Corporate credentials (including multifactor authentication) can then be used to log in. After authentication, the identity provider redirects the user back to the vSphere Client via a cryptographic token used for authorization.

Identity Federation utilizes OAuth 2 and OIDC protocols for information exchange. Identity Federation replaces traditional Active Directory, integrated Windows authentication, and LDAP/LDAPS authentication for vCenter. vSphere Single Sign-On is not replaced, however, to allow additional administration or emergency access.



The following requirements apply to Active Directory Federation Services:

- AD Federation Services for Windows Server 2016 R2 or higher must be deployed.
- AD Federation Services must be connected to Active Directory.
- An application group for vCenter is created in AD Federation Services for configuration (see <https://kb.vmware.com/s/article/78029>).
- The AD Federation Services server certificate must be signed by a certificate authority. If self-signed certificates are used, the root CA certificate has to be imported to the vCenter JRE truststore.

The following requirements apply to vSphere to support Identity Federation:

- vSphere 7.0 or later is required.
- Communication is possible between the vCenter Server and AD Federation Services endpoint, authorization, token, logout, JWKS, and other advertised endpoints.
- The VcIdentityProviders.Manage privilege is required in vCenter to create, update, or delete a vCenter Server identity provider for Federation Services authentication. If a user should be limited to viewing the identity provider configuration and not changing it, you need to use the VcIdentityProviders.Read privilege.

Follow these steps to enable Identity Federation:

- Step 1.** In the vSphere Client, navigate to **Home > Administration > Configuration > Single Sign On.**
- Step 2.** Select the **Identity Provider** tab.
- Step 3.** Click the **i** next to Change Identity Provider.
- Step 4.** Copy the URIs in the pop-up (as they will be needed to configure the AD Federation Services server) and then close the pop-up.
- Step 5.** Create an OpenID Connect configuration in AD Federation Services and configure it for the vCenter Server. A shared secret and identifying information must be established between the identity provider and vCenter. An OpenID Connect configuration is created in AD Federation Services and is known as an application group; it contains a server application and a web API. These components define the information that vCenter uses to trust and communicate with the AD Federation Services server. To create the AD Federation Services application group, the two redirect URIs from step 4 are used. It is important to copy or record the following information:
  - a.** Client identifier
  - b.** Shared secret
  - c.** OpenID address of the AD Federation Services server
- Step 6.** On the **Identity Provider** tab, click **Change Identity Provider**, select **Microsoft ADFS**, and click **Next**.
- Step 7.** Enter the following information:
  - a.** Client identifier
  - b.** Shared secret
  - c.** OpenID address of the AD Federation Services serverClick **Next**.
- Step 8.** Enter the user and group information for Active Directory over LDAP to search for users and groups and click **Next**.
- Step 9.** Review the configuration information and click **Finish**.
- Step 10.** Go to **Home > Administration > Single Sign On > Users and Groups**.

**Step 11.** On the **Groups** tab, click **Administrators (group)** and click **Add Members**.

**Step 12.** From the drop-down, select **Microsoft ADFS**, and in the text box under the drop-down menu, enter **vcenter** and wait for the drop-down to show a selection. Then select **vCenter Admins** and add it to the group.

**Step 13.** Click **Save**.

**Step 14.** Log in to vCenter with an AD user's credentials to verify functionality.

## Initial vSphere Configuration

This section covers some of the vSphere settings, components, and features that are typically configured in conjunction with a new vSphere deployment.

### Implementing vSphere Client

The vSphere Client, which is the primary vSphere GUI, is HTML5 based and uses the Clarity style user interface. The vSphere Client is a service that is installed automatically when you install vCenter Server.

### Implementing VMware vSphere Lifecycle Manager

In vSphere 7.0, **vSphere Lifecycle Manager** replaces VMware Update Manager from prior versions. Lifecycle Manager adds to the functionality of Update Manager to include features and capabilities for ESXi lifecycle management at the cluster level. Lifecycle Manager operates as a service that runs on the vCenter Server. This service is available via the vSphere Client immediately after the vCenter Server deployment; no special steps are required to install vSphere Lifecycle Manager unless you need to install the optional module VMware vSphere Update Manager Download Service (UMDS).

In scenarios where vCenter Server is installed in a secured network with no Internet access, you can install UMDS and use it to download updates. You can use UMDS to export the updates to a portable media drive that you then present to vSphere Lifecycle Manager. Or, if network connectivity exists between the vCenter Server and UMDS, you can automate the export process by leveraging the web server on the UMDS machine.

## Note

See [Chapter 13, "Managing vSphere and vCenter Server,"](#) for more details on vSphere Lifecycle Manager.

## Configuring the vCenter Server Inventory

This section describes the procedures you use to create an inventory hierarchy that includes data centers, folders, clusters, and resource pools.

To create a data center, you can follow these steps:

- Step 1.** Log on to the vSphere Client as a user with the Datacenter. Create datacenter privilege.
- Step 2.** Right-click the vCenter Server in the inventory and select **New Datacenter**.
- Step 3.** Provide a name for the data center and click **OK**.

To create a folder, you can follow these steps:

- Step 1.** In the vSphere Client, right-click the appropriate parent object (a data center or another folder).
- Step 2.** Select **New Folder**. If the parent object is a folder, the new folder type is automatically set to match the parent; otherwise, specify the folder type as **Host and Cluster** folder, **Storage** folder, or **VM and Template** folder.
- Step 3.** Provide a name for the folder and click **OK**.

To add an ESXi host, you must address the following prerequisites:

- Ensure that the appropriate data center and (optionally) folder objects are created in the vCenter Server inventory.
- Obtain the root account credentials for the host.
- Verify that the host and vCenter Server can communicate via port 902 or a custom port.
- Verify that any NFS mounts on the host are active.

- For a host with more than 512 LUNs and 2048 paths, verify that the vCenter Server instance is set to support a Large or X-Large environment.

To add an ESXi host, you can follow these steps:

- Step 1.** Log on to the vSphere Client as a user with the Host.Inventory.Add standalone host privilege.
- Step 2.** Right-click the data center or folder in the inventory and select **Add Host**.
- Step 3.** Provide the host's FQDN (or IP address) and credentials.
- Step 4.** Provide licensing information (using a new or existing license).
- Step 5.** Optionally, set the lockdown mode, remote access, and virtual machine folder information and continue navigating to the final wizard page.
- Step 6.** Click **OK**.

## Implementing vCenter HA



You can use the following procedure to configure vCenter HA:

- Step 1.** In the vSphere Client, select the vCenter Server in the inventory pane.
- Step 2.** Select **Configure > Select vCenter HA > Set Up vCenter HA**.
- Step 3.** If your vCenter Server is managed by another vCenter Server in a different SSO domain, complete the following steps:
  - a.** Click **Management vCenter Server Credentials**. Provide the FQDN and Single Sign-On credentials and click **Next**.
  - b.** If you see a certificate warning displayed, review the SHA1 thumbprint and click **Yes** to continue.
- Step 4.** In the Resource Settings section, select the vCenter HA network for the active node from the drop-down menu.
- Step 5.** Click the checkbox if you want to automatically create clones for Passive and Witness nodes.
- Step 6.** For the Passive node, follow these steps:

- Click **Edit** and provide details for the Passive node virtual machine, such as the name, compute resources, and datastore.
- Select the Management (NIC 0) and vCenter HA (NIC 1) networks.
- Complete the settings and click **Finish**.

**Step 7.** For the Witness node, follow these steps:

Click **Edit** and provide details for the passive node virtual machine, such as the name, compute resources, and datastore.

- Select the vCenter HA (NIC 1) networks.
- Complete the settings and click **Finish**.

**Step 8.** Click **Next**.

**Step 9.** In the IP Settings section, provide the IP address details.

**Step 10.** Click **Finish**.

## Using Host Profiles

This section provides details on using host profiles to configure ESXi hosts.

### Configuring ESXi Using Host Profiles

During the implementation of vSphere, you can use host profiles to efficiently deploy a standard configuration to a set of ESXi hosts. To do so, you can configure a single ESXi host, create a host profile from that host, and apply the profile to other recently deployed hosts. This process reduces the time required to configure ESXi hosts and minimizes the risk of misconfigured hosts. The host profile contains all the networking, storage, security, and other host-level settings. The host from which the profile is created is known as the reference host. You can attach a host profile to individual hosts, a cluster, or all the hosts and clusters managed by a vCenter Server.

### Applying a Host Profile

After attaching a host profile, you can check compliance of the associated hosts and remediate as necessary.

You can use this procedure to create a host profile from a reference host:

- Step 1.** Navigate to the host profiles main view and click **Extract Host Profile**.
- Step 2.** Select the host that acts as the reference host and click **Next**.
- Step 3.** Enter the name and description for the new profile and click **Next**.
- Step 4.** Review the summary information for the new profile and click **Finish**.

You can use this procedure to attach a profile to ESXi hosts and clusters:

- Step 1.** From the host profiles main view, select the host profile to be applied to the host or cluster.
- Step 2.** Click **Attach/Detach Hosts and Clusters**, select the host or cluster from the expanded list, and click **Attach**.
- Step 3.** Optionally, click **Attach All** to attach all listed hosts and clusters to the profile.
- Step 4.** Optionally, enable **Skip Host Customization**; if you do, you do not need to customize hosts during this process.
- Step 6.** Click **Next**.
- Step 7.** Optionally, update or change the user input parameters for the host profiles policies by customizing the host.
- Step 8.** Click **Finish** to finish attaching the host or cluster to the profile.

You can use this procedure to remediate an ESXi host:

- Step 1.** Navigate to the host profiles main view.
- Step 2.** Right-click a host profile and select **Remediate**.
- Step 3.** Select the host or hosts you want to remediate with the host profile.
- Step 4.** Optionally, enter the host customizations to specify host properties or browse to import a host customization file.
- Step 5.** Click **Pre-check Remediation** to check whether the selected hosts are ready for remediation.
- Step 6.** Select the checkbox to reboot the host if a reboot is required in order to complete the remediation process. To manually reboot the host after the process, do not select the checkbox.



**Step 7.** Review the tasks that are necessary to remediate the host profile and click **Finish**.

## **Editing Host Profiles**

To edit a host profile, you can use this procedure:

**Step 1.** Navigate to the host profiles main view.

**Step 2.** Select the host profile that you want to edit and click the **Configure** tab.

**Step 3.** Click **Edit Host Profile**.

**Step 4.** Optionally, click the **Name and Description** tab to change the profile name and description.

**Step 5.** In the Edit Host Profile page, expand each category to view or edit a specific policy or setting.

**Step 6.** Select **All** to view all host profile configurations or select **Favorites** to view only those configurations.

**Step 7.** Optionally, in the search field, filter the configuration names and values you want to view. For example, enter **SNMP**, and all configurations that relate to SNMP are displayed.

**Step 8.** Optionally, customize the hosts. Make any changes to the available configuration values for this profile and click **Save**.

## **Applying Permissions to ESXi Hosts by Using Host Profiles**

You can use host profiles to apply ESXi host permissions to be used when users access a host directly. To configure the host profile with the appropriate permissions, you can use the vSphere Client (not the vSphere Web Client) and follow this procedure:

**Step 1.** Select **View > Management > Host Profiles**.

**Step 2.** Select an existing profile and click **Edit Profile**.

**Step 3.** In the profile tree, locate and expand **Security Configuration**.

**Step 4.** Right-click on the **Permission Rules** folder and click **Add Profile**.

**Step 5.** Expand **Permissions Rules** and select **Permission**.

- Step 6.** On the **Configuration Details** tab, click the **Configure Permission** drop-down menu and select **Require a Permission Rule**.
- Step 7.** Enter the name of a user or group, in the format *domain\name*, where *domain* is the domain name, and *name* is the user or group name.
- Step 8.** If a group name is used, select the **Name Refers to a Group of Users** checkbox.
- Step 9.** Enter the assigned role name, which is case sensitive. This can be the name of a built-in role on the host or a custom role that you have created on the host. For system roles, use the non-localized role name, such as Admin for the administrator role or ReadOnly for the read-only role.
- Step 10.** Optionally, select **Propagate Permission**.
- Step 11.** Click **OK**.

After configuring the host profile, you can use it to apply the permissions to new or existing ESXi hosts.

## VMware Tools

Ideally, you should install VMware Tools in all your virtual machines. When deploying a new vSphere environment, you should install VMware Tools in any virtual machines deployed as part of the virtual infrastructure and management. For example, if you use virtual machines to run Active Directory domain controllers, DNS servers, or DHCP servers, consider installing VMware Tools.

VMware Tools is a suite of utilities that you install in the operating system of a virtual machine. VMware Tools enhances the performance and management of the virtual machine. You can use the following procedure to install VMware Tools in a virtual machine using the VMware Host Client:

- Step 1.** Click **Virtual Machines** in the VMware Host Client inventory.
- Step 2.** Select a powered-on virtual machine from the list. (The virtual machine must be powered on to install VMware Tools.)
- Step 3.** Open a console to the virtual machine and log in with administrator or root privileges.
- Step 4.** Click **Actions**, select **Guest OS** from the drop-down menu, and select

## Install VMware Tools.

**Step 5.** Use the guest OS to complete the installation.

This procedure is useful for installing VMware Tools in a DNS, Active Directory domain controller, database server, or other virtual machine that you may deploy prior to deploying vCenter Server.

## Advanced ESXi Host Options

This section describes how to configure some advanced options for ESXi hosts.

### ESXi Advanced System Settings

You can use the vSphere Host Client or the vSphere Client to set advanced attributes on ESXi hosts. You should change the advanced options only when you get specific instructions to do so from VMware technical support or a knowledge base article.

To change a host's advanced settings using the vSphere Host Client, you can navigate to Manage > System > Advanced Settings. To change a host's advanced settings using the vSphere Client, you can select the host and navigate to Configure > System > Advanced System Settings.

### ESXi Kernel Options

Disk partitioning in ESXi 7.0 boot devices has changed from previous ESXi versions. The changes include the following:

- The system boot partition is now larger.
- The boot banks are now larger.
- Coredump, tools, and scratch have been consolidated into a single VMFS-L-based ESX-OSData volume.
- Coredumps default to a file in the ESX-OSData volume.

You need to issue boot options at the time of boot, either by defining the kernel options in the ESXi boot.cfg file or by manually entering the boot options after pressing Shift+O in the ESXi boot loader. [Table 8-11](#) lists the boot options.

#### **Table 8-11** ESXi 7.0 Kernel Options

---

<b>Kernel Option</b>	<b>Description</b>
----------------------	--------------------

---

<b>autoPartition=TRUE/FALSE (default FALSE)</b>	This option, if set to TRUE, defines automatic partitioning of the unused local storage devices at boot time. The boot disk gets partitioned with boot bands, ESXi-OSData, and, if the disk is larger than 128 GB, a VMFS partition. Any new empty device discovered will be auto-partitioned as well. Auto-partitioning can be set for only the first unused device with the setting <b>autoPartitionOnlyOnceAndSkipSsd=TRUE</b> . On hosts with USB boot and VMFS-L, ESX-OSData does not exist on other local disks.
---	--

If a storage device has both a scratch partition and a coredump partition, the scratch partition is converted to ESX-OSData; otherwise, the first unused disk identified is partitioned with ESX-OSData as well.

---

<b>skipPartitioningSsds=TRUE/FALSE (default FALSE)</b>	If this option is set to TRUE, local SSDs are excluded from automatic partitioning.
--	---

---

<b>autoPartitionOnlyOnceAndSkipSsd=TRUE/FALSE (default FALSE)</b>	If this option is set to TRUE, SSD/NVMe devices are excluded, and the ESXi host automatically partitions the first unused local disk if there is no VMFS-L ESX-OSData volume.
---	---

E)

---

**allowCoreDumpOnUSB=TRUE/FALSE**  
(default FALSE)  
E)

If this option is set to TRUE, ESXi can write kernel crash core dumps to the VMFS-L Locker volume on a USB boot device.

---

**dumpSize**  
(default:0  
(automatically sized))

This option sets the size of the core dump file (in megabytes) created on the system VMFS-L volume. This is limited to one-half of the space available on the VMFS-L volume.

---

**autoCreatedDumpFile=TRUE/FALSE**  
(default TRUE)  
)

This option, when set to TRUE, automatically creates a core dump file. This is attempted in the following order:

- VMFS-L ESX-OSData
- USB VMFS-L
- Local VMFS

---

**Note**

The following kernel boot options have been deprecated and are no longer supported in ESXi 7.0:

- **--no-auto-partition**
- **autoPartitionCreateUSBCoreDumpPartition**
- **autoPartitionDiskDumpPartitionSize**

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

## Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 8-12](#) lists these key topics and the page number on which each is found.



**Table 8-12** Key Topics for [Chapter 8](#)

<b>Key Topic Element</b>	<b>Description</b>	<b>Page Number</b>
List	Installing ESXi interactively	<a href="#">286</a>
List	First-time boot sequence using Auto Deploy	<a href="#">296</a>
List	Deploying VCSA using the GUI installer	<a href="#">299</a>
List	Using VMCA in a subordinate CA manner	<a href="#">305</a>
List	Adding an Active Directory over LDAP identity source	<a href="#">308</a>
List	Enabling Identity Federation	<a href="#">313</a>
List	Implementing vCenter HA	<a href="#">316</a>

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

[VMware Lifecycle Manager](#)  
[vSphere Client](#)  
[VMware Enhanced Authentication Plug-in](#)  
[VMware Certificate Authority \(VMCA\)](#)  
[VMware Endpoint Certificate Service \(VECS\)](#)  
[stateless caching](#)

## Review Questions

1. You are using the GUI installer for vCenter Server 7.0. Which of the following statements is true?
  - a. In the first stage, you choose the deployment type. In the second state, you navigate through the installation wizard.
  - b. In the first stage, you provide the appliance settings. In the second stage, you navigate through the installation wizard.
  - c. In the first stage, you choose the deployment type. In the second stage, you deploy the OVA.
  - d. In the first stage, you provide the appliance settings. In the second stage, you configure SSO.
2. You are adding an Active Directory over LDAP identity source. Which setting must you provide?
  - a. UPN
  - b. SPN

- c. Use machine account
  - d. Base DN for users
- 3. You are implementing a new vSphere environment and want to install services for updating the ESXi hosts. What should you do?
  - a. Deploy a VMware Update Manager appliance.
  - b. Deploy a vSphere Lifecycle Manager appliance.
  - c. Deploy vCenter Server with Update Manager embedded.
  - d. Nothing; the software service is included in vCenter Server.
- 4. You are implementing Auto Deploy and want to control its behavior with rules. Which of the following is not a means by which a rule can identify target hosts?
  - a. MAC address
  - b. Model
  - c. Serial number
  - d. BIOS UUID
- 5. You are using host profiles to deploy a standard configuration to ESXi hosts. Which of the following provides the proper order of operation?
  - a. Click Attach Host Profile, click Pre-check Remediation, click Remediate.
  - b. Click Attach Host Profile, click Remediate, click Pre-check Remediation, click Finish.
  - c. Click Pre-check Remediation, click Pre-check Remediation, click Remediate.
  - d. Click Pre-check Remediation, click Remediate, click Attach Host Profile, click Finish.



## Chapter 9

# Configuring and Managing Virtual Networks

This chapter covers the following topics:

- [vSphere Standard Switches \(vSS\)](#)
- [vSphere Distributed Switches \(vDS\)](#)
- [VMkernel Networking](#)
- [Configuring and Managing Networking Features](#)
- [Managing Host Networking with vDS](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.4, 1.7, 1.7.1, 1.7.2, 1.7.3, 1.7.4, 4.2, 4.14, and 5.4.

This chapter concentrates on how to configure and manage different network objects in vSphere 7. The procedures described here can be useful in creating the various networking objects that you will encounter on the 2V0-21.20 exam.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 9-1](#) 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 and Review Questions.](#)”

**Table 9-1** “Do I Know This Already?” Section-to-Question Mapping

---

---

vSphere Standard Switches (vSS)

---

vSphere Distributed Switches (vDS)

---

VMkernel Networking

---

Configuring and Managing Networking Features

---

Managing Host Networking with vDS

---

- 1.** You want to connect physical adapters to an existing vSphere Standard Switch. Which of the following steps should you take?
  - a.** Select Manage Physical Adapters and click the Add Adapters button.
  - b.** Select VMkernel Adapters.
  - c.** Create a port group and then add the adapter to it.
  - d.** Create an uplink port group and then add the adapter to it.
  
- 2.** You are assigning a VLAN ID to a standard port group. What is the acceptable range?
  - a.** 1–4094
  - b.** 0–4095
  - c.** 0–4094
  - d.** 1–1095
  
- 3.** You are modifying an existing vDS. Which of the following options is not available on the General settings page?
  - a.** Name
  - b.** Number of Uplinks
  - c.** Network I/O Control
  - d.** VLAN ID
  
- 4.** You are modifying an existing vDS. Which of the following options is not available on the Advanced settings page?
  - a.** Traffic Filtering and Marking

- b.** MTU
  - c.** Multicast Filtering Mode
  - d.** Discovery Protocol
- 5.** You are adding a VMkernel adapter to a vDS. Which one of the following settings is not available?
  - a.** SR-IOV
  - b.** MTU
  - c.** TCP/IP Stack
  - d.** Available Services
- 6.** You enabled NIOC for a distributed switch, and you want to change shares for the system traffic. Which of the following is not an available system traffic type?
  - a.** Fault Tolerance
  - b.** vSAN
  - c.** vSphere HA Heartbeat
  - d.** NFS
- 7.** You are configuring port mirroring for a distributed switch. Which of the following is not an available session type?
  - a.** Distributed port monitoring
  - b.** Port group
  - c.** Remote mirroring destination
  - d.** Encapsulated remote mirroring (L3) source
- 8.** You want to implement LAGs to support your vSphere 7.0 networking. Which of the following steps should you take to prepare?
  - a.** Ensure that all the NICs in a LAG are configured with the same speed and duplexing.
  - b.** Ensure that the number of ports in a single port channel on the switch is equivalent to the total number of participating NICs for ESXi hosts.

- c. Ensure that the number of ports in a single port channel on the switch is equivalent to or greater than the number of participating NICs from a specific ESXi host.
  - d. Ensure that the number of participating NICs on each host is greater than the number of ports in the port channel.
9. You want to enable vDS health checks in your vSphere 7.0 environment. Which of the following are valid health check services that you can enable or disable in the vSphere 7.0 Client?
- a. MTU and Failover and xxx
  - b. VLAN and MTU
  - c. VLAN and Teaming
  - d. MTU and Teaming
10. You are configuring rules to mark network packets. Which of the following is not a valid option for qualifying packets?
- a. Destination IP address
  - b. Source IP address
  - c. Source and destination MAC addresses
  - d. Destination port number

## Foundation Topics

### vSphere Standard Switches (vSS)

This section addresses the creation and configuration of vSS and standard port groups.

#### Creating and Configuring vSphere Standard Switches



You can use the following procedure to create a vSS that provides network connectivity for hosts and virtual machines:

**Step 1.** In the vSphere Client, select an ESXi host in the inventory pane and

navigate to **Configure > Networking > Virtual Switches**.

**Step 2.** Click **Add Networking**.

**Step 3.** Select a connection type (VMkernel Network Adapter, Physical Network Adapter, or Virtual Machine Port Group for a Standard Switch) for which you want to use the new standard switch and click **Next**.

**Step 4.** Select **New Standard Switch**, optionally change the MTU setting (from the default 1500), and click **Next**.

**Step 5.** Carry out the appropriate steps for the selection you made in Step 3:

- On the Create a Standard Switch page, to add physical network adapters to the standard switch, do the following steps. Otherwise, click **Next**.
  - a. In the Assigned Adapters window, click the **Add Adapter** (green plus sign) button.
  - b. From the list of available network adapters, select one or more adapters.
  - c. Click **OK** and then click **Next**.
- If you created a new standard switch with a VMkernel adapter, use the Port Properties page to configure the adapter:
  - a. Provide a network label that indicates its purpose, such as vMotion or NFS.
  - b. Optionally, set a VLAN ID.
  - c. Select **IPv4**, **IPv6**, or **IPv4 and IPv6**.
  - d. Set MTU to a custom size for the VMkernel adapter or choose to get the MTU from a switch.
  - e. Select a TCP/IP stack for the VMkernel adapter. (You cannot change it later.)
  - f. If you use the default TCP/IP stack, select from the available services (vMotion, Provisioning, Fault Tolerance logging, Management, vSphere Replication, vSphere Replication NFC, vSAN).

- g. Click **Next**.
- h. Configure the IP settings and click **Next**.
- If you created a new standard switch with a virtual machine port group, use the Connection Settings page to configure the port group:
  - a. Provide a network label for the port group.
  - b. Optionally, assign a VLAN ID.
  - c. Click **Next**.

**Step 6.** On the Ready to Complete page, click **Finish**.

### Note

If you create a standard switch without physical network adapters, all traffic on that switch is confined to that switch. You can create a standard switch without physical network adapters if you want a group of virtual machines to be able to communicate with each other but with nothing else.

You can make vSS configuration settings, including settings that control switch-wide defaults for ports. Such settings can be overridden by port group settings. To modify the settings of a vSS, select the host in the vSphere Client inventory pane and click Edit. You can then adjust the following settings:

- **Number of vSS ports:** The number of ports on a standard switch is dynamically scaled up and down. You cannot explicitly set the number of ports in a vSS. A vSS can expand up to the maximum number of ports supported on the host, which is based on the maximum number of virtual machines the host can handle.
- **Maximum transmission unit (MTU):** You can enable jumbo frames on a vSS by increasing the MTU setting from the standard 1500 bytes. You can set the MTU between 1280 bytes and 9000 bytes.
- **Physical network adapter:** Virtual machines connected to a vSS can only reach the physical networks through uplink physical network adapters. If you can connect two or more adapters to a vSS, they are transparently teamed.

To change or add physical adapters that are assigned to a vSS, you can use the following procedure:

- Step 1.** In the vSphere Client, select the ESXi host in the inventory pane and navigate to **Configure > Networking > Virtual Switches**.
- Step 2.** Navigate to the appropriate standard switch and select **Manage Physical Adapters**.
- Step 3.** In the Manage Physical Adapters window, click the **Add Adapter** (green plus sign) button.
- Step 4.** In the Add Physical Adapters to Switch window, select one or more adapters to assign to the vSS and click **OK**.
- Step 5.** In the Manage Physical Adapters window, use the up and down buttons to set each assigned vSS adapter to **Active**, **Standby**, or **Unused**. Click **OK**.

To view the MAC address and other characteristics of a host's physical NICs, you can select the host and navigate to **Configure > Networking > Physical Adapters**. To change the speed and duplexing of an adapter, select the adapter, click **Edit**, and make the change.

When configuring networks for the virtual machines in your vSphere environment, consider whether you want to migrate the virtual machines among a set of hosts. If so, be sure that the hosts are in the same broadcast domain (that is, the same Layer 2 subnet). ESXi does not support migration of virtual machines between hosts in different broadcast domains as a virtual machine may lose access to required resources in the destination network. Even if your network provides high availability or includes intelligent switches that can resolve the virtual machine's needs across different networks, you may experience significant lag times as the Address Resolution Protocol (ARP) table updates.

## **Creating and Configuring Standard Port Groups**

You can use the following procedure to add a virtual machine port group to a standard virtual switch:

- Step 1.** In the vSphere Client, right-click a host in the inventory pane and select **Add Networking**.

**Step 2.** For Select Connection Type, select **Virtual Machine Port Group for a Standard Switch** and click **Next**.

**Step 3.** For **Select Target Device**, choose whether to create a new standard switch or to use an existing switch:

- If the new port group is for an existing standard switch, click **Browse**, select the standard switch, click **OK**, and click **Next**.
- If you are creating a new standard switch, you can assign physical network adapters to the standard switch or you can choose to create the standard switch with no assigned physical network adapters:
  - a. Click **Add Adapters**.
  - b. Select an adapter from the Network Adapters list.
  - c. Use the Failover Order Group drop-down menu to assign the adapter to **Active Adapters**, **Standby Adapters**, or **Unused Adapters** and click **OK**.
  - d. Use the up and down arrows in the Assigned Adapters list to change the position of the adapter, if needed.
  - e. Click **Next**.

**Step 4.** On the Connection Settings page, set a network label for the port group and, optionally, set a VLAN ID. Click **Next**.

**Step 5.** On the Ready to Complete page, click **Finish**.

On a standard switch port group, the VLAN ID reflects the VLAN tagging mode in the port group, as shown in [Table 9-2](#).

**Table 9-2** VLAN ID Details

<b>VLAN ID</b>	<b>VLAN Tagging Mode</b>	<b>Description</b>
0	External switch tagging (EST)	The virtual switch does not pass traffic associated with a VLAN.
1 to 4094	Virtual switch tagging (VST)	The virtual switch tags traffic with the entered tag.
409	Virtual guest	Virtual machines handle VLANs. The virtual



You can edit an existing standard switch port group by using the following procedure.:

- Step 1.** In the vSphere Client, select a host in the inventory pane and navigate to **Configure > Networking > Virtual Switches**.
- Step 2.** Select the appropriate standard switch and navigate to the switch's topology diagram.
- Step 3.** In the topology diagram, click on the name of the port group and click the **Edit Settings** icon.
- Step 4.** On the Properties page, optionally change the port group's network label and VLAN ID.
- Step 5.** On the Security page, optionally override the switch settings concerning MAC address impersonation and using promiscuous mode.
- Step 6.** On the Traffic Shaping page, optionally override the switch settings for throttling network traffic based on average and peak bandwidth.
- Step 7.** On the Teaming and Failover page, optionally override the teaming and failover settings inherited from the standard switch. Optionally configure traffic distribution across the physical adapters and the failover order.
- Step 8.** Click **OK**.

To remove a port group from a standard switch, navigate to the switch's topology, select the port group, and click the Remove Selected Port Group icon.

## vSphere Distributed Switches (vDS)

This section addresses the creation and configuration of vSphere Distributed Switches (vDS) and distributed port groups.

### Creating and Configuring vSphere Distributed Switches

You can use the vSphere Client to create a vDS:

- Step 1.** In the vSphere Client, right-click a data center in the inventory pane and select **Distributed Switch > New Distributed Switch**.

- Step 2.** On the Name and Location page, enter a name for the new distributed switch or accept the generated name and click **Next**.
- Step 3.** On the Select Version page, select a distributed switch version (7.0, 6.6, or 6.5) and click **Next**. Features released with vSphere versions later than the version selected are not supported for the distributed switch.
- Step 4.** On the Configure Settings page, provide the following vDS settings:
- Select the number of uplinks. The number of uplink ports determines how many physical NICs are connected per ESXi host.
  - Enable or disable Network I/O Control. If it is enabled, it prioritizes network traffic.
  - (Optional) Select the **Create a Default Port Group** checkbox to create a new distributed port group with default settings for this switch. Enter a port group name or accept the generated name.
  - If your system has custom port group requirements, create distributed port groups that meet those requirements after you add the distributed switch.
- Click **Next**.
- Step 5.** On the Ready to Complete page, review the settings you selected and click **Finish**.

If you plan on using NSX-T, set the vDS version to 7.0 and use NSX-T 3.0 or later.

## Upgrading a vDS

You can upgrade a vDS from Version 6.x to a later version, but you cannot revert a vDS to an earlier version. As a rollback plan, you should export the distributed switch configuration prior to upgrading. In the export wizard, choose the option to include the distributed port groups. If an issue emerges, you can re-create the vDS by importing the switch configuration file and choosing the Preserve Original Distributed Switch and Port Group Identifiers option.

The following list provides information for exporting and importing vDS configuration.:

- **Exporting a vDS configuration:** To export a vDS configuration, select it in the inventory pane, select Settings > Export Configuration, and use the wizard. In the wizard, select whether you want to include the configuration of the distributed port groups in the export. Optionally, you can provide a description for the export. The file is saved to your local system.
- **Importing a vDS configuration:** To import a vDS configuration, right-click a data center in the inventory pane, select Distributed Switch > Import Distributed Switch, and use the wizard. In the wizard, to assign the keys from the configuration file to the switch and its port groups, select the Preserve Original Distributed Switch and Port Group Identifiers checkbox.

Upgrading a distributed switch causes the hosts and virtual machines attached to the switch to experience brief downtime. VMware recommends performing the upgrade during a maintenance window and changing the DRS mode to manual (and ignoring DRS recommendations) during the upgrade.

You can use the following procedure to upgrade a vDS:

- Step 1.** In the vSphere Client, navigate to **Networking**, right-click the distributed switch in the inventory pane, and select **Upgrade > Upgrade Distributed Switch**.
- Step 2.** Select the vSphere Distributed Switch version (7.0, 6.6, or 6.5) that you want to upgrade the switch to and click **Next**.
- Step 3.** Complete the wizard and click **Finish**.

### Note

If some ESXi hosts are incompatible with the selected target version, you should upgrade (or remove) the incompatible hosts or select another distributed switch version.

## Modifying a vSphere Distributed Switch

To use the vSphere Client to configure general properties on an existing vDS, you can use the following procedure:

- Step 1.** Select **Home > Networking**.

- Step 2.** Select the vDS in the navigation pane and navigate to **Configure > Settings > Properties**.
- Step 3.** Click **Edit**.
- Step 4.** Click **General**.
- Step 5.** Change the general setting of the distributed switch, including name, number of uplinks, whether Network I/O Control is enabled or disabled, and description.

## **Configure NetFlow on a vDS**

To configure NetFlow on a vDS, you can use the following procedure:

- Step 1.** In the vSphere Client home page, select the distributed switch in the inventory pane.
- Step 2.** Select **Actions > Settings > Edit NetFlow**.
- Step 3.** Provide the collector IP address and collector port of the NetFlow collector.
- Step 4.** Set an observation domain ID that identifies information related to the switch.
- Step 5.** Optionally, set the switch IP address and provide an IP address if you want to see the information from the distributed switch in the NetFlow collector under a single network device (IP address) instead of under a separate device for each host.
- Step 6.** Optionally, set the Active Flow Export Timeout and Idle Flow Export Timeout options to time values, in seconds, to wait before sending information after the flow is initiated.
- Step 7.** Optionally, to change the portion of data that the switch collects, configure the sampling rate.
- Step 8.** Optionally, to collect data on network activity between virtual machines on the same host, enable **Process Internal Flows Only**.
- Step 9.** Click **OK**.

## **Configuring Advanced vDS Settings**

To configure advanced properties for an existing vDS, you can use the following procedure:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane and navigate to **Configure > Settings > Properties**.
- Step 3.** Click **Edit**.
- Step 4.** Click **Advanced**.
- Step 5.** Set the MTU (in bytes), multicast filtering mode (basic or IGMP/MLD snooping), discovery protocol, and administrator contact.

## Creating and Configuring Distributed Port Groups



You can use the following procedure to add a distributed port group to a vDS to create a network for connecting virtual machines and VMkernel adapters:

- Step 1.** In the vSphere Client, right-click a distributed switch in the inventory pane and select **Distributed Port Group > New Distributed Port Group**.
- Step 2.** In the wizard, provide a name for the new distributed port group or accept the generated name and click **Next**.
- Step 3.** On the **Configure Settings** page, optionally change any of the following properties:
  - **Port Binding:** Choose **Static** or **Ephemeral**.
  - **Port Allocation:** Choose **Elastic** or **Fixed**.
  - **Number of Ports:** Increase or decrease the value from the default (which is 8).
  - **Network Resource Pool:** Select an available pool.
  - **VLAN:** Set VLAN Type to **none**, **VLAN**, **VLAN Trunking**, or **Private VLAN** and provide the corresponding settings.
  - **Advanced:** Select the **Customize Default Policy Configuration** checkbox.

Click **Next**.

**Step 4.** If you selected the **Customize Default Policy Configuration** checkbox in step 4, you can use the following pages to customize policies:

- On the Security page, provide your choices for accepting or rejecting Promiscuous Mode, MAC Address Changes, and Forged Transmits and click **Next**.
- On the Traffic Shaping page, enable ingress traffic shaping or egress traffic shaping, or both. If you enable traffic shaping, you can set the average bandwidth, peak bandwidth, and burst size. Click **Next**.
- On the Teaming and Failover page, optionally set the Load Balancing, Network Failure Detection, Notify Switches, Failback, and Failover Order options. Click **Next**.

**Step 5.** On the Monitoring page, enable or disable NetFlow and click **Next**.

**Step 6.** On the Miscellaneous Settings page, click **Next**.

**Step 7.** On the Ready to Complete page, review the settings and click **Finish**.

## VMkernel Networking

This section describes the procedures for configuring VMkernel networking.

### Configuring and Managing VMkernel Adapters

To create a VMkernel virtual network adapter on a standard switch, you can use the Add Networking wizard, as previously described in the “Creating and Configuring vSphere Standard Switches” section. To open the wizard, you can right-click a host in the inventory pane and select Add Networking. In the wizard, you can choose whether to add the adapter to a new standard switch or to an existing standard switch.

To add a VMkernel adapter to a distributed port group, you can use the following procedure:

**Step 1.** In the vSphere Client, right-click a distributed port group in the inventory pane and select **Add VMkernel Adapters**.

**Step 2.** On the Select Hosts page, click **Attached Hosts** (green plus sign), select the appropriate hosts, and click **Next**.

**Step 3.** Configure the VMkernel adapter IP, MTU, Stack and Available Services settings, as previously described.

**Step 4.** Complete the wizard and click **Finish**.

To view information about the VMkernel adapters in a host, select the host in the inventory pane and select **Configure > Networking > VMkernel Adapters**. To view details, select a specific adapter and examine the **All, Properties, IP Settings, and Policies** tabs. To modify a VMkernel adapter, select the adapter and click **Edit**.

## Configuring TCP/IP Stacks



To view and edit the configuration of existing TCP/IP stacks on a host, you can use the following procedure:

**Step 1.** In the vSphere Client, select the host in the inventory pane and navigate to **Configure > Networking > TCP/IP Configuration**.

**Step 2.** Select any of the stacks in the table, such as **Default, vMotion, Provisioning, or a custom stack**.

**Step 3.** Examine the details pane, which may include **DNS, routing, IPv4/IPv6 routing tables, a control algorithm, and the maximum number of allowed connections**.

**Step 4.** Click **Edit** and use the following pages to modify the selected stack.

**Step 5.** On the **DNS Configuration** page, choose one of the following methods:

- **Obtain Settings Automatically from a VMkernel Network Adapter:** Select an existing VMkernel adapter.
- **Enter Settings Manually:** Provide the host name, domain name, preferred DNS server, alternate DNS server, and search domains.

**Step 6.** On the **Routing** page, edit the VMkernel gateway settings.

**Step 7.** ON the **Name** page, edit the name of the stack.

**Step 8.** On the **Advanced** page, edit the maximum number of connections and the congestion control algorithm.

**Step 9.** Click **OK**.

To create a custom TCP/IP stack, you can use the following command in the ESXi Shell:

[Click here to view code image](#)

```
esxcli network ip netstack add -N="stack_name"
```

After creating a custom stack, you can use the previous procedure to configure the stack. When creating a VMkernel virtual network adapter, you can select any existing custom stack or predefined stack (default, vMotion, or provisioning).

## Configuring and Managing Networking Features

This section describes the procedures for implementing networking features supported by vSphere.

### Configuring Network I/O Control (NIOC)



To guarantee minimum bandwidth for system traffic and for virtual machines, you can enable and configure NIOC. You can enable NIOC on a distributed switch by using the following procedure:

**Step 1.** In the vSphere Client, select the distributed switch in the inventory pane and navigate to **Actions > Settings > Edit Settings**.

**Step 2.** From the Network I/O Control drop-down menu, select **Enable**.

**Step 3.** Click **OK**.

By default, NIOC applies shares to each network traffic type as shown in the following list:

- **Management traffic:** 50 shares
- **Fault Tolerance (FT) traffic:** 50 shares
- **NFS traffic:** 50 shares
- **vSAN traffic:** 50 shares
- **vMotion traffic:** 50 shares



- **vSphere Replication (VR) traffic:** 50 shares
- **vSphere Data Protection backup traffic:** 50 shares
- **Virtual machine traffic:** 100 shares

To configure resource allocation for system traffic, you can use the following procedure:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane.
- Step 2.** On the Configure tab, expand **Resource Allocation**.
- Step 3.** Click **System Traffic**.
- Step 4.** Select the appropriate traffic type and click **Edit**.
- Step 5.** Set the desired values for Shares, Reservation, and Limit.
- Step 6.** In the Reservation text box, enter a value for the minimum bandwidth that must be available for the traffic type.
- Step 7.** In the Limit text box, set the maximum bandwidth that system traffic of the selected type can use.
- Step 8.** Click **OK** to apply the allocation settings.

## Creating a Network Resource Pool

If you enabled NIOC on a distributed switch and reserved bandwidth for the virtual machine feature, then you can create a set of network resource pools and allocate the reserved bandwidth among the pools. The total reservation from the virtual network adapters of the powered-on associated VMs must not exceed the quota of the pool. You can create a network pool by using the following procedure:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane.
- Step 2.** On the **Configure** tab, expand **Resource Allocation**.
- Step 3.** Click **Network Resource Pools**.
- Step 4.** Click the **Add** button.

**Step 5.** Provide a name and a description for the pool.

**Step 6.** Set the reservation quota (in Mbps).

**Step 7.** Click **OK**.

### **Note**

The maximum quota that you can assign to the pool is equal to the aggregated reservation for virtual machine system traffic minus the quotas of the other network resource pools.

After creating a network resource pool, you can assign a distributed port group to the resource pool by using the following procedure:

**Step 1.** In the vSphere Client, right-click a distributed port group in the inventory pane and select **Edit Settings**.

**Step 2.** In the settings, click **General**.

**Step 3.** In the Network Resource Pool drop-down menu, select the network resource pool and click **OK**.

Finally, you can set the network shares, reservation, and limit settings for individual virtual machines that are connected to the distributed port group in a network resource pool by using the following procedure:

**Step 1.** In the vSphere Client, select a virtual machine in the inventory pane and navigate to **Actions > Edit Settings**.

**Step 2.** Expand the **Network Adapter** section of the VM network adapter.

**Step 3.** Either add and configure a new network adapter or select an existing network adapter.

**Step 4.** Configure the network adapter's Shares, Reservation, and Limit settings.

**Step 5.** Click **OK**.

## **Using Private VLANs**

To use private VLANs (PVLANS), you must first define the PVLANS on a vDS.

You can use the following procedure to do so:

- Step 1.** In the vSphere Client, select a distributed switch in the inventory pane and navigate to **Configure > Settings > Private VLAN**.
- Step 2.** Click **Edit**.
- Step 3.** To add a primary VLAN, above Primary VLAN ID click the **plus sign (+)** button.
- Step 4.** To add a secondary VLAN, in the right pane click the **plus sign (+)** button.
- Step 5.** For Secondary VLAN type, select either **Isolated** or **Community**.
- Step 6.** Click **OK**.

After creating the PVLANS, you can use them when assigning the VLAN network policies for distributed port groups and distributed ports, just as you do with standard VLANs.

## Using DirectPath I/O

To allow virtual machines to access a physical NIC using DirectPath I/O, you must first enable DirectPath I/O for the NIC. You can do so with this procedure:

- Step 1.** In the vSphere Client, select the ESXi host in the inventory.
- Step 2.** On the **Configure** tab, expand **Hardware** and click **PCI Devices**.
- Step 3.** Click **Edit**.
- Step 4.** Select the NIC that has a green icon, indicating that it is active and ready. (An orange icon indicates that the device state has changed, and you must reboot the host prior to using the NIC.)
- Step 5.** Click **OK**.

Now you are ready to use the following procedure to configure a virtual machine for DirectPath access to the passthrough NIC:

- Step 1.** In the vSphere Client, locate the virtual machine in the inventory pane.
- Step 2.** Power off the virtual machine.
- Step 3.** Select **Actions > Edit Settings**.

- Step 4.** Select the **Virtual Hardware** tab.
- Step 5.** Select **Memory** and set Limit to **Unlimited**.
- Step 6.** Click **Add New Device** and select **Other Devices > PCI Device**.
- Step 7.** From the New PCI Device drop-down menu, select the appropriate passthrough device and click **OK**.
- Step 8.** Power on the virtual machine.

## Single Root I/O Virtualization (SR-IOV)

To enable a virtual machine to use the capabilities of SR-IOV, you must enable the SR-IOV virtual functions on the host and connect a virtual machine. To enable SR-IOV on a host, you can use the following procedure:

- Step 1.** In the vSphere Client, select the host in the inventory pane.
- Step 2.** Select the **Configure** tab and select **Networking > Physical Adapters**. Examine the SR-IOV property to see which NIXs supports SR-IOV.
- Step 3.** Select the NIC and click **Edit Adapter Settings**.
- Step 4.** Under SR-IOV, set Status to **Enabled**.
- Step 5.** In the Number of Virtual Functions text box, type the number of virtual functions that you want to configure for the adapter.
- Step 6.** Click **OK**.
- Step 7.** Restart the host.

You can use the following vCLI command to examine a host's virtual functions:

```
esxcli network sriovnic
```

To associate a virtual machine with one or more virtual functions as SR-IOV passthrough network adapters, you should first verify the following:

- Virtual functions exist on the host.
- Passthrough networking devices for the virtual functions are active in the host's PCI Devices list.
- The virtual machine compatibility is ESXi 5.5 and later.

- The guest OS is Red Hat Enterprise Linux 6 or later or Windows.

You can use the following procedure to implement SR-IOV for a virtual machine:

- Step 1.** In the vSphere Client, select the virtual machine in the inventory pane.
- Step 2.** Power off the virtual machine.
- Step 3.** Select **Actions > Edit Settings**.
- Step 4.** Select the **Virtual Hardware** tab.
- Step 5.** From the Add New Device drop-down menu, select **Network Adapter**.
- Step 6.** Expand the **New Network** section and connect the virtual machine to a port group. (The virtual NIC does not use this port group for data traffic. The port group is used to identify the networking properties, such as VLAN tagging, to apply on the data traffic.)
- Step 7.** Select **Adapter Type > SR-IOV Passthrough**.
- Step 8.** From the Physical Function drop-down menu, select the physical NIC.
- Step 9.** To allow changes in the MTU of packets from the guest operating system, use the Guest OS MTU Change drop-down menu.
- Step 10.** Expand the **Memory** section, select **Reserve All Guest Memory (All Locked)** and click **OK**.
- Step 11.** Power on the virtual machine.

The host selects a free virtual function from the physical adapter and maps it to the SR-IOV passthrough adapter. The host validates all properties of the virtual machine adapter and the underlying virtual function against the settings of the port group to which the virtual machine belongs.

## Note

The step to reserve all the guest memory is required to allow the I/O memory management unit (IOMMU) and the passthrough device to access the memory using direct memory access (DMA).

Optionally, you can use the virtual switch, port group, or port to set the MTU size, security policy for VF traffic, and VLAN tagging mode.

You can enable SR-IOV with host profiles. In a host profile, expand General System Settings > Kernel Module. Select the appropriate physical function driver and provide the number of virtual functions that you want to provide for each physical function. The details depend on your hardware.

## Configuring and Managing Port Mirroring



To configure port mirroring, you can use the following procedure:

**Step 1.** In the vSphere Client, select a vDS in the inventory pane and navigate to **Configure > Settings > Port Mirroring**.

**Step 3.** Click **New**.

**Step 4.** Select one of the following options for Session Type:

- **Distributed Port Monitoring**
- **Remote Mirroring Source**
- **Remote Mirroring Destination**
- **Encapsulated Remote Mirroring (L3) Source**

Click **Next**.

**Step 5.** Provide the following applicable session properties:

- **Name, Status, Description, and Sampling Rate**
- **Normal I/O on Destination Ports**
- **Mirrored Packet Length**

Click **Next**.

**Step 6.** Identify the traffic source by using the following options:

- **Add Existing Ports from a List:** Click **Select Distributed Ports**, select each port, and click **OK**.
- **Add Existing Ports by Port Number:** Click **Add Distributed Ports**,

enter the port number, and click **OK**.

- **Set Traffic Direction:** Select ingress, egress, or ingress/egress.
- **Specify the Source VLAN:** If you selected a remote mirrored destination, click **Add** and provide a VLAN ID.

Click **Next**.

**Step 7.** Select the destination by using the following information:

- **Select a destination distributed port:** Click either **Select Distributed Ports** or **Add Distributed Ports** to add by port number.
- **Select an uplink:** Select an uplink and click **Add**.
- **Select ports or uplinks:** Select distributed ports and uplinks.
- **Specify IP address:** Click **Add** and provide an IP address.

Click **Next**.

You can use a similar procedure to edit port mirroring sessions; however, in step 2, you should select a session and click **Edit**. To remove a session, click **Remove**.

## **Configuring and Managing Link Aggregation Groups (LAGs)**

This section provides information on configuring ESXi hosts to connect to physical switches using dynamic link aggregation. It involves creating link aggregation groups (LAGs) on distributed switches using host NICs that are connected to LACP port channels on physical switches.

To get started, you must create a LACP port channel on physical switches for each LAG on each participating ESXi host. The steps are hardware dependent and not covered here. Before you create the port channel, you should consider the following requirements:

- The number of ports in each port channel must match the number of physical NICs that will be aggregated on the host (a minimum of two).
- The same hashing algorithm must be used for the port channel and the associated LAG on the vDS.
- All the NICs in a LAG must be configured with the same speed and

duplexing.

Before creating the LAG on a vDS in vSphere 7.0, you should address the following requirements:

- An LACP port channel must be available on a physical switch and is configured to support the host, including the appropriate number of ports, speed, duplexing, and hashing (load balancing) algorithm.
- The vDS must be Version 6.5 or later.
- Enhanced LACP must be supported on the distributed switch, so use the vSphere Client to select the vDS in the inventory pane, navigate to Summary > Features, and verify that Link Aggregation Control Protocol is set to Enhanced Support.

Enhanced LACP support for vDS supports the following load-balancing modes (hashing algorithms):

- Destination IP address
- Destination IP address and TCP/UDP port
- Destination IP address and VLAN
- Destination IP address, TCP/UDP port, and VLAN
- Destination MAC address
- Destination TCP/UDP port
- Source IP address
- Source IP address and TCP/UDP port
- Source IP address and VLAN
- Source IP address, TCP/UDP port, and VLAN
- Source MAC address
- Source TCP/UDP port
- Source and destination IP address
- Source and destination IP address and TCP/UDP port
- Source and destination IP address and VLAN



- Source and destination IP address, TCP/UDP port, and VLAN
- Source and destination MAC address
- Source and destination TCP/UDP port
- Source port ID
- VLAN

To change the LACP feature state from Basic Support to Enhanced Support, you can use the following procedure:

- Step 1.** In the vSphere Client, select the vDS in the inventory pane and navigate to **Summary > Features**.
- Step 2.** Verify that Link Aggregation Control Protocol is set to **Basic Support**.
- Step 3.** Select **Actions > Upgrade**.
- Step 4.** Select **Enhance LACP Support**.
- Step 5.** Navigate through the wizard to verify the port group accessibility and LACP configuration prerequisites.
- Step 6.** If the prerequisite verifications passed, complete the wizard and click **Finish**.
- Step 7.** Verify that Link Aggregation Control Protocol is set to **Enhanced Support**.
- Step 8.** Navigate to the **Configure** tab and verify that **LACP** appears in the Settings section.

You can use the following procedure to create a LAG:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane, navigate to **Configure > Settings**, and select **LACP**.
- Step 2.** Click the **New Link Aggregation Group** icon.
- Step 3.** Provide the following information:
  - a. Name for the LAG
  - b. Number of ports to include in the LAG (which must match the number of ports in the LACP port channel)
  - c. LACP negotiating mode (If the corresponding LACP-enabled

physical switch ports are set to **Active**, then set the LAG's mode to **Passive** and vice versa.)

- d. Load-balancing mode (which must match the load-balancing mode in the LACP port channel)

**Step 4.** If you want to override the VLAN and NetFlow policies for each individual uplink port, set the VLAN and NetFlow policies for the LAG.

**Step 5.** Click **OK**.

The LAG is now available for use. It appears as unused in the teaming and failover settings of the distributed port groups. To use the LAG as the uplink for a distributed port group, you can use the following procedure:

**Step 1.** In the vSphere Client, select the distributed switch in the inventory pane.

**Step 2.** Use the following steps to set the LAG as **Standby** for the appropriate distributed port groups:

- a. Select **Actions > Distributed Port Group > Manage Distributed Port Groups**.
- b. Select **Teaming and Failover** and click **Next**.
- c. Select the port groups where you want to use the LAG.
- d. In Failover Order, select the LAG and use the arrow keys to move it to the **Standby Uplinks** list.
- e. Complete the wizard and click **Finish**.

**Step 3.** Use the following steps to assign the host's physical NICs to the LAG:

- a. Select **Actions > Add and Manage Hosts > Manage Host Networking**.
- b. Select the host and click **Next**.
- c. On the Select Network Adapter Tasks page, select **Manage Physical Adapters** and click **Next**.
- d. On the Manage Physical Adapters page, select a NIC and click **Assign an Uplink**.
- e. Select a LAG port and click **OK**.

- f. Repeat the steps d and e for each NIC you want to include in the LAG.
  - g. Complete the wizard and click **Finish**.
- Step 4.** Use the following steps to activate the LAG for the appropriate distributed port groups:
- a. Select **Actions > Distributed Port Group > Manage Distributed Port Groups**.
  - b. Select **Teaming and Failover** and click **Next**.
  - c. Select the port groups where you previously set the LAG for standby.
  - d. In Failover Order, select the LAG and use the arrow keys to move it to the **Active Uplinks** list.
  - e. Likewise, move each standalone uplink to the **Unused** list.
  - f. Complete the wizard and click **Finish**.

## Managing Host Networking with vDS

This section describes the steps for managing host networking using vDS.

### Adding Hosts to a vDS



To prepare for adding ESXi hosts to a vDS, you should do the following.:

- Step 1.** Create distributed port groups for virtual machines.
- Step 2.** Create distributed port groups for VMkernel networking, such as management, vMotion, and Fault Tolerance.
- Step 3.** Configure uplinks on the distributed switch for physical NICs that you want to connect to the switch.
- Step 4.** Configure the vDS to support the hosts' VMs. For example, set the vDS's MTU and discovery protocols.

You can use the Add and Manage Hosts wizard to add multiple hosts at a time. To do so, follow these steps:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane and navigate to **Actions > Add and Manage Hosts**.
  - Step 2.** On the Select Task page, select **Add Hosts** and click **Next**.
  - Step 3.** On the Select Hosts page, click **New Hosts**.
  - Step 4.** Select the appropriate hosts in your data center, click **OK**, and then click **Next**.
  - Step 5.** On the next page, select the tasks for configuring network adapters to the distributed switch and click **Next**.
  - Step 6.** On the Manage Physical Network Adapters page, do the following:
    - a.** From the On Other Switches/Unclaimed list, select an unclaimed physical NIC or select a NIC to migrate from another virtual switch.
    - b.** Click **Assign Uplink**.
    - c.** Select an uplink and click **OK**.
  - Step 7.** Click **Next**.
  - Step 8.** On the next page, do the following:
    - a.** Select a VMkernel adapter and click **Assign Port Group**.
    - b.** Select a distributed port group and click **OK**.
  - Step 9.** Click **Next**.
  - Step 10.** Optionally, on the Migrate VM Networking page, select the checkbox **Migrate Virtual Machine Networking** to configure virtual machine networking and then complete the following steps:
    - a.** To connect all network adapters of a virtual machine to a distributed port group, select the virtual machine or select an individual network adapter to connect only that adapter.
    - b.** Click **Assign Port Group**.
    - c.** Select a distributed port group from the list, click **OK**, and then click **Next**.
  - Step 11.** Click **Finish**
-

## Note

In the vSphere inventory, the hosts that you add must reside in the same data center as the vDS.

## Managing Host Physical Network Adapters on a vDS

You can configure the physical NICs for multiple hosts connected to a single vDS at the same time. For consistent network configuration, you can use the following procedure to assign the same physical NIC from each host to the same vDS uplink:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane and navigate to **Actions > Add and Manage Hosts**.
- Step 2.** In the wizard, select **Manage Host Networking** and click **Next**.
- Step 3.** On the Select Hosts page, click **Attached Hosts** (green plus sign) and select the appropriate hosts. Click **Next**.
- Step 4.** When you see all the ESXi hosts you selected on the Select Hosts page, click **Next**.
- Step 5.** On the Manage Physical Adapters page, select a physical NIC from the On Other Switches/Unclaimed list to assign an uplink to the adapter.
- Step 6.** Click **Assign Uplink**.
- Step 7.** Select an uplink or select **Auto-assign** and click **OK**.
- Step 8.** Click **Next**.
- Step 9.** Continue through the wizard. Optionally, you can use the wizard to migrate VMkernel adapters and virtual machines.
- Step 10.** Click **Finish**.

## Migrating VMkernel Network Adapters to a vDS

You can use the following procedure to migrate VMkernel network adapters to a vDS:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane and navigate to **Actions > Add and Manage Hosts**.

- Step 2.** In the wizard, select **Manage Host Networking** and click **Next**.
- Step 3.** On the Select Hosts page, click the **Attached Hosts** button (green plus sign) and select the appropriate hosts. Click **Next**.
- Step 4.** Click **Next**.
- Step 5.** Continue through the wizard. Optionally, you can use the wizard to make changes and migrate virtual machines.
- Step 6.** On the Manage VMkernel Adapters page, select a VMkernel adapter from each host, click **Assign Port Group**, click **OK**, and click **Next**.
- Step 7.** Click **Finish**.

### Note

If you migrate or create VMkernel adapters for iSCSI, verify that the teaming and failover policy of the target distributed port group meets the requirements for iSCSI:

- Verify that only one uplink is active, the standby list is empty, and the rest of the uplinks are unused.
- Verify that only one physical NIC per host is assigned to the active uplink.

## Removing Hosts from a vDS

Prior to removing an ESXi host for a vDS, you should migrate all the host's virtual machines, VMkernel adapters, and physical NICs from the vDS:

- Step 1.** In the vSphere Client, select the vDS in the inventory pane and navigate to **Actions > Add and Manage Hosts**.
- Step 2.** Select **Remove Hosts** and click **Next**.
- Step 3.** Select the hosts you want to remove and click **Next**.
- Step 4.** Click **Finish**.

## Migrating Virtual Machines to a vDS

If you want to relocate a virtual machine to a vDS, you can edit each virtual machine and change the network to which the vNICs are connected. To migrate a set of virtual machines from multiple networks to the distributed port groups of a specific vDS, you can use the following procedure:

- Step 1.** In the vSphere Client, select the distributed switch in the inventory pane and navigate to **Actions > Add and Manage Hosts**.
- Step 2.** In the wizard, select **Manage Host Networking** and click **Next**.
- Step 3.** On the Select Hosts page, click the **Attached Hosts** button (green plus sign) and select the appropriate hosts. Click **Next**.
- Step 4.** Click the **Next** button to navigate through the wizard until you reach the Migrate VM Networking page.
- Step 5.** Check the **Migrate Virtual Machine Networking** box and examine the list of virtual machines that appear.
- Step 6.** Select a virtual machine, click **Assign Port Group**, and select the distributed port group where the virtual machine should connect.
- Step 7.** Repeat steps 5 and 6 for each virtual machine that you want to migrate and then click **Next**.
- Step 8.** On the Ready to Complete page click **Finish**.

To migrate a set of virtual machines from one network (distributed port group or standard port group) to another, you can use the following procedure:

- Step 1.** In the vSphere Client, select the network in the inventory pane and navigate to **Actions > Migrate VMs to Another Network**.
- Step 2.** In the wizard, select **Destination Network**, click **OK**, and click **Next**.
- Step 3.** On the next page, select the virtual machines that you want to migrate and click **Next**.
- Step 4.** On the Ready to Complete page, click **Finish**.

## **Monitoring the State of Ports in a Distributed Port Group**

You can use the following procedure to examine the ports in a distributed port group:

- Step 1.** In the vSphere Client, select a distributed port group in the inventory

pane.

- Step 2.** Click the **Ports** tab and examine the list of ports that exist in the port group.
- Step 3.** To sort the rows by the data in a specific column (such as Port ID), click on the header for the appropriate column.
- Step 4.** To filter the list using data in the column, click on the **Filter** icon in the column and enter a value. For example, to search for a port for a virtual machine with MAC address ending with 83, click on **Runtime MAC Address** and enter **83**.
- Step 5.** To examine details for a port, click on the port and examine the details pane beneath the list of ports. Click the **Network Connection**, **Policies**, and **Traffic Filtering and Marking** tabs to see related details. Click on **Statistics** to view statistics on the amount of data and number of packets flowing into and out of the port.

## Using the vDS Health Check

You can choose whether and when to use the vDS Health Check feature. For example, you could choose to enable it only temporarily, as you plan and make vDS configuration changes. By default it is disabled. You can use the following procedure to enable or disable Health Check for a vDS:

- Step 1.** In the vSphere Client, select the vDS in the inventory pane and navigate to **Configure > Setting > Health Check**.
- Step 2.** Click the **Edit** button.
- Step 3.** For each for the following available health checks, choose **Enabled** or **Disabled** and set the interval, in minutes:
  - VLAN and MTU
  - Teaming and Failover
- Step 4.** Click **OK**.

You can view a vDS health check by using this procedure:

- Step 1.** In the vSphere Client, select the vDS in the inventory pane and navigate to **Monitor > Health**.



- Step 2.** On the Host Member Health Status page, examine the overall, VLAN, MTU and teaming and failover health statuses of the hosts connected to the switch.
- Step 3.** For more detail, select any host in the list and examine the details pane. In the details pane, select the **VLAN**, **MTU**, or **Teaming and Failover** tabs to get details on specific health checks.

## Networking Policies and Advanced Features

To view the policies applied to a distributed port group, select the port group in the inventory pane, select **Configure > Policies**, and examine the applied security, ingress traffic shaping, egress traffic shaping, VLAN, and teaming and failover policies. To change the policies applied to a distributed port group, you can use the following procedure:

- Step 1.** In the vSphere Client, select the distributed port group in the inventory pane, navigate to **Configure > Policies**, and click the **Edit** button.
- Step 2.** On the left side, select any of the following policies and use the associated page to make changes:
- **Security:** On the Security page, provide your choices for accepting or rejecting promiscuous mode, MAC address changes, and forged transmits.
  - **Traffic Shaping:** On the Traffic Shaping page, you can enable Ingress Traffic Shaping or Egress Traffic Shaping or both. If you enable traffic shaping, you can set the Average Bandwidth, Peak Bandwidth, and Burst Size options.
  - **Teaming and Failover:** On the Teaming and Failover page, you can set the Load Balancing, Network Failure Detection, Notify Switches, Failback, and Failover Order options.
  - **Monitoring:** On the Monitoring page, you can enable or disable NetFlow.
- Step 3.** Select any of the following categories and use the associated page to modify switch settings:
- **General:** On the General page, modify settings such as Name, Port Binding, Port Allocation, Number of Ports, Network Resource Pool,

and Description.

- **Advanced:** On the Advanced page, choose **Allowed** or **Disabled** for each policy to indicate whether the policy can be overridden at the port level.
- **Miscellaneous:** On the Miscellaneous page, you can set Block All Ports to **yes** or **no**.

**Step 4.** Click **OK**.

For example, to configure a traffic filtering and marking policy for a vDS, you can use the following procedure:

**Step 1.** In the vSphere Client, select a distributed port group or uplink port group in the inventory pane and navigate to **Configure > Settings > Traffic Filtering and Marking**.

**Step 2.** Click **Enable and Reorder**.

**Step 3.** Click **Enable All Traffic Rules**.

**Step 4.** Click **OK**.

**Step 5.** To create a rule to mark traffic, use the following steps:

- a. Click **Add**.
- b. Select **Action > Tag** and configure the tag either for **CoS** or **DSCP**.
- c. Set the traffic direction and specify the traffic qualifiers (system, MAC, or IP).
- d. Optionally, click the **Enable Qualifier** checkbox and provide information for qualifying the packets to mark. You can use the following tabs to qualify data:
  - **IP:** Identify packets by source and destination addresses and ports.
  - **MAC:** Identify packets by source and destination addresses and by VLAN.
  - **System Traffic:** Identify data by system traffic type (such as management or vMotion).
- e. Click **OK**.

**Step 6.** To create a rule to filter traffic, repeat step 5 but configure the action to either allow traffic to pass or to block the traffic.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

## Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 9-3](#) lists these key topics and the page number on which each is found.



**Table 9-3** Key Topics for [Chapter 9](#)

Key Topic Element	Description	Page Number
List	Creating a vSphere Standard Switch	<a href="#">330</a>
List	Creating and configuring distributed port groups	<a href="#">337</a>
List	Configuring TCP/IP stacks	<a href="#">339</a>
List	Configuring Network I/O Control	<a href="#">340</a>
List	Configuring and managing port mirroring	<a href="#">345</a>
Section	Adding hosts to a vDS	<a href="#">350</a>

## Complete Tables and Lists from Memory

Print a copy of [Appendix B](#), “[Memory Tables](#)” (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C](#), “[Memory Tables Answer Key](#)” (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

No new terms are defined in this chapter.

## Review Questions

1. You want to use VLAN guest tagging with your vSphere Standard Switch. What setting should you make on the standard port group?
  - a. Set VLAN ID to 0.
  - b. Set VLAN ID to 4095.
  - c. Set VLAN Type to Trunking.
  - d. Set VLAN Type to Guest Tagging.
2. You are preparing to upgrade a vDS from Version 6.x to 7.0. What step should you take prior to upgrading?
  - a. Copy the vDS.
  - b. Back up vCenter Server.
  - c. Export the vDS configuration, including the distributed port group configuration.
  - d. Export the vDS configuration, excluding the distributed port group configuration.
3. You enabled NIOC, reserved virtual machine system traffic, and created a network resource pool. Which of the following steps do you need to take to allow a virtual machine to use the network resource pool?
  - a. Edit the virtual machine and set the network resources allocation policy.
  - b. Add the virtual machine to the resource pool.
  - c. Assign the network resource pool to the distributed port group where the virtual machines are connected.
  - d. In the inventory pane, drag and drop the virtual machine onto the network resource pool.
4. You are creating a VMkernel virtual adapter for vMotion traffic.

Which of the following is not a valid option?

- a.** In a standard switch, assign the adapter to the vMotion stack.
  - b.** In a distributed switch, assign the adapter to the vMotion stack.
  - c.** In a standard switch, assign the adapter to the default stack.
  - d.** In a distributed switch, assign the adapter to the provisioning stack.
5. You want to enable NetFlow in a distributed port group. Which of the following steps should you take?
- a.** Change the distributed port group's monitoring policy.
  - b.** Change the distributed port group's security policy.
  - c.** In the distributed port group's Advanced settings, set NetFlow to Enable.
  - d.** Enable port mirroring on the distributed port group.

# **Part III**

# **vSphere Management and Optimization**

## Chapter 10

# Managing and Monitoring Clusters and Resources

This chapter covers the following topics:

- [Creating and Configuring a vSphere Cluster](#)
- [Creating and Configuring a vSphere DRS Cluster](#)
- [Creating and Configuring a vSphere HA Cluster](#)
- [Monitoring and Managing vSphere Resources](#)
- [Events, Alarms, and Automated Actions](#)
- [Logging in vSphere](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.6, 1.6.2, 1.6.3, 1.6.4, 1.6.4.1, 4.9, 5.1, 5.1.1, 5.2, 5.3, 7.5, and 7.12.

This chapter introduces vSphere 6.7, describes its major components, and identifies its requirements.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 10-1](#) 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 and Review Questions.”](#)

**Table 10-1** “Do I Know This Already?” Section-to-Question Mapping

---

<b>Foundation Topics Section</b>	<b>Questions</b>
Creating and Configuring a vSphere Cluster	1
Creating and Configuring a vSphere DRS Cluster	2, 3
Creating and Configuring a vSphere HA Cluster	4, 5
Monitoring and Managing vSphere Resources	6–8
Events, Alarms, and Automated Actions	9
Logging in vSphere	10

1. In a cluster that you initially created using Quickstart and for which you chose the option Configure Network Settings Later, you now want to add a host. Which of the following is a true statement?
  - a. You cannot use Quickstart to add more hosts to the cluster.
  - b. You can use Quickstart to add hosts to the cluster and configure the host networking.
  - c. You can use Quickstart to add hosts to the cluster but must manually configure the host networking.
  - d. You can edit the cluster and change the Configure Networking Settings Later option.
  
2. You are creating a resource pool in a DRS cluster. Which of the following statements is not true?
  - a. When you create a child resource pool, the system applies admission control.
  - b. If you choose Scale Descendant's Shares, child pools use scalable share.
  - c. The default CPU reservation is 0.
  - d. The default memory limit is 0.
  
3. You are configuring a vSphere HA cluster. Which of the following is not a valid setting for Define Host Failover Capacity?
  - a. Standby
  - b. Cluster Resource Percentage
  - c. Slot Policy (powered-on VMs)



- d.** Dedicated Failover Hosts
4. You want to configure VMCP in a vSphere cluster. Which of the following settings is valid?
- a.** In the vSphere HA settings, select Failures and Responses > Datastore with PDL and choose Power Off and Restart VMs–Conservative Restart Policy.
  - b.** In the vSphere HA settings, select Failures and Responses > Datastore with PDL and choose Power Off and Restart VMs.
  - c.** In the vSphere DRS settings, select Failures and Responses > Datastore with APD and choose Power Off and Restart VMs–Conservative Restart Policy.
  - d.** In the vSphere DRS settings, select Failures and Responses > Datastore with APD and choose Power Off and Restart VMs–Aggressive Restart Policy.
5. You are configuring a vSphere HA cluster and want to configure proactive HA. Which of the following is not a requirement?
- a.** Host.Config.Quarantine and Host.Config.Maintenance privileges
  - b.** A vendor-supplied vSphere Client plug-in
  - c.** A VMware-supplied plug-in
  - d.** vSphere DRS
6. You are experiencing poor performance for an application in a virtual machine. You learn from guest OS software and the vSphere client performance charts that the guest OS is paging. Which of the following is likely to fix the problem?
- a.** Increase the memory in the ESXi host.
  - b.** Increase the memory size of the virtual machine.
  - c.** Migrate the virtual machine to a host that has plenty of free memory.
  - d.** Reserve all of the virtual machine’s memory.
7. You are configuring virtual disks for the virtual machines in your vSphere environment. Which provisioning type is the best choice

when you care more about optimizing the space usage than about performance or availability risk?

- a. Thin
  - b. Thick eager zeroed
  - c. Thick lazy zeroed
  - d. Sparse
8. You are using ESXTOP to analyze vSphere performance. Which of the following statistics is the best indicator of some resource contention?
- a. %USED
  - b. %DRPTX
  - c. OVHD
  - d. READ/s
9. You are configuring alarms in your vSphere environment. Which of the following is not a valid event type?
- a. Error
  - b. Warning
  - c. Information
  - d. Audit
10. You are examining vSphere logs. Which of the following logs contains data about the agent that manages and configures the ESXi host?
- a. /var/log/vmkernel.log
  - b. /var/log/vpxa.log
  - c. /var/log/hostd.log
  - d. /var/log/vmksummary.log

## Foundation Topics

# Creating and Configuring a vSphere Cluster

By using the vSphere Client, you can create a vSphere cluster and use its Quickstart feature to configure the cluster. You can configure DRS, vSphere HA, and EVC on the cluster, as described in this chapter. You can also configure vSAN on the cluster, as described in [Chapter 11, “Managing Storage.”](#)

## Creating a Cluster

To create a vSphere cluster that you plan to configure using Quickstart, you should ensure that the hosts have the same ESXi version and patch level. If you are adding hosts to the vCenter Server inventory, you need the credentials for the root user account for the hosts. You must have the Host.Inventory.Create Cluster privilege. To create a cluster that you manage with a single image, verify that you have a supported ESXi 7.0 or later image available in the vSphere Lifecycle Manager depot. You can use the following procedure to create the cluster:

**Step 1.** In the vSphere Client, right-click a data center in the inventory pane and select **New Cluster**.

**Step 2.** Enter a name for the cluster.

**Step 3.** Optionally, for each of the following services, slide the switch to the right to enable the service:

- DRSv
- Sphere HA
- vSAN

If you enable DRS, you can optionally change its automation setting (default is Fully Automated using Threshold Level 3).

**Step 4.** Optionally, to create a cluster that you manage with a single image, select **Manage All Hosts in the Cluster with a Single Image**, and then do the following:

- Select an ESXi version (image) from the drop-down menu.
- Optionally, select options from the Vendor Addon and Vendor Addon Version drop-down menus.

**Step 5.** Click **OK**.

## Configuring a Cluster with Quickstart

To modify an existing cluster, you can select the cluster in the inventory pane and click **Configure > Configuration > Quickstart**. On the Quickstart page are three cards: Cluster Basics, Add Hosts, and Configure Cluster. To change the name and the enabled cluster services, click **Cluster Basics > Edit**. To add a host to a cluster, you can use the following procedure:



- Step 1.** In the vSphere Client, select a cluster in the inventory pane and navigate to **Configure > Configuration > Quickstart > Add Hosts > Add**.
- Step 2.** Click **New Hosts > Add** and provide the name (or IP address) and credentials for each host that you want to add that is not already in the vCenter Server inventory.
- Step 3.** Optionally, select the **Use the Same Credentials for All Hosts** option.
- Step 4.** Click **Existing Hosts > Add** and select each host that you want to add that is already in the vCenter Server inventory.
- Step 5.** Click **Next**.
- Step 6.** On the Host Summary page, click **Next**.
- Step 7.** On the Ready to Complete page, click **Finish**.
- Step 8.** Monitor the progress under Recent Tasks, where you can see any errors.
- Step 9.** When the task is complete, you can view the number of hosts and the health on the **Add Hosts** card. Optionally, select **Re-validate**.
- Step 10.** Use the inventory pane to verify that the hosts are attached to the cluster and are in maintenance mode.

To configure cluster settings and host networking in a cluster, you can use the following procedure:

- Step 1.** In the vSphere Client, select a cluster in the inventory pane and navigate to **Configure > Configuration > Quickstart**.
- Step 2.** Optionally, if you want to configure the cluster manually, click **Skip Quickstart**, which is irreversible. Otherwise, continue with the following steps to use Quickstart to configure the cluster.

**Step 3.** Click **Configure Cluster > Configure**.

**Step 4.** On the Distributed Switches page, you can either select the irreversible option **Configure Networking Settings Later** or use the following steps to configure the cluster networking:

- a. Specify the number of distributed switches to create (up to three).
- b. Enter a unique name for each distributed switch. Alternatively, click **Use Existing** and select an existing compatible distributed switch and distributed port group.
- c. To set up the vMotion network, select a distributed switch and assign a new port group to it.
- d. In the Physical Adapters section, for each physical network adapter, assign a distributed switch name from the drop-down menu. Ensure that each new distributed switch is assigned to at least one physical adapter. For any existing distributed switch, to avoid an error, select the physical adapter that is currently mapped to the switch.
- e. Click **Next**.
- f. If you enabled the vSphere DRS feature during cluster creation, in the vMotion Traffic page that appears, provide the VLAN ID, protocol type, and IP configuration.

**Step 5.** Click **Next**.

**Step 6.** In the Advanced Options page, configure the following options:

- a. If you enabled vSphere HA during cluster creation, use the options in the High Availability section to enable or disable host failure monitoring, virtual machine monitoring, and admission control. For admission control, you can specify the number of hosts for failover capacity.
- b. If you enabled vSphere DRS during cluster creation, use the options in the Distributed Resource Scheduler section to set the automation level and migration threshold.
- c. In the Host Options section, set the lockdown mode and enter an NTP server address.
- d. Optionally, in the Enhanced vMotion Capability section, use the options to enable EVC and select a mode.

**Step 7.** Click **Next**.

**Step 8.** On the Ready to Complete page that appears, review the settings and click **Finish**.

You can extend a cluster by adding more hosts. If you initially selected the Skip Quickstart option, then you should add hosts manually. If you previously used Quickstart but selected Configure Networking Settings Later, you can add hosts by using Quickstart but must manually configure the host networking. To extend a cluster, you can use the following procedure:

**Step 1.** In the vSphere Client, right-click a configured cluster in the inventory pane and select **Add Hosts**.

**Step 2.** In the wizard, select hosts from the vCenter Server inventory and add new hosts (by providing names and credentials) to include in the cluster.

**Step 3.** On the Ready to Complete page, click **Finish**.

**Step 5.** On the Configure Hosts card of the Extend Cluster Guide page that appears, select **Configure**. If you previously used Quickstart to configure the host networking, the vMotion Traffic page appears. Provide the VLAN ID, protocol type, and IP configuration. Then a pop-up window appears, informing you that the configuration for the hosts that exist in the cluster is applied to the newly added hosts.

**Step 6.** Click **Continue**.

After successful validation, the Configure button in the Configure Hosts card becomes inactive, and the Re-validate button is available.

If you enable DRS, the default Automation Level setting is Fully Automated and the default Threshold setting is 3. If you enable HA, the default values are Host Monitoring and Admission Control Are Enabled and VM Monitoring Is Disabled. You can override the default values later in the workflow.

If you select an image for managing all the hosts in the cluster, you can later edit the image specification on the Updates tab. If you do not choose an image to manage hosts, you must manage the cluster by using baselines and baseline groups. You can switch from using baselines to using images later.

Starting with vSphere 7.0, you can use vSphere Lifecycle Manager to upgrade and update the hosts in a cluster. A vSphere Lifecycle Manager image is a

combination of vSphere software, driver software, and firmware for specific host hardware. You can assign an image to a cluster used to control the software set to be installed on the hosts, including the ESXi version, additional VMware-provided software, and vendor software, such as firmware and drivers.

The image that you define during cluster creation is not immediately applied to the hosts. If you do not set up an image for a cluster, the cluster uses baselines and baseline groups. For more information about using images and baselines to manage hosts in clusters, see the *Managing Host and Cluster Lifecycle* documentation.

## **EVC Mode**

As previously described, you can configure EVC by using Quickstart > Configure Cluster. You can also configure EVC directly in the cluster settings. You can set VMware EVC to Disable EVC, Enable EVC for AMD Hosts, or Enable EVC for Intel Hosts.

If you choose Enable EVC for AMD Hosts, you can set the mode to one of the options listed in [Table 4-3](#) in [Chapter 4](#), “[Clusters and High Availability](#).”

If you choose Enable EVC for Intel Hosts, you can set the mode to one of the options listed in [Table 4-2](#) in [Chapter 4](#).

To view the EVC modes for all of a cluster’s virtual machines in the vSphere Client, you can select a cluster, navigate to its VMs tab, and select Show/Hide Columns > EVC Mode.

## **Creating and Configuring a vSphere DRS Cluster**

This section describes how to create and configure a vSphere DRS cluster.

### **Creating a vSphere DRS Cluster**

To create a vSphere DRS cluster, follow the procedure in the section “Creating a Cluster,” earlier in this chapter, and ensure that you choose to enable the DRS service. Use the information in the rest of this section to configure the DRS cluster.

### **Creating a Resource Pool**

You can use the following procedure to create a child resource pool in a DRS

cluster:

- Step 1.** In the vSphere Client, navigate to **Hosts and Clusters**, right-click a DRS cluster in the inventory, and select **New Resource Pool**.
- Step 2.** Provide a name for the resource pool.
- Step 3.** Optionally, select the **Scale Descendant's Shares** checkbox to enable scalable shares. (Enabling this option causes any child resource pools to use scalable shares, which scale dynamically when virtual machines are added and removed.)
- Step 4.** Optionally, set CPU and Memory Shares to **Low**, **Normal**, **High**, or **Custom**. If you select custom, enter a numeric value.
- Step 5.** Optionally, set CPU and Memory Reservation to a numeric value (the default is 0) and a unit of measure (MB, GB, MHz, or GHz).
- Step 6.** Optionally, set CPU and Memory Limit to a numeric value (the default is Unlimited) and a unit of measure (MB, GB, MHz, or GHz).
- Step 7.** Optionally, set CPU and Memory Reservation Type to **Expandable**.
- Step 8.** Click **OK**.

### Note

When you create a child resource pool, the vSphere Client prompts you for resource pool attribute information. The system uses admission control to ensure that you do not allocate resources that are not available. If you choose Scale Descendant's Shares, each descendant pool will also use scalable shares. You cannot change this behavior for each child pool.

## Configuring Advanced DRS Options

This section describes how to configure some advanced options for vSphere DRS.

### Creating Affinity/Anti-Affinity Rules

[Table 10-2](#) provides some common use cases for VM-to-VM affinity and anti-affinity rules.





**Table 10-2** Use Cases for VM-to-VM Rules

Use Case	Rule Details
To improve application and communication performance for a multinode application.	Use VM-to-VM affinity rules to ensure that sets of virtual machines that engage in significant data exchange reside on the same host, such that the data transfer occurs within the host system hardware and does not traverse the physical network infrastructure.
To improve application availability for a multinode application	Use VM-to-VM anti-affinity rules to ensure that sets of peer virtual machines reside on separate hosts, such that the failure of a single host does not result in the failure of all the peer application nodes.

You can use the following procedure to create a VM-to-VM affinity or anti-affinity rule:

- Step 1.** Browse to the cluster in the vSphere Client.
- Step 2.** Navigate to **Configure > VM/Host Rules** and click **Add**.
- Step 3.** In the Create VM/Host Rule dialog box, type a name for the rule.
- Step 4.** From the Type drop-down menu, select either **Keep Virtual Machines Together** (affinity) or **Separate Virtual Machines** (anti-affinity).
- Step 5.** Click **Add**.
- Step 6.** Select at least two virtual machines to which the rule will apply and click **OK**.
- Step 7.** Click **OK**.

## Configuring Predictive DRS

To configure predictive DRS, you can use the following procedure:

- Step 1.** In the vRealize Operations (vROps) GUI, select the appropriate vCenter Server adapter instance and choose **Advanced Settings**.

- Step 2.** Set Provide Data to vSphere Predictive DRS to **True**.
- Step 3.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Cluster > Services > vSphere DRS > Edit**.
- Step 4.** Check the **Predictive DRS** checkbox.

## Creating and Configuring a vSphere HA Cluster

This section describes how to create and configure a vSphere HA cluster.

### Creating a vSphere HA Cluster

To create a vSphere HA cluster, follow the procedure in the section “Creating a Cluster,” earlier in this chapter, and ensure that you choose to enable the vSphere HA service. Use the information in the rest of this section to configure the vSphere HA cluster.

### Configuring Advanced vSphere HA Options

You can use the following procedure to add vSphere HA advanced options, as described in [Table 4-9](#):

- Step 1.** In the vSphere Client, select a vSphere HA cluster in the inventory pane and navigate to **Configure > vSphere Availability > Edit > Advanced Options**.
- Step 2.** Click **Add**.
- Step 3.** Enter the name of the advanced option and the value.
- Step 4.** Click **OK**.

### Configuring vSphere HA Admission Control



To configure admission control for a vSphere HA cluster, you can use the following procedure:

- Step 1.** In the vSphere Client, select the vSphere HA cluster in the inventory pane and navigate to **Configure > vSphere Availability > Edit**.
- Step 2.** Click **Admission Control** and set Host Failures Cluster Tolerates to the

maximum number of host failures you want the cluster to support.

**Step 3.** Select one of the following options for Define Host Failover Capacity By, as described in [Table 4-8](#) in [Chapter 4](#):

- **Cluster Resource Percentage**
- **Slot Policy** (powered-on VMs)
- **Dedicated failover hosts**
- **Disabled** (disables admission control)

**Step 4.** Optionally, set Performance Degradation VMs Tolerate to a percentage.

**Step 5.** Click **OK**.

## Configuring VMCP

To configure Virtual Machine Component Protection (VMCP) in a vSphere HA cluster, you can use the following procedure:

**Step 1.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Configure > vSphere Availability > Edit**.

**Step 2.** Select **Failures and Responses > Datastore with PDL** and choose one of the following:

- **Issue Events**
- **Power Off and Restart VMs**

**Step 3.** Select **Failures and Responses > Datastore with APD** and choose one of the following:

- **Issue Events**
- **Power Off and Restart VMs–Conservative Restart Policy**
- **Power Off and Restart VMs–Aggressive Restart Policy**

## Configuring Virtual Machine and Application Monitoring

You can use the following procedure to turn on and configure virtual machine and application monitoring in a vSphere HA cluster:

**Step 1.** In the vSphere Client, select the vSphere HA cluster in the inventory

pane and navigate to **Configure > vSphere Availability > Edit**.

**Step 2.** Select **Failures and Responses > VM Monitoring**.

**Step 3.** Select **VM Monitoring** to turn on VMware Tools heartbeats.

**Step 4.** Select **Application Monitoring** to turn on application heartbeats.

**Step 5.** To set the heartbeat monitoring sensitivity, move the slider between Low and High or select **Custom** and provide a custom value.

**Step 6.** Click **OK**.

## Configuring Proactive HA



To get started with implementing Proactive HA, you need to install a supported vendor-supplied vSphere Client plug-in and register the proactive HA provider. When you turn on proactive HA in a cluster, you can select from the list of providers for installed plug-ins that are monitoring every host in the cluster. You can use the following procedure to configure proactive HA in a cluster:

**Step 1.** Ensure that the following prerequisites are met:

- vSphere HA and DRS are enabled.
- To allow remediation actions, ensure that you have the Host.Config.Quarantine and Host.Config.Maintenance privileges.

**Step 2.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Configure > vSphere Availability > Edit**.

**Step 3.** Select **Turn on Proactive HA**.

**Step 4.** Click **Proactive HA Failures and Responses**.

**Step 5.** Set Automation Level to **Manual** or **Automated**.

**Step 6.** Set Remediation to one of the following:

- **Quarantine Mode for All Failures**
- **Quarantine Mode for Moderate and Maintenance Mode for Severe Failure (Mixed)**
- **Maintenance Mode for All Failures**

## Configuring vSphere Fault Tolerance

Before enabling vSphere Fault Tolerance (FT) for a virtual machine, you must prepare the hosts and cluster by doing the following:

- Step 1.** Configure vSphere HA on the cluster.
- Step 2.** On each participating host, configure a vMotion port group, a VMkernel adapter enabled for vMotion, a Fault Tolerance logging network, and a VMkernel adapter enabled for FT Logging.

To turn on FT for a virtual machine, you can use the following procedure:

- Step 1.** In the vSphere Client, right-click the virtual machine in the inventory pane and select **Fault Tolerance > Turn On Fault Tolerance**.
- Step 2.** Click **Yes**.
- Step 3.** Select a datastore on which to place the secondary VM configuration files and click **Next**.
- Step 4.** Select a host on which to place the secondary VM and click **Next**.
- Step 5.** Review your selections and then click **Finish**.

Before FT is turned on for a virtual machine, FT performs several validation steps related to the FT requirements listed in [Chapter 4](#). The virtual machine datastores and memory are replicated as FT is turned on. This may take several minutes, during which the virtual machine status does not appear as protected. When the replication completes and the state of the secondary VM is synchronized with that of the primary VM, the status changes to Protected.

To test FT failover for a virtual machine, right-click the virtual machine and select Fault Tolerance > Test Failover. Likewise, you can select Fault Tolerance > Test Restart Secondary to restart the Secondary VM.

## Monitoring and Managing vSphere Resources

You can use the vSphere [client performance charts](#) to view compute, storage, and network resource usage for virtual machines, hosts, and clusters. For a more granular look from the host perspective, you can use the [ESXTOP](#) utility. You can use vCenter Server alarms to bring attention to conditions and events that may call for human intervention, such as low resource availability on a cluster, host, or data-store. To bring multi-vCenter Server monitoring, predictive

analysis, and intelligent operations to an environment, you can consider integrating vRealize Operations (vROps).

## Metrics

Performance metrics are organized into logical groups based on the object or object device, as shown in [Table 10-3](#).

**Table 10-3** Metrics

<b>Metric Group</b>	<b>Description</b>
Cluster Services	Performance metrics on vSphere host clusters.
CPU	CPU utilization metrics for hosts, virtual machines, resource pools, or compute resources.
Datastore	Datastore utilization metrics.
Disk	Disk utilization metrics for hosts, virtual machines, or datastores.
Memory	Memory utilization metrics for hosts, virtual machines, resource pools, or compute resources.
Network	Network utilization metrics for physical NICs, virtual NICs, and other network devices.
Power	Energy and power utilization metrics for hosts.
Storage Adapter	Data traffic metrics for host bus adapters (HBAs).
Storage Path	Data traffic metrics for paths.
System	Overall system availability metrics, such as the system heartbeat and uptime. These counters are available directly from hosts and from vCenter Server.
Virtual Disk	Disk utilization and disk performance metrics for virtual machines.
Virtual Flash	Virtual flash metrics.

Virtual Machine Operations Virtual machine power and provisioning operations metrics in a cluster or data center.

---

vSphere Replication Virtual machine replication metrics.

---

Disk metrics include I/O performance, such as latency and read/write speeds, and utilization metrics for storage as a finite resource.

The value obtained for memory utilization is one of the following:

- For virtual machines, memory refers to the guest physical memory, which is the virtual memory the hypervisor presents to the guest as physical memory.
- For hosts, memory refers to the machine memory, which is the physical memory in the host system.

## vSphere Client Performance Charts

The vSphere client performance charts enable you to view performance metrics in different types of charts, depending on the selected object and metric type, as described in [Table 10-4](#).

**Table 10-4** Performance Chart Types

Chart Type	Description	Example
Line	Displays metrics for a single inventory object, where data for each metric is represented by a separate line.	For example, A network chart for a host can contain one line showing the number of packets received and another line showing the number of packets transmitted.
Bar	Displays metrics for	A bar chart can display metrics for datastores,

arrt objects, where each bar represents metrics for an object. where each datastore is represented as a bar. Each bar displays metrics based on the file type, such as virtual disk or snapshot.

---

Prta Displays metrics for a single object, where each slice represents a category or child object. A datastore pie chart can display the amount of storage space occupied by each virtual machine or by each file type.

---

Stacked Displays metrics for child objects. A host's stacked CPU usage chart displays metrics for the 10 virtual machines on the host that are consuming the most CPU. The Other amount displays the total CPU usage of the remaining virtual machines.

---

Overview and advanced performance charts are available for data center, cluster, host, resource pool, vApp, and virtual machine objects. Overview performance charts are also available for datastores and datastore clusters. Performance charts are not available for network objects. Charts are organized into views, which you can use to see related data together on one screen. You can specify the time range or data collection interval. Advanced charts contain more information than overview charts. You can print, configure, and export advanced charts (PNG, JPEG, or CSV formats).

## Overview Performance Charts



You can use the vSphere Client to examine the overview performance charts for



data centers, clusters, datastores (and datastore clusters), hosts, resource pools, vApps, and virtual machines.

To view a performance chart, you can use the following procedure:

**Step 1.** In the vSphere Client, select an appropriate object in the inventory pane and navigate to **Monitor > Performance**.

**Step 2.** Select a view.

**Step 3.** Select a predefined or custom time range.

[Table 10-5](#) lists the available performance chart views by object type.

**Table 10-5** Views by Object Type

<b>Object Type</b>	<b>View List Items</b>
Data center	<b>Clusters:</b> Thumbnail CPU and memory charts for each cluster and stacked charts for total data center CPU and memory.  <b>Storage:</b> Space utilization charts for each datastore by file type.
Datastore and datastore cluster	<b>Space:</b> Space utilization charts by datastore, by file type, and by virtual machine.  <b>Performance:</b> Disk performance (latency, throughput, and queuing) charts for the datastore (or datastore cluster, when Storage DRS is enabled) by virtual machine, by virtual disk, and by file type.
Cluster	<b>Home:</b> CPU and memory charts for the cluster.  <b>Resource pools and virtual machines:</b> Thumbnail charts for resource pools and virtual machines and stacked charts for total cluster CPU and memory usage.  <b>Hosts:</b> Thumbnail charts for each host and stacked charts for total cluster CPU, memory, disk usage, and network usage.
Host	<b>Home:</b> CPU, memory, disk, and network charts for the host.

**Virtual machines:** Thumbnail charts for virtual machines, and stacked charts for total CPU usage and total host memory usage.

---

Resource Pool **Home:** CPU and memory charts for the resource pool.

**Resource pools and virtual machines:** Thumbnail charts for resource pools and virtual machines and stacked charts for total resource CPU and memory usage.

---

vApps **Home:** CPU and memory charts for the resource pool.

**Resource pools and virtual machines:** Thumbnail charts for resource pools and virtual machines and stacked charts for total vApp CPU and memory usage.

---

Virtual Machine **Storage:** Space utilization charts for the virtual machine by file type and by datastore.

**Fault tolerance:** CPU and memory charts that display metrics for the fault-tolerant primary and secondary virtual machines.

**Home:** CPU, memory, network, and host thumbnail charts and disk performance charts for the virtual machine.

---

### Note

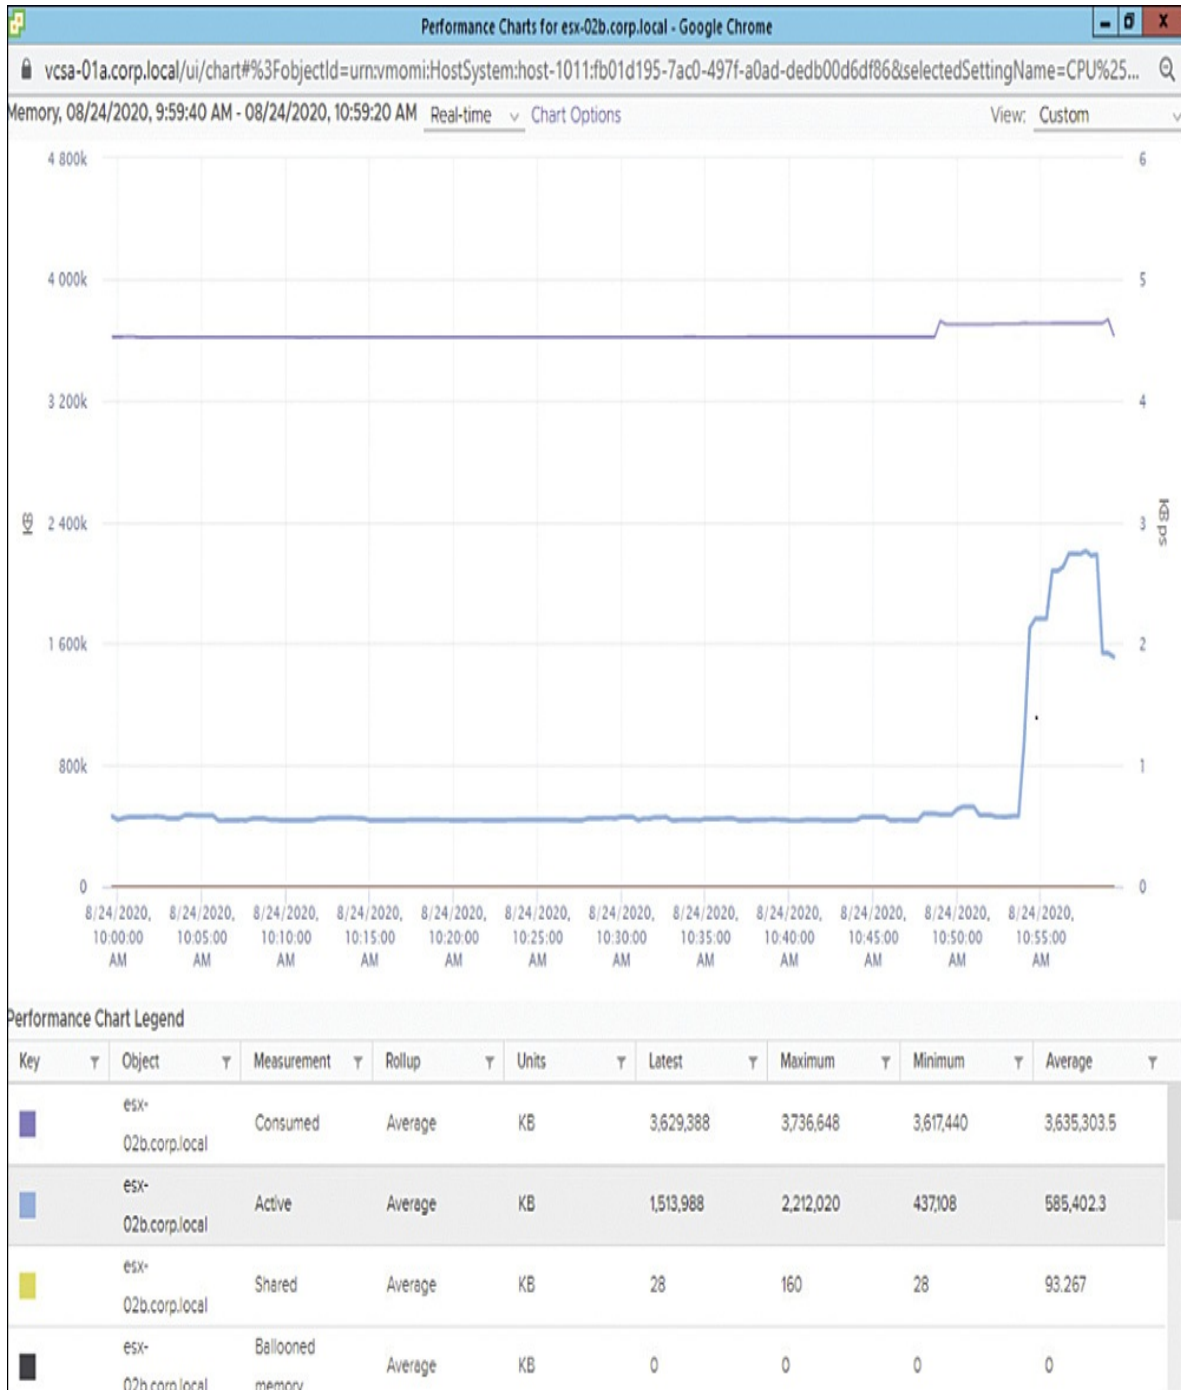
When Storage I/O Control is disabled, the values for the Storage I/O Normalized Latency metrics are zeros.

## Advanced Performance Charts

For more granularity, you can use advanced performance charts or create your own custom charts. Advanced performance charts are especially useful when overview performance charts do not provide sufficient data for troubleshooting a specific issue. Advanced performance charts include the following features:

- **Customizable:** You can change and save chart settings.
- **More information:** You can include data counters that are not supported in other performance charts. You can hover over a data point to see details at that point.
- **Exportable:** You can save an image to a file or spreadsheet. You can export the data to a spreadsheet.

Figure 10-1 shows an example of an advanced performance chart that includes memory metrics for a virtual machine.



**FIGURE 10-1** Sample Advanced Performance Chart

You can use the following procedure to access an advanced performance chart:

- Step 1.** In the vSphere Client, select an appropriate object in the inventory pane and navigate to **Monitor > Performance**.

- Step 2.** Click **Advanced**.
- Step 3.** Optionally, select an appropriate view from View drop-down list.
- Step 4.** Optionally, click the **Popup Chart** icon to open the chart in a separate window.
- Step 5.** Click **Chart Options**.
- Step 6.** Under Chart Metrics, select an appropriate metric group.
- Step 7.** Select a timespan. If you choose **Custom Interval**, then select one of the following:
- **Last:** Specify the number of hours, days, weeks, or months.
  - **From:** Specify beginning and ending times.
- Step 8.** Under Target Objects, select the appropriate inventory objects. (Optionally, use the **All** or **None** buttons.)
- Step 9.** Select an appropriate chart type.
- Step 10.** Under Counters, select the data counters to display in the chart. (Optionally, use the **All** or **None** buttons.)
- Step 11.** Optionally, click **Save Options As** and save your settings as a custom chart.

## Note

Pop-up charts are useful for maximizing the available real estate for a chart and for comparing two separate charts side by side.

## Note

For the stacked graph type, you can use only one measurement unit. In addition, per-virtual-machine stacked graphs are available only for hosts. You can click on a counter's description name to display details, such as whether the selected metric can be stacked for each virtual machine.

After you create a custom chart, the chart is added to the View drop-down list. You can then use the chart in the same manner that you would any prebuilt view.

You can use the following procedure to delete a custom chart:

**Step 1.** In the vSphere Client, select an appropriate object in the inventory pane and navigate to **Monitor > Performance**.

**Step 2.** Select **Advanced > Chart Options**.

**Step 3.** Select the chart and click **Delete Options**.

You can use the following procedure to save data from an advanced performance chart to a file either in a graphic format or in a comma-separated values (CSV) format:

**Step 1.** In the vSphere Client, select an object in the inventory pane and navigate to **Monitor > Performance**.

**Step 2.** Click **Advanced**.

**Step 3.** Optionally, select a view or change chart options until you are satisfied with the chart.

**Step 4.** Click the **Export** icon.

**Step 5.** Select one of the following options:

- **To PNG:** Exports a bitmap image to PNG format.
- **To JPEG:** Exports a bitmap image to JPEG format.
- **To CSV:** Exports text data to CSV format.
- **To SVG:** Exports a vector image to SVG format.

**Step 6.** Provide a filename and location.

**Step 7.** Click **Save**.

## Troubleshooting and Optimizing Performance

Table 10-6 provides the likely causes and potential solutions for some sample symptoms, based on vSphere performance metrics.

**Table 10-6** CPU Performance Analysis

---

<b>Symptoms</b>	<b>Likely Causes</b>	<b>Potential Solutions</b>
Host: CPU usage is consistently high.	The host has insufficient CPU resources to meet the demand.	Add the host to a DRS cluster.
Virtual machine: CPU usage is above 90%. CPU ready is above 20%. Application performance is poor.	Too many virtual CPUs are running on the host.	Increase the number of hosts in the DRS cluster.
	Storage or network operations are placing the CPU in a wait state.	Migrate one or more virtual machines to other hosts.
	The guest OS generates too much load for the CPU.	Upgrade the physical CPUs of the host.
		Upgrade ESXi to the latest version.
		Enable CPU-saving features such as TCP segmentation offload, large memory pages, and jumbo frames.
		Increase the amount of memory allocated to the virtual machines, which may improve

cached I/O and reduce CPU utilization.

Reduce the number of virtual CPUs assigned to virtual machines.

Ensure that VMware Tools is installed.

Compare the CPU usage of troubled virtual machines with that of other virtual machines on the host or in the resource pool. (Hint: Use a stacked graph.)

Increase the CPU limit, shares, or reservation on the troubled virtual machine.

---

Host: Memory usage is consistently 94% or higher. Free memory is 6% or less.

The host has insufficient memory resources to meet the demand.

Ensure that VMware Tools is installed and that the balloon driver is enabled for all virtual machines.



Virtual machine: Swapping is occurring. (Memory usage may be high or low.)

---

Reduce the memory size on oversized virtual machines.

Reduce the memory reservation of virtual machines where it is set higher than needed.

Add the host to a DRS cluster.

Increase the number of hosts in the DRS cluster.

Migrate one or more virtual machines to other hosts.

Add physical memory to the host.

---

Virtual machine:

The guest OS is not provided sufficient memory by the virtual machine.

Increase the memory size of

Memory usage is high.

the virtual machine.

Guest OS: Memory usage is high. Paging is occurring.

---

Virtual machine: CPU ready is low.

The guest OS is not provided sufficient CPU resources by the virtual machine.

Increase the number of CPUs for the virtual machine.

Guest OS: CPU utilization is high.

Migrate the virtual machine to a host with faster CPUs.

---

Datastore: Space utilization is high.

Snapshot files are consuming a lot of datastore space.

Delete or consolidate virtual machine snapshots.

Some virtual machines are provisioned with more storage space than required.

The datastore has insufficient storage space to meet the demand.

Convert some virtual disks to thin provisioned.

Migrate one or more virtual machines (or virtual disks) to other datastores.

Add the datastore

to a Storage DRS  
datastore cluster.

Add datastores  
with available  
space to the  
datastore cluster.

Add more storage  
space to the  
datastore.

---

Disk: Device latency is greater than 15 ms.	Problems are occurring with the storage array.	Migrate the virtual machines to datastores backed by other storage arrays.
---	---	--

---

Disk: VMkernel latency is greater than 4 ms. Queue latency is greater than zero.	The maximum throughput of a storage device is not sufficient to meet the demand of the current workload.	Migrate the virtual machines to datastores backed by storage devices (LUNs) with more spindles.
--	--	--

Balance virtual  
machines and their  
disk I/O across the  
available physical  
resources. Use  
Storage DRS I/O  
balancing.

Add more disks  
(spindles) to the  
storage device  
backing the

datastore.

Configure the queue depth and cache settings on the RAID controllers. Adjust the Disk.SchedNumReqOutstanding parameter.

Configure multipathing.

Increase the memory size of the virtual machine to eliminate any guest OS paging. Increase the guest OS caching of disk I/O.

Ensure that no virtual machine swapping or ballooning is occurring.

Defragment guest file systems.

Use eager zeroed

thick provisioned virtual disks.

---

Network:  
The number of packets dropped is greater than zero. Latency is high. The transfer rate is low.

The maximum throughput of a physical network adapter is not sufficient to meet the demand of the current workload.

Virtual machine network resource shares are too few.

Network packet size is too large, which results in high network latency. Use the VMware AppSpeed performance monitoring application or a third-party application to check network latency.

Network packet size is too small, which increases the demand for the CPU resources needed for processing each packet. Host CPU, or possibly virtual machine CPU, resources are not enough to handle the load.

Install VMware Tools on each virtual machine and configure the guest OS to use the best-performing network adapter driver (such as vmxnet3).

Migrate virtual machines to other hosts or to other physical network adapters.

Verify that all NICs are running in full duplex mode.

Implement TCP Segmentation Offload (TSO) and jumbo frames.

Assign additional physical adapters as uplinks for the associated port groups.

Replace physical network adapters with high-bandwidth adapters.

Place sets of virtual machines that communicate with each other regularly on the same ESXi host.

---

Performance charts are empty.

Some metrics are not available for pre-ESXi 5.0 hosts.

Upgrade hosts to a later version of ESXi.

Data is deleted when you remove objects to vCenter Server or remove them.

Allow time for data collection on objects that were recently added, migrated, or recovered to the vCenter Server.

Performance chart data for inventory objects that were moved to a new site by VMware vCenter Site Recovery Manager is deleted from the old site and not copied to the new site.

Performance charts data is deleted when you use VMware vMotion across vCenter Server instances.

Power on all hosts and allow time for real-time statistics to collect.

Real-time statistics are not available for disconnected hosts or powered-off virtual machines.

Allow time for the required roll-ups for non-real-time statistics.

---

Non-real-time statistics are rolled up at specific intervals. For example, 1-day

statistics might not be available for 30 minutes after the current time, depending on when the sample period began.

The 1-day statistics are rolled up to create one data point every 30 minutes. If a delay occurs in the roll-up operation, the 1-week statistics might not be available for 1 hour after the current time. It takes 30 minutes for the 1-week collection interval, plus 30 minutes for the 1-day collection interval.


The 1-week statistics are rolled up to create one data point every two hours. If a delay occurs in the roll-up operations, the 1-month statistics might not be available for 3 hours. It takes 2 hours for the 1-month collection interval, plus 1 hour for the 1-week collection interval.

The 1-month statistics are rolled up to create one data point every day. If a delay occurs in the roll-up operations, the statistics might not be available for 1 day and 3 hours. It takes 1 day for the past year collection interval, plus 3 hours for the past month collection interval. During this time, the charts are empty.

---

## **Monitoring and Managing Cluster Resources**

On a Summary tab for a vSphere DRS cluster, you can see the Capacity, Used, and Free metrics for CPU, Memory, and Storage resources in the upper-right corner, as shown in [Figure 10-2](#).

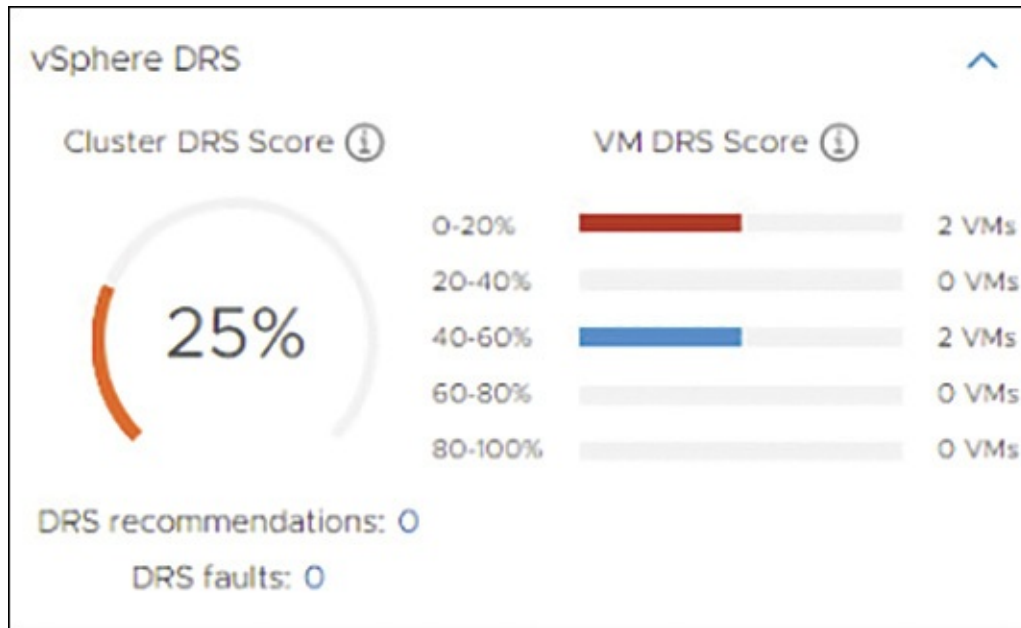


**FIGURE 10-2** DRS Cluster Resource Usage

To examine the CPU and memory usage more closely, you can navigate to Monitor > vSphere DRS and select CPU Utilization or Memory Utilization. Each of these pages shows a bar graph, where each bar represents the total resource (CPU or memory) usage of a specific host and each bar is split into sections representing the resource usage of individual virtual machines. Likewise, you can select Monitor > vSphere DRS > Network Utilization to examine the network utilization of each host in the cluster.

The Summary tab shows the vSphere DRS score, the number of DRS recommendations, and the number of DRS faults, as shown in [Figure 10-3](#).





**FIGURE 10-3** Sample DRS Score

If DRS is in manual mode, you can click on the number of DRS recommendations on the summary tab, which is a link that takes you to the DRS Recommendations page. On the DRS Recommendations page, you can view the current recommendations, select those that you want to apply, and click the Apply Recommendations button. Each recommendation includes a description, such as which virtual machine to migrate to which host, and a reason, such as balance average memory loads.

Optionally, you can click the Run DRS Now button to make DRS perform its analysis and potentially generate new recommendations.

## Monitoring and Managing Resource Pool Resources

To view resource pool configuration details, you can select a DRS cluster in the inventory pane and navigate to Hosts > Resource Pools. On that page, you can see all the resource pools that are direct children of the cluster. For each pool, you see the CPU and memory resource settings, including Reservation, Limit, Shares Setting (such as Low or Custom), Shares Value (numeric share value), and Allocation Type (Expandable or non-Expandable). You can click on the name of a resource pool, which is a link to the pool's Summary page, which shows the current capacity, usage, and free compute resources for the resource pool. The summary page also shows the number of virtual machines, powered-

on virtual machines, child resource pools, and vApps in the pool.

For more detail, you can navigate to Monitor > Utilization or to Monitor > Resource Allocation and select CPU, Memory, or Storage. For both CPU and memory resources, the Utilization page shows the resource configuration and the consumed, active, and worst-case allocation. The Utilization page also shows a breakdown of guest memory, including the Active Guest Memory, Swapped, Compressed, and Ballooned metrics.

You can use overview and advanced performance charts with resource pools. When you see undesired behavior, you can edit the settings for an existing resource pool to change the pool's CPU and memory shares, reservations, and limits. For example, consider a scenario where you configure two resource pools in a cluster with 100 GHz CPU capacity. In a pool with 40 virtual machines, you set CPU Shares to High. In another pool, which has 8 virtual machines, you set CPU Shares to Normal. You see in the performance charts that the virtual machines in the pool with the 40 virtual machines have greater CPU Ready values than the virtual machines in the other pool. You realize that although you used higher CPU shares for the first pool, the virtual machines are experiencing more CPU contention than are virtual machines in the second pool. To correct this, you could take one of the following actions:

- Increase the CPU shares on the first pool by using a custom value.
- Change the CPU shares on the second pool to Low.
- Set an appropriate CPU reservation on the first pool.
- Set an appropriate CPU limit on the second pool.
- Change the configuration to use scalable shares.

## **Monitoring and Managing Host Resources and Health**

You can use the vSphere Client to monitor the state of host hardware components, such as CPUs, memory, fans, temperature, voltage, power, network, battery, storage, cable (interconnect), software components, watchdog, and PCI devices. To view the host hardware health status with the vSphere Client, you can use the following procedure:

**Step 1.** In the vSphere Client, select host in the inventory pane and navigate to **Monitor > Hardware Health**.

**Step 2.** Select the type of information to view:

- **Sensors**
- **Storage Sensors**
- **Alerts and Warnings**
- **System Event Log**

The host health monitoring tool presents data gathered using Systems Management Architecture for Server Hardware (SMASH) profiles. The information displayed depends on the sensors available on the server hardware.

### Note

You can also set alarms to trigger when the host health status changes.

If you participate in the Customer Experience Improvement Program (CEIP), you can configure Skyline Health to perform online health checks. If CEIP is not enabled, the Internet connectivity check is unavailable. You can use the following procedure to configure Skyline Health:



- Step 1.** In the vSphere client, select a vCenter Server or a host in the inventory pane and navigate to **Monitor > Skyline Health**.
- Step 2.** Expand the **Online Health Connectivity** category and select one of the following options:
- **CEIP:** Verifies whether CEIP is enabled for the vCenter Server.
  - **Online Health Connectivity (Internet check):** Verifies vCenter Server to [vmware.com](https://www.vmware.com) connectivity via HTTPS/443.
  - **Advisor:** Provides additional features, such as automatic support log bundle transfer with Log Assist. (Advisor is included in Production and Premier Support contracts.)
  - **Audit CEIP Collected Data:** Allows you to view data collected and sent for CEIP.

- **Number of Online Health Checks Performed Successfully:** Indicates, as it says, how many checks have been performed successfully.

**Step 3.** Expand the following categories and examine the related health warnings:

- **Compute Health Checks**
- **Network Health Checks**
- **Security Health Check**
- **Storage Health Checks**
- **General Health Checks**

**Step 4.** Click **Retest** to run the health checks immediately.

**Step 5.** Optionally, if issues are discovered, click the **Ask VMware** button to request a knowledge base article that describes how to resolve the issue.

## Monitoring and Managing Virtual Machine Resources

Table 10-7 lists and describes some of the key metrics for monitoring virtual machines.

**Table 10-7** Virtual Machine Metrics

Metric	Unit	Description
CPU Usage	%	Indicates the CPU workload for the virtual machine.
<i>CPU Ready Time</i>	ms	Indicates the amount of time a virtual CPU is ready to work (that is, has a workload and is ready to be scheduled) but is waiting to be scheduled on hardware. High CPU Ready Time is a sign of CPU contention.
Memory Consumed	KB	Indicates the amount of physical memory currently backing the virtual machine.

Memory Active	KB	Indicates the amount of consumed memory that is actively being read or written by the guest OS.
Memory Swap In Rate	Kbps	Indicates the amount of memory read from the virtual machine's swap file over time.
Disk Usage	Kbps	Indicates the disk throughput.
Virtual Disk Read Latency	ms	Indicates the average amount of time for a read operation to complete.
Network Usage	Kbps	Indicates the amount of data transmitted and received over time.
Network Transmit Packets Dropped	number	Indicates the number of packets transmitted to the network that were dropped.

## Shares, Limits, and Reservations

You can set the CPU and memory shares, reservation, and limit on a virtual machine by using the following procedure:

- Step 1.** In the vSphere Client, right-click your virtual machine in the inventory and select **Edit Settings**.
- Step 2.** Edit the **Shares**, **Reservation**, and **Limit** values under CPU Resources.
- Step 3.** Set the **Shares**, **Reservation**, and **Limit** values under Memory Resources.
- Step 4.** Click **OK**.

The relative priority represented by each share changes whenever additional sibling virtual machines are powered on or powered off. Likewise, each share's relative priority changes whenever the shares on siblings are increased or decreased. This affects all virtual machines in the same resource pool.

For example, consider the following scenario:

- All virtual machines have the same number of vCPUs.
- Two virtual machines are run in a resource pool with CPU Limit set to 8 GHz.
- The virtual machines are CPU bound (that is, they are demanding more CPU resources than they are receiving).
- The virtual machines' CPU shares are set to Normal.
- You should expect each virtual machine's performance chart to show CPU Utilization as 4 GHz.
- When you power on a third CPU-bound sibling virtual machine with CPU Shares value set to High, you should expect to see that the new virtual machine uses 4 GHz and the first two machines drop to 2 GHz each.

To understand the impact of shares, consider another scenario, where a set of sibling virtual machines are frequently CPU bound and are using all the resources in their parent resource pool. During these periods of CPU contention in the resource pool, you see significantly high CPU Ready Time values on each of the virtual machines. You are only concerned about improving the performance of one specific virtual machine, so you increase its CPU Shares value. The CPU's Ready Time setting for that machine should decrease during periods of CPU contention, and the CPU Ready Time settings of its siblings should rise.

To guarantee that a specific amount of resources are always available to a running virtual machine, even when the physical server is heavily loaded, you can set its CPU or memory reservation. The vCenter Server or ESXi host allows you to power on a virtual machine only if there are enough unreserved resources to satisfy the virtual machine's reservation. Likewise, your attempts to increase the reservation on a running virtual machine (or a resource pool) succeed only if there are enough unreserved resources to satisfy the request. In the previous scenario, if you want to ensure that a virtual machine always has access to at least 1 GHz, regardless of the number or resource settings of siblings, you

should set its CPU reservation to 1 GHz.

## Note

The default CPU and memory reservation for a virtual machine is zero, meaning that its guest OS is not guaranteed any specific amount of either resource. Instead, with default settings, shares would be applied during periods of compute resource contention.

You can set limits for CPU, memory, and storage I/O for a virtual machine to establish an upper bound (maximum) amount of resources that can be allocated to the virtual machine. The host never allocates more than the limit, even when there are unused resources on the system. By default, the limits are set to Unlimited, which means the virtual machine's configured memory becomes its effective limit. Using limits has both benefits and drawbacks:

- **Benefits:** If you are concerned that the performance of a virtual machine may deteriorate as you add virtual machines to the cluster, you could set limits on the virtual machine to simulate having fewer available resources and measure its performance.
- **Drawbacks:** You might be wasting idle resources because the system prevents virtual machines from exceeding the limits that you set, even when the system is underutilized and idle resources are available.

## Note

If you want to reduce the risk that a virtual machine may consume excessive resources and impact the performance of other virtual machines, you can consider setting low shares on the virtual machine. Low shares decrease the virtual machine's access to the resource during periods of resource contention but also do not prevent the virtual machine from using idle resources.

## Admission Control

When you power on a virtual machine, the system checks the amount of

available unreserved CPU and memory resources. The system determines whether it can guarantee the reservation for the virtual machine. This process is called *admission control*. If enough unreserved CPU and memory are available (or if there is no reservation), the virtual machine is powered on. Otherwise, an “Insufficient Resources” warning appears.

### **Note**

Each virtual machine, including those with no user-specified memory reservation, may have some reservation for its memory overhead. The memory overhead reservation is considered by admission control.

### **Note**

When the vSphere Distributed Power Management (DPM) feature is enabled and some hosts are in standby mode, their unreserved resources are considered available for admission control. If a virtual machine cannot be powered on without these resources, vSphere DPM makes a recommendation to power on one or more standby hosts.

## **VMware Tools and Microsoft Windows Perfmon**

When VMware Tools is installed, VMware provides performance counters that enable you to view data within a Windows guest OS by using the Microsoft Windows Perfmon utility. VMware provides virtual machine-specific performance counter libraries for the Windows Perfmon utility, which enables administrators to accurately examine virtual machine usage data and guest OS usage data within the same pane of glass.

For a Windows virtual machine where VMware Tools is installed, you can use the following procedure to examine VMware-specific statistics in the Windows Perfmon utility:

- Step 1.** Log on to Windows and click **Start > Run**.
- Step 2.** Enter **Perfmon** and press **Enter**.
- Step 3.** In the Performance dialog box, click **Add**.



- Step 4.** In the Add Counters dialog box, select **Use Local Computer Counters**.
- Step 5.** Select a performance object whose name begins with **VM** (that is, a virtual machine performance object).
- Step 6.** Select the counters that you want to display for that object.
- Step 7.** If the performance object has multiple instances, select the instances you want to display.
- Step 8.** Click **Add**.
- Step 9.** Examine the data for the selected performance object.
- Step 10.** Click **Close**.

## Latency Sensitivity

If you have a latency-sensitive application, such as voice over IP (VOIP) or a media player application, you can edit the virtual machine's settings and set VM Options > Advanced > Latency Sensitivity to High. With this setting, you should ensure that all the virtual machine's configured CPU and memory are reserved. With this setting, the system effectively gives exclusive physical CPU access to each virtual CPU. If the virtual machine is in a DRS cluster, DRS automatically creates a VM-Host soft affinity rule.

## The Impact of Virtual Machine Configurations



The specific settings you make for a virtual machine can impact its performance, as summarized in [Table 10-8](#).

**Table 10-8** The Impact of Virtual Machine Configurations

Configur ation	Impact
-------------------	--------

Comput e oversiz e/under	An oversized compute size for a virtual machine may result in wasted resources. With an undersized compute size, the virtual machine may experience poor performance.
-----------------------------------	---

size

---

Virtual disk oversize/undersize	An oversized virtual disk may result in wasted resources. With an undersized virtual disk, the virtual machine may experience denial of service.
---------------------------------	--

---

VMDK provisioning types	If a virtual disk is thin provisioned, then you may be maximizing the use of your storage space while decreasing the virtual machine's performance and increasing its risk of denial of service.
-------------------------	--

---

Resource reservations	If a resource is reserved, you may be improving and guaranteeing the guest OS performance while reducing the density of virtual machines on the resource.
-----------------------	---

---

Independent disks	If a virtual disk is set to independent mode, then you are prevented from taking snapshots of it. If it is set to Independent–Nonpersistent, all changes are discarded when you power off or reset the virtual machine.
-------------------	---

---

Guest OS type	The choice for the guest OS type during virtual machine creation directly impacts the type of virtual devices that are used in the virtual machine.
---------------	---

---

VMware Tools version	The VMware Tools version impacts the set of device drivers it provides to the guest OS.
----------------------	---

---

Permissions	The permissions set on a virtual machine impacts who can use (power on, open a console), who can modify (change virtual hardware settings), and who can manage (set permissions, migrate) the virtual machine.
-------------	--

---

## Other Virtual Machine Resource Management Features

You can configure virtual machines to support SRIOV, VGPU, RDMA, and DirectPath I/O passthrough, as discussed in [Chapter 14, “Managing Virtual Machines.”](#)

## ESXTOP

ESXTOP is a utility that provides a detailed real-time look at resource usage from the ESXi Shell. You can run ESXTOP in Interactive, Batch, or Replay mode. You must have root user privileges. RESXTOP is a similar tool that can be installed and run from a Linux server and connected to ESXi hosts.

By default, when you issue the command **esxtop**, the utility opens in interactive mode to show the CPU panel, where statistics for each virtual machine and other groups are displayed in separate rows. To see just virtual machines statistics, you can press Shift+V. Each column provides CPU statistics, such as %USED, %WAIT, %RDY, %CSTP, and %SWPWT. To see statistics for the multiple worlds (processes) that comprise a virtual machine, you can press the E key and enter the virtual machine's ID. [Figure 10-4](#) shows an example of an ESXTOP CPU panel, displaying virtual machine statistics with one virtual machine (GID 33791) expanded.

```

1:24:45pm up 1:31, 562 worlds, 3 VMs, 4 vCPUs; CPU load average: 0.12, 0.14, 0.19
PCPU USED(%): 11 11 AVG: 11
PCPU UTIL(%): 11 11 AVG: 11

```

GID NAME	NWLD	%USED	%RUN	%SYS	%WAIT	%VMWAIT	%RDY	%IDLE	%OVRLP	%CSTP	%MLMTD	%SWPWT
24851 kms-01b	10	12.86	12.92	0.06	984.53	0.31	2.84	185.43	0.07	0.00	0.00	0.00
33791 vmx	1	0.08	0.01	0.07	99.97	-	0.01	0.00	0.00	0.00	0.00	0.00
33791 NetWorld-VM-135	1	0.00	0.00	0.00	99.98	-	0.00	0.00	0.00	0.00	0.00	0.00
33791 vmast.135915	1	0.09	0.10	0.00	99.35	-	0.54	0.00	0.00	0.00	0.00	0.00
33791 vmx-vthread-135	1	0.00	0.00	0.00	99.99	-	0.00	0.00	0.00	0.00	0.00	0.00
33791 vmx-filtPoll:wi	1	0.00	0.00	0.00	99.98	-	0.00	0.00	0.00	0.00	0.00	0.00
33791 vmx-mks:win10	1	0.01	0.01	0.00	99.98	-	0.00	0.00	0.00	0.00	0.00	0.00
33791 vmx-svga:win10	1	0.00	0.00	0.00	99.99	-	0.00	0.00	0.00	0.00	0.00	0.00
33791 vmx-vcpu-0:win1	1	1.55	1.55	0.00	96.95	0.17	1.50	96.78	0.02	0.00	0.00	0.00
33791 LSI-135915:0	1	0.00	0.00	0.00	99.99	-	0.00	0.00	0.00	0.00	0.00	0.00
33780 app-01b	9	1.04	1.00	0.03	897.71	0.80	1.19	97.11	0.00	0.00	0.00	0.00

**FIGURE 10-4** Sample ESXTOP CPU Panel

You can change the view from the CPU panel to other panels by using keystrokes. For example, you can press *M* Key for the memory panel, *V* Key for the virtual machine storage panel, or *N* Key for the network panel. [Table 10-9](#) describes some of the key statistics available for each panel.

**Table 10-9** Key ESXTOP Panels and Metrics

Panel	Static	Description
-------	--------	-------------

CPU	% U S E D	Percentage of physical CPU core cycles used by the virtual machine.
CPU	% R U N	Percentage of total time scheduled for the virtual machine without accounting for hyperthreading, system time, co-stopping, and waiting:  $\%RUN = 100\% - \%RDY - \%CSTP - \%WAIT$
CPU	% R D Y	Percentage of time the virtual machine was ready to run but was not provided CPU resources on which to execute. Indicator of CPU contention on the host.
CPU	% W A I T	Percentage of time the virtual machine spent in the blocked or busy wait state, including idle time. %WAIT includes %SWPWT.
CPU	% C S T P	Percentage of time a virtual machine spends in a ready, co-deschedule state. A high value indicates that the virtual machine's multiple CPUs are in contention.
CPU	% S W P W T	Percentage of time a virtual machine spends waiting for the host to swap memory.
Memory	M E M S Z	Amount of physical memory allocated to a virtual machine:  $MEMSZ = GRANT + MCTLSZ + SWCUR + \text{"Never Touched"}$
Memory	G R A N T	Amount of guest physical memory mapped to a virtual machine

Memory	CNSM	Amount of the memory consumed by the virtual machine:  CNSM = GRANT – Shared Memory
--------	------	---

Memory	SWCUR	Amount of memory swapped by the virtual machine.
--------	-------	--

Memory	SWR/s	Rate at which the host swaps in memory from disk for the virtual machine.
--------	-------	---

Memory	OVHD	Amount of memory used for virtual machine overhead, which is memory charged to the virtual machine that is not used by the guest OS.
--------	------	--

Virtual Machine Storage	READ S/s	Number of read commands issued per second.
-------------------------	----------	--

Virtual Machine Storage	WRITE S/s	Number of write commands issued per second.
-------------------------	-----------	---

Virtual Machine Storage	MBR/s	Megabytes read per second.
-------------------------	-------	----------------------------

Virtual Machine	LAT/rd	Average latency (in milliseconds) per read.
-----------------	--------	---

## Storage

---

Network	PKRSX/s	Number of packets received per second.
Network	MbTxs	Megabits transmitted per second.
Network	%D RPTX	Percentage of transmit packets dropped. Indicates that the physical network adapter cannot meet the demand, perhaps due to load from other virtual machines.
Network	%D RPTX	Percentage of receive packets dropped. Indicates that insufficient CPU resources are available for network processing.

---

### Note

The Network panel contains a row for each NIC in a virtual machine rather than a row for each virtual machine. The E and Shift+V keystrokes are not applicable to the Network panel.

You can use the **-b** argument to run ESXTOP in batch mode, in which you collect statistics in a CSV file. You can later manipulate this file with other tools, such as Microsoft Perfmon or Excel. For example, you can use the following command to collect statistics in a file named mydata.csv:

```
esxtop -b > mydata.csv
```

You can use ESXTOP in Replay mode, where it uses pre-collected data rather

than real-time data. To collect the data, you should run **vm-support** in Snapshot mode, specifying the data collection interval and duration (in seconds), as shown in the following example:

```
vm-support -S -d 3600 -I 5
```

After collecting the data, you must unpack and decompress the resulting tar file. Then you can run ESXTOP in Replay mode, providing the path to the data file, as shown here:

```
esxtop -R vm-support_dir_path
```

## VIMTOP

**VIMTOP** is a tool you can run in vCenter Server Appliance to see resource usage for the services that are running. It is like ESXTOP but displays services, such as vCenter Server, Certificate Manager, vPostgres, and ESXi Agent Manager, rather than virtual machines and ESXi worlds (processes). You can use VIMTOP to identify which service is using the most compute, disk, or network resources whenever vCenter Server is running poorly.

## vCenter Server Appliance Management Interface (VAMI)

In the vCenter Server Appliance Management Interface (VAMI), you select Monitor to view the resource usage of the vCenter Server. To see compute usage graphs, select Monitor > CPU and Memory. To see the usage of each storage partition, select Monitor > Disks. To use a graph where you can select and view specific network metrics, select Monitor > Network.

You can navigate to Monitor > Database to view database utilization of alarms, events, tasks, and statistics. You can also view the overall space utilization of the database and database log.

## Events, Alarms, and Automated Actions

A configurable events and alarms subsystem exists in vSphere that tracks events throughout vSphere and stores the data in log files and in the vCenter Server database. It enables you to specify the conditions under which alarms are triggered. Alarms can change state from normal (green) to warning (yellow) to alert (red), depending on changing conditions. Triggered alarms can automatically launch alarm actions.

## Events

Events are simply recorded incidents, such as user actions or system actions, that occurred involving a host or any object managed by vCenter Server. The following are a few examples:

- A license key expires.
- A virtual machine is migrated.
- A virtual machine is powered on.
- A host connection is lost.

Event data includes details such as who generated it, when it occurred, and what type of event it was. [Table 10-10](#) describes the types of events

**Table 10-10** Event Types

Event Type	Description
Audit	Provides data concerning events that are tracked because that data is crucial for the security framework. The data includes action details, such as who did it, when it occurred, and the IP address of the user.
Information	Indicates that the operation completed successfully.
Warning	Indicates a potential risk to the system that needs to be addressed. This event does not terminate the process or operation.
Fatal	Indicates that a fatal problem has occurred in the system and terminates the process or operation.



## Viewing Events in the vSphere Client

You can use the following procedure to view events in the vSphere Client:

- Step 1.** In the vSphere Client, select an object in the inventory pane and navigate to **Monitor > Events**.
- Step 2.** Select an event to see its details.
- Step 3.** Use the column headings to sort the events, show columns, hide columns, and filter the events.

## Viewing the System Event Log

To view system events that are recorded in the vCenter Server database, you can use the following procedure:

- Step 1.** Log on to the vSphere Client as a user with the Global.Diagnostics privilege.
- Step 2.** Select the vCenter Server in the inventory pane.
- Step 3.** Navigate to **Monitor > Hardware Health**.
- Step 4.** Click **System Event Log**.
- Step 5.** Optionally, click **Export**.

## Streaming Events to a Remote Syslog Server

You can enable remote streaming, such that the vCenter Server streams newly generated events to a remote syslog server. In the syslog server, the events have the following format:

[Click here to view code image](#)

```
<syslog-prefix> : Event [eventId] [partInfo] [createdTime]  
[eventType] [severity] [user] [target] [chainId] [desc]
```

Messages that are longer than 1024 characters are split into multiple syslog messages.

### Note

In an environment with no more than 30 hosts, you can configure hosts to

send log files to a vCenter Server rather than store them to a local disk. This option is intended for smaller environments with stateless ESXi hosts. For all other cases, VMware recommends that you use a dedicated log server.

As an alternative to streaming events, you can forward events. When you forward events, the events are sent to a remote server rather than recorded.

You can use the following procedure to forward vCenter Server logs to a remote syslog server:

- Step 1.** Log on to VAMI as root.
- Step 2.** Select **Syslog**.
- Step 3.** In the Forwarding Configuration section, click **Configure**.
- Step 4.** In the Create Forwarding Configuration pane, enter the server address of the destination host. The maximum number of supported destination hosts is three.
- Step 5.** Select a protocol (TLS, TCP, RELP, or UDP) to use.
- Step 6.** Provide a port number.
- Step 7.** Optionally, add more destination servers.
- Step 8.** Click **Save**.
- Step 9.** Optionally, click **Send Test Message**.

You can configure the writing of events to the vCenter Server streaming facility. Streaming events is disabled by default. You can use the following procedure to stream events to a remote syslog server:

- Step 1.** In the vSphere Client, select the vCenter Server in the inventory pane and navigate to **Configure > Settings > Advanced Settings**.
- Step 2.** Click **Edit**.
- Step 3.** Enable the **vpxd.event.syslog** option.

## Alarms

An alarm is a notification that is activated in response to an event, a set of

conditions, or the state of an inventory object. [Table 10-11](#) describes the elements that are used in an alarm definition

**Table 10-11** Alarm Definition Elements

Element	Description
Name	A name (label) that is used to identify the alarm
Description	Text that is useful for understanding the purpose of the alarm
Targets	The type of object that is monitored by the alarm
Alarm Rules	A set of rules that define the alarm's triggers, severity, and actions
Last Modified	The date of the most recent change to the alarm definition

For example, you might want to monitor the memory usage of all virtual machines in a specific vSphere cluster. In the vSphere Client, you can select the cluster in the inventory, create an alarm for the cluster, set the alarm's Targets value to virtual machine, and configure rules with triggers based on memory usage.

### Note

You can enable, disable, and modify alarms only from the object at which the alarm is defined. For example, if you define a virtual machine memory alarm on a cluster, you cannot change the alarm at the individual virtual machine level.

## Viewing and Acknowledging Triggered Alarms

To view triggered alarm, you can use the following procedure:

- Step 1.** In the vSphere Client, select an object in the inventory pane and navigate to **Monitor > Issues and Alarms**.
- Step 2.** Click **Triggered Alarms**.
- Step 3.** Optionally, select an alarm and click **Acknowledge**.

You can acknowledge an alarm to let other users know that you are taking

ownership of the issue and to prevent the alarm from sending more email messages. The alarm, however, is still visible in the system.

## Note

After you acknowledge an alarm in the vSphere Client, its alarm actions are discontinued. Alarms are not cleared or reset when acknowledged.

To clear an alarm (that is, reset its state to normal), you need the Alarm.Set Alarm Status privilege. You can select a triggered alarm and choose Reset to Green.

## Creating Alarm Definitions

To create or configure an alarm, you must use a user account with the Alarms.Create Alarm or Alarms.Modify Alarm privilege. To create an alarm, you can use the following procedure:

- Step 1.** In the vSphere client, select an object in the inventory pane and navigate to **Configure > More > Alarm Definitions**.
- Step 2.** Click **Add**.
- Step 3.** Provide the name, description, target type, and target.
- Step 4.** Click **Next**.
- Step 5.** Create an alarm rule by specifying the following:
  - **Conditions:** Trigger, Arguments, Operator, Threshold
  - **Severity:** Warning or Critical
  - **Actions:** Send Email Notifications, Send SNMP Traps, Run Script
- Step 6.** Optionally click **Add Another Rule**, **Duplicate Rule**, or **Remove Rule**.
- Step 7.** Click **Next**.
- Step 8.** Specify alarm reset rules by enabling the **Reset the Alarm to Green** option and providing details, such as arguments, operators, and actions.

**Step 9.** Click **Next**.

**Step 10.** Click **Enable this Alarm**.

## Alarm Actions

Alarm actions are operations that are automatically triggered by alarms. [Table 10-12](#) provides details on available alarm actions.

**Table 10-12** Alarm Actions

Action Details	
Send Email Notification	Indicates the recipient email address.  Requires that you first configure the mail settings for your vCenter Server. You must set the primary receiver URL to the DNS name or IP address of your SNMP receiver. You should set the receiver port to an appropriate value between 1 and 65535 and set the community string to an appropriate community identifier.
Send SNMP Traps	Requires that you first configure the SNMP Receivers settings for your vCenter Server. You must set Mail.Server to the DNS name or IP address of your SMTP gateway. You must set Mail.Sender to the email address of the sender.
Run Scripts	Provides the full pathname of the command or script, formatted into a single string. The execution occurs on the vCenter Server Appliance.
Advanced	Only applicable to alarms that target virtual machines and hosts. Examples of host actions include Enter Maintenance Mode and Exit Maintenance Mode. Examples of virtual machine actions include Migrate VM and Reboot Guest on VM.

## Advanced Use Cases for Alarms

You can create custom alerts with notifications for many purposes, such as the following:

- Something has failed or disconnected (such as host connection failure or VASA provider disconnected).
- Something is not performing well (such as excessive CPU ready time, memory swapping, disk latency, or packets dropped).
- Health is poor (such as vSAN health, key management server health, vCenter HA cluster health).

## Logging in vSphere

It is important that you understand logging in vSphere components and related products and that you be prepared to implement logging.

### ESXi Logs

[Table 10-13](#) provides details on most of the ESXi log files, including the location and purpose of each of them. You should become familiar with each of them and learn which logs are useful for various troubleshooting scenarios. For example, when troubleshooting virtual machine issues, the only directly useful logs are vmkernel, vmkwarning, hostd, and the specific virtual machine's log files. When troubleshooting issues related to the connection between an ESXi host and the vCenter Server, the vpxa log is most useful.

**Table 10-13** ESXi Log Files

<b>Component</b>	<b>Location</b>	<b>Description</b>
VMkernel	/var/log/vmkernel.log	Data related to virtual machines and ESXi
VMk	/var/log/vmkwarni	Data related to virtual machines

ernel warnings	ng.log	
VMkernel summary	/var/log/vmksummary.log	Data related to uptime and availability statistics for ESXi
ESXi host agent	/var/log/hostd.log	Data related to the agent that manages and configures the ESXi host and its virtual machines
vCenter agent	/var/log/vpxa.log	Data related to the agent that communicates with vCenter Server
ESXi Shell	/var/log/shell.log	Data related to each command typed into the ESXi Shell as well as shell events
Authentication	/var/log/auth.log	Data related to event authentication for the local system
System messages	/var/log/syslog.log	General log messages that can be used for troubleshooting
Virtual machines	vmware.log located in the same folder as the virtual machine configuration file.	Data related to virtual machine power events, system failure information, tool status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and more
Trusted infrastructure	/var/run/log/kmxa.log	Data related to the client service on the ESXi trusted host

agen  
t

---

Key provi der servi ce	/var/run/log/kmxd.l og	Data related to the vSphere Trust Authority key provider service
Attes tatio n servi ce	/var/run/log/attestd .log	Data related to the vSphere Trust Authority attestation service
ESX toke n servi ce	/var/run/log/esxtok end.log	Data related to the vSphere Trust Authority ESXi token service
ESX API forw arder	/var/run/log/esxapi adapter.log	Data related to the vSphere Trust Authority API forwarder
Quic k Boot	/var/log/loadESX.l og	Data related to restarting an ESXi host through Quick Boot

---

You can use the ESXi host client to examine the logs on a specific ESXi host by navigating to Monitor > Logs and selecting a specific log file. You can scroll through the log and search for specific text. You can select a log, click Actions, and choose Open in New Window or Generate a Support Bundle.

Likewise, you can use the ESXi Direct Console User Interface (DCUI) to view system logs. In the DCUI, after you click View System Logs and select the log you want, you can use the Enter key (or Spacebar) to scroll through the log messages and press the forward slash (/) key to begin a search.

If you have the Global.Diagnostics privilege, you can also use the vSphere Client to export a host's system logs by following these steps:



- Step 1.** In the vSphere Client, right-click an ESXi host in the inventory pane.
- Step 2.** Click **Export System Logs**.
- Step 3.** Select the appropriate objects.
- Step 4.** Optionally, click **Gather Performance Data**.
- Step 5.** Optionally, provide a password for encrypted coredumps.
- Step 6.** Click **Export Logs**.
- Step 7.** Monitor the status of the Downloading Log Bundles task in the Recent Tasks pane.

When you finish this process, the file is located in the default location. On a Windows desktop, the location is the Downloads folder, and the filename begins with VMware-vCenter-support.

### Note

In step 3, you can select or deselect entire categories, such as System, Virtual Machines, and Storage. You can also select or deselect specific objects within each category, such as logs and coredumps.

You can collect ESXi log files by using the `/usr/bin/vm-support` command, which generates a file named using the following format:

```
esx-date-unique-xnumber.tgz
```

## vCenter Server Logs

The main logs in a vCenter Server appliance are located in `/var/log/vmware`. The most important logs are in the `vpzd` subdirectory. Some other sibling subdirectories include `vsan-health`, `vsphere-ui`, and `vpostgres`.

## Uploading System Logs to VMware

To export system logs from the vCenter Server and all its hosts, you can use the procedure outlined in the section “ESXi Logs” but begin by selecting the vCenter Server instead of a specific host. In the wizard, you can select which hosts to include, and you can optionally select Include vCenter Server and

vSphere UI Client Logs.

You can export a vCenter Server instance's support bundle by using the URL shown on the DCUI home screen (<https://FQDN:443/appliance/support-bundle>).

Alternatively, you can run the **vc-support.sh** script in the vCenter Server appliance Bash Shell to collect the support bundle.

You can directly upload a log package to an open VMware service request by using the following procedure:

- Step 1.** In the vSphere Client, navigate to **Administration > Support**.
- Step 2.** Click **Upload File to Service Request**.
- Step 3.** Provide a service request ID.
- Step 4.** Click **Choose File**, select the appropriate log bundle, and click **OK**.

## Log Levels

The default log level setting is Info; this is where errors, warnings, and informational level are logged. You can change the log level to lower levels, such as Verbose, which is useful for troubleshooting and debugging but is not recommended for normal use in production environments. You can use the vSphere Client to change the logging level by selecting the vCenter Server, selecting **Configure > Settings > General > Edit**, and setting the logging settings to the appropriate level, as described in [Table 10-14](#).

**Table 10-14** vCenter Server Logging Options

Logging Option	Description
None (Disable Logging)	No vCenter Server logging occurs.
Error (Errors Only)	The vCenter Server collects only error entries in its log files.
Warning (Warning and Errors)	The vCenter Server collects warning and error entries in its log files.
Info (Normal Logging)	The vCenter Server collects information, warning, and error entries in its log files.

Verbose (Verbose)	The vCenter Server collects verbose, information, warning, and error entries in its log files.
Trivia (Extended Verbose)	The vCenter Server collects trivia, verbose, information, warning, and error entries in its log files.

Although setting the logging level to Verbose or Trivia may be beneficial for troubleshooting, doing so for long durations may cause noticeable vCenter Server performance degradation. VMware recommends that you use these levels in rare cases, such as while actively troubleshooting, and that you reset the logging level immediately afterward. Changes to the logging level are saved in the vCenter Server configuration file `/etc/vmware-vpx/vpxd.cfg`. You can make additional changes to logging behavior by editing the advanced settings of a vCenter Server. For example, you can use the vSphere Client to edit the following settings, which impact log size, retention, rotation, and compression:

- `config.log.level`
- `config.log.maxFileNum`
- `config.log.maxFileSize`
- `config.log.compressOnRoll`

### Note

By default, vCenter Server vpxd log files are rolled up and compressed into `.gz` files. You can turn off compression for vpxd log files by adding the `log.compressOnRoll` key with the value `false` to the vCenter Server advanced settings.

## Configuring Syslog on ESXi Hosts

You can use the following procedure to configure the syslog service for a host:

- Step 1.** In the vSphere Client, select a host in the inventory pane and navigate to **Configure > System > Advanced System Settings**.
- Step 2.** Click **Edit**.
- Step 3.** Filter for **syslog**.

**Step 4.** To set up logging globally for the following options, select the appropriate option and enter the appropriate value:

- **Syslog.global.defaultRotate:** Maximum number of logs to keep when rotating logs.
- **Syslog.global.defaultSize:** Size of log (in KB) before triggering a log rotation.
- **Syslog.global.LogDir:** Directory in a VMFS or NFS datastore to store logs specified in the format *[datastore] /path*. For example, to store logs in the */vmfs/volumes/VMFS-01/systemlogs* folder, specify *[VMFS-01] /systemlogs*.
- **Syslog.global.logDirUnique:** A subdirectory for the host at the specified path, which is useful when multiple hosts use the same shared datastore for logging.
- **Syslog.global.LogHost:** Remote syslog host and port to which messages are forwarded. For example, to forward to a server named *syslogsvr-1* using port 1514, specify *ssl://syslogsvr-1:1514*.

**Step 5.** Optionally, select specific log names and change the number of rotations and log size for just that specific log.

**Step 6.** Click **OK**.

You can control how log files are maintained for virtual machines. A new log file is created each time you power on or resume a virtual machine or whenever the file size exceeds the *vmx.log.rotateSize* value, unless the value is 0 (default). VMware recommends saving 10-log files, each one limited to no less than 2 MB. If you need logs for a longer time span, you can set *vmx.log.keepOld* to 20.

You can use the following procedure to change the number of log files for a single virtual machine:

**Step 1.** In the vSphere Client, right-click a host or a virtual machine in the inventory pane and click **Edit Settings**.

**Step 2.** Select **VM Options > Advanced**.

**Step 3.** Click **Edit Configuration**.

**Step 4.** Add or edit the *vmx.log.keepOld* parameter, set to the appropriate number.

**Step 5.** Click **OK**.

### Note

To set the `vmx.log.keepOld` value for all virtual machines on a specific host, edit the `/etc/vmware/config` file and add or edit a line like the following:

```
vmx.log.keepOld = "10"
```

You can modify the `/etc/vmware/logfilters` file on a host to change its logging behavior. In this file you can add an entry specifying the following options:

- Add **numLogs** to specify the maximum number of log entries before the specified log messages are filtered and ignored. Use 0 to filter and ignore all the specified log messages.
- Add **Ident** to specify one or more system components to apply the filter.
- Add **logRegexp** to specify a case-sensitive phrase to filter the log messages by their content.
- Add the following line to the `/etc/vmsyslog.conf` file: **enable\_logfilters = true**.
- Run the command **esxcli system syslog reload**.

## vRealize Log Insight (vRLI)

To collect and analyze vSphere data using vRLI, after you deploy vRLI, follow these steps:

**Step 1.** In the vRLI web interface, navigate to the **Administration** tab.

**Step 2.** Click **Integration > vSphere**.

**Step 3.** Provide the hostname and credentials to connect to a vCenter Server.

**Step 4.** Click **Test Connection**.

**Step 5.** If you use an untrusted SSL certificate, click **Accept** in the dialog.

**Step 6.** Click **Save**.

**Step 7.** Use the vRLI web interface to configure the data collection from vCenter.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

### Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 10-15](#) lists these key topics and the page number on which each is found.



**Table 10-15** Key Topics for [Chapter 10](#)

Key Topic Element	Description	Page Number
List	Add a host using Quickstart	<a href="#">365</a>
<a href="#">Table 10-2</a>	Use cases for VM-to-VM rules	<a href="#">369</a>
List	Configuring vSphere HA admission control	<a href="#">371</a>
Section	Configuring proactive HA	<a href="#">372</a>
Section	Performance graph overview	<a href="#">375</a>
List	Configuring Skyline Health	<a href="#">387</a>
Section	The impact of virtual machine configuration	<a href="#">392</a>

### Complete Tables and Lists from Memory

Print a copy of [Appendix B](#), “[Memory Tables](#)” (found on the companion website), or at least the section for this chapter, and complete the tables and lists

from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

[CPU Ready Time](#)  
[ESXTOP](#)  
[VIMTOP](#)  
[client performance charts](#)

## Review Questions

1. You are creating a resource pool in a vSphere DRS cluster. Which of the following is a default setting?
  - a. The Memory Limit is disabled.
  - b. CPU Shares is 0.
  - c. Memory Reservation is 0.
  - d. The CPU Reservation is normal.
2. You want to configure predictive DRS in your vSphere cluster. Which of the following is a requirement?
  - a. Set DRS to Fully Automated.
  - b. In the cluster, set Provide Data to vSphere Predictive DRS to True.
  - c. In the vRealize Operations, set Provide Data to vSphere Predictive DRS to True.
  - d. In the vRealize Automation, set Provide Data to vSphere Predictive DRS to True.
3. You are configuring a vSphere HA cluster and do not want it to automatically reserve resources for failure. What setting should you use?
  - a. Set Cluster Resource Percentage to 0.

- b.** Set Cluster Resource Percentage to 100.
  - c.** Set Define Host Failover Capacity to Dedicated Host Failures.
  - d.** Set Define Host Failover Capacity to Disabled.
- 4.** You want to use a command-line tool that shows real-time CPU statistics for the services running in the vCenter Server. Which should you choose?
  - a.** VIMTOP
  - b.** ESXTOP
  - c.** Performance charts
  - d.** vCenter Server Management Interface
- 5.** You are examining vSphere logs. Which of the following logs is in the same folder as the virtual machine configuration file?
  - a.** vpxa.log
  - b.** vmksummary.log
  - c.** auth.log
  - d.** vmware.log



# Chapter 11

## Managing Storage

This chapter covers the following topics:

- [Configuring and Managing vSAN](#)
- [Managing Datastores](#)
- [Storage DRS and SIOC](#)
- [NVMe and PMem](#)
- [Multipathing, Storage Policies, and vVols](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20), exam objectives 1.3.1, 1.3.2, 1.3.3, 1.3.4, 1.3.5, 1.4, 1.6.5, 1.9.1, 5.5, 7.4, 7.4.1, 7.4.2, and 7.4.3.

This chapter provides information on configuring and managing storage in a vSphere environment.

### “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 11-1](#) 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 and Review Questions.](#)”

**Table 11-1** “Do I Know This Already?” Section-to-Question Mapping

<b>Foundation Topics Section</b>	<b>Questions</b>
Configuring and Managing vSAN	1, 2

Managing Datastores	3, 4
Storage DRS and SIOC	5, 6
NVMe and PMem	7, 8
Multitasking, Storage Policies, and vVols	9, 10

1. You are configuring a hybrid vSAN cluster in a vSphere 7.0 environment. By default, what percentage of the flash space is used as a write buffer?

  - a. 100%
  - b. 70%
  - c. 30%
  - d. 0%
2. You are configuring vSAN in a vSphere 7.0 environment. Which of the following is supported when using Quickstart to configure a vSAN cluster?

  - a. ESXi 6.5.0
  - b. Network I/O Control (NIOC) Version 2
  - c. Hosts with dissimilar network configurations
  - d. Fibre Channel storage
3. You want to increase the size of a VMFS 6 datastore. Which one of the following statements is true?

  - a. You can extend the datastore by using available space on the storage device that is backing the datastore.
  - b. You can expand the datastore by using a separate storage device.
  - c. If the datastore is 100% full, you cannot increase its capacity.
  - d. You can expand the datastore by using available space on the storage device that is backing the datastore.
4. You are configuring NFS datastores for your vSphere 7.0 environment. Which one of the following statements is true?

  - a. You can use multiple IP addresses with any NFS datastore.

- b. You can use multiple IP addresses with NFS Version 4.1 but not with NFS Version 3.
  - c. You can use multiple IP addresses with NFS Version 3 but not with NFS Version 4.1.
  - d. You cannot use multiple IP addresses with any version of NFS.
5. You are configuring SIOC and want to change the threshold it uses to begin prioritizing I/O based on shares. Which of the following options is the acceptable range?
- a. 1 to 100 ms
  - b. 10 to 100 ms
  - c. 30 to 100 ms
  - d. 10 to 50 ms
6. You want to perform maintenance on a datastore that is a member of a data-store cluster. Which of the following actions should you take?
- a. Right-click the host and choose Enter Maintenance Mode.
  - b. Right-click the datastore and choose Enter Maintenance Mode.
  - c. Right-click the host and choose Enter SDRS Maintenance Mode.
  - d. Right-click the datastore and choose Enter SDRS Maintenance Mode.
7. You need to configure an ESXi 7.0 host to access shared NVMe devices using RDMA over Converged Ethernet (RoCE) Version 2. Which steps should you take? (Choose three.)
- a. Configure a VMkernel network adapter.
  - b. Add a software adapter to the host's network adapters.
  - c. Navigate to **Storage Adapters** > **RDMA Adapters** and verify the VMkernel adapter bindings.
  - d. Navigate to **Networking** > **RDMA Adapters** and verify the VMkernel adapter bindings.
  - e. Add a software adapter to the host's storage adapters.
8. In a vSphere 7.0 environment, you want to allow a virtual machine to

use NVDIMMs as standard memory. What should you configure?

- a. vPMemDisk
  - b. vPMem
  - c. NVMe-oF
  - d. RDMA
9. You want to set the path selection policy for a storage device managed by NMP such that it uses a preferred path. Which of the following policies should you choose?
- a. FIXED
  - b. LB\_RR
  - c. VMW\_PSP\_FIXED
  - d. VMW\_PSP\_RR
10. You are preparing to configure vVols in a vSphere 7.0 environment. Which of the following components should you configure in the storage system? (Choose two.)
- a. Protocol endpoints
  - b. Storage containers
  - c. LUNs
  - d. Virtual volumes

## Foundation Topics

### Configuring and Managing vSAN

This section provides information on configuring and managing vSAN clusters and vSAN datastores.

#### Preparing for vSAN

Before creating and configuring vSAN clusters, you should be aware of the following vSAN characteristics:

**Key  
Topic**

- Multiple vSAN clusters can be configured in a single vCenter Server instance.
- vSAN does not share devices with other vSphere features.
- At a minimum, a vSAN cluster must include three hosts with capacity devices. In addition, it can include hosts with or without capacity devices.
- For best results, use uniformly configured hosts in each vSAN cluster.
- If a host contributes capacity, it must have at least one flash cache device and one capacity device.
- In hybrid clusters, magnetic disks are used for capacity, and flash devices serve as a read cache and a write buffer. In a hybrid cluster, 70% of the flash space is used for the read cache, and 30% is used for the write buffer.
- In all-flash clusters, one designated flash device is used as a write cache, and additional flash devices are used for capacity. No read cache is used. All read requests come directly from the flash pool capacity.
- Only local (or directly attached) devices can participate in a vSAN cluster.
- Only ESXi 5.5 Update 1 or later hosts can join a vSAN cluster.
- Before you move a host from a vSAN cluster to another cluster, you need to make sure the destination cluster is vSAN enabled.
- To use a vSAN datastore, an ESXi host must be a member of the vSAN cluster.

It is important to ensure that you meet all the vSAN hardware, cluster, software, and network requirements described in [Chapter 2, “Storage Infrastructure.”](#)

## **Creating a vSAN Cluster with Quickstart**

Quickstart, which is described in [Chapter 10, “Managing and Monitoring Clusters and Resources,”](#) allows you to quickly create, configure, and expand a vSAN cluster, using recommended default settings for networking, storage, and services. It uses the vSAN health service to help you validate and correct configuration issues using a checklist consisting of green messages, yellow warnings, and red failures.

To use Quickstart to configure a vSAN cluster, the hosts must use ESXi 6.0 Update 2 or later. The hosts must have a similar network configuration to allow Quickstart to configure network settings based on cluster requirements. You can use Quickstart to configure vSAN on an existing cluster by using the following procedure:

- Step 1.** In the vSphere Client, select the cluster in the Hosts and Clusters inventory and click **Configure** > **Configuration** > **Quickstart**.
- Step 2.** On the Cluster Basics card, click **Edit**, select the vSAN service, and optionally select other services, such as DRS and vSphere HA. Then click **Finish**.
- Step 3.** Click **Add Hosts** > **Add** and use the wizard to add hosts to the cluster.
- Step 4.** On the Cluster Configuration card, click **Configure** and use the wizard to configure the following:
  - On the Configure the Distributed Switches page, enter networking settings, including distributed switches, port groups, and physical adapters.
  - On the vMotion Traffic page, enter the vMotion IP address information.
  - On the Storage Traffic page, enter the storage IP address information.
  - On the Advanced Options page, provide vSAN cluster settings. Optionally, provide settings for DRS, HA, and EVC.
  - On the Claim Disks page, select disks on each host to claim for vSAN cache and capacity.
  - Optionally, on the Create Fault Domains page, define fault domains for hosts that can fail together.
  - On the Ready to Complete page, verify the cluster settings and click **Finish**.

## Note

If you are running vCenter Server on a host, the host cannot be placed into Maintenance Mode as you add it to a cluster using the Quickstart workflow. The same host also can be running a Platform Services

Controller. All other virtual machines on the host must be powered off.

## Note

Distributed switches with Network I/O Control (NIOC) 2 cannot be used with vSAN Quickstart.

## Manually Enabling vSAN

While it is recommended that all of the ESXi hosts in a vSAN cluster contribute storage to that vSAN cluster, it is not required, and ESXi hosts without any capacity can be added to and make use of the vSAN cluster. This is possible provided that the following requirements are met:

- There must be at least three ESXi hosts in the vSAN cluster, and they must all contribute storage, or the cluster cannot tolerate host and/or device failures.
- ESXi 5.5 Update 1 or higher must be used on all the hosts in the cluster.
- If a host is being moved from one vSAN cluster to another, vSAN must be enabled on the destination cluster.
- ESXi hosts must be members of the vSAN cluster to access the vSAN data-store (regardless of whether they are contributing storage to the vSAN cluster).

You can use the following procedure to manually enable vSAN:

**Step 1.** Prepare a VMkernel network adapter on each participating host.

- In the vSphere Client, select a host in the inventory pane and navigate to **Networking > VMkernel Adapters**.
- Click the **Add Networking** icon.
- Use the wizard to configure the adapter's network settings and to enable vSAN.

**Step 2.** In the inventory pane, right-click a data center and select **New Cluster**.

- Step 3.** Provide a name for the cluster.
- Step 4.** Optionally, configure other cluster settings, such as DRS, vSphere HA, and EVC.
- Step 5.** Add hosts to the cluster.
- Step 6.** Navigate to **Configure > vSAN > Services** and click **Configure**.
- Step 7.** Select one of the following configuration types:
- **Single Site Cluster**
  - **Two Host Cluster**
  - **Stretched Cluster**
- Click **Next**.
- Step 8.** In the next wizard page, optionally configure the following:
- Enable **Deduplication and Compression** on the cluster.
  - Enable **Encryption** and select a KMS.
  - Select the **Allow Reduced Redundancy** checkbox to enable encryption or deduplication and compression on a vSAN cluster that has limited resources.
- Click **Next**.
- Step 9.** On the Claim Disks page, select the disks for use by the cluster and click **Next**.
- Step 10.** Follow the wizard to complete the configuration of the cluster, based on the fault tolerance mode:
- **For a two-host vSAN cluster:** Choose a witness host for the cluster and claim disks for the witness host.
  - **For a stretched cluster:** Define fault domains for the cluster, choose a witness host, and claim disks for the witness host.
  - **If you selected fault domains:** Define the fault domains for the cluster.
- Step 11.** On the Ready to Complete page, click **Finish**.



## Note

When claiming disks for each host that contributes storage to a vSAN cluster, select one flash device for the cache tier and one or more devices for the capacity tier.

## Editing vSAN Settings

You can modify the settings of an existing vSAN cluster by using the following procedure:

**Step 1.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Configure > vSAN > Services**.

**Step 2.** Click **Edit**.

**Step 3.** Optionally modify the following:

- Deduplication and compression
- vSAN encryption
- vSAN performance service
- iSCSI target
- **Advanced Settings > Object Repair Timer**
- **Advanced Settings > Site Read Locality for stretched clusters**
- Thin swap provisioning
- Large cluster support for up to 64 hosts
- Automatic rebalance

**Step 4.** Click **Apply**.

## Licensing vSAN

You need a vSAN license to use it beyond the evaluation period. The license capacity is based on the total number of CPUs in the hosts participating in the cluster. The vSAN license is recalculated whenever ESXi hosts are added to or removed from the vSAN cluster.

The Global.Licenses privilege is required on the vCenter Server. You can use the following procedure to assign a vSAN license to a cluster:

**Step 1.** In the vSphere Client, select the vSAN cluster in the inventory pane.

**Step 2.** On the **Configure** tab, right-click the vSAN cluster and choose **Assign License**.

**Step 3.** Select an existing license and click **OK**.

### Note

You can use vSAN in Evaluation Mode to explore its features for 60 days. To continue using vSAN beyond the evaluation period, you must license the cluster. Some advanced features, such as all-flash configuration and stretched clusters, require a license that supports the feature.

## Viewing a vSAN Datastore

When you enable vSAN on a cluster, a vSAN datastore is created. You can use the following procedure to review the capacity and other details of a vSAN datastore:

**Step 1.** In the vSphere Client, navigate to **Home > Storage**.

**Step 2.** Select the vSAN datastore.

**Step 3.** On the Configure tab, review the following:

- Capacity (total capacity, provisioned space, and free space)
- Datastore capabilities
- Policies

A vSAN datastore's capacity depends on the capacity devices per host and the number of hosts in the cluster. For example, if a cluster includes eight hosts, each having seven capacity drives, where each capacity drive is 2 TB, then the approximate storage capacity is  $8 \times 7 \times 2 \text{ TB} = 112 \text{ TB}$ .

Some capacity is allocated for metadata, depending on the on-disk format version:

- On-disk format Version 1.0 adds approximately 1 GB overhead per capacity device.
- On-disk format Version 2.0 adds overhead that is approximately 1% to 2% of the total capacity.
- On-disk format Version 3.0 and later adds overhead that is approximately 1% to 2% of the total capacity plus overhead for checksums used by deduplication and compression (approximately 6.2% of the total capacity).

## Configuring vSAN and vSphere HA

You can enable vSphere HA and vSAN on the same cluster but with some restrictions. The following ESXi requirements apply when using vSAN and vSphere HA together:

- ESXi Version 5.5 Update 1 or later must be used on all participating hosts.
- The cluster must have a minimum of three ESXi hosts.

The following networking differences apply when using vSAN and vSphere HA together:

- The vSphere HA traffic flows over the vSAN network rather than the management network.
- vSphere HA uses the management network only when vSAN is disabled.

Before you enable vSAN on an existing vSphere HA cluster, you must first disable vSphere HA. After vSAN is enabled, you can re-enable vSphere HA.

[Table 11-2](#) describes the vSphere HA networking differences between clusters where vSAN is enabled and is not enabled.

**Table 11-2** Network Differences in vSAN and non-vSAN Clusters

Factor	vSAN Is Enabled	vSAN Is Not Enabled
Network used by vSphere HA	vSAN network	Management network
Heartbeat	Any datastore, other than a vSAN	Any datastore that is

datastores	datastore, that is mounted to multiple hosts in the cluster	mounted to multiple hosts in the cluster
Host isolation criteria	Isolation addresses not pingable and vSAN storage network inaccessible	Isolation addresses not pingable and management network inaccessible

You must account for the vSAN's rule set's Primary Level of Failures to Tolerate setting when configuring the vSphere HA admission control policy. The vSAN primary level of failures must not be lower than the capacity reserved by the vSphere HA admission control setting. If vSphere HA reserves less capacity, failover activity might be unpredictable. For example, for an eight-host cluster, if you set the vSphere HA admission control to more than 25% of the cluster resources, then you should not set the vSAN rule's Primary Level of Failures to Tolerate setting higher than two hosts.

In response to events involving the failure of multiple hosts in a cluster where vSphere HA and vSAN are enabled, vSphere HA may not be able to restart some virtual machines where the most recent copy of an object is inaccessible. For example, consider the following scenario:

- In a three-host cluster with vSAN and vSphere HA enabled, two hosts fail.
- A VM continues to run on the third host.
- The final host fails.
- The first two hosts are recovered.
- vSphere HA cannot restart the VM because the most recent copy of its object is on the third host, which is still unavailable.

## Setting the HA Capacity Reservation Setting

Capacity can be reserved for failover in the vSphere HA admission control policies. Such a reservation must be coordinated with the vSAN policy Primary Level of Failures to Tolerate. The HA reserved capacity cannot be higher than the vSAN Primary Level of Failures to Tolerate setting.

For example, if you set vSAN Primary Level of Failures to Tolerate to 1, the HA admission control policy must reserve resources equal to those of one host. If you set the vSAN Primary Level of Failures to Tolerate to 2, the HA admission control policy must reserve resources equal to those of two ESXi hosts.

## Disabling vSAN

You can use the following procedure to disable vSAN for a host cluster, which causes all virtual machines located on the vSAN datastore to become inaccessible:

- Step 1.** In the vSphere Client, select the cluster in the inventory pane.
- Step 2.** Verify that the host in the cluster is in Maintenance Mode.
- Step 3.** Select **Configure > vSAN > Services**.
- Step 4.** Click **Turn Off vSAN**.
- Step 5.** In the dialog box that appears, confirm your selection.

### Note

If you intend to use virtual machines while vSAN is disabled, make sure you migrate the virtual machines from a vSAN datastore to another datastore before disabling the vSAN cluster.

## Shutting Down and Restarting vSAN

To shut down an entire vSAN cluster prior to performing some maintenance activities, vSAN does not have to be disabled. The following procedure details how you can shut down a vSAN cluster:



- Step 1.** Power off all virtual machines in the vSAN cluster except for the vCenter Server, if it is running in the cluster.
- Step 2.** In the vSphere Client, select the cluster and navigate to **Monitor > vSAN > Resyncing Objects**.
- Step 3.** When all resynchronization tasks are complete, on the Configure tab, turn off DRS and HA.
- Step 4.** On each host, use the following command to disable cluster member updates:

[Click here to view code image](#)

```
esxcfg-advcfg -s 1 /VSAN/IgnoreClusterMemberListUpdates
```

**Step 5.** If vCenter Server runs in the vSAN cluster, shut it down. (This makes the vSphere Client unavailable.)

**Step 6.** On each host, use the following command to place the hosts in Maintenance Mode with no data migration:

[Click here to view code image](#)

```
esxcli system maintenanceMode set -e true -m noAction
```

**Step 7.** Shut down each host.

### Note

When you plan to shut down a vSAN cluster, you do not need to disable vSAN on the cluster.

After you perform maintenance activities, you can restart a vSAN cluster by using the following procedure:

**Step 1.** Power on the hosts.

**Step 2.** Use the hosts' consoles to monitor the ESXi startup

**Step 3.** Optionally, use a web browser to connect directly to the ESXi host client to monitor the host's status, events, and logs. You can ignore misconfiguration status messages that appear temporarily when fewer than three hosts have come online and joined the cluster.

**Step 4.** On each host, use the following commands to exit Maintenance Mode and to ensure that each host is available in the cluster:

[Click here to view code image](#)

```
esxcli system maintenanceMode set -e false  
esxcli vsan cluster get
```

**Step 5.** Restart the vCenter Server VM.

**Step 6.** On each host, use the following command to re-enable updates:

[Click here to view code image](#)

```
esxcfg-advcfg -s 0 /VSAN/IgnoreClusterMemberListUpdates
```

**Step 7.** In the vSphere Client, select the vSAN cluster in the inventory pane.

**Step 8.** On the Configure tab, re-enable DRS and HA.

You can now start virtual machines in the cluster and monitor the vSAN health service.

## Deploying vSAN with vCenter Server

You can simultaneously deploy a VCSA and create a vSAN cluster by using the vCenter Server installer and following these steps:

**Step 1.** Create a single-host vSAN cluster.

**Step 2.** Place the vCenter Server on the host in the cluster.

**Step 3.** Choose **Install** on a new vSAN cluster containing the target host.

This process deploys a one-host vSAN cluster. After the deployment, you can use the vSphere Client to configure the vSAN cluster and add additional nodes to the cluster.

## Expanding a vSAN Cluster

You can expand a vSAN cluster by adding to the cluster ESXi hosts with storage. Keep in mind that ESXi hosts without local storage can also be added to a vSAN cluster. You can use the following procedure to expand a vSAN cluster by adding hosts:

**Step 1.** In the vSphere Client, right-click a cluster in the inventory pane and select **Add Hosts**

**Step 2.** Using the wizard, add hosts using one of the following options:

- **New Hosts:** Provide the host name and credentials.
- **Existing Hosts:** Select a host in the inventory that is not yet in the cluster.

**Step 3.** Complete the wizard and click **Finish** on the final page.

You can also use the following procedure to move multiple existing ESXi hosts into a vSAN cluster by using host profiles:

**Step 1.** In the vSphere Client, navigate to **Host Profiles**.

- Step 2.** Click the **Extract Profile from a Host** icon.
- Step 3.** Select a host in the vSAN cluster that you want to use as the reference host and click **Next**.
- Step 4.** Provide a name for the new profile and click **Next**.
- Step 5.** On the next wizard page, click **Finish**.
- Step 6.** In the Host Profiles list, select the new host profile and attach multiple hosts to the profile.
- Step 7.** Click the **Attach/Detach Hosts and Clusters to a Host Profile** icon.
- Step 8.** Detach the reference vSAN host from the host profile.
- Step 9.** In the Host Profiles list, select the new host profile and click the **Check Host Profile Compliance** icon.
- Step 10.** Select the **Monitor > Compliance**.
- Step 11.** Right-click the host and select **All vCenter Actions > Host Profiles > Remediate**.
- Step 12.** When prompted, provide appropriate input parameters for each host and click **Next**.
- Step 13.** Review the remediation tasks and click **Finish**.

The hosts and their resources are now part of the vSAN cluster.

You can use the following procedure to add hosts to a vSAN cluster by using Quickstart:

- Step 1.** Verify that no network configuration that was previously performed through the Quickstart workflow has been modified from outside the Quickstart workflow.
- Step 1.** In the vSphere Client, select the vSAN cluster in the inventory and click **Configure > Configuration > Quickstart**.
- Step 2.** Click **Add hosts > Launch**.
- Step 3.** Use the wizard to provide information for new hosts or to select existing hosts from the inventory.
- Step 4.** Complete the wizard and click **Finish** on the last page.



- Step 5.** Click **Cluster Configuration > Launch**.
- Step 6.** Provide networking settings for the new hosts.
- Step 7.** On the Claim Disks page, select disks on each new host.
- Step 8.** On the Create Fault Domains page, move the new hosts into their corresponding fault domains.
- Step 9.** Complete the wizard and click **Finish**.

### Note

When adding a host to a vSAN cluster by using Quickstart, the vCenter Server must not be running on the host.

## Working with Maintenance Mode

Before shutting down, rebooting, or disconnecting a host that is a member of a vSAN cluster, you must put the ESXi host in Maintenance Mode. Consider the following guidelines for using Maintenance Mode for vSAN cluster member hosts:

- When entering host Maintenance Mode, you must select a data evacuation mode, such as Ensure Accessibility or Full Data Migration.
- When a vSAN cluster member host enters Maintenance Mode, the cluster capacity is automatically reduced.
- Each impacted virtual machine may have compute resources, storage resources, or both on the host entering Maintenance Mode.
- Ensure Accessibility Mode, which is faster than Full Data Migration Mode, migrates only the components from the host that are essential for running the virtual machines. It does not reprotect your data. When in this mode, if you encounter a failure, the availability of your virtual machine is affected, and you might experience unexpected data loss.
- When you select Full Data Migration Mode, your data is automatically reprotected against a failure (if the resources are available and Primary Level of Failures to Tolerate is set to 1 or more). In this mode, your virtual machines can tolerate failures, even during planned maintenance.

- When working with a three-host cluster, you cannot place a server in Maintenance Mode with Full Data Migration Mode.

Prior to placing a vSAN cluster member host in Maintenance Mode, you must do the following:

- If using Full Data Migration Mode, ensure that the cluster has enough hosts and available capacity to meet the requirements of the Primary Level of Failures to Tolerate policy.
- Verify that remaining hosts have enough flash capacity to meet any flash read cache reservations. To analyze this, you can run the following VMware Ruby vSphere Console (RVC) command:

```
vsan.whatif_host_failures
```

- Verify that the remaining hosts have devices with sufficient capacity to handle stripe width policy requirements, if selected.
- Make sure that you have enough free capacity on the remaining hosts to handle the data that must be migrated from the host entering Maintenance Mode.

You can use the Confirm Maintenance Mode dialog box to determine how much data will be moved, the number of objects that will become noncompliant or inaccessible, and whether sufficient capacity is available to perform the operation. You can use the Data Migration Pre-check button to determine the impact of data migration options when placing a host into Maintenance Mode or removing it from the cluster.

To place a vSAN cluster member host in Maintenance Mode, you can use the following procedure:

**Step 1.** In the vSphere Client, select the cluster in the inventory pane.

**Step 2.** Optionally, use the following steps to run Data Migration Pre-check:

- Click **Data Migration Pre-check**.
- Select a host and a data migration option and click **Pre-check**.
- View the test results and decide whether to proceed.

**Step 3.** Right-click the host and select **Maintenance Mode > Enter**

## Maintenance Mode.

**Step 4.** Select one of the following data evacuation modes:

- **Ensure Accessibility:** If hosts are powered off or removed from a vSAN cluster, vSAN makes sure the virtual machines on the ESXi host that is removed can still run those virtual machines. This moves some of the virtual machine data off the vSAN cluster, but replica data remains. If you have a three-host cluster, this is the only evacuation mode available.
- **Full Data Migration:** As its name implies, this mode moves all the VM data to other ESXi hosts in the cluster. This option makes sense if you are removing the host from the cluster permanently. If a virtual machine has data on the host and that data is not migrated off, the host cannot enter this mode.
- **No Data Migration:** If this option is selected, vSAN does not move any data from this ESXi host.

Click **OK**.

## Managing vSAN Fault Domains

Fault domains provide additional protection against outage in the event of a rack or blade chassis failure. A **vSAN fault domain** contains at least one vSAN host, depending on the physical location of the host. With fault domains, vSAN can withstand rack, blade chassis, host, disk, or network failure within one fault domain, as the replica and witness data are stored in a different fault domain.

You can use the following procedure to create a new fault domain in a vSAN cluster:

- Step 1.** In the vSphere Client, examine each host in a vSAN cluster.
- Step 2.** Verify that each host is running ESXi 6.0 or later (to support fault domains) and is online.
- Step 3.** Select the vSAN cluster in the inventory pane and click **Configure > vSAN > Fault Domains**.
- Step 4.** Click the **Add** (plus sign) icon.
- Step 5.** In the wizard, provide a name for the fault domain.

**Step 6.** Select one or more hosts to add to the fault domain.

**Step 7.** Click **Create**.

You can use the vSphere Client to add hosts to an existing fault domain by selecting **Configure > vSAN > Fault Domains** and dragging the host to the appropriate fault domain. Likewise, you can drag a host out of a fault domain to remove the host from the fault domain and create a single-host fault domain.

## Extending a vSAN Datastore Across Two Sites

A vSAN stretched cluster extends across two physical data center locations to provide availability in the event of site failure as well as provide load balancing between sites. With a stretched vSAN cluster, both sites are active, and if either site fails, vSAN uses storage on the site that is still up. One site must be designated as the preferred site, which makes the other site the secondary, or nonpreferred, site.

You can use the following procedure to leverage Quickstart to create a stretched cluster across two sites:



**Step 1.** Ensure that the following prerequisites are met:

- A host is deployed outside any cluster for the witness host.
- ESXi 6.0 Update 2 or later is used on each host.
- The hosts in the cluster do not have any existing vSAN or networking configuration.

**Step 2.** Click **Configure > Configuration > Quickstart**.

**Step 3.** Click **Cluster Configuration > Edit**.

**Step 4.** In the wizard, provide a cluster name, enable vSAN, and optionally enable other features, such as DRS or vSphere HA.

**Step 5.** Click **Finish**.

**Step 6.** Click **Add Hosts > Add**.

**Step 7.** In the wizard, provide information for new hosts or select existing hosts from the inventory. Click **Finish**.

**Step 8.** Click **Cluster Configuration > Configure**.

**Step 9.** In the wizard, configure the following:

- Configure settings for distributed switch port groups, physical adapters, and the IP configuration associated with vMotion and storage.
- Set vSAN Deployment Type to **Stretched Cluster**.
- On the Claim Disk page, select disks on each host for cache and capacity.
- On the Create Fault Domains page, define fault domains for the hosts in the preferred site and the secondary site.
- On the Select Witness Host page, select a host to use as a witness host. This host cannot be part of the cluster and can have only one VMkernel adapter configured for vSAN data traffic.
- On the Claim Disks for Witness Host page, select disks on the witness host for cache and capacity.
- On the Ready to Complete page, verify the cluster settings and click **Finish**.

When creating a vSAN stretched cluster, DRS must be enabled on the cluster. There are also several DRS requirements for stretched vSAN clusters:

- Two host groups must be created: one for the preferred site and another for the secondary site.
- Two VM groups must be created: one for the preferred site VMs and one for the VMs on the secondary site.
- Two VM-to-host affinity rules must be created for the VMs on the preferred site and VMs on the secondary site.
- VM-to-host affinity rules must be used to define the initial placement of virtual machines on ESXi hosts in the cluster.

In addition to the DRS requirements, there are also HA requirements for stretched vSAN clusters:

- HA must be enabled.

- HA rules should allow the VM-to-host affinity rules in the event of a failover.
- HA datastore heartbeats should be disabled.

vSAN has numerous requirements for implementing stretched clusters:

- Stretched clusters must use on-disk format Version 2.0 or higher. If your vSAN cluster is not using on-disk format Version 2.0, it must be upgraded before you configure the stretched vSAN cluster.
- Failures to Tolerate must be set to 1.
- Symmetric Multiprocessing Fault Tolerance (SMP-FT) VMs are supported only when PFFT is at 0 and Data Locality is either Preferred or Secondary. SMP-FT VMs with PFFT set to 1 or higher are not supported.
- If hosts are disconnected or fail in a not responding state, the witness cannot be added or removed.
- Adding ESXi hosts via **esxcli** commands on stretched clusters is not supported.

## Managing Devices in a vSAN Cluster

You can use the following procedure to create a disk group on a vSAN cluster member host:

- Step 1.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Configure > vSAN > Disk Management**.
- Step 2.** Select the host and click **Create Disk Group**.
- Step 3.** Select the flash device to be used for the cache.
- Step 4.** Select the type of capacity disks to use (**HDD** for hybrid or **Flash** for all-flash).
- Step 5.** Select the devices you want to use for capacity.
- Step 6.** Click **Create** or **OK**.

You can use the following procedure to claim storage devices for a vSAN cluster:

- Step 1.** In the vSphere Client, select the cluster in the inventory pane and

navigate to **Configure > vSAN > Disk Management > Claim Unused Disks**.

**Step 2.** Select a flash device to be used for the cache and click **Claim** for the cache tier.

**Step 3.** Select one or more devices (**HDD** for hybrid or **Flash** for all-flash) to be used as capacity and click **Claim** for the capacity tier.

**Step 4.** Click **Create** or **OK**.

To verify that the proper role (cache or capacity) has been assigned to each device in an all-flash disk group, examine the Disk Role column at the bottom of the Disk Management page. If the vSAN cluster is set to claim disks in manual mode, you can use the following procedure to add additional local devices to an existing disk group. The additional devices must be the same type (flash or HDD) as existing devices in the disk group:

**Step 1.** In the vSphere Cluster, select the vSAN cluster in the inventory pane and navigate to **Configure > vSAN > Disk Management**.

**Step 2.** Select the disk group and click **Add Disks**.

**Step 3.** Select the device and click **Add**.

### Note

If you add a used device that contains residual data or partition information, you must first clean the device. For example, you can run the RVC command **host\_wipe\_vsan\_disks**.

You can use the following procedure to remove specific devices from a disk group or remove an entire disk group. However, you should typically do so only when you are upgrading a device, replacing a failed device, or removing a cache device. Deleting a disk group permanently deletes the data stored on the devices. Removing one flash cache device or all capacity devices from a disk group removes the entire disk group. Follow these steps to remove specific devices from a disk group or remove an entire disk group:

**Step 1.** In the vSphere Cluster, select the vSAN cluster in the inventory pane.

**Step 2.** Click **Configure > vSAN > Disk Management**.

**Step 3.** To remove a disk group, select the disk group, click **Remove**, and select a data evacuation mode.

**Step 4.** To remove a device, select the disk group, select the device, click **Remove**, and select a data evacuation mode.

**Step 5.** Click **Yes** or **Remove**.

If ESXi does not automatically identify your devices as being flash devices, you can use the following procedure to manually mark them as local flash devices. For example, flash devices that are enabled for RAID 0 Mode rather than Passthrough Mode may not be recognized as flash. Marking these devices as local flash makes them available for use as vSAN cache devices. Before starting this procedure, you should verify that the device is local and not in use.

**Step 1.** In the vSphere Cluster, select the vSAN cluster in the inventory pane and navigate to **Configure > vSAN > Disk Management**.

**Step 2.** Select a host to view the list of available devices.

**Step 3.** In the **Show** drop-down menu, select **Not in Use**.

**Step 4.** Select one or more devices from the list and click **Mark as Flash Disk**.

**Step 5.** Click **Yes**.

Likewise, you can use this procedure in other scenarios where you want to change how a device is identified. In step 4, you can choose Mark as HDD Disk, Mark as Local Disk, or Mark as Remote.

## Increasing Space Efficiency in a vSAN Cluster

To increase space efficiency in a vSAN cluster, you can use SCSI Unmap, deduplication, compression, RAID5 erasure coding, and RAID 6 erasure encoding.

Unmap capability is disabled by default. To enable SCSI Unmap on a vSAN cluster, use the RVC command **vsan.unmap\_support --enable**.

### Note

Unmap capability is disabled by default. When you enable Unmap on a vSAN cluster, you must power off and then power on all VMs. VMs must use virtual hardware Version 13 or above to perform Unmap operations.



---

When you enable or disable deduplication and compression, vSAN performs a rolling reformat of every disk group on every host. Depending on the data stored on the vSAN datastore, this process might take a long time. Do not perform such operations frequently. If you plan to disable deduplication and compression, you must first verify that enough physical capacity is available to place your data.

You should consider the following when managing disks in a vSAN cluster where deduplication and compression are enabled:

- For efficiency, consider adding a disk group to cluster capacity instead of incrementally adding disks to an existing disk group.
- When you add a disk group manually, add all the capacity disks at the same time.
- You cannot remove a single disk from a disk group. You must remove the entire disk group in order to make modifications.
- A single disk failure causes an entire disk group to fail.

To enable deduplication and compression for an existing vSAN cluster, you can use the following procedure:

**Step 1.** Verify that the cluster is all-flash.

**Step 2.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Configure > vSAN > Services**.

**Step 3.** Click **Edit**.

**Step 4.** Enable **Deduplication and Compression**.

**Step 5.** Optionally, select **Allow Reduced Redundancy**.

**Step 6.** Click **Apply** or **OK**.

When you enable deduplication and compression, vSAN updates the on-disk format of each disk group of the cluster by evacuating data from the disk group, removing the disk group, and re-creating it with a new format. This operation does not require virtual machine migration or DRS. If you choose the Allow Reduced Redundancy option, the virtual machines may continue to keep running even if the cluster does not have enough resources for the disk group to be fully evacuated. In this case, your virtual machines might be at risk of experiencing

data loss during the operation.

You can use the vSphere Client to check the storage savings provided by deduplication and compression. To do so, select a cluster, navigate to Monitor > Capacity, and examine Capacity Overview.

To use RAID 5 erasure coding in a vSAN cluster, set the following options:

- Set Failure Tolerance Method to RAID-5/6 (Erasure Coding)–Capacity.
- Set Primary Level of Failures to Tolerate to 1.

To use RAID 6 erasure coding in a vSAN cluster, set the following options:

- Set Failure Tolerance Method to RAID-5/6 (Erasure Coding)–Capacity.
- Set Primary Level of Failures to Tolerate to 2.

To use RAID 1, set Failure Tolerance Method to RAID-1 (Mirroring)–Performance.

### **Note**

RAID 5 and RAID 6 erasure coding do not support Primary Level of Failures set to a value higher than 2.

## **Using Encryption in a vSAN Cluster**

When planning to implement vSAN encryption, you should consider the following:

- To provide the encryption keys for the vSAN datastore, you must implement a key management server (KMS) cluster server that is KMIP 1.1 compliant and is in the vSphere compatibility matrices.
- You should not deploy the KMS server on the same vSAN datastore that it will help encrypt.
- Encryption is CPU intensive. Enable AES-NI in your BIOS.
- In a stretched vSAN cluster, the Witness host only stores metadata and does not participate in encryption.

- You should establish a policy regarding the encryption of coredumps because they contain sensitive information such as keys for hosts. In the policy, consider the following:
  - You can use a password when you collect a vm-support bundle.
  - The password re-encrypts coredumps that use internal keys based on the password.
  - You can later use the password to decrypt the coredumps in the bundle.
  - You are responsible for keeping track of the password. It is not saved anywhere in vSphere.

To use encryption in a vSAN datastore, you must add a KMS to the vCenter Server and establish trust with the KMS. You can use the following procedure to add a KMS to vCenter Server:

- Step 1.** Ensure that the user has the `Cryptographer.ManageKeyServers` privilege.
- Step 2.** In the vSphere Client, select the vCenter Server in the inventory pane and navigate to **Configure > Key Management Servers**.
- Step 3.** Click **Add** and specify the following KMS information in the wizard:
  - For KMS Cluster, select **Create New Cluster**.
  - Specify the cluster name, alias, and address (FQDN or IP address).
  - Specify the port, proxy, and proxy port.
- Step 4.** Click **Add**.

### Note

Connecting to a KMS through a proxy server that requires a username or password is not supported. Connecting to a KMS by using only an IPv6 address is not supported.

You can use the following procedure to establish a trusted connection for a KMS:

- Step 1.** In the vSphere Client, select the vCenter Server in the inventory pane and navigate to **Configure > Key Management Servers**.
- Step 2.** Select the KMS instance and click **Establish Trust with KMS**.
- Step 3.** Select one of the following options, as appropriate for the selected KMS instance:
- **Root CA Certificate**
  - **Certificate**
  - **New Certificate Signing Request**
  - **Upload Certificate and Private Key**

When multiple KMS clusters are used, you can use the following procedure to identify a default KMS cluster:

- Step 1.** In the vSphere Client, select the vCenter Server in the inventory pane and navigate to **Configure > Key Management Servers**.
- Step 2.** Select the KMS cluster and click **Set KMS Cluster as Default**.
- Step 3.** Click **Yes**.
- Step 4.** Verify that the word default appears next to the cluster name.

You can make vCenter Server trust the KMS by using the following procedure:

- Step 1.** In the vSphere Client, select the vCenter Server in the inventory pane and navigate to **Configure > Key Management Servers**.
- Step 2.** Select the KMS instance and do one of the following:
- Select **All Actions > Refresh KMS Certificate > Trust**.
  - Select **All Actions > Upload KMS Certificate > Upload File**.

If you want to enable encryption on a vSAN cluster, you need the following privileges:

- Host.Inventory.EditCluster
- Cryptographer.ManageEncryptionPolicy
- Cryptographer.ManageKMS

- Cryptographer.ManageKeys

You can use the following procedure to enable encryption on a vSAN cluster:

- Step 1.** In the vSphere Client, select the cluster in the inventory pane and navigate to **vSAN > Services**.
- Step 2.** Click the **Edit** button.
- Step 3.** In the vSAN Services dialog, enable **Encryption** and select a KMS cluster.
- Step 4.** Optionally, select the **Erase Disks Before Use** checkbox, based on the following:
  - If this is a new cluster with no virtual machines, you can deselect the checkbox.
  - If it is an existing cluster with unwanted data, select the checkbox, which increases the processing time for each disk.
- Step 5.** Click **Apply**.

To generate new encryption keys, you can use the following procedure:

- Step 1.** Log on to the vSphere Client as a user with Host.Inventory.EditCluster and Cryptographer.ManageKeys privileges.
- Step 2.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Configure > vSAN > Services**.
- Step 3.** Click **Generate New Encryption Keys**.
- Step 4.** To generate a new KEK, click **Apply**. Each host's DEK is re-encrypted with the new KEK.
- Step 5.** Optionally, select **Also Re-encrypt All Data on the Storage Using New Keys**.
- Step 6.** Optionally, select the **Allow Reduced Redundancy** checkbox, which may put your data at risk during the disk reformatting operation.

If a host member of a vSAN cluster that uses encryption has an error, the resulting coredump is encrypted. Coredumps that are included in the vm-support package are also encrypted.

## Using vSAN Policies

Virtual machine performance and availability requirements can be defined for vSAN, if required. Once virtual machines are created, their storage policy is enforced on the vSAN datastore. Underlying components of virtual disks are spread across the vSAN datastore to meet the requirements defined in the storage policy. Storage providers provide information about the physical storage to vSAN to assist with placement and monitoring.

## Creating a vSAN Storage Policy

The following procedure can be used to create a vSAN storage policy:

- Step 1.** In the vSphere Client, go to **Policies and Profiles > VM Storage Policies**.
- Step 2.** Click on the **Create a New VM Storage Policy** icon.
- Step 3.** On the Name and Description page, select an appropriate vCenter Server, enter a name and description for the policy, and click **Next**.
- Step 4.** On the Policy Structure page, select **Enable Rules for “vSAN” Storage** and click **Next**.
- Step 5.** On the vSAN page, set the policy:
  - a.** On the Availability tab, set Site Disaster Tolerance and Failures to Tolerate.
  - b.** On the Advanced Policy Rules tab, set Disk Stripes per Object and IOPS Limit.
  - c.** On the Tags tab, click **Add Tag Rule** and configure its options. Click **Next**.
- Step 6.** On the Storage Compatibility page, review the list of compatible datastores and click **Next**.
- Step 7.** On the Review and Finish page, review all the settings and click **Finish**.

## Changing the vSAN Default Storage Policy

vSAN datastore default policies can be changed, if desired, using the following procedure:

- Step 1.** In the vSphere Client storage inventory view, right-click the vSAN data-stores and select **Configure**.
- Step 2.** Select **General**, click **Edit** next to the default storage policy, and select a storage policy to be defined as the new default.
- Step 3.** Click **OK**.

## Viewing vSAN Storage Providers

vSAN 6.7 and above register one storage provider for all vSAN clusters managed by vCenter. To access the storage providers, use the URL <https://VCfqdn:VCport/vsanHealth/vsanvp/version.xml>.

To view the vSAN storage providers, in the vSphere client, select a vCenter Server and navigate to **Configure > Storage Providers**.

Each ESXi host has a vSAN storage provider, but only one is active. Storage providers on other ESXi hosts are in standby. If an ESXi host with an active storage provider fails, a storage provider from another host activates.

## Using vSAN File Service

You can use the following procedure to configure (enable) the **vSAN file service** on a vSAN cluster, which enables you to create file shares:

- Step 1.** Address the following prerequisites:
  - Identify a set of available IPv4 addresses, preferably one per host (for best performance), that are from the same subnet and are part of the forward and reverse lookup zones in the DNS server.
  - Create a dedicated distributed port group.
  - Ensure that vDS 6.6.0 or higher is in use.
  - Promiscuous Mode and forged transmits are enabled during file services configuration. If an NSX-based network is used, you must provide similar settings.
- Step 2.** In the vSphere Client, select the vSAN cluster and select **Configure > vSAN > Services**.
- Step 3.** In the File Service row, click **Enable**.

**Step 4.** In the wizard, click **Next**.

**Step 5.** On the next page, select either of the following options:

- **Automatic:** Automatically searches for and downloads the OVF
- **Manual:** Allows you to manually select an OVF and associated files (CERT, VMDK, and so on)

**Step 6.** Continue the wizard to provide file service domain, DNS, and networking information.

**Step 7.** On the IP Pool page, enter the set of available IPv4 addresses and assign one as the primary IP address. To simplify this process, you can use the Auto Fill or Look Up DNS options.

## Note

vSAN stretched clusters do not support the file service.

You can use the following procedure to create a vSAN file service:

**Step 1.** In the vSphere Client, select the vSAN cluster in the inventory pane and navigate to **Configure > vSAN > File Service Shares**.

**Step 2.** Click **Add**.

**Step 3.** In the wizard, enter the following general information:

- **Protocol:** Select either NFS Version 3 or NFS Version 4.1.
- **Name:** Specify a name.
- **Storage Policy:** Select the vSAN default storage policy.
- **Storage space quotas:** Set the share warning threshold and the share hard quota.
- **Labels:** Specify up to 50 labels (key-value pairs) per share, a label key (up to 250 characters), and a label value (fewer than 1000 characters).

Click **Next**.



**Step 4.** In the Net Access Control page, select one of the following options:

- **No Access:** Use this option to prevent access to the file share.
- **Allow Access from Any IP:** Use this to allow access from any IP address.
- **Customize Net Access:** Use this to control whether specific IP addresses can access, read, or modify the file share. You can configure Root Squash based on IP address.

Click **Next**.

**Step 5.** In the Review page, click **Finish**.

## Managing Datastores

This section provides information on managing datastores in a vSphere 7.0 environment.

### Managing VMFS Datastores

You can set up VMFS datastores on any SCSI-based storage device that is discovered by a host, such as a Fibre Channel device, an iSCSI device, or a local device. To view a host's SCSI devices, you can use the following procedure:

**Step 1.** In the vSphere Client, select an ESXi host in the inventory pane and navigate to **Configure > Storage > Storage Adapters**.

**Step 2.** Select a storage adapter.

**Step 3.** Optionally, click the **Rescan Adapter** or **Rescan Storage** button.

**Step 4.** In the details pane, select the **Devices** tab and examine the details for each discovered SCSI device, including type, capacity, and assigned datastores.

**Step 5.** Optionally, to manipulate a specific device, select the device and click the **Refresh**, **Attach**, or **Detach** button.

To create a VMFS 6 datastore on a SCSI device, you can use the following procedure:

**Step 1.** In the vSphere Client, right-click a host in the inventory pane and select **Storage > New Datastore**.

- Step 2.** For datastore type, select **VMFS** and click **Next**.
- Step 3.** Provide a name for the datastore, select an available SCSI device, and click **Next**.
- Step 4.** Select **VMFS 6** and click **Next**.
- Step 5.** Keep the default Partition Configuration setting **Use All Available Partitions**. Alternatively, set the datastore size, block size, space reclamation granularity, and space reclamation priority.
- Step 6.** Click **Next**.
- Step 7.** On the Ready to Complete page, click **Finish**.

You can increase the size of a VMFS datastore by adding an extent or by expanding the datastore within its own extent. A VMFS datastore can span multiple devices. Adding an extent to a VMFS datastore means adding a storage device (LUN) to the datastore. A spanned VMFS datastore can use any extent at any time. It does not require filling up a specific extent before using the next one.

A datastore is expandable when the backing storage device has free space immediately after the datastore extent. You can use the following procedure to increase the size of a datastore:



- Step 1.** In the vSphere Client, right-click the datastore in the inventory pane and select **Increase Datastore Capacity**.
- Step 2.** Select a device from the list of storage devices, based on the following.
  - To expand the datastore, select a storage device whose Expandable column contains YES.
  - To add an extent to the datastore, select a storage device whose Expandable column contains NO.
- Step 3.** Review the available configurations in the partition layout.
- Step 4.** In the menu, select one of the following available configuration options, depending on your previous selections:
  - **Use Free Space to Expand the Datastore:** Select this option to

expand the existing datastore and disk partition to use the adjacent disk space.

- **Use Free Space:** Select this option to deploy an extent in the remaining free space.
- **Use All Available Partitions:** Select this option to reformat a disk and deploy an extent using the entire disk. (This option is available only for non-blank disks.)

**Step 5.** Set the capacity. (The minimum extent size is 1.3 GB.) Click **Next**.

**Step 6.** Click **Finish**.

### Note

If a shared datastore becomes 100% full and has powered-on virtual machines, you can increase the datastore capacity—but only from the host where the powered-on virtual machines are registered.

Each VMFS datastore is assigned a universally unique ID (UUID). A storage device operation, such as a LUN snapshot, LUN replication, or a LUN ID change, might produce a copy of the original datastore such that both the original and a copy device contain a VMFS datastore with identical signatures (UUID). When ESXi detects a VMFS datastore copy, it allows you to mount it with the original UUID or mount it with a new UUID. The process of changing the UUID is called *resignaturing*.

To allow a host to use the original datastore and the copy, you can choose to resignature the copy. If the host will only access the copy, you could choose to mount the copy without resignaturing.

You should consider the following:

- When resignaturing a datastore, ESXi assigns a new UUID to the copy, mounts the copy as a datastore that is distinct from the original, and updates all corresponding UUID references in the virtual machine configuration files.
- Datastore resignaturing is irreversible.

- After resignaturing, the storage device is no longer treated as a replica.
- A spanned datastore can be resignatured only if all its extents are online.
- The resignaturing process is fault tolerant. If the process is interrupted, you can resume it later.
- You can mount the new VMFS datastore without risk of its UUID conflicting with UUIDs of any other datastore from the hierarchy of device snapshots.

To mount a VMFS datastore copy on an ESX host, you can use the following procedure:

- Step 1.** In the vSphere Client, select the host in the inventory page and navigate to **Configure > Storage Adapters**.
- Step 2.** Rescan storage.
- Step 3.** Unmount the original VMFS datastore, which has the same UUID as the VMFS copy.
- Step 4.** Right-click the host and select **Storage > New Datastore**.
- Step 5.** Select **VMFS** as the datastore type.
- Step 6.** Enter the datastore name and placement (if necessary).
- Step 7.** In the list of storage devices, select the device that contains the VMFS copy.
- Step 8.** Choose to mount the datastore and select one of the following options:
  - a. Mount Options > Assign a New Signature**
  - b. Mount Options > Keep Existing Signature**
- Step 9.** Click **Finish**.

Beginning with vSphere 7.0, you can use clustered virtual disks (VMDKs) on a VMFS 6 datastore to support Windows Server Failover Clustering (WSFC). To enable support for clustered VMDK, you should set Clustered VMDK Support to Yes when creating a VMFS 6 datastore. The datastores must only be used by ESXi 7.0 or later and must be managed by the same vCenter Server 7.0 or later. For a datastore that supports clustered VMDK, you must also enable clustered VMDK. In the vSphere Client, select the VMFS 6 datastore in the inventory

pane, and set Datastore Capabilities > Clustered VMDK to Enable. After enabling this setting, you can place the clustered virtual disks on the datastore. To disable the setting, you need to first power off the virtual machines with clustered virtual disks.

Table 11-3 provides details for other administration operations that you can perform on VMFS datastores.

**Table 11-3** VMFS Datastore Operations

<b>Operation</b>	<b>Notes</b>
<p>Change datastore name</p> <ol style="list-style-type: none"> <li>1. Right-click the datastore and select Rename.</li> <li>2. Provide a name (42-character limit).</li> </ol>	<p>If the host is managed by vCenter Server, you must rename the datastore from vCenter Server, not from the vSphere Host Client.</p> <p>You can successfully rename a datastore that has running virtual machines.</p>
<p>Unmount datastore</p> <ol style="list-style-type: none"> <li>1. Ensure that Storage DRS and Storage I/O Control are not enabled for the datastore.</li> <li>2. Right-click the datastore and select Unmount Datastore.</li> <li>3. Select the hosts.</li> </ol>	<p>When you unmount a datastore, it remains intact but can no longer be seen from the specified hosts.</p> <p>Do not perform other configuration operations on the datastore during the unmount operation.</p> <p>Ensure that the datastore is not used by vSphere HA heartbeating (which could trigger host failure event and restart virtual machines).</p>

---

M o u n t d a t a s t o r e	<ol style="list-style-type: none"> <li>1. Right-click the datastore and select Mount Datastore (if the datastore is inactive) or Mount Datastore on Additional Hosts (if the datastore is active).</li>   <li>2. Select the host(s) and click OK.</li> </ol>	<p>A VMFS datastore that is unmounted from all hosts is marked as inactive. You can mount the unmounted VMFS datastore.</p> <p>If you unmount an NFS or a vVols datastore from all hosts, the datastore disappears from the inventory. To mount the NFS or vVols datastore that has been removed from the inventory, use the New Datastore wizard.</p>
--	--	--

---

R e m o v e d a t a s t o r e	<ol style="list-style-type: none"> <li>1. Move or migrate virtual machines from the datastore. Unmount the datastore.</li>   <li>2. Right-click the datastore and select Delete Datastore.</li>   <li>3. Click OK to confirm the deletion.</li> </ol>	<p>Deleting a datastore permanently destroys it and all its data, including virtual machine files. You are not required to unmount the datastore prior to deletion, but you should.</p>
---	---	---

---

In the vSphere Client, you can use the Datastore Browser to examine and manage the datastore contents. To get started, right-click the datastore in the inventory pane and select Browse Files. In the Datastore browser, select any of the options listed in [Table 11-4](#).

**Table 11-4** Datastore Browser Options

---

<b>Option</b>	<b>Description</b>
Upload Files	Upload a local file to the datastore.
Upload Folder	Upload a local folder to the datastore.
Download	Download a file from the datastore to the local machine.
New Folder	Create a folder on the datastore.
Copy to	Copy selected folders or files to a new location on the datastore or on another datastore.
Move to	Move selected folders or files to a new location on the datastore or on another datastore.
Rename to	Rename selected files.
Delete	Delete selected folders or files.
Inflate	Convert a selected thin virtual disk to thick.

When you use the vSphere Client to perform VMFS datastore operations, vCenter Server uses default storage protection filters. The filters help you avoid data corruption by displaying only the storage devices that are suitable for an operation. In the rare scenario in which you want to turn off the storage filters, you can do so using the following procedure:

- Step 1.** In the vSphere Client, select the vCenter Server instance in the inventory pane and navigate to **Configure > Settings > Advanced Settings > Edit Settings**.
- Step 2.** Specify one of the filter names described in [Table 11-5](#) and set its value to False.

**Table 11-5** Storage Filters

<b>Filter</b>	<b>Description</b>
config.vpxd.fil	Hides storage devices (LUNs) that are used by a VMFS datastore on any host managed by vCenter Server.

ter.vmf  
sFilter

(VMFS  
filter)

---

config. Hides storage devices (LUNs) that are used by an RDM on any  
vpxd.fil host managed by vCenter Server.  
ter.rdm  
Filter

(RDM  
filter)

---

config. Hides storage devices (LUNs) that are ineligible for use as VMFS  
vpxd.fil datastore extents because of incompatibility with the selected  
ter.sam datastore. Hides LUNs that are not exposed to all hosts that share  
eHosts the original datastore. Hides LUNs that use a storage type (such as  
AndTra Fibre Channel, iSCSI, or local) that is different from the original  
nsports datastore.  
Filter

(Same  
Hosts  
and  
Transp  
orts  
filter)

---

config. Automatically rescans and updates VMFS datastores following  
vpxd.fil datastore management operations. If you present a new LUN to a  
ter.host host or a cluster, the hosts automatically perform a rescan,  
Rescan regardless of this setting.  
Filter

(Host  
Rescan  
filter)

---



## Note

You should consult the VMware support team prior to changing device filters.

## Managing Raw Device Mappings (RDMs)

You can use the following procedure to add an RDM to a virtual machine:

- Step 1.** In the vSphere Client, open the settings for a virtual machine.
- Step 2.** Click **Add New Device** and select **RDM Disk**.
- Step 3.** Select a LUN and click **OK**.
- Step 4.** Click the **New Hard Disk** triangle to expand the RDM properties.
- Step 5.** Select a datastore to place the RDM, which can be the same as or different from where the virtual machine configuration file resides.
- Step 6.** Select either **Virtual Compatibility Mode** or **Physical Compatibility Mode**.
- Step 7.** If you selected Virtual Compatibility Mode, select a disk mode: **Dependent**, **Independent–Persistent**, or **Independent–Nonpersistent**.
- Step 8.** Click **OK**.

You can use the following procedure to manage paths for the storage devices used by RDMs:

- Step 1.** In the vSphere Client, right-click the virtual machine in the inventory pane and select **Edit Settings**.
- Step 2.** Select **Virtual Hardware > Hard Disk**.
- Step 3.** Click the device ID that appears next to Physical LUN to open the Edit Multipathing Policies dialog box.
- Step 4.** Use the Edit Multipathing Policies dialog box to enable or disable paths, set multipathing policy, and specify the preferred path.

If the guest OS in your virtual machine is known to have issues using the SCSI INQUIRY data cached by ESXi, you can either modify the virtual machine or

the host to ignore the cached data. To modify the virtual machine, you can edit its VMX file and add the following parameter, where `scsiX:Y` represents the SCSI device:

[Click here to view code image](#)

```
scsiX:Y.ignoreDeviceInquiryCache = "true"
```

To modify the host, you can use the following command, where `deviceID` is the device ID of the SCSI device:

[Click here to view code image](#)

```
esxcli storage core device inquirycache set --device deviceID
--ignore true
```

## Managing NFS Datastores

NFS Version 3 and Version 4.1 are supported by ESXi, which uses a different client for each protocol. When mounting NFS datastores on an ESXi host, the following best practices should be observed:

- On ESXi, the NFS Version 3 and NFS Version 4.1 clients use different locking mechanisms. You cannot use different NFS versions to mount the same data-store on multiple hosts.
- ESXi hosts can make use of both NFS Version 3 and Version 4.1 if the previous rule is observed.
- ESXi hosts cannot automatically upgrade NFS Version 3 to NFS Version 4.1.
- NFS datastores must have folders with identical names mounted on all ESXi hosts, or functions such as vMotion may not work.
- If an NFS device does not support internationalization, you should use ASCII characters only.

How you configure an NFS storage device to use with VMware varies by vendor, so you should always refer to the vendor documentation for specifics.

The following is the procedure to configure an NFS server (but refer to vendor documentation for specifics on how to carry out this procedure):

**Step 1.** Use the VMware Hardware Compatibility List to ensure that the NFS

server is compatible. Pay attention to the ESXi version, the NFS server version, and the server firmware version.

**Step 2.** Configure the NFS volume and export it (by adding it to `/etc/exports`) using the following details:

- NFS Version 3 or Version NFS 4.1 (only one protocol per share)
- NFS over TCP

**Step 3.** For NFS Version 3 or non-Kerberos NFS Version 4.1, ensure that each host has root access to the volume. The typical method for this is to use the **`no_root_squash`** option.

**Step 4.** If you are using Kerberos, ensure that the NFS exports provide full access to the Kerberos user. In addition, if you are going to use Kerberos with NFS Version 4.1, you need to enable either AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 on the NFS storage device.

To prepare an ESXi host to use NFS, you must configure a VMkernel virtual adapter to carry NFS storage traffic. If you are using Kerberos and NFS Version 4.1, you should take the following additional steps:

**Step 1.** Ensure that the DNS settings on the ESXi hosts are pointing to the DNS server that is used for DNS records for Kerberos Key Distribution Center (KDC). This will most likely be the Active Directory server if that is being used for name resolution.

**Step 2.** Configure NTP because Kerberos is sensitive to time drift.

**Step 3.** Configure Active Directory for Kerberos.

To create (mount) an NFS datastore in vSphere, you need the IP address or DNS name of the NFS server as well as the path to the share (folder name). When using Kerberos, you need to configure the ESXi hosts for Kerberos authentication prior to creating the NFS datastore.

## Note

Multiple IP addresses or DNS names can be used with NFS Version 4.1 multipathing.

You can use the following procedure to create an NFS datastore:



- Step 1.** In the vSphere Client, right-click a data center, cluster, or ESXi host object in the inventory pane and select **Storage > New Datastore**.
- Step 2.** Select **NFS** as the new datastore type.
- Step 3.** Select the correct NFS version (Version 3 or Version 4.1). Be sure to use the same version on all ESXi hosts that are going to mount this datastore.
- Step 4.** Define the datastore name (with a maximum of 42 characters).
- Step 5.** Provide the appropriate path for the folder to mount, which should start with a forward slash (/).
- Step 6.** Set **Server** to the appropriate IPv4 address, IPv6 address, or server name.
- Step 7.** Optionally, select the **Mount NFS Read Only** checkbox. (This can only be set when mounting an NFS device. To change it later, you must unmount and remount the datastore from the hosts.)
- Step 8.** If using Kerberos, select **Kerberos** and define the Kerberos model as one of the following:
  - a. Use Kerberos for Authentication Only (krb5):** This method supports identity verification only.
  - b. Use Kerberos for Authentication and Data Integrity (krb5i):** This method supports identity verification and also ensures that data packets have not been modified or tampered with.
- Step 9.** If you selected a cluster or a data center object in step 1, then select the ESXi hosts to mount this datastore.
- Step 10.** Verify the configuration and click **Finish**.

To rename or unmount an NFS datastore, you can use the same procedure as described for VMFS datastores in [Table 11-3](#). To remove an NFS datastore from the vSphere inventory, you should unmount it from every host.

## Storage DRS and SIOC

This section provides details on configuring and managing Storage DRS and Storage I/O Control (SIOC).

## Configuring and Managing Storage DRS

To create a datastore cluster using the vSphere Client, you can right-click on a data center in the inventory pane, select New Datastore Cluster, and complete the wizard. You can use the following procedure to enable Storage DRS (SDRS) in a data-store cluster:

**Step 1.** In the vSphere Client, select the datastore cluster in the inventory pane and navigate to **Configure > Services > Storage DRS**.

**Step 2.** Click **Edit**.

**Step 3.** Select **Turn ON vSphere DRS** and click **OK**.

You can use similar steps to set the SDRS Automation Mode to No Automation, Partially Automated, or Fully Automated. You can set Space Utilization I/O (SDRS Thresholds) Latency. You can select or deselect Enable I/O Metric for SDRS Recommendations. You can configure the advanced options, which are Space Utilization Difference, I/O Load Balancing Invocation Interval, and I/O Imbalance Threshold.

You can add datastores to a datastore cluster by using drag and drop in the vSphere Client. Each datastore can only be attached to hosts with ESXi 5.0 or later. The datastores must not be associated with multiple data centers.

If you want to perform a maintenance activity on an SDRS cluster member datastore or its underlying storage devices, you can place it in Maintenance Mode. (Standalone datastores can be placed in Maintenance Mode.) SDRS has recommendations for migrating the impacted virtual machine files, including virtual disk files. You can let SDRS automatically apply the recommendations, or you can manually make recommendations. To place a datastore in Maintenance Mode using the vSphere Client, right-click the datastore in the inventory pane, select Enter SDRS Maintenance Mode, and optionally apply any recommendations.

The Faults tab displays a list of the disks that cannot be migrated and the reasons.

If SDRS affinity or anti-affinity rules prevent a datastore from entering Maintenance Mode, you can select an option to ignore the rules. To do so, edit

the settings of the datastore cluster by selecting SDRS Automation > Advanced Options and setting IgnoreAffinityRulesForMaintenance to 1.

When reviewing each SDRS recommendation on the Storage SDRS tab in the vSphere Client, you can examine the information described in [Table 11-6](#) and use it when deciding which recommendations to apply.

**Table 11-6** SDRS Recommendations

<b>Recommendations</b>	<b>Details</b>
Priority	Priority level (1–5) of the recommendation. (This is hidden by default.)
Recommendation	Recommended action.
Reason	Why the action is needed.
Space Utilization % Before (source) and (destination)	Percentage of space used on the source and destination datastores before migration.
Space Utilization % After (source) and (destination)	Percentage of space used on the source and destination datastores after migration.
I/O Latency Before (source)	Value of I/O latency on the source datastore before migration.
I/O Latency Before (destination)	Value of I/O latency on the destination datastore before migration.

You can use the following procedure to override the SDRS datastore cluster automation level per virtual machine:

**Step 1.** In the vSphere Client, right-click a datastore cluster in the inventory pane and select **Edit Settings**.

**Step 2.** Select **Virtual Machine Settings**.

**Step 3.** Select one of the following automation levels:

- **Default (Manual)**
- **Fully Automated**
- **Disabled**

**Step 4.** Optionally select or deselect the **Keep VMDKs Together** option.

**Step 5.** Click **OK**.

You can use the following procedure to create an inter-VM anti-affinity rule (that is, a rule specifying that two or more virtual machines are placed on separate datastores):

**Step 1.** In the vSphere Client, right-click a datastore cluster in the inventory pane and select **Edit Settings**.

**Step 2.** Select **Rules > Add**.

**Step 3.** Provide a name and set Type to **VM Anti-affinity**.

**Step 4.** Click **Add**.

**Step 5.** Click **Select Virtual Machine**.

**Step 6.** Select at least two virtual machines and click **OK**.

**Step 7.** Click **OK** to save the rule.

To create an intra-VM anti-affinity rule (that is, a rule which says that virtual disks for a specific virtual machine are placed on separate datastores), you use a similar procedure but set Type to VMDK-Affinity and select the appropriate virtual machine and virtual disks.

## **Configuring and Managing SIOC**

*Storage I/O Control (SIOC)* allows you to prioritize storage access during periods of contention, ensuring that the more critical virtual machines obtain more I/O than less critical VMs. Once SIOC has been enabled on a datastore, ESXi hosts monitor the storage device latency. If the latency exceeds a predetermined threshold, the datastore is determined to be under contention, and the virtual machines that reside on that datastore are assigned I/O resources based on their individual share values. You can enable SIOC as follows:

**Step 1.** In the vSphere Client, select a datastore in the Storage inventory view and select **Configuration > Properties**.

**Step 2.** Click the **Enabled** checkbox under Storage I/O Control and click **Close**.

### **Note**

SIOC is enabled automatically on Storage DRS–enabled datastore clusters.

In addition to share values, which are similar to shares defined for CPU and memory, storage I/O limits can be defined on individual virtual machines to limit the number of I/O operations per second (IOPS). By default, just as with CPU and memory resources, there are no limits set for virtual machines. In a virtual machine with more than one virtual disk, limits must be set on all of the virtual disks for that VM. If you do not set a limit on all the virtual disks, the limit won't be enforced. To view the shares and limits assigned to virtual machines, you can use the vSphere Client. To select a datastore, select the Virtual Machines tab and examine the associated virtual machines. The details for each virtual machine include its respective shares, the IOPS limit, and the percentage of shares for that datastore.

## Setting SIOC Shares and Limits

As with CPU and memory shares, SIOC shares establish a relative priority in the event of contention. In the event of storage contention, virtual machines with more shares will observe more disk I/O than will a virtual machine with fewer shares. The following procedure outlines how you configure SIOC shares and limits for virtual machines:



- Step 1.** In the vSphere Client, right-click a virtual machine in the inventory pane and select **Edit Settings**.
- Step 2.** Expand one of the hard disks (for example, **Hard disk 1**).
- Step 3.** From the Shares drop-down menu, select **High**, **Normal**, **Low**, or **Custom** to define the share value.
- Step 4.** Set the Limit–IOPS drop-down to **Low** (500), **Normal** (1000), **High** (2000), or **Custom** (and enter a custom value for the IOPS limit).
- Step 5.** Click **OK** to save your changes.

## Monitoring SIOC Shares

To view the impact of shares on individual datastores, in the vSphere Client,



select a datastore in the inventory pane, select the Performance tab, and select View > Performance. Here, you can observe the following data:

- Average latency and aggregated IOPS
- Host latency
- Host queue depth
- Host read/write IOPS
- Virtual machine disk read/write latency
- Virtual machine disk read/write IOPS

## SIOC Threshold

The default threshold for SIOC to begin prioritizing I/O based on shares is 30 ms and typically does not need to be modified. However, you can modify this threshold if you need to. Be aware that SIOC will not function properly unless all the data-stores that share drive spindles have the same threshold defined. If you set the value too low, shares will enforce priority of resources sooner but could decrease aggregated throughput, and if you set it too high, the result might be higher aggregated throughput but less prioritization of disk I/O.

The following procedure allows you to modify the threshold:

- Step 1.** In the vSphere Client Storage Inventory view, select a datastore and select the **Configuration** tab.
- Step 2.** Select **Properties** and under Storage I/O Control, select **Enabled** if it is not already.
- Step 3.** Click **Advanced** to modify the threshold for contention; this value must be between 10 ms and 100 ms.
- Step 4.** Click **OK** and then click **Close**.

The procedure to reset the threshold to the default is similar:

- Step 1.** In the vSphere Client Storage Inventory view, select a datastore and select the **Configuration** tab.
- Step 2.** Select **Properties** and under Storage I/O Control, select **Advanced**.
- Step 3.** Click **Reset**.

**Step 4.** Click **OK** and then click **Close**.

## NVMe and PMem

This section provides details on configuring and managing Non-Volatile Memory Express (NVMe).

### Managing VMware NVMe

As described in [Chapter 2, \*Non-Volatile Memory Express \(NVMe\)\*](#) devices are a high-performance alternative to SCSI storage. There are three mechanisms for NVMe:

- **NVMe over PCIe:** NVMe over PCIe is for local storage, and NVMe over fabrics (NVMe-oF) is for connected storage.
- **NVMe over Remote Direct Memory Access (RDMA):** NVMe over RDMA is shared NVMe-oF storage using RDMA over Converged Ethernet (RoCE) Version 2 transport.
- **NVMe over Fibre Channel (FC-NVMe):** FC-NVMe is shared NVMe-oF storage using Fibre Channel transport.

[Chapter 2](#) describes the requirements for each of these mechanisms.

After you install the hardware for NVMe over PCIe, ESXi detects it as a storage adapter that uses PCIe. You can use the vSphere Client to view the storage adapter and storage device details. No other configuration is needed.

### NVMe over Fabrics Shared Storage

When using NVMe storage for shared storage, you must not mix transport types on the same namespace. You should ensure that the active paths are presented because they cannot be registered until the path has been discovered.

[Table 11-7](#) identifies additional information about NVMe over Fabric versus SCSI over Fabric storage.

**Table 11-7** SCSI over Fabric and NVMe over Fabric Comparison

---

<b>Shared Storage Capability</b>	<b>SCSI over Fabric</b>	<b>NVMe over Fabric</b>
----------------------------------	-------------------------	-------------------------

---

RDM	Supported	Not supported
Coredump	Supported	Not supported
SCSI-2 reservations	Supported	Not supported
Shared VMDK	Supported	Not supported
vVols	Supported	Not supported
Hardware acceleration with VAAI plug-ins	Supported	Not supported
Default MPP	NMP	HPP (NVMe-oF targets cannot be claimed by NMP.)
Limits	LUNs=1024, paths=4096	Namespaces=32, paths=128 (maximum 4 paths per namespace in a host)

To use FC-NVMe, you must add an appropriate supported adapter and use the following procedure to add the controller to the host:

- Step 1.** In the vSphere Client, select the host in the inventory pane and navigate to **Configure > Storage > Storage Adapters**.
- Step 2.** Click **Controllers > Add Controller**.
- Step 3.** Select one of the following options:
  - a. Automatically Discover Controllers:** Click **Discover Controllers** and select a controller.
  - b. Enter Controller Details Manually:** Provide the subsystem NQN, the worldwide node name, and the worldwide port name. Optionally, provide an admin queue size and keepalive timeout.

## Configuring ESXi to Support RDMA

You can configure an ESXi 7.0 host to access shared NVMe devices using RDMA over Converged Ethernet (RoCE) Version 2. The host must have a network adapter that supports RoCE Version 2, and you must configure a software NVMe over RDMA adapter.

For hosts with a NIC that supports RoCE Version 2, the vSphere Client shows both the network adapter component and the RDMA component. You can select

the host in the inventory pane and navigate to **Configure > Networking > RDMA Adapters**. Here you can see the unique names assigned to the RDMA devices, such as `vmrdma0`. For each device, you can see its paired uplink (that is, its integrated NIC, such as `vmnic9`). To complete the host configuration, you can use the following procedure:

- Step 1.** Create a new VMkernel virtual network adapter on a vSphere standard or distributed switch and configure its uplink to use the RDMA paired uplink (for example, `vmnic9`).
- Step 2.** Select the host and navigate to **Configure > Networking > RDMA Adapters**.
- Step 3.** Select the appropriate RDMA device (for example, `vmrdma0`) and select **VMkernel Adapters Bindings** in the details pane.
- Step 4.** Verify that the new VMkernel adapter (for example, `vmk2`) appears.
- Step 5.** Select the host and navigate to **Configure > Storage > Storage Adapters**.
- Step 6.** Click the **Add Software Adapter** button.
- Step 7.** Select **Add Software NVMe over RDMA Adapter**.
- Step 8.** Select the appropriate RDMA adapter (for example, `vmrdma0`) and click **OK**.
- Step 9.** In the list of storage adapters, identify the new adapter in the category **VMware NVME over RDMA Storage Adapter** and make note of its assigned device number (for example, `vmhba71`).
- Step 10.** To identify the available storage devices, select the storage adapter (for example, `vmhba71`) and select **Devices** in the details pane. You can use these devices to create VMFS datastores.

## Configuring HPP

As described in [Chapter 2, \*High-Performance Plug-in \(HPP\)\*](#) is the default plug-in that claims NVMe-oF targets. NVMe over PCIe targets default to the VMware Native Multipathing Plug-in (NMP). You can use the **esxcli storage core claimrule add** command to change the claiming plug-in in your environment. For example, to set a local device to be claimed by HPP, use the **--pci-vendor-id** parameter and set the **--plugin** parameter to **HPP**. To change the

claim rule based on an NVMe controller model, use the **--nvme-controller-model** parameter.

To assign a specific HPP Path Selection Scheme (PSS) to a specific device, you can use the **esxcli storage hpp device set** command with **-pss** parameter to specify the scheme and the **--device** parameter to specify the device. The available HPP PSS options are explained in [Table 2-6](#) in [Chapter 2](#). To create a claim rule that assigns the HPP PSS by vendor and model, you can use **esxcli storage core claimrule add** with the **-V** (vendor), **-M** (model), **-P** (plug-in), and **--config-string** parameters. In the value for **--config-string**, specify the PSS name and other settings, such as “**pss=LB-Latency,latency-eval-time=40000.**”

### Note

Enabling HPP on PXE-booted ESXi hosts is not supported.

After using these commands, you should reboot the hosts to apply the changes.

## Managing PMem

**PMem devices** are non-volatile dual in-line memory modules (NVDIMMs) on the ESXi host that reside in normal memory slots. They are non-volatile and combine the performance of volatile memory with the persistence of storage. PMem devices are supported on ESXi 6.7 and later.

ESXi hosts detect local PMem devices and expose the devices as host-local PMem datastores to virtual machines. Virtual machines can directly access and utilize them as either memory (virtual NVDIMM) or storage (PMem hard disks). An ESXi host can have only one PMem datastore, but it can be made up of multiple PMem modules.

## Virtual PMem (vPMem)

In vPMem mode, a virtual machine can directly access PMem resources and use the resources as regular memory. The virtual machine uses NVDIMMs that represent physical PMem regions. Each virtual machine can have up to 64 virtual NVDIMM devices, and each NVDIMM device is stored in the host-local PMem datastore. Virtual machines must be at hardware Version 14, and the guest OS must be PMem aware.

## Virtual PMem Disks (vPMemDisk)

In **vPMemDisk** mode, a virtual machine cannot directly access the PMem resources. You must add a virtual PMem disk to the virtual machine. A virtual PMem disk is a regular virtual disk that is assigned a PMem storage policy, forcing it to be placed on a host-local PMem datastore. This mode has no virtual machine hardware or operating system requirements.

## PMem Datastore Structure

The following are components of the PMem structure on an ESXi host:

- **Modules:** These are the physical NVDIMMs that reside on the motherboard.
- **Interleave sets:** These are logical groupings of one or more modules. ESXi hosts read from an interleave set in turns, so if there are two modules on an ESXi host, they will be read in parallel. You can identify the way the NVDIMMs are grouped into interleave sets via the vSphere Client.
- **Namespaces:** PMem datastores are built on top of namespaces, which are regions of contiguously addressed memory ranges.

To view information about the PMem modules, interleave sets, and namespaces, you can follow this procedure:

**Step 1.** In the vSphere Host Client, select **Storage** from the inventory pane.

**Step 2.** Click on the **Persistent Memory** tab.

**Step 3.** Click **Modules** to see the NVDIMMs that contribute to the PMem datastore.

**Step 4.** Click **Namespaces** to see namespace information.

**Step 5.** Click **Interleave Sets** to see how the modules are grouped into interleave sets.

To delete namespaces that were created by an operating system that was previously installed on the host machine, you can use this procedure to navigate to Namespaces, select the namespace, and click Delete. This frees up the PMem space, but you must reboot the host to access it.

# Multipathing, Storage Policies, and vVols

This section provides information on managing storage multipathing, storage policies, and Virtual Volumes (vVOLS) in vSphere 7.0.

## Managing Multipathing

As explained in [Chapter 2](#), ESXi uses the Pluggable Storage Architecture (PSA), which allows plug-ins to claim storage devices. The plug-ins include the Native Multipathing Plug-in (NMP), the High-Performance Plug-in (HPP), and third-party multipathing modules (MPPs).

## esxcli Commands for Multipathing

You can use **esxcli** commands to manage the PSA plug-ins. For example, you can use the following command from an ESXi shell to view the multipathing modules (plug-ins):

[Click here to view code image](#)

```
esxcli storage core plugin list -plugin-class=MP
```

You can use the following command to list all devices controlled by the NMP module. For each device, you will find details, such as assigned storage array type (SATP) and the path selection policy (PSP):

```
esxcli storage nmp device list
```

To see details for a specific device, you can provide the **--device** option with the previous command. For example, if you have a device that is identified by **mpx.vmbha0:C0:T0:L0**, you can use the following command to retrieve details for just that device:

[Click here to view code image](#)

```
esxcli storage nmp device list --device=mpx.vmbha0:C0:T0:L0
```

[Table 11-8](#) provides information on some other **esxcli** commands that you can use with NMP.

**Table 11-8** ESXLI Commands for NMP

Command	Description
---------	-------------

<b>esxcli storage nmp satp list</b>	Provides information for each available SATP, including the default PSP
<b>esxcli storage nmp psp list</b>	Provides a description for each available PSP
<b>esxcli storage nmp satp set --default- psp=policy -- satp=satpname</b>	Changes the default PSP policy for an SATP named <i>satpname</i> , where the policy is VMW_PSP_MRU, VMW_PSP_FIXED, or VMW_PSP_RR, as explained in <a href="#">Table 2-11</a> in <a href="#">Chapter 2</a>

### Note

In many cases, the storage system provides ESXi with the storage device names and identifiers, which are unique and based on storage standards. Each identifier uses a *naa.xxx*, *eui.xxx*, or *t10.xxx* format. Otherwise, the host generates an identifier in the form *mpx.path*, where *path* is the first path to the device, such as *mpx.vmhba1:C0:T1:L3*.

[Table 11-9](#) provides information on some **esxcli** commands that you can use with HPP.

**Table 11-9** ESXCLI Commands for HPP

Command	Description	Options
<b>esxcli storage hpp path list</b>	Lists which paths are claimed by HPP	<b>-d   --device</b>  <b>-p   --path=&lt;path&gt;</b>
<b>esxcli storage hpp device list</b>	Lists devices controlled by HPP	<b>-d   --device=&lt;device&gt;</b>
<b>esxcli storage hpp device set</b>	Configures HPP settings	<b>-B   --bytes=</b> <b>&lt;max_bytes_on_path&gt;</b>  <b>-g   --cfgfile</b>



**-d | --device=<device>**

**-I | --iops=  
<max\_iops\_on\_path>**

**-T | --latency-eval-time=  
<interval\_in\_ms>**

**-M | --mark-device-ssd=  
<value>**

**-p | --path=<path>**

**-S | --sampling-ios-per-  
path=<value>**

**-P | --pss=<FIXED|LB-  
Bytes|LB-IOPs|LB-  
Latency|LB-RR>**

---

<b>esxcli storage hpp device usermarkedssd list</b>	Lists devices that someone marked as SSD
---	--

---

**-d | --device=<device>**

## Managing Paths with the vSphere Client

This section provides information on using the vSphere Client to manipulate the path selection policy and available paths for NMP. For information on selecting PSS for HPP, see the “Configuring NVMe and HPP” section, earlier in this chapter.

A path to a storage device is represented as the storage adapter, storage channel number, target number, and LUN number (the LUN position within the target) set that is used to connect to the device. For example, vmhba1:C0:T1:L3 indicates that the path uses storage adapter vmhba1, channel 0, target 1, and LUN 3. To view the storage paths for a specific device, you can use the

following procedure:

- Step 1.** In the vSphere Client, select a host in the inventory pane and navigate to **Configure > Storage > Storage Devices**.
- Step 2.** Select the storage device.
- Step 3.** Click the **Properties** tab and review the details. For NMP devices, the details include the assigned SATP and PSP.
- Step 4.** Click **Paths** and review the available paths to the device. The status for each path can be Active (I/O), Standby, Disabled, or Dead. For devices using the Fixed path policy, an asterisk (\*) represents the preferred path.

To disable a path to a storage device, you can follow this procedure, select a path, and choose Disable.

In the vSphere Client, you can select a VMFS datastore and navigate to Configure > Connectivity and Multipathing to review information on the paths to the storage devices backing the datastore.

To change the PSP that is assigned to a storage device, you can navigate to the device's Properties page (see the previous procedure) and click Edit Multipathing. On the multipathing page, you can choose a policy, such as VMW\_PSP\_FIXED, VMW\_PSP\_RR, or VMW\_PSP\_MRU, as described in [Table 2-11](#) in [Chapter 2](#).

## Managing Claim Rules

ESXi uses claim rules to determine which multipathing module owns the paths to a specific storage device and the type of multipathing support that is applied. Core claim rules determine which multipathing module (NMP, HPP, or a third-party module) claims a device. For NMP, SATP claim rules determine which SATP submodule claims the device. [Table 11-10](#) describes a few commands for claim rules.



**Table 11-10** Sample Claim Rules Commands

Command	Description
<code>esxcli storage</code>	Lists the claim rules on the host.

**core  
claimrule list  
--claimrule-  
class=MP**

---

**esxcli storage  
core  
claimrule add** Defines a new claim rule. The rule may contain multiple options, such as plug-in (-P), model (-M), and vendor (-V). The value for the plug-in can be NMP, HPP, MASK\_PATH, or a third-party plug-in name.

---

**esxcli storage  
core  
claimrule  
load** Loads new claim rules into the system.

---

**esxcli storage  
core  
claimrule run** Applies the loaded claim rules.

---

## Managing Storage Policies

In vSphere, Storage Policy Based Management (SPBM) can be used to align storage with the application demands of your virtual machines. With SPBM, you can assign a storage policy to a virtual machine to control the type of storage that can be used by the virtual machine and how the virtual machine is placed on the storage. You can apply a storage policy as you create, clone, or migrate a virtual machine.

Prior to creating virtual machine storage policies, you must populate the VM Storage Policy interface with information about storage entities and data services in your storage environment. When available, you can use a vSphere APIs for Storage Awareness (VASA) provider to provide the information. Alternatively, you can use datastore tags.

To use storage policies, multiple steps are required, and the particular steps depend on the type of storage or services you need. This section describes the major steps.



In many cases, VM storage policies populate automatically. However, you can manually assign tags to datastores as follows:

**Step 1.** Create a category for the storage tags:

- a. In the vSphere Client, select **Home > Tags & Custom Attributes**.
- b. Click **Tags > Categories**.
- c. Click **Add Category**.
- d. Define the following:
  - Category Name
  - Description
  - Tags per Object
  - Associable Object Types.
- e. Click **OK**.

**Step 2.** Create a storage tag:

- a. Click **Tags** on the Tags tab.
- b. Click **Add Tag**.
- c. Define the tag properties:
  - Name
  - Description
  - Category
- d. Click **OK**.

**Step 3.** Apply the tag to the datastore:

- a. In the storage inventory view, right-click the datastore and select **Tags & Custom Attributes > Assign Tag**.
- b. Select a tag from the list and click **Assign**.

To create a VM storage policy, follow this procedure:

**Step 1.** For host-based services:

- a. In the vSphere Client, select **Home > Policies and Profiles**.
- b. Click **VM Storage Policies**.
- c. Select **Create VM Storage Policy**.

- d. Provide the following information:
  - vCenter Server
  - Name
  - Description
- e. On the Policy Structure page:
  - Select the tab for the data service category.
  - Define custom rules.
- f. Review the datastores that match the policy.
- g. Review the settings and click **Finish**.

**Step 2.** For vVols:

- a. In the vSphere Client, select **Home > Policies and Profiles**.
- b. Click **VM Storage Policies**.
- c. Select **Create VM Storage Policy**.
- d. Provide the following information:
  - vCenter Server
  - Name
  - Description
- e. On the Policy Structure page, enable rules.
- f. On the Virtual Volumes Rules page, set storage placement rules:
  - Select **Placement > Add Rule**.
  - Use the **Add Rule** drop-down to define the available capacity.
  - Click **Tags** to define tag-based rules
- g. Optionally, set rules for datastore-specific services.
- h. On the Storage Compatibility page, review the datastores matching the policy.
- i. On the Review and Finish page, verify the settings and click **Finish**.

**Step 3.** For tag-based placement:

- a. In the vSphere Client, select **Home > Policies and Profiles**.
- b. Click **VM Storage Policies**.
- c. Select **Create VM Storage Policy**.
- d. Provide the following information:
  - vCenter Server
  - Name
  - Description
- e. On the Policy Structure page, select **Add Tag Rule** and define the tag category, usage option, and tags. Repeat as needed.
- f. Review the datastores that match the policy.
- g. Verify the policy settings on the Review and Finish page and click **Finish**.

## **VASA: Registering a Storage Provider**

To collect storage entities and data services information from a VASA storage provider, you can use the following procedure:

- Step 1.** In the vSphere Client, select **vCenter** in the inventory pane and navigate to **Configure > Storage Providers**.
- Step 2.** Click the **Add** icon.
- Step 3.** Provide the connection information for the provider, including the name, URL, and credentials.
- Step 4.** Select one of the following security methods:
  - a. Select the **Use Storage Provider Certificate** option and define the location of the certificate.
  - b. Review and accept the displayed certificate thumbprint that is displayed.
- Step 5.** Storage provider adds the vCenter certificate to the truststore when the vCenter server initially connects to the storage device.
- Step 6.** Click **OK**.

## VASA: Managing Storage Providers

To perform management operations involving a storage provider, such as rescanning, you can use the following procedure:

- Step 1.** In the vSphere Client, select **vCenter** in the inventory pane and navigate to **Configure > Storage Providers**.
- Step 2.** Select a storage provider and choose one of the following options:
  - **Synchronize Storage Providers:** Synchronizes vCenter Server with information for all storage providers.
  - **Rescan:** Synchronizes vCenter Server with information from a specific storage provider.
  - **Remove:** Unregisters a specific storage provider, which is useful whenever upgrading a storage provider to a later VASA version requires you to unregister and reregister.
  - **Refresh Certificate:** Refreshes a certificate before it retires.

## Applying Storage Policies to VMs

When performing a provisioning, cloning, or migration operation for a virtual machine, you can assign a storage policy. For example, when using the New Virtual Machine wizard, on the Select Storage page, you can select a policy in the VM Storage Policy drop-down menu to assign a storage policy to the entire virtual machine. After selecting the policy, you should choose a datastore from the list of compatible datastores. If you use the replication service provided with vVols, you should either specify a preconfigured replication group or choose to have vVols create an automatic replication group.

Optionally, you can set different storage policies for each virtual disk. For example, on the Customize Hardware page, you can select New Hard Disk and set a VM storage policy for that disk.

To change the storage policy assigned to a virtual machine, you can use the following procedure:

- Step 1.** In the vSphere Client, navigate to **Menu > Policies and Profiles > VM Storage Policies**.
- Step 2.** Select a storage policy and click **VM Compliance**.

- Step 3.** Select a virtual machine that is currently assigned the selected policy.
- Step 4.** Click **Configure > Policies**.
- Step 5.** Click **Edit VM Storage Policies**.
- Step 6.** Assign the appropriate policy to the virtual machine or assign separate policies to different virtual disks.
- Step 7.** If you use the replication service provided with vVols, configure the replication group.
- Step 8.** Click **OK**.

## Configuring and Managing vVols

To use vVols, you must ensure that your storage and vSphere environment are properly configured.

To work with vVols, the storage system must support vVols and must integrate with vSphere using VASA. You should consider the following guidelines:

- The storage system must support thin provisioning and snapshotting.
- You need to deploy the VASA storage provider.
- You need to configure the following components on the storage side:
  - Protocol endpoints
  - Storage containers
  - Storage profiles
  - Replication configurations if you plan to use vVols with replication
- You need to follow appropriate setup guidelines for the type of storage you use (Fibre Channel, FCoE, iSCSI, or NFS). If necessary, you should install and configure storage adapters on your ESXi hosts
- You need to use NTP to ensure time synchronization among the storage system components and vSphere.

To configure vVols, you can use the following procedure:

- Step 1.** Register the storage providers for vVols, as described previously.
- Step 2.** Create a vVols datastore by using the following steps:



- a. In the vSphere Client, right-click a host, cluster, or data center in the inventory pane and select **Storage > New Datastore**.
- b. Select **vVol** as the datastore type
- c. Select the hosts that will access the datastore.
- d. Click **Finish**.

**Step 3.** Navigate to **Storage > Protocol Endpoints** to examine and manage the protocol endpoints. Optionally, you can take the following steps:

- a. Use the **Properties** tab to modify the multipathing policy.
- b. Use the **Paths** tab to change the path selection policy, disable paths, and enable paths.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15, “Final Preparation,”](#) and the exam simulation questions on the companion website.

## Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 11-12](#) lists these key topics and the page number on which each is found.



**Table 11-12** Key Topics for [Chapter 11](#)

Key Topic Element	Description	Page Number
List	vSAN characteristics	<a href="#">414</a>
List	Shutting down a vSAN cluster	<a href="#">421</a>
List	Creating a stretched vSAN cluster	<a href="#">427</a>
List	Increasing the size of a datastore	<a href="#">439</a>
List	Creating an NFS datastore	<a href="#">446</a>

List	Setting SIOC shares and limits	450
Table 11-10	Sample claim rules commands	458
List	Implementing storage policies	459

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

- virtual PMem disk (vPMemDisk)
- PMem device
- vSAN fault domain
- vSAN file service
- Storage I/O Control (SIOC)
- Non-Volatile Memory Express (NVMe) device
- High-Performance Plug-in (HPP)

## Review Questions

1. You are implementing encryption for a vSAN cluster in vSphere 7.0. Which of the following options is a requirement?
  - a. Deploy KMIP 1.0.
  - b. Deploy the KMS as a virtual machine in the vSAN datastore.
  - c. Ensure that the KMS is in the vSphere compatibility matrices.
  - d. Ensure that the witness host participates in encryption.
2. You want to save space in your vSAN cluster by removing redundant data blocks. Which of the following steps should you take?
  - a. Enable Compression.

- b.** Enable Deduplication.
  - c.** Enable Deduplication and Compression.
  - d.** Enable Allow Reduced Redundancy.
- 3.** In your vSphere 7.0 environment, you are using the Datastore Browser to perform administrative tasks. Which of the following options is not available in the Datastore Browser?
  - a.** Upload files
  - b.** Download
  - c.** Mount
  - d.** Inflate
- 4.** For your vSphere 7.0 environment, you are comparing NVMe-oF with SCSI over Fibre Channel. Which one of the following statements is true?
  - a.** Virtual volumes are supported with NVMe-oF.
  - b.** SCSI-2 reservations are supported with NVMe-oF.
  - c.** RDMS are supported with NVMe-oF.
  - d.** HPP is supported with NVMe-oF.
- 5.** You are using the vSphere Client to manage the storage providers. Which one of the following is not an option?
  - a.** Replace
  - b.** Synchronize Storage Providers
  - c.** Rescan
  - d.** Refresh Certificate

## Chapter 12

# Managing vSphere Security

This chapter covers the following topics:

- [Configuring and Managing Authentication and Authorization](#)
- [Configuring and Managing vSphere Certificates](#)
- [General ESXi Security Recommendations](#)
- [Configuring and Managing ESXi Security](#)
- [Additional Security Management](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.10, 1.11, 4.1, 4.1.1, 4.1.2, 4.3, 4.3.1, 4.3.2, 4.3.3, 4.10, 4.11, 4.11.1, 4.13, 7.7, and 7.x.

This chapter covers the procedures for managing security in a vSphere environment.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 12-1](#) 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 and Review Questions.](#)”

**Table 12-1** “Do I Know This Already?” Foundation Topics Section-to-Question Mapping

---

Foundations Topics Section	Questions Covered in This Section
----------------------------	-----------------------------------

---

Configuring and Managing Authentication and Authorization	1, 2
Configuring and Managing vSphere Certificates	3, 4
General ESXi Security Recommendations	5, 6
Configuring and Managing ESXi Security	7, 8
Additional Security Management	9, 10

1. You are responsible for multiple vSphere environments. What must you do to enable the use of Enhanced Linked Mode in vSphere 7.0?
  - a. Associate two vCenter Servers with the same external PSC.
  - b. Map the external PSC of one vCenter Server to the embedded PSC of another vCenter Server.
  - c. Configure vCenter Server HA.
  - d. Connect two vCenter Servers to the same SSO domain.
2. You are configuring permissions in a vSphere environment. When editing an existing permission, which of the following properties can you change?
  - a. Role
  - b. Privilege
  - c. User
  - d. Object
3. You are managing certificates in your vSphere environment. By default, what types of certificates are in VECS? (Choose two.)
  - a. ESXi certificates
  - b. Machine SSL certificates
  - c. Trusted root certificates
  - d. vCenter Server Certificate
4. You are responsible for performing certificate management for your ESXi hosts. Which of the following privileges do you need?
  - a. Certificates.Manage Certificates

- b. Host.Manage Certificates
  - c. Manage.Certificates
  - d. Certificates.Manage.Host
5. You are enabling direct ESXi access using local accounts. To change the password requirements, such as minimum length, which of the following steps should you take?
- a. Select Single Sign On > Configuration.
  - b. Configure Lockdown Mode.
  - c. Use the Set-PasswordControl cmdlet.
  - d. Configure Security.PasswordQualityControl.
6. You want to enable passthrough for a network device on your ESXi host. You see that an orange icon is associated with a device. Which of the following actions should you take?
- a. Reboot the host.
  - b. Ignore the icon, select the device, and click OK.
  - c. Navigate to Configure > Services and restart a specific service.
  - d. Give up. The device is not compatible with passthrough.
7. You want to configure your ESXi host's acceptance level such that you cannot install VIBs signed at or below the PartnerSupported level but you can install VIBs signed at higher levels. Which option should you choose?
- a. VMwareCertified
  - b. VMwareAccepted
  - c. PartnerSupported
  - d. CommunitySupported
8. You want to enable UEFI Secure Boot. To determine whether your ESXi host supports Secure Boot, which of the following steps should you take?
- a. Use the following command:  
`/usr/lib/vmware/secureboot/bin/secureBoot.py -c`

- b. Check for compliance by using a host profile.
  - c. Check for compliance by using Life Cycle Manager.
  - d. Use the Security Profile section in the vSphere Client.
9. You need to use encryption in your vSphere environment. Which of the following should you use to configure a trust relationship between a KMS and vCenter?
- a. In the vSphere Appliance Management Interface, choose Configuration > Key Management Servers.
  - b. In the vSphere Client, select the vCenter Server and choose Configuration > Key Management Servers.
  - c. In the vSphere Appliance Management Interface, choose Configuration > Encryption.
  - d. In the vSphere Client, select the vCenter Server and choose Configuration > Encryption.
10. You want to configure vSphere Trust Authority. Which of the following is a necessary step?
- a. Create the trusted key provider on the trusted cluster.
  - b. Import the trusted key provider to the trusted authority cluster.
  - c. Configure the trusted key provider for the trusted hosts on the trusted cluster.
  - d. Configure the trusted key provider for the hosts on the trusted authority cluster.

## Foundation Topics

### Configuring and Managing Authentication and Authorization

This section describes configuration and management tasks related to vSphere authentication and authorization. Authentication tasks involve vCenter Single Sign-On (SSO), and authorization involves permissions.

## Managing SSO

As explained in previous chapters, you can use the built-in identity provider vCenter SSO and external identity providers for vSphere authentication. SSO includes the Security Token Service (STS), an administration server, the vCenter Lookup Service, and the VMware Directory Service (vmdir). The VMware Directory Service is also used for certificate management.

Chapter 8, “vSphere Installation,” discusses the following procedures:

- Adding and editing identity sources
- Adding the vCenter Appliance to an Active Directory domain
- Configuring SSO password, lockout, and token policies

This section describes the procedures for enabling Windows session authentication (SSPI) and managing STS. This section also describes how to implement and use Enhanced Linked Mode.

### Note

The lockout policy applies only to user accounts and not to system accounts such as [administrator@vsphere.local](mailto:administrator@vsphere.local).

## Enabling SSO with Windows Session Authentication

To enable Windows session authentication (SSPI), you can use the following procedure:

- Step 1.** Prepare an Active Directory domain and its trusts in an SSO-trusted manner, as described at <https://kb.vmware.com/s/article/2064250>.
- Step 2.** Join the vCenter Server to the Active Directory domain, as described in [Chapter 8](#).
- Step 3.** Install the Enhanced Authentication Plug-In.
- Step 4.** Instruct vSphere Client users to select the **Use Windows Session Authentication** checkbox during login.



## Note

If you use federated authentication with Active Directory Federation Services, the Enhanced Authentication Plug-in applies only if vCenter Server is the identity provider.

## Managing Service Token Service (STS)

The *vCenter Single Sign-On Security Token Service (STS)* is a web service that issues, validates, and renews security tokens. It uses a private key to sign tokens and publishes the public certificates. SSO manages the certificates that are used by STS for signing and stores the certificates (the signing certificates) in VMware Directory Service (vmdir). You can use the following procedure to generate a new STS signing certificate:

- Step 1.** Create a top-level directory.
- Step 2.** Copy the certool.cfg file into the new directory.
- Step 3.** Modify the certool.cfg file to use the local vCenter Server IP address and host name.
- Step 4.** Generate the key by running `/usr/lib/vmware-vmca/bin/certool --genkey`.
- Step 5.** Generate the certificate by running `/usr/lib/vmware-vmca/bin/certool --gencert`.
- Step 6.** Create a PEM file with the certificate chain and private key.

## Note

The certificate is not external facing, and it is valid for 10 years. You should replace this certificate only if required by your company's security policy.

To use a company-required certificate or to refresh a certificate that is near expiration, you can use the PEM file from the previous procedure and the **sso-config** utility command to refresh the STS certificate, as in the following

example:

[Click here to view code image](#)

```
/opt/vmware/bin/sso-config.sh -set_signing_cert -t vsphere.local  
~/newsts/newsts.pem
```

## Enhanced Linked Mode

To join vCenter Server systems in Enhanced Linked Mode, you need to connect them to the same SSO domain. For example, during the deployment of a vCenter Server, you can choose to join the SSO domain of a previously deployed vCenter Server.

When implementing Enhanced Linked Mode, you should ensure that you properly synchronize the time settings of the new appliance to match those of the previously deployed appliance. For example, if you are using the vCenter Server GUI installer to deploy two new vCenter Server systems joined in Enhanced Linked Mode to the same ESXi host, you can configure them to synchronize the time settings with the host.

You can use a single vSphere Client window to manage multiple vCenter Server systems that are joined with Enhanced Linked Mode. Enhanced Linked Mode provides the following features for vCenter Server:

- You can log in to all linked vCenter Server systems simultaneously.
- You can view and search the inventories of all linked vCenter Server systems.
- Roles, permission, licenses, tags, and policies are replicated across linked vCenter Server systems.

### Note

Enhanced Linked Mode requires the vCenter Server Standard licensing level.

## Users and Groups

By default, immediately following installation, only the localos and the SSO domain (which is sphere.local by default) identity sources are available. [Chapter](#)

8 describes how to add identity sources, such as a native Active Directory (integrated Windows authentication) domain, an OpenLDAP directory service, or Active Directory as an LDAP server. It also describes how to create users and groups in the SSO domain.

### Note

After creating a user or group, you cannot change its name.

When using the procedure in [Chapter 8](#) to add members to a group in the SSO domain, you can add users from identity sources.

In some cases, you might want to manage multiple independent vSphere environments that have similar but separate SSO domains and users. In such scenarios, you can export SSO users by using this procedure:

- Step 1.** Log on to the vSphere Web Client.
- Step 2.** Select **Home > Administration**.
- Step 3.** Select **Single Sign On > Users and Groups**.
- Step 4.** Click the **Users** tab.
- Step 5.** Click the **Export List** icon in lower-right corner.

You can use a similar procedure to export SSO groups except that you choose the Groups tab in step 4 instead of the Users tab.

## Privileges and Roles

To create a role in vCenter Server using the vSphere Client, you can use this procedure:

- Step 1.** Click **Menu > Administration > Roles**.
- Step 2.** Click the **Create Role Action (+)** button.
- Step 3.** Provide a name for the role.
- Step 4.** Select the desired privileges.
- Step 5.** Click **OK**.

After you create custom roles, you can use these roles when assigning permissions in the same manner as you use the vCenter Server system roles and sample roles.

To clone a sample role or custom role in the vSphere Client, navigate to Administration > Roles and select the role, click the Clone role action icon, and provide a name for the new role. To edit a sample role or custom role in the vSphere Client, navigate to Administration > Roles and select the role, click the Edit Role Action icon, and modify the set of privileges in the role.

## Permissions

To set a permission using the vSphere Client, you can use the following steps:

- Step 1.** Select the object in the inventory.
- Step 2.** Click the **Permissions** tab.
- Step 3.** Click the **Add Permission** icon.
- Step 4.** Select a user or group from the **User** drop-down menu.
- Step 5.** Select a role from the **Role** drop-down menu.
- Step 6.** Optionally, select **Propagate to Children**.
- Step 7.** Click **OK**.

By assigning a different role to a group of users on different objects, you control the tasks that those users can perform in your vSphere environment. For example, to allow a group to configure memory for the host, select that host and add a permission that grants a role to that group that includes the Host.Configuration.Memory Configuration privilege.

## Global Permissions

In some cases, you might assign a global permission and choose not to propagate to child objects. This may be useful for providing a global functionality, such as creating roles. To assign a global permission, you should use the vSphere Client with a user account that has the Permissions.Modify privilege on the root object of all inventory hierarchies. Select Administration > Global Permissions > Manage and use the Add Permission icon (plus sign). Then use the dialog that appears to select the desired user group (or user) and role.

## Note

By default, the administrator account in the SSO domain, such as [administrator@vsphere.local](mailto:administrator@vsphere.local), can modify global permissions, but the vCenter Server Appliance root account cannot.

## Note

Be careful when applying global permission. Decide whether you genuinely want a permission to apply to all solutions and to all objects in all inventory hierarchies.

## Editing Permissions

To modify an existing permission, you can edit the permission and change role assignment. You cannot change the object, user, or user group in the permission, but you can change the role. If this is not adequate, you need to remove the permission and create a new permission with the correct settings. You must do this work as a user with sufficient privileges to change permissions on the associated object.

The biggest challenge in editing permissions may be locating the permission in order to modify it. If you know the object on which a permission was created, then you can select the object in the vSphere Client inventory, select **Configure > Permissions**, right-click the permission, and choose **Change Role**. Then you select the appropriate role and click **OK**.

If you do not already know which permission to modify or on which object the permission is assigned, you may need to investigate. Begin by selecting an object in the inventory on which you know the applied user permissions are incorrect. Select **Manage > Permissions** to discover all the permissions that apply to the object. Use the **Defined In** column to identify where each applied permission is defined. Some of the permissions may be assigned directly on the object, and some may be assigned to ancestor objects. Determine which permissions are related to the issue and where they are assigned.

## Configuring and Managing vSphere Certificates

You can use the vSphere Client and vSphere *Certificate Manager* to view and manage certificates. With the vSphere Client, you can perform the following tasks:

- View trusted root certificates and machine SSL certificates.
- Renew or replace existing certificates.
- Generate a custom certificate signing request (CSR) for a machine SSL certificate.

For each certificate management task, you should use the administrator account in the SSO domain (which is vsphere.local by default).

## Managing vSphere Client Certificates

You can use the following procedure to explore and take actions on the certificate stored in a VMware Endpoint Certificate Service (VECS) instance:

- Step 1.** In the vSphere Client, navigate to **Home > Administration > Certificates > Certificate Management**.
- Step 2.** If the system prompts you to do so, enter the credentials for your vCenter Server. The Certificate Management page shows the certificate types in the VECS. By default, the types are
  - machine SSL Certificates and Trusted Root certificates.
- Step 3.** For more details, click **View Details** for the certificate type.
- Step 4.** For the machine SSL certificates, optionally choose from the following actions:
  - **Renew**
  - **Import and Replace Certificate**
  - **Generate CSR**
- Step 5.** For trusted root certificates, optionally choose **Add**.

### Note

To replace all VMCA-signed certificates with new VMCA-signed

certificates, choose the Renew action for the machine SSL certificates.

If you replace an existing certificate, you can remove the old root certificate (as long as you are sure it is no longer in use).

By default, vCenter Server monitors all certificates in VECS and raises an alarm for any certificate that will expire in 30 days or less. You can change the 30-day threshold by modifying vCenter Server's advanced setting `vpxd.cert.threshold`.

## Using Custom Certificates

To set up your environment to use custom certificates, you need to generate a CSR for each machine and each solution user and replace certificates when you receive them. You can generate the CSRs by using the vSphere Client or Certificate Manager. You can use the vSphere Client to upload both the root certificate and the signed certificates that are returned from the CA.

You can use the following procedure to generate a CSR for custom certificates:

- Step 1.** Verify that you meet the certificate requirements described in [Chapter 7, “vSphere Security.”](#)
- Step 2.** In the vSphere Client, navigate to **Home > Administration > Certificates > Certificate Management.**
- Step 3.** If prompted to do so, enter the credentials for your vCenter Server.
- Step 4.** In the Machine SSL Certificate section, for the certificate you want to replace, click **Actions > Generate Certificate Signing Request (CSR).**
- Step 5.** Enter your certificate information and click **Next.**
- Step 6.** Copy or download the CSR.
- Step 7.** Click **Finish.**
- Step 8.** Provide the CSR to your certificate authority.

Alternatively, you can use the vSphere Certificate Manager utility from the vCenter Server shell to generate the CSR, by using the following command, selecting option 1, and providing the certificate information:

[Click here to view code image](#)

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

After your CA processes the CSR, you can use the following procedure to add the custom certificates:

- Step 1.** In the vSphere Client, navigate to **Home > Administration > Certificates > Certificate Management**.
- Step 2.** If the system prompts you to do so, enter the credentials for your vCenter Server.
- Step 3.** In the Machine SSL Certificate section, for the certificate you want to replace, click **Actions > Import and Replace Certificate**.
- Step 4.** Select the **Replace with External CA Certificate (requires private key)** option and click **Next**.
- Step 5.** Upload the certificates and click **Replace**.
- Step 6.** Wait for the vCenter Server services to restart.

## Managing ESXi Certificates

In vSphere 6.0 and later, ESXi hosts initially boot with an autogenerated certificate. When a host is added to a vCenter Server system, it is provisioned with a certificate signed by the VMware Certificate Authority (VMCA). You can view and manage ESXi certificates by using the vSphere Client or the vim.CertificateManager API in the vSphere Web Services SDK. You cannot use the vCenter Server certificate management CLIs to view or manage ESXi certificates.

## Changing the Certificate Mode

You can change the ESXi certificate mode from VMCA Mode to Custom Certificate Authority Mode or to Thumbprint Mode. In most cases, mode switches are disruptive and not necessary. If you require a mode switch, be sure to review the potential impact before you start. You should use Thumbprint Mode only for debugging.

### Note

Thumbprint Mode was used in vSphere 5.5 and should not be used in later



versions unless it is necessary because some services may not work. Also, in Thumbprint Mode, vCenter Server checks only the certificate format and not its validity. Even expired certificates are accepted.

To perform certificate management for ESXi, you must have the Certificates.Manage Certificates privilege.

For example, if you want to use custom certificates instead of using VMCA to provision ESXi hosts, you need to edit the vCenter Server `vpdx.certmgmt.mode` advanced option. In the vSphere client, you can use this procedure to change the certificate mode:

- Step 1.** Select the vCenter Server and click **Configure**.
- Step 2.** Click **Advanced Settings** and then click **Edit**.
- Step 3.** In the Filter box, enter **certmgmt** to display only certificate management keys.
- Step 4.** Change the value of `vpdx.certmgmt.mode` to **custom** and click **OK**.
- Step 5.** Restart the vCenter Server service.

## Using Custom ESXi Certificates



You can switch the certificate mode from VMCA to a different root CA by using these steps:

- Step 1.** Obtain the certificates from the trusted CA.
- Step 2.** Place the host or hosts into Maintenance Mode and disconnect them from vCenter Server.
- Step 3.** Add the custom CA's root certificate to VECS.
- Step 4.** Deploy the custom CA certificates to each host and restart services on that host.
- Step 5.** Change Certificate Mode to **Custom CA Mode** (as described in the previous section).
- Step 6.** Connect the host or hosts to the vCenter Server system.

## Switching Back to VMCA Mode

If you are using the Custom CA Mode, you can switch back to VMCA Mode by using this procedure:

- Step 1.** Remove all hosts from the vCenter Server system.
- Step 2.** On the vCenter Server system, remove the third-party CA's root certificate from VECS.
- Step 3.** Change Certificate Mode to **VMCA Mode**. (See the [“Changing the Certificate Mode”](#) section.)
- Step 4.** Add the hosts to the vCenter Server system.

## Certificate Expiration

For ESXi 6.0 and later, you can use the vSphere Client to view information, including expiration, for all certificates that are signed by VMCA or a third-party CA. In the vSphere Client, select the host and navigate to **Configure > System > Certificate**. Here you can examine the Issuer, Subject, Valid From, Valid To, and Status fields. The value of the Status field may be Good, Expiring, Expiring Shortly, Expiration Imminent, or Expired.

A yellow alarm is raised if a certificate's status is Expiring Shortly (that is, if it expires in less than eight months). A red alarm is raised if the certificate's status is Expiration Imminent (that is, if it expires in less than two months).

By default, each time a host reconnects to vCenter Server, it renews any host certificates whose status is Expired, Expiring Immediately, or Expiring. If a certificate is already expired, you must disconnect the host and reconnect it. To renew or fresh the certificates, you can use the following procedure:

- Step 1.** In the vSphere Client, select the host in the navigation pane.
- Step 2.** Navigate to **Configure > System > Certificate**.
- Step 3.** Click one of the following options:
  - **Renew:** Retrieves a fresh signed certificate for the host from VMCA.
  - **Refresh CA Certificates:** Pushes all certificates in the VECS TRUSTED\_ROOTS store to the host.
- Step 4.** Click **Yes**.

# General ESXi Security Recommendations

In [Chapter 7](#), you learned that vSphere has built-in security features and that you can take additional steps to harden ESXi. The following items are additional security measures recommended by VMware:

- Limit access to the Direct Console User Interface (DCUI), the ESXi Shell, and Secure Shell (SSH). If you allow access to these items, which have privileged access to certain ESXi components, you need to ensure that only trusted users have access and that timeouts are set.
- Do not directly access ESXi hosts that are managed by vCenter Server. Although it may be possible to access a host via DCUI, SSH, ESXi Shell, API, or vSphere Host Client, you should not normally do so. Instead, you should use the vSphere Client (or vSphere Web Client) or API connected to vCenter Server to manage the ESXi host.
- Use the DCUI only for troubleshooting. Likewise, use root access to the ESXi Shell only for troubleshooting.
- When upgrading ESXi components, use only VMware sources. Although a host runs several third-party packages, VMware supports upgrades to those packages only from VMware sources. Check third-party vendor sites and the VMware knowledge base for security alerts.
- You should follow the VMware security advisories at <http://www.vmware.com/security/>.
- Configure ESXi hosts with host profiles, scripts, or some other automation.

## Configuring ESXi Using Host Profiles

You can use host profiles to set up standard secured configurations for your ESXi hosts and automate compliance.

You can consider any setting that is applied by a host profile to be important to ensuring that your hosts are secured. Some settings, such as direct ESXi permissions, may be obvious. Other settings, such as NTP settings, may not be obvious, but time synchronization issues impact integration with Active Directory, which impacts user authentication. Network settings, such as physical NIC speed, could impact the ability of a host to connect to the proper management network.

As discussed in [Chapter 8](#), host profiles can be used to apply many host configuration settings, including security measures, such as ESXi-level permissions. You can use the vSphere Client to configure a host profile for a reference host and apply the host profile to a set of hosts. You can also use host profiles to monitor hosts for host configuration changes. You can attach a host profile to a cluster to apply it to all hosts in the cluster. These are the high-level steps:

- Step 1.** Set up the reference host to specification and create a host profile.
- Step 2.** Attach the profile to a host or cluster.
- Step 3.** Apply the host profile from the reference host to other hosts or clusters.

To ensure that an ESXi host is properly configured according to your standards, you can check for its compliance to its attached host profile. You can use the results to identify noncompliant settings on the host and remediate with the host profiles settings. You can use these steps to check compliance:

- Step 1.** Navigate to the Host Profiles main view.
- Step 2.** Right-click a host profile.
- Step 3.** Click **Check Host Profile Compliance**

The compliance status for each ESXi host is Compliant, Unknown, or Noncompliant. Noncompliant status indicates a specific inconsistency between the profile and the host, which you should remediate. Unknown status indicates that the compliance of the host is not known because it could not be verified. A common root cause is that the host is disconnected. You should resolve the issue and recheck compliance.

## Using Scripts to Manage Host Configuration Settings

Another means to establish a standard secured configuration for ESXi hosts in your vSphere environment is to use vSphere PowerCLI, ESXCLI, or custom code leveraging the vSphere API.

### Note

Starting with vSphere 7.0, the vSphere CLI package is end of life. Its capabilities are supported with more API-centric tools such as ESXCLI and Perl SDK.

From the ESXi Shell, you can use the ESXCLI command set to configure the host and to perform administrative tasks. ESXCLI provides a collection of namespaces that allow an administrator to quickly discover the precise command necessary for a specific task. For example, all the commands to configure networking exist in the `esxcli network` namespace, and all the commands to configure storage exist in the `esxcli storage` namespace. Each namespace is further divided into child namespaces that comprise various functions performed under the parent namespace. For example, the `esxcli storage` parent namespace contains a `core` namespace that deals with storage adapters and devices and an `nmp` namespace that deals with path selection and storage array types. Therefore, a typical ESXCLI command is composed of multiple namespaces, where each additional namespace is used to narrow the scope of the command, ending with the actual operation to be performed.

To identify the proper ESXCLI command to perform a specific task, you can begin by entering `esxcli` at the command prompt in the ESXi Shell. Because it is not a command by itself, just the entry point to the namespace hierarchy, the results will show the first level of the namespace hierarchy. The first level of available namespaces includes `device`, `esxcli`, `fcoe`, `graphics`, `hardware`, `iscsi`, `network`, `nvme`, `rdma`, `sched`, `software`, `storage`, `system`, `vm`, and `vsan`. You can use the brief description of each namespace shown in the results to identify which namespace is most likely to serve your need. You can press the up-arrow key on the keyboard to retrieve the last entered namespace and add the name for the next namespace. You can continue reviewing namespaces until you discover the command you need.

For example, if you are seeking a command to list all standard virtual switches, you could enter `esxcli network` to learn that it contains several namespaces, including one named `vswitch`. You could then enter `esxcli network vswitch` and learn that its namespaces are `standard` and `dvs`. Going further, you could learn that the `esxcli network vswitch standard` namespace contains the list. You can conclude that the command you need is `esxcli network vswitch standard list`. [Table 12-2](#) lists a few other examples of ESXCLI commands.

**Table 12-2** Sample ESXCLI Commands

Command	Description
<code>esxcli system account add</code>	Creates an ESXi host local

	user account
<b>esxcli system account set</b>	Configures an ESXi host local user account
<b>esxcli system account list</b>	Lists ESXi host local user accounts
<b>esxcli system account remove</b>	Deletes ESXi host local user accounts
<b>esxcli network ip dns server list</b>	Lists the host's DNS servers
<b>esxcli network nic list</b>	Lists the ESXi host's physical network adapters
<b>esxcli system settings advanced get /UserVars/ESXiShellTimeout</b>	Displays the shell interactive timeout for the host

Likewise, you can use PowerCLI to manage and configure a vSphere environment. When connecting to a vCenter Server environment, the functionality scope of PowerCLI is similar to the functionality scope of using the vSphere Client with the vCenter Server. [Table 12-3](#) describes a few popular PowerCLI commands.

**Table 12-3** Sample PowerCLI Commands

<b>Command</b>	<b>Description</b>	<b>Example</b>
<b>Connect-Viserver</b>	Connects to a vCenter Server	<b>Connect-Viserver vc01 -User administrator@vsphere.local</b>  Connects to a vCenter Server named vc01 as user administrator@vsphere.local
<b>Get-VMHost</b>	Retrieves information about one or more ESXi hosts	<b>Get-VMHost -Location MyDC</b>  Retrieves details about all ESXi hosts in a data center named MyDC
<b>Set-VMHost</b>	Changes a setting or state of an ESXi host	<b>Set-VMHost -VMHost Host -State "Disconnected"</b>

t

Disconnects the host from vCenter Server

---

If you want to develop code using other tools, you may want to get familiar with vSphere REST APIs. To do so, you can browse to the FQDN of your vCenter Server and select Browse vSphere REST APIs. In vCenter Server 7.0, this link takes you to the API Explorer section of the Developer Center in the vSphere Client. Here you can learn how to make GET and POST calls to query and modify the state and configuration of your ESXi hosts and other vSphere objects.

## ESXi Passwords and Account Lockout

For direct ESXi host access, you can use the local root account and additional user accounts that you create directly on the host. When setting a password on these accounts, you must comply with or modify the predefined requirements. ESXi uses the Linux PAM module `pam_passwdqc` for password management and control. You can change the required length, change character class requirement, and allow passphrases by using the `Security.PasswordQualityControl` advanced option.

### Note

The default requirements for ESXi passwords can change from one release to the next. You can check and change the default password restrictions by using the `Security.PasswordQualityControl` advanced option.



One step in hardening an ESXi host is to harden the password required to use its predefined local administrator account, which is called root. By default, the ESXi host enforces passwords for its local user accounts, which may be used to access the host via the DCUI, the ESXi Shell, SSH, or the vSphere Client. Starting with ESXi 6.0, the default password policy must contain characters from at least three character classes (of the four character classes, which are lowercase letters, uppercase letters, numbers, and special characters) and must be at least seven characters long.

An uppercase character that begins a password and a number that ends a password do not count toward the number of character classes used. A password cannot contain a dictionary word or part of a dictionary word. For example, xQaT3!A is an acceptable password because it contains four character classes and seven characters. However, Xqate!3 is not an acceptable password because it contains only two character classes; the leading X and ending 3 do not count toward the number of used character classes. You can modify the ESXi password requirements by using the ESXi host Security.PasswordQualityControl advanced option. You can set Security.PasswordQualityControl to configure the ESXi host to accept passphrases, which it does not accept by default. The key to changing the password and passphrase requirements is understanding the syntax and functionality of the Security.PasswordQualityControl parameter, which has the following default value:

[Click here to view code image](#)

```
retry=3 min=disabled,disabled,disabled,7,7
```

The first part of the value used for this parameter identifies the number of retries allowed for the user following a failed logon attempt. In the default value, **retry=3** indicates that three additional attempts are permitted following a failed logon. The remainder of the value can be abstracted as follows:

```
min=N0, N1, N2, N3, N4
```

The values in this parameter are as follows:

- **N0:** This is the minimum number of accepted characters for passwords that contain characters from only one class; it can be disabled to disallow passwords that contain characters from only one class.
- **N1:** This is the minimum number of accepted characters for passwords that contain characters from only two classes; it can be disabled to disallow passwords that contain characters from only two classes.
- **N2:** This is the minimum number of accepted characters for passphrases, and it can be disabled to disallow passphrases. In addition, to require a passphrase, you can append `passphrase=N` to the end of the value, where *N* specifies the minimum number of words, separated by spaces, in the passphrase.
- **N3:** This is the minimum number of accepted characters for passwords that contain characters from only three classes; it can be disabled to disallow



passwords that contain characters from only three classes.

- **N4:** This is the minimum number of accepted characters for passwords that contain characters from all four classes.

For example, to require a passphrase with a minimum of 16 characters and 3 words, you can set `Security.PasswordQualityControl` as follows:

[Click here to view code image](#)

```
retry=3 min=disabled,disabled,16,7,7,passphrase=3
```

The password requirements in ESXi 6.0 are implemented by `pam_passwdqc`. For more details, see the man pages for `pam_passwdqc`.

In vSphere 6.x, account locking is supported for access through SSH and through the vSphere Web Services SDK. The DCUI and the ESXi Shell do not support account lockout. By default, a maximum of 10 failed attempts is allowed before an account is locked. The account is unlocked after two minutes by default. You can modify the lockout behavior by using the host's advanced options:

- **Security.AccountLockFailures:** The maximum number of failed login attempts before a user's account is locked. Zero disables account locking.
- **Security.AccountUnlockTime:** The number of seconds that a user is locked out.

## SSH and ESXi Shell Security

You can use SSH to remotely log in to the ESXi Shell and perform troubleshooting tasks for the host. SSH configuration in ESXi is enhanced to provide a high security level. VMware does not support SSH Version 1 and uses Version 2 exclusively. SSH supports only 256-bit and 128-bit AES ciphers for connections.

The ESXi Shell is disabled by default on ESXi hosts. If necessary, you can enable local and remote access to the shell. However, to reduce the risk of unauthorized access, you should enable the ESXi Shell only when troubleshooting. If the ESXi Shell or SSH is enabled and the host is placed in Lockdown Mode, accounts on the exception users list who have administrator privileges can use these services. For all other users, ESXi Shell or SSH access is disabled. Starting with vSphere 6.0, ESXi or SSH sessions for users who do

not have administrator privileges are closed.

If the ESXi Shell is enabled, you can still log in to it locally, even if the host is running in Lockdown Mode. To enable local ESXi Shell access, enable the ESXi Shell service. To enable remote ESXi Shell access, enable the SSH service.

## Note

The root user and users with the administrator role can access the ESXi Shell. Users who are in the Active Directory group ESX Admins are automatically assigned the administrator role. By default, only the root user can run system commands (such as **vmware -v**) by using the ESXi Shell.

You can use the following procedure to enable the ESXi Shell:

- Step 1.** In the vSphere Client, select the host in the navigation pane.
- Step 2.** Navigate to **Configure > Services**.
- Step 3.** Select **ESXi Shell** and click **Start**.
- Step 4.** Optionally, click **Edit Startup Policy** and then select one of the following options:
  - **Start and Stop Manually**
  - **Start and Stop with Host**
  - **Start and Stop with Port Usage**
- Step 5.** Click **OK**.

You can use a similar procedure to control local and remote access to the ESXi Shell by configuring the startup policy for DCUI and SSH services.

To increase security, you can set ESXi Shell timeout. The Availability Timeout setting is the amount of time that can elapse before you must log in after the ESXi Shell is enabled. After the timeout period, the service is disabled, and users are not allowed to log in. To set the timeout, you can use this procedure:

- Step 1.** Browse to the host in the vSphere Web Client inventory.
- Step 2.** Click **Configure**.

- Step 3.** Under System, select **Advanced System Settings**.
- Step 4.** Select **UserVars.ESXiShellTimeOut** and click **Edit**.
- Step 5.** Enter the idle timeout setting.
- Step 6.** Click **OK**.
- Step 7.** Restart the SSH service and the ESXi Shell service for the timeout to take effect.

Likewise, you can set a timeout for idle ESXi Shell sessions. The idle timeout is the amount of time that can elapse before a user is logged out of an idle interactive session. You can control the amount of time for both local and remote (SSH) sessions from the DCUI or from the vSphere Web Client by using the following procedure:

- Step 1.** Browse to the host in the vSphere Web Client inventory.
- Step 2.** Click **Configure**.
- Step 3.** Under System, select **Advanced System Settings**.
- Step 4.** Select **UserVars.ESXiShellInteractiveTimeOut**, click the **Edit** icon, and enter the timeout setting.
- Step 5.** Restart the ESXi Shell service and the SSH service so that the timeout takes effect. Now, if the session is idle, users are logged out after the timeout period elapses.

An SSH key can allow a trusted user or script to log in to a host without specifying a password. You can upload the authorized keys file for the root user, an RSA key, or an RSA public key to a host. To upload an RSA public key to a host, you can use the following **vifs** command:

[Click here to view code image](#)

```
vifs --server hostname --username username --put filename /host/  
ssh_host_dsa_key_pub
```

To upload an RSA key or root user authorized key files, use the same command but change the target to **/host/ssh\_host\_rsa\_key** or **/host/ssh\_root\_authorize\_keys**, respectively.

## **PCI and PCIe Devices and ESXi**

You can use the VMware DirectPath I/O feature to pass through a PCI or a PCIe device to a virtual machine, but this results in a potential security vulnerability that could be triggered when buggy or malicious code, such as a device driver, is running in Privileged Mode in the guest OS. Therefore, you should use PCI or PCIe passthrough to a virtual machine *only* if a trusted entity owns and administers the virtual machine. Otherwise, you risk the host being compromised in several ways:

- The guest OS might generate an unrecoverable PCI or PCIe error.
- The guest OS might generate a Direct Memory Access (DMA) operation that causes an IOMMU page fault on the ESXi host.
- If the operating system on the ESXi host is not using interrupt remapping, the guest OS might inject a spurious interrupt into the ESXi host on any vector.

To enable passthrough for a network device on a host, you can use the following procedure:

**Step 1.** In the vSphere Client, select the host in the navigation pane.

**Step 2.** Navigate to **Configure > Hardware > PCI Devices** and click **Edit**.

**Step 3.** Select a device with a green icon and click **OK**.

### Note

An orange icon indicates that the status of the device has changed, and you must reboot the host before you can use the device.

## Disabling the Managed Object Browser

The *managed object browser (MOB)* provides a means to explore the VMkernel object model. Starting with vSphere 6.0, the MOB is disabled by default to avoid malicious configuration changes or actions. You can enable and disable the MOB manually. VMware recommends that you not enable MOB in production systems.

To enable the MOB by using the vSphere Client, you can use the following procedure:

- Step 1.** In the vSphere Client, select the host in the inventory.
- Step 2.** In the right pane, click the **Configuration** tab.
- Step 3.** Select **System > Advanced Settings** and click **Edit**.
- Step 4.** Select **Config.HostAgent.plugins.solo.enableMob** and set its value to **true**.

## ESXi Networking Security Recommendations

[Chapter 7](#) provides VMware general network security recommendations for vSphere. Concerning each ESXi host, you can summarize the network isolation into the following categories.

- **vSphere infrastructure networks:** Isolate these networks for their specific functions, such as vMotion vSphere Fault Tolerance, storage, and vSAN. In many cases, you may not need to route these networks outside a single rack.
- **Management network:** This network carries client API and third-party software traffic. Isolate this network such that it is accessible only by the appropriate administrators and systems. Consider using a jump box or a virtual private network (VPN).
- **Virtual machine networks:** This network may involve many networks, each with unique isolation requirements. Consider using virtual firewall solutions, such as NSX.

## ESXi Web Proxy Settings

If you configure a web proxy, consider the following:

- Do not use certificates that use passwords or passphrases. ESXi does not support web proxies with passwords or passphrases, also known as encrypted keys.
- If you want to disable SSL for vSphere Web Services SDK connections, you can change the connection from HTTPS to HTTP. You should consider doing this only if you have a fully trusted environment, where firewalls are in place and transmissions to and from the host are fully isolated.
- Most internal ESXi services are accessible only through port 443. Port 443 acts as a reverse proxy for ESXi. You can change the configuration to allow

direct HTTP connections but should consider this only for a fully trusted environment.

- During upgrades, the certificate remains in place.

## **vSphere Auto Deploy Security Considerations**

If you use vSphere Auto Deploy, consider the networking security, boot image security, and potential password exposure through host profiles.

You should secure an Auto Deploy network as you would secure the network for any PXE-based deployment method. Auto Deploy transfers data over SSL, but it does not check the authenticity of the client or of the Auto Deploy server during a PXE boot.

The boot image includes the host's public and private SSL key and certificate. If Auto Deploy rules are set up to provision the host with a host profile or host customization, the boot image includes the host profile and host customization. The root password and user passwords in the host profile and host customization are hashed with SHA-512. Other passwords, such as those to set up Active Directory using the host profile, are not protected. You can use vSphere Authentication Proxy to avoid exposing Active Directory passwords.

Ideally, you should completely isolate an Auto Deploy network.

## **Controlling CIM Access**

***Common Information Model (CIM)*** is an open standard that defines a framework for agentless, standards-based monitoring of ESXi host hardware resources. The framework consists of a CIM broker and a set of CIM providers. Hardware vendors, including server manufacturers and hardware device vendors, can write providers that monitor and manage their devices. VMware has CIM providers that monitor server hardware, ESXi storage infrastructure, and virtualization-specific resources. These lightweight providers run inside the ESXi host and perform specific management tasks.

Instead of using root credentials, you can create a less-privileged vSphere user account to provide to remote applications that access the CIM interface. The required privilege for the user account is Host.CIM.Interaction.

You can use the VIM API ticket function to issue a session ID (ticket) to a user account to authenticate to CIM. You need to ensure that the account is granted permission to obtain CIM tickets.

When you install a third-party CIM VIB, the CIM service starts. To manually enable the CIM service, you can use the following command:

```
esxcli system wbem set -e true
```

You can use the API SDK of your choice to call `AcquireCimServicesTicket` to return a ticket that you can use to authenticate the user with vCenter Server using CIM-XML port 5989 or WS-Man port 433 APIs.

## Configuring and Managing ESXi Security

This section describes procedures for securing ESXi.

### Configuring the ESXi Firewall

By default, the ESXi firewall is configured to block incoming and outgoing traffic, except traffic for services that are enabled in the hosts' security profile. The firewall also allows Internet Control Message Protocol (ICMP) pings. Prior to opening any ports on the firewall, you should consider the impact it may have for potential attacks and unauthorized user access. You can reduce this risk by configuring the firewall to only allow communication on the port with authorized networks. To modify the firewall's rule set, you can use the vSphere Client to modify the host's security profile, using this procedure:



- Step 1.** In the vSphere Client, select the host in the inventory pane and navigate to **Configure > System > Firewall**.
- Step 2.** Select the appropriate service name, such as the incoming SSH server (TCP 22) or the outgoing DNS client (TCP/UDP 53), and click **Edit**.
- Step 3.** Examine the rule set. Change the state of any rule by selecting the rule (placing a check in the rule's box) to enable the rule or deselecting the rule to disable it.
- Step 4.** Optionally, for some services, you can deselect the **Allow Connections from Any IP Address** box and enter specific IP addresses in the accompanying text box to restrict use to only those IP addresses.
- Step 5.** Click **OK**.

When specifying particular IP addresses in the firewall settings, you can use the

formats used in the following examples:

- 192.168.10.0/24
- 192.168.11.2, 2001::1/64
- fd3e:29a6:0a79:e462::/64

The NFS Client firewall rule set behaves differently than other rule sets. ESXi configures NFS Client settings when you mount or unmount an NFS datastore. When you mount an NFS Version 3 datastore, the following events occur:

- If the nfsClient rule set is disabled, ESXi enables the rule set, sets allowedAll to FALSE, and adds the NFS server IP address to the list of allowed IP addresses.
- If the nfsClient rule set is enabled, ESXi adds the NFS server IP address to the list of allowed IP addresses but does not change the state of the rule set or allowedAll.
- When you mount an NFS Version 4.1 datastore, ESXi enables the nfs41client rule set and sets allowedAll to TRUE.

When you remove or unmount an NFS Version 3 datastore from a host, ESXi removes the IP address from the list of allowed IP addresses. When you remove or unmount the last NFS Version 3 datastore, ESXi stops the nfsClient rule set. Unmounting an NFS Version 4.1 datastore does not impact the firewall.

The ESXi software firewall is enabled by default. It should never be disabled while running production virtual machines. In rare cases, such as temporarily during troubleshooting, you can disable the ESXi firewall by using the **esxcli network firewall set --enabled false** command.

## Customizing ESXi Services

Several optional services that are provided in an ESXi host are disabled by default. VMware disables these services to provide strong security out of the box. In a default installation, you can modify the status of the following services from the vSphere Client:

- **Running services:** DCUI, Load-Based Teaming Daemon, CIM Server, and VMware vCenter Agent



- **Stopped services:** ESXi Shell, SSH, attestd, kexd, Active Directory service, NTP daemon, PC/SC smart card daemon, SNMP server, syslog server, X.Org server

In some circumstances, you might want to configure and enable these services. A good example of an optional service that you might decide to configure and enable in most environments is NTP because solid time synchronization is vital for many services. As another example, you might want to temporarily enable SSH while troubleshooting. To enable, disable, and configure services, you can use the following procedure:

- Step 1.** In the vSphere Client, select the host in the navigation pane and navigate to **Configure > Services**.
- Step 2.** Select a service that you want to modify and click **Start**, **Stop**, or **Restart** to immediately change the state of the service.
- Step 3.** To change the behavior permanently, click **Edit Startup Policy** and then choose one of the following options:
  - **Start and Stop with Port Usage**
  - **Start and Stop with Host**
  - **Start and Stop Manually**
- Step 4.** Click **OK**.

## Using Lockdown Mode

In vSphere 5.0 and earlier, only the root account can log in to the DCUI on an ESXi host that is in Lockdown Mode. In vSphere 5.1 and later, you can add a user to the DCUI. Access advanced system setting to grant the user access to the DCUI on a host that is in Lockdown Mode, even if the user is not granted the administrator role on the host. The main purpose of this feature is to prepare for catastrophic vCenter Server failures.

vSphere 6.0 and later include an Exception Users list, whose main purpose is to support the use of Lockdown Mode but still support service accounts, which must log on directly to the ESXi host. User accounts in the Exception Users list who have administrator privileges can log on to the DCUI and ESXi Shell.

As described in [Chapter 7](#), you can place a host in normal Lockdown Mode,

Strict Lockdown Mode, or Normal Mode.

To change the Lockdown Mode setting, you can use the followings procedure:

**Step 1.** In the vSphere Client, select an ESXi host in the navigation pane and navigate to **Configure > Security Profile**.

**Step 2.** In the Lockdown Mode panel, click the **Edit** button.

**Step 3.** Click **Lockdown Mode** and choose either **Normal** or **Strict**.

**Step 4.** Click **OK**.

By default, the root account is included in DCUI.Access. You could consider removing the root account from DCUI.Access and replacing it with another account for better auditability.

[Table 12-4](#) provides details on the behavior of an ESXi host in Lockdown Mode.

**Table 12-4** ESXi Lockdown Mode Behavior

Service	Normal Mode	Normal Lockdown Mode	Strict Lockdown Mode
vSphere Web Services API	All users, based on permissions	vCenter (vpxuser)	vCenter (vpxuser)
		Exception users, based on permissions	Exception users, based on permissions
		vCloud Director (vslauser, if available)	vCloud Director (vslauser, if available)
CIM providers	Users with administrator privileges on the host	vCenter (vpxuser)	vCenter (vpxuser)
		Exception users, based on permissions	Exception users, based on permissions

		vCloud Director (vslouser, if available)	vCloud Director (vslouser, if available)
DCUI	Users with administrator privileges on the host and users defined in the DCUI.Access advanced option	Users defined in the DCUI.Access advanced option  Exception users with administrator privileges on the host	DCUI service is stopped
ESXi Shell (if enabled)	Users with administrator privileges on the host	Users defined in the DCUI.Access advanced option  Exception users with administrator privileges on the host	Users defined in the DCUI.Access advanced option  Exception users with administrator privileges on the host
SSH (if enabled)	Users with administrator privileges on the host	Users defined in the DCUI.Access advanced option  Exception users with	Users defined in the DCUI.Access advanced option  Exception users with

administrator  
privileges on  
the host

administrator  
privileges on  
the host

---

## Managing the Acceptance Levels of Hosts and VIBs

A [vSphere Installation Bundle \(VIB\)](#) is a software package that includes a signature from VMware or a VMware partner. To protect the integrity of the ESXi host, do not allow users to install unsigned (community-supported) VIBs. An unsigned VIB contains code that is not certified by, accepted by, or supported by VMware or its partners. Community-supported VIBs do not have digital signatures. The host's acceptance level must be the same or less restrictive than the acceptance level of any VIB you want to add to the host. For example, if the host acceptance level is `VMwareAccepted`, you cannot install VIBs at the `PartnerSupported` level. You should use extreme caution when allowing community-supported VIBs. The following list provides details on defined VIB acceptance levels:

### Key Topic

- **VMwareCertified:** These VIBs go through thorough testing equivalent to VMware in-house quality assurance testing for the same technology. Only I/O Vendor Partner (IOVP) program drivers are published at this level. VMware takes support calls for VIBs with this acceptance level.
- **VMwareAccepted:** These VIBs go through testing that is run by a partner and verified by VMware. CIM providers and PSA plug-ins are among the VIBs published at this level. VMware directs support calls for VIBs with this acceptance level to the partner's support organization.
- **PartnerSupported:** These VIBs are published by a partner that VMware trusts. The partner performs all testing, but VMware does not verify it. VMware directs support calls for VIBs with this acceptance level to the partner's support organization.
- **CommunitySupported:** These VIBs have not gone through any VMware-approved testing program and are not supported by VMware Technical Support or by a VMware partner.

To change the host acceptance level, you can use the following command:

[Click here to view code image](#)

```
esxcli --server=<server_name> software acceptance set
```

## Assigning Privileges for ESXi Hosts

Typically, users should access vSphere via vCenter Server, where privileges are managed centrally. For use cases where some users access ESXi hosts directly, you can manage privileges directly on the host. The following roles are predefined in ESXi:

- **Read Only:** Ability to view but not change assigned objects
- **Administrator:** Ability to change assigned objects
- **No Access:** No access to assigned objects (This role is the default role, and you can override it.)

In vSphere 6.0 and later, you can use ESXCLI to manage local user accounts and to configure permissions on local accounts and on Active Directory accounts. You can connect directly to an ESXi host by using the vSphere Host Client and navigate to Manage > Security & Users > Users to create, edit, and remove local user accounts.

The following user accounts exist on an ESXi host that is not added to a vCenter System:

- **root:** A user account that is created and assigned the administrator role by default on each ESXi host.
- **vpxuser:** A local ESXi user account that is created, managed, and used for management activities by vCenter Server.
- **dcui:** A user account that acts as an agent for the DCUI and cannot be modified or used by interactive users.

### Note

You can remove the access privileges for the root user. But you should first create another user account at the root level and assign it the administrator role.

Much as with vCenter Server, each ESXi host uses role-based permissions for users who log on directly to the ESXi host rather than accessing the host through vCenter Server. ESXi allows the creation of custom roles, but these roles are applied only when a user directly logs on to the host, such as when the user uses the vSphere Host Client to connect to the host directly. In most cases, managing roles and permissions at the host level should be avoided or minimized. To create, edit, and remove roles, you can connect directly to an ESXi host by using the vSphere Host Client and navigate to **Manage > Security & Users > Roles**.

## Using Active Directory to Manage ESXi Users

In scenarios where multiple users need to access multiple ESXi hosts directly (rather than accessing vCenter Server), you face challenges in synchronizing usernames and passwords. To address the challenges, consider joining the hosts to Active Directory and assigning roles to specific Active Directory users and groups. Then require users to provide their Active Directory credentials when logging in directly to the host.

In scenarios where Active Directory users need to access an ESXi host directly, you need to add the host to a directory service and apply permissions to those users. You can use the following steps to add an ESXi host to an Active Directory domain:

- Step 1.** Verify that an Active Directory domain is available.
- Step 2.** Ensure that the host name of the ESXi host is fully qualified with the domain name that matches the domain name of the Active Directory forest. For example, if the Active Directory domain name is [mydomain.com](#) and the ESXi host name is host-01, then the host's fully qualified name is [host-01.mydomain.com](#).
- Step 3.** Synchronize time between the ESXi host and domain controllers by using NTP.
- Step 4.** Ensure that DNS is configured and that the ESXi host can resolve the host names of the Active Directory domain controllers.
- Step 5.** In the vSphere Client, select the ESXi host in the inventory pane and navigate to **Configure > Authentication Services**.
- Step 6.** Click **Join Domain**.
- Step 7.** In the dialog box, specify the domain and user credentials. Optionally,

specify a proxy server.

**Step 8.** Enter a domain, either in the form *name.tld* or in the form *name.tld/container/path*, where *name.tld* is the domain name, and */container/path* is an optional path to an organization unit where the host computer object should be created. For example, you can use [domain.com/ou01/ou02](#). to add the host to an organization unit named ou02 that resides in an organization unit named ou01 in a domain named [domain.com](#).

**Step 9.** Click **OK**.

## Configuring vSphere Authentication Proxy

You can use vSphere Authentication Proxy to add hosts to an Active Directory domain instead of adding hosts to the domain explicitly. When vSphere Authentication Proxy is enabled, it automatically adds hosts that are being provisioned by Auto Deploy to the Active Directory domain. You can also use vSphere Authentication Proxy to add hosts that are not provisioned by Auto Deploy.

To start the vSphere Authentication Proxy service and add a domain, you can use the following procedure:

**Step 1.** Log on to the vCenter Server Appliance Management Interface (VAMI) as root.

**Step 2.** Select **Services > VMware vSphere Authentication Proxy**.

**Step 3.** Click **Start**.

**Step 4.** In the vSphere Client, select the vCenter Server in the inventory pane and navigate to **Configure > Authentication Proxy > Edit**.

**Step 5.** Enter the domain name and credentials of a user who can add hosts to the domain

**Step 6.** Click **Save**.

Now you can add a host to an Active Directory domain by using the procedure outlined in the section “Use Active Directory to Manage ESXi Users,” but this time, you select the **Using Proxy Server** option in step 8.

## Configuring Smart Card Authentication for ESXi

As an alternative to specifying a username and password, you can use smart card authentication to log in to the ESXi DCUI by using a Personal Identity Verification (PIV) credential, a Common Access Card (CAC), or an SC650 smart card. In this case, the DCUI prompts for a smart card and PIN combination. To configure smart card authentication, you should set up the smart card infrastructure (Active Directory domain accounts, smart card readers, smart card, and so on), configure ESXi to join an Active Directory domain that supports smart card authentication, use the vSphere Client to add root certificates, and follow these steps:



- Step 1.** In the vSphere Client, select the host in the inventory pane and navigate to **Configure > Authentication Services**.
- Step 2.** In the Smart Card Authentication panel, click **Edit**.
- Step 3.** In the dialog box, select the **Certificates** page.
- Step 4.** Add trusted certificate authority (CA) certificates, such as root and intermediary CA certificates, in the PEM format.
- Step 5.** Open the Smart Card Authentication page, select the **Enable Smart Card Authentication** checkbox, and click **OK**.

## Configuring UEFI Secure Boot for ESXi Hosts

Starting with vSphere 6.5, ESXi supports UEFI Secure Boot, which you can enable in the UEFI firmware. With Secure Boot enabled, a machine refuses to load any UEFI driver or app unless the operating system bootloader is cryptographically signed. In vSphere 6.5 and later, the ESXi bootloader contains and uses a VMware public key to verify the signature of the kernel and a small subset of the system that includes a Secure Boot VIB verifier that verifies each VIB packages installed on the host.

### Note

You cannot perform a Secure Boot on ESXi servers that were upgraded by using ESXCLI commands because the upgrade does not update the bootloader.



You can use the following command to run the Secure Boot validation script on an upgraded ESXi host to determine if it supports Secure Boot:

[Click here to view code image](#)

```
/usr/lib/vmware/secureboot/bin/secureBoot.py -c
```

The output is either “Secure boot can be enabled” or “Secure boot CANNOT be enabled.”

To resolve issues with Secure Boot, you can follow these steps:

- Step 1.** Reboot the host with Secure Boot disabled.
- Step 2.** Run the Secure Boot verification script.
- Step 3.** Examine the information in the `/var/log/esxupdate.log` file.

## Securing ESXi Hosts with Trusted Platform Module

ESXi 6.7 can use Trusted Platform Module (TPM) Version 2.0 chips, which are secure cryptoprocessors that enhance host security by providing trust assurance rooted in hardware. A TPM 2.0 chip attests to an ESXi host’s identity. Host attestation is the process of authenticating and attesting to the state of the host’s software at a given point in time. UEFI Secure Boot, which ensures that only signed software is loaded at boot time, is a requirement for successful attestation. The TPM 2.0 chip securely stores measurements of the software modules loaded in the ESXi host, and vCenter Server remotely verifies the measurements. The automated high-level steps of the attestation process are as follows:

- Step 1.** Establish the trustworthiness of the remote TPM chip and create an attestation key (AK) on it.
- Step 2.** Retrieve the attestation report from the host.
- Step 3.** Verify the host’s authenticity.

To use TPM 2.0 chips, you should ensure that your vSphere environment meets these requirements:

- vCenter Server 6.7
- ESXi 6.7 host with TPM 2.0 chip installed and enabled in UEFI
- UEFI Secure Boot enabled

In addition, you should ensure that the TPM chip is configured in the ESXi host's BIOS to use the SHA-256 hashing algorithm and the TIS/FIFO (first-in, first-out) interface and not CRB (Command Response Buffer).

During the boot of an ESXi host with an installed TPM 2.0 chip, vCenter Server monitors the host's attestation status. The vSphere Client displays the hardware trust status in the vCenter Server's Summary tab under Security with the following alarms:

- **Green:** Normal status, indicating full trust
- **Red:** Attestation failed

If the "Host secure boot was disabled" message appears in the vSphere Client, you must re-enable Secure Boot to resolve the problem. If the "No cached identity key loading from DB" message appears, you must disconnect and reconnect the host.

## Securing ESXi Log Files

To increase the security of a host, you can take the following measures:

- Configure persistent logging to a datastore. By default, ESXi logs are stored in an in-memory file system that keeps only 24 hours' worth of data and loses data during host reboot.
- Configure syslog to use remote logging from ESXi hosts to a central host, where you can monitor, search, and analyze logs from all hosts with a single tool.
- Query the syslog configuration to ensure that the syslog server and port are valid.

For details see [Chapter 10, "Managing and Monitoring Clusters and Resources."](#)

## Additional Security Management

Managing vSphere security can involve other tasks, such as those described in this section.

### Key Management Server



In order to use encryption in vSphere, you must be running a key management server (KMS) that has a trust relationship with vCenter Server. To add a KMS to vCenter Server, you can use the following procedure:

- Step 1.** In the vSphere Client, select the vCenter Server in the inventory pane and navigate to **Configuration > Key Management Servers**.
- Step 2.** Click **Add**.
- Step 3.** Provide the server name, server address (FQDN), and server port.
- Step 4.** Optionally, provide other appropriate details, such as proxy details and user credentials.
- Step 5.** If you are adding the first KMS in a cluster, provide a cluster name.
- Step 6.** Click the radio button next to the KMS name.
- Step 7.** In the Make vCenter Trust KMS window, click **TRUST**.
- Step 8.** Click **Make KMS Trust vCenter**.
- Step 9.** Select **KMS Certificate and Private Key** and click **Next**.
- Step 10.** In the next window, next to KMS Certificate, click **Upload File** and open an available certificate PEM file.
- Step 11.** In the same window, next to KMS Private Key, click **Upload File** and open an available certificate PEM file.
- Step 12.** Click the **Establish Trust** button.

## Changing Permission Validation Settings

Periodically, vCenter Server validates its user and group lists against the users and groups in the Windows Active Directory domain. It removes users and groups that no longer exist in the domain. You can change the behavior of this validation by using the vSphere Client to edit the general settings of the vCenter Server and change the Validation and Validation Period options. If you want to disable the validation, deselect the Enable checkbox under Validation. If you want to adjust the frequency at which this validation is performed, enter a value in the Validation Period text box to specify a time, in minutes, between validations.

## Configuring and Managing vSphere Trust Authority (vTA)

With vSphere Trust Authority (vTA), you can do the following.

- Provide a hardware root of trust and remote attestation to ESXi hosts.
- Restrict the release of encryption keys to only attested ESXi hosts.
- Centralize and secure the management of multiple key servers.
- Enhance the level of encryption key management that is used to perform cryptographic operations on virtual machines.

With vTA, you can run workloads in a secure environment where you detect tampering, disallow unauthorized changes, prevent malware, and verify the hardware and software stacks.

When you configure vTA, you enable the Attestation service and the Key Provider service on the ESXi host in the Trust Authority cluster. The Attestation service attests to the state of the trusted ESXi hosts, using a TPM 2.0 chip as the basis for software measurement and reporting. The Attestation service verifies that the software measurement signature can be attributed to a previously configured trusted TPM endorsement key (EK). The Key Provider service removes the need for the vCenter Server and the ESXi hosts to require direct key server credentials. The Key Provider service acts as a gatekeeper for the key servers, releasing keys only to trusted ESXi hosts.

A trusted ESXi host must contain a TPM chip. A TPM chip is manufactured with an EK, which is a public/private key pair that is built into the hardware. You can configure the Attestation service to trust all CA certificates where the manufacturer signed the TPM chip (the EK public key) or to trust the host's TOM CA certificate and EK public key.

### Note

If you want to trust individual ESXi hosts, the TPM chip must include an EK certificate. Some TPM chips do not.

You can use VMware PowerCLI to configure and manage vSphere Trust Authority. Alternatively, you can use vSphere APIs or the vSphere Client for at least some of the activities. To configure vTA, you can perform the following

high-level tasks:

- Step 1.** On a Windows system with access to the vTA environment, install PowerCLI 12.0.0 and Microsoft .NET Framework 4.8 or greater, and create a local folder.
- Step 2.** Add your user account to the TrustedAdmins groups on the vCenter Server managing the Trust Authority cluster and on the vCenter Server of the trusted cluster.
- Step 3.** Enable **Trust Authority State**.
- Step 4.** Collect information about the trusted hosts in the trusted cluster (using **Export-Tpm2CACertificate**).
- Step 5.** Import the trusted host data to the trust authority cluster (**New-TrustAuthorityPrincipal**).
- Step 6.** Create the trusted key provider on the Trust Authority cluster (using **New-TrustAuthorityKeyProvider**).
- Step 7.** Export the Trust Authority cluster information from the Trust Authority cluster (using **Export-TrustAuthorityServicesInfo**).
- Step 8.** Import the Trust Authority cluster data to the trusted cluster (using **Import-TrustAuthorityServicesInfo**).
- Step 9.** Configure the trusted key provider for the trusted hosts on the trusted cluster (using **Register-KeyProvider** and **Set-KeyProvider**).

After configuring vTA, you can perform management operations, including those summarized in [Table 12-5](#).

**Table 12-5** vTA Operations

<b>Operation</b>	<b>Key Steps</b>
Start, stop, and restart vTA services	In the vSphere Client, select the host, navigate to Configure > Services > System, and select Restart, Start, or Stop.
View Trust Authority hosts	In the vSphere Client, select the trusted cluster's vCenter Server and select Configure > Security > Trust Authority.
View vTA cluster state	In the vSphere Client, select the Trust Authority cluster's vCenter Server and select Configure > Trust Authority >

## Trust Authority Cluster.

Restart the Trusted Host service	In an SSH session, enter <b>/etc/init.d/kmxa restart</b> .
Add a Trust Authority host	Use PowerCLI to run <b>Add-TrustAuthorityVMHost</b> .
Add a trusted host	Use PowerCLI to run <b>Add-TrustedVMHost</b> .
Change the master key of a key provider	Use PowerCLI to run <b>Set-TrustAuthorityKeyProvider</b> .

Most of the vTA configuration and state information is stored on the ESXi hosts in the ConfigStore database. Backups of vCenter Server do not include vTA configuration. You can leverage the files that you exported during the configuration of vTA vSphere as your backup. If you need to restore vTA, use the exported files to reconfigure vTA.

## Securing Virtual Machines with Intel Software Guard Extensions (SGX)

You can enable vSGX on a virtual machine on an ESXi host that has compatible CPUs and SGX enabled in the BIOS. The virtual machine must use hardware version 17 or later (with Compatibility set to ESXi 7.0 and Later) and a supported guest OS (Linux, 64-bit Windows 10 or later, or 64-bit Windows Server 2016 or later). To enable vSGX, configure the following hardware settings:

- Select the Security Devices > SGX > Enable checkbox.
- Set VM Options > Boot Options > Firmware to EFI.
- Set the Enter Enclave Page Cache (EPC) size and select Flexible Launch Control (FLC) Mode.

To enable vSGX, the virtual machine must be powered off. You can enable vSGX as you provision a new virtual machine. To remove vSGX from a virtual machine, uncheck the Security Devices > SGX > Enable checkbox.

## Encrypting a Virtual Machine

You can use the following procedure to create a new encrypted virtual machine:

- Step 1.** Establish a trusted connection with the KMS and select a default KMS.
- Step 2.** Create an encryption storage policy or plan to use the bundled sample, VM Encryption Policy.
- Step 3.** Ensure that you have the Cryptographic Operations.Encrypt New privileges.
- Step 4.** If the host encryption mode is not enabled, ensure that you have the Cryptographic Operations.Register Host privilege.
- Step 5.** In the vSphere Client, launch the New Virtual Machine wizard.
- Step 6.** In the wizard, provide the following settings to encrypt the virtual machine:
  - **Compute Resource Settings:** Select a compatible cluster or host. ESXi 6.5 or later is required.
  - **Select Storage:** Select **Encrypt This Virtual Machine**, select the storage policy (from step 2), and select an appropriate datastore.
  - **Virtual Machine Hardware Compatibility:** Select **ESXi 6.5 and Later**.
  - **Customize Hardware Settings:** Optionally, select **VM Options** > **Encryption** and select virtual disks to exclude from encryption.
- Step 7.** Complete the wizard and click **Finish**.

To encrypt an existing virtual machine, you can use the following procedure:

- Step 1.** Establish a trusted connection with the KMS and select a default KMS.
- Step 2.** Create an encryption storage policy or plan to use the bundled sample, VM Encryption Policy.
- Step 3.** Ensure that you have the Cryptographic Operations.Encrypt New privileges.
- Step 4.** If the host encryption mode is not enabled, ensure that you have the Cryptographic Operations.Register Host privilege.

- Step 5.** Ensure the virtual machine is powered off.
- Step 6.** In the vSphere Client, right-click the virtual machine and select **VM Policies > Edit VM Storage Policies**.
- Step 7.** Select the storage policy (from step 2).
- Step 8.** Optionally, select **Configure per Disk** and set encryption as needed for each virtual disk.
- Step 9.** Click **OK**.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15, “Final Preparation,”](#) and the exam simulation questions on the companion website.

## Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 12-6](#) lists these key topics and the page number on which each is found.



**Table 12-6** Key Topics for [Chapter 12](#)

Key Topic Element	Description	Page Number
List	Change the certificate mode	<a href="#">480</a>
Paragraph	Harden the ESXi password	<a href="#">485</a>
List	Modifying an ESXi firewall rule set	<a href="#">492</a>
List	VIB acceptance levels	<a href="#">496</a>
List	Configure smart card authentication	<a href="#">499</a>
List	Adding a KMS to vCenter Server	<a href="#">502</a>



## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

- [vCenter Single Sign-On Security Token Service \(STS\)](#)
- [Certificate Manager](#)
- [managed object browser \(MOB\)](#)
- [Common Information Model \(CIM\)](#)
- [vSphere Installation Bundle \(VIB\)](#)

## Review Questions

1. You want to add a global permission. Which of the following privileges do you need?
  - a. Permissions.Modify Permission privilege on the vCenter root object
  - b. Permissions.Modify Permission privilege on the global root object
  - c. Permissions.Add Permission privilege on the vCenter root object
  - d. Permissions.Add Permission privilege on the global root object
2. A yellow alarm is raised due to a host’s certificate expiration date. Which of the following is a true statement concerning the state of the certificate?
  - a. The certificate is expired.
  - b. The certificate will expire in less than two months.
  - c. The certificate will expire in more than two and in less than six months.

- d.** The certificate will expire in more than two and in less than eight months.
- 3.** You set the Security.PasswordQualityControl parameter to retry=3 min=disabled,disabled,disabled,7,7. With this setting, which of the following statements is true?
  - a.** You cannot use passphrases.
  - b.** Your password can use just a single character class.
  - c.** Your password must include at least two character classes and seven letters.
  - d.** Vmware1 is an acceptable password.
- 4.** You configured an ESXi host with a TPM 2.0 chip and enabled UEFI Secure Boot. During the boot, you get the message “No cached identity key, loading from DB.” What should you do?
  - a.** Reinstall ESXi.
  - b.** Reboot ESXi.
  - c.** Re-enable Secure Boot.
  - d.** Disconnect the host from the vCenter Server and reconnect.
- 5.** You want to have a backup in case you ever need to restore vSphere Trusted Authority. What should you do?
  - a.** Keep a copy of the files that you exported while configuring vTA.
  - b.** In the vSphere Client, choose Backup vTA Configuration.
  - c.** Clone the vCenter Server.
  - d.** Use the vCenter Server File Backup feature.

## Chapter 13

# Managing vSphere and vCenter Server

This chapter covers the following topics:

- [vCenter Server Backup](#)
- [Upgrading to vSphere 7.0](#)
- [Using vSphere Lifecycle Manager](#)
- [Managing ESXi Hosts](#)
- [Monitoring and Managing vCenter Server](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 1.8, 4.4, 4.6, 4.8, 4.11, 4.12, 4.16.1, 5.2, 5.7, 7.10, 7.11, 7.11.1, 7.11.2, 7.11.3, 7.11.4, and 7.11.5.

This chapter covers topics related to managing vCenter Server and vSphere components.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 13-1](#) 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 and Review Questions.”](#)

**Table 13-1** “Do I Know This Already?” Section-to-Question Mapping

---

Foundation Topics Section	Questions
---------------------------	-----------

vCenter Server Backup	1, 2
Upgrading to vSphere 7.0	3, 4
Using vSphere Lifecycle Manager	5–7
Managing ESXi Hosts	8
Monitoring and Managing vCenter Server	9, 10

- 1.** You want to back up your vCenter Server. Which of the following approaches is valid?
  - a.** Use the vSphere Client to perform a file-based backup.
  - b.** Use the vSphere Client to perform an image-based backup.
  - c.** Use the vCenter Server Appliance Management Interface to perform a file-based backup.
  - d.** Use the vCenter Server Appliance Management Interface to perform an image-based backup.
  
- 2.** You want to back up your vCenter Server. Which of the following options is not valid?
  - a.** SCP
  - b.** FTP
  - c.** HTTP
  - d.** HTTPS
  
- 3.** You want to upgrade a vSphere 6.7 environment to vSphere 7.0. Which of the following is the appropriate order?
  - a.** Virtual machines, ESXi hosts, vCenter Server
  - b.** ESXi hosts, vCenter Server, virtual machines
  - c.** vCenter Server, ESXi hosts, virtual machine hardware, VMware Tools
  - d.** vCenter Server, ESXi hosts, VMware Tools, virtual machine hardware
  
- 4.** You plan to upgrade a Windows-based vCenter Server to vCenter Server Appliance 7.0 and want to transfer data in the background.

Which of the following types of data can be transferred in the background following the upgrade?

- a. Configuration data only
  - b. Configuration data and performance data
  - c. Historical and performance data
  - d. Data from the embedded database
5. You are preparing to use Lifecycle Manager. Which of the following is the smallest installable software package (metadata and binary payload) for ESXi?
- a. An update
  - b. An upgrade
  - c. A patch
  - d. A VIB
6. You want to enable Quick Boot for all the hosts in your vSphere cluster. In the vSphere Client, where should you navigate?
- a. Menu > Lifecycle Manager
  - b. Menu > Host and Clusters > DRS Cluster Settings
  - c. Menu > Host and Clusters > HA Cluster Settings
  - d. Menu > Host and Clusters > ESXi Host Settings
7. You want to use Lifecycle Manager to update ESXi firmware. Which of the following is a requirement?
- a. Firmware baselines
  - b. VMware-provided add-on
  - c. Vendor-provided plug-in
  - d. Vendor-provided baselines
8. You want to manage the services running in an ESXi host. Which of the following actions are not available using the vSphere Host Client?
- a. Starting a service
  - b. Stopping a service

- c. Removing a service
  - d. Changing a service's policy
9. You are examining the health state in the vCenter Server Appliance Management Interface (VAMI). What color indicates an alert, where one or more components may be degraded?
- a. Red
  - b. Orange
  - c. Yellow
  - d. Gray
10. You are repointing a vCenter Server to an existing domain. Which of the following is not a valid resolution setting involving conflicts among the settings between the old and new domains?
- a. Delete
  - b. Copy
  - c. Skip
  - d. Merge

## Foundation Topics

### vCenter Server Backup

To provide backup and recovery protection for your vCenter Server, you can use the integrated file-based backup feature. Alternatively, you can perform image-based backups using the vSphere API.

The vCenter Server Appliance Management Interface (VAMI) provides a file-based backup feature for the vCenter Server. If you need to restore the vCenter Server, you can use the vCenter Server installer to deploy a new vCenter Server Appliance and to restore the database and configuration from the file-based backup. You can configure the backup to stream the data to a target using FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB.

When planning a vCenter Server backup, you should consider the following details:

- When a vCenter Server high-availability cluster is involved in a backup, only the primary vCenter Server is backed up.
- Also consider the following protocol considerations:
  - FTP and HTTP are not secure protocols.
  - Backup servers must support at least 10 simultaneous connections for each vCenter Server.
  - You must have read and write permissions on the backup target.
  - FTPS supports only Explicit Mode.
  - HTTP or HTTPS requires WebDAV on the backup web server.
  - To support an HTTP proxy server, you can use only FTP, FTPS, HTTP, or HTTPS.
  - You can use IPv4 or IPv6 URLs for the vCenter Server and file backup, but mixing IP versions between the backup server and the vCenter Server is unsupported.
  - The vCenter Server Appliance Management API is required to restore from NFS- or SMB-based backups.
- After a restore completes, the following configurations are restored:
  - Virtual machine resource settings
  - Resource pool hierarchy and setting
  - Cluster-host membership
  - DRS configuration and rules
- The state of various vSphere components may change following a restore, depending on the changes made since the backup. For example, the following items may be impacted:
  - Storage DRS datastore cluster configuration and membership.
  - Standby Mode for some ESXi hosts where DPM is used.
  - Distributed switch configuration. (Consider exporting the switch configuration prior to restoring vCenter Server.)
  - Content libraries and library items.
  - The registration of virtual machines on ESXi hosts. (Some virtual

machines may be orphaned or missing from the vCenter Server inventory.)

- The host membership of a vSphere HA cluster. (The vCenter Server may be out of sync with the actual current membership.)
- Security patches, which may be missing following a restore. (You should re-apply the missing patches.)

If you have prepared a supported target server, you can use the following procedure to schedule a file-based backup of the vCenter Server:

**Step 1.** Log on to the VAMI (<https://vCenterFQDN:5480>) as root.

**Step 2.** Click **Backup > Configure**.

**Step 3.** Enter the backup location details:

- **Backup Location:** Provide the protocol, port, server address, and folder.
- **Backup Server Credentials:** Provide the username and password with write privileges.

**Step 4.** Configure the schedule and time.

**Step 5.** Optionally, provide an encryption password.

**Step 6.** Provide a number of backups to retain or select **Retain All Backups**.

**Step 7.** Optionally, select **Stats, Events, and Tasks** to back up historical data.

**Step 8.** Click **Create**.

You can use the VAMI to manually back up a vCenter Server by selecting **Backup > Backup Now**.



To restore a vCenter Server, launch the vCenter Server installer (described in [Chapter 8, “vSphere Installation”](#)) on your desktop (Windows, Linux, or Mac) and use the following procedure:

**Step 1.** On the Home page, click **Restore**.

**Step 2.** On the next page, click **Next**.



- Step 3.** Accept the license agreement and click **Next**.
- Step 4.** Provide the backup location and credentials for the backup file to be restored and click **Next**.
- Step 5.** Review the backup information and click **Next**.
- Step 6.** Continue using the wizard to provide connection details (FQDN, credentials, and certificate information) for the ESXi host or vCenter Server to which the appliance will be restored.
- Step 7.** When prompted in the wizard, provide a name and a root password for the vCenter Server Appliance.
- Step 8.** Select a deployment size (from Tiny to X-Large).
- Step 9.** Select the storage size (from Default to X-Large).
- Step 10.** Select a datastore and provide virtual disk and network settings for the appliance.
- Step 11.** On the Ready to Complete Stage 1 page, click **Finish**.
- Step 12.** When the OVA deployment finishes, click **Continue** to proceed to Stage 2.
- Step 13.** Continue navigating the wizard, and when it prompts you for Single Sign-On credentials, provide the credentials and click **Validate and Recover**.
- Step 14.** On the Ready to Complete page, review the details, click **Finish**, and click **OK**.

### **Note**

If a restore fails, power off and delete the partially restored VM. Then try to restore the VM again.

### **Note**

You must power off the active, passive, and witness nodes in a vCenter Server high-availability cluster prior to restoring. You must reconstruct the

cluster after a restore operation completes successfully.

If you prefer to use an image-based backup, you can leverage the vSphere API. For image-based backups, you should consider the following:

- You must ensure that the vCenter Server uses a fully qualified domain name (FQDN) with correct DNS resolution or configure its hostname to be an IP address.
- If DHCP is used, you must configure the restored vCenter Server's IP address back to the original value.
- Ensure that all vSphere component clocks are synchronized.
- The set of restored configurations for image-based restoration is identical to the file-based restoration.
- The impact to the state of vSphere components following an image-based restoration is nearly identical to the impact of a file-based restoration.

To perform an image-based backup or recovery of vCenter Server, you must use a third-party product or custom code.

## Upgrading to vSphere 7.0

This section provides details on upgrading a vSphere environment to vSphere 7.0. For information on installing a new vSphere 7.0 environment, see [Chapter 8](#).

To upgrade a vSphere 6.5 or 6.7 environment to vSphere 7.0, you should upgrade the major components in the following order:

1. vCenter Server
2. ESXi hosts
3. VMware Tools on the virtual machines
4. Virtual machine hardware

### Note

For vCenter Server 6.0 and earlier, you should upgrade to vSphere 6.5 or 6.7 and then upgrade to vSphere 7.0.

You should back up vCenter Server prior to upgrading it. For details, see the section “[vCenter Server Backup](#),” later in this chapter.

Upgrading your environment to use vCenter Server 7.0 requires you to either upgrade an existing vCenter Server Appliance or migrate from an existing Windows-based vCenter Server. When you upgrade or migrate a vCenter Server that uses an external Platform Services Controller (PSC), you converge the PSC into a vCenter Server Appliance.

Prior to upgrading to vCenter Server 7.0, you should consider its compatibility with other vSphere components, as summarized in [Table 13-2](#).



**Table 13-2** vCenter Server 7.0 Compatibility

---

**Component Compatibility**

---

vCenter Server Appliance	You can upgrade vCenter Server Appliance 6.5 and 6.7 to 7.0, except with specific build combinations that violate the back-in-time restrictions identified at <a href="https://kb.vmware.com/s/article/67077">https://kb.vmware.com/s/article/67077</a> .
vCenter Server for Windows	You can migrate vCenter Server for Windows 6.5 and 6.7 (with or without an embedded PSC) to a vCenter Server Appliance 7.0.
vCenter Server database	vCenter Server 7.0 uses PostgreSQL for the embedded database. It does not support external databases.
ESXi hosts	vCenter Server 7.0 requires ESXi host Version 6.5 or later.
Host profiles	vCenter Server 7.0 requires host profiles Version 6.0 or later.
VMFS	vCenter Server 7.0 supports VMFS 3 and later but can only create VMFS 5 and VMFS 6 datastores.
Virtual machines and	Review the ESXi upgrade documentation for specific upgrade options, which are dependent on the current versions.

## VMware Tools

---

Auto Deploy	If you currently use Auto Deploy, when you upgrade to vCenter Server 7.0, VMware recommends that you use Auto Deploy to upgrade hosts to ESXi 7.0.
-------------	--

---

vSphere Distributed Switch (vDS)	Upgrade to vDS 6.0 before upgrading vCenter Server.
----------------------------------	---

---

Network I/O Control (NIOC)	Upgrade to NIOC Version 3 before upgrading vCenter Server.
----------------------------	--

---

vSAN	VMware recommends that you synchronize versions of vCenter Server and ESXi to avoid potential faults.
------	---

---

vSAN disk version	Supported versions and paths may be impacted by the current version and upgrade history. See <a href="https://kb.vmware.com/articleview?docid=2145267">https://kb.vmware.com/articleview?docid=2145267</a> .
-------------------	--

---

Legacy Fault Tolerance (FT)	If you use Legacy FT on any virtual machines, you must turn off or upgrade the Legacy FT feature prior to vCenter Server upgrade or migration.
-----------------------------	--

---

## vCenter Server Data Transfer

If you migrate a Windows-based vCenter Server or upgrade a vCenter Server with an external PSC, you need to transfer data to the embedded PostgreSQL database in the target vCenter Server Appliance. At a minimum, you must transfer configuration data. You can choose whether you want to transfer historical data and performance metrics data. Specifically, you can choose one of the following options:

- **Configuration data:** Transferring just configuration data minimizes downtime during the upgrade.
- **Configuration and historical data:** You can choose to transfer historical data (usage statistics, tasks, and events) during an upgrade (impacting the

downtime) or in the background following the upgrade.

- **Configuration, historical, and performance data:** You can transfer the configuration data during the upgrade and transfer the remaining data in the background following the upgrade.

## Note

The option to transfer data in the background following an upgrade is applicable only in scenarios where the source vCenter Server uses an external database.

You can monitor the background data transfer by using the VAMI. You can pause and cancel the data transfer.

## Upgrading vCenter Server Appliance

You should address the following prerequisites prior to upgrading a vCenter Server Appliance to Version 7.0:

### Key Topic

- Ensure that the clocks of all the vSphere components are synchronized.
- Ensure that the system has the minimum hardware and software components.
- Ensure that the target ESXi host is not in Lockdown, Maintenance, or Standby Mode.
- Ensure that the target ESXi host is not part of a fully automated DRS cluster.
- Verify that port 22 is open on the source vCenter Server Appliance and that port 443 is open on the ESXi host on which the source vCenter Server Appliance is running.
- Verify that the source appliance has sufficient free space to accommodate data that is used for the upgrade.
- If the source vCenter Server uses an external database, determine its size and ensure that you account for it in the size of the new appliance.

- Ensure that network connectivity exists between the vCenter Server or ESXi that hosts the source vCenter Server Appliance and the new vCenter Server Appliance.
- If you plan to set the system name to a FQDN, ensure that forward and reverse DNS records are created.

Upgrading a vCenter Server Appliance is a two-stage process. The first stage is to deploy the OVA. The second phase is to transfer the data and configure the vCenter Server Appliance. For a vCenter Server with an external PSC, you can use the following procedure for the first stage:

**Step 1.** Launch the vCenter Server (GUI) installer and select **Upgrade**.

**Step 2.** Review the upgrade process on the first wizard page and click **Next**.

**Step 3.** Accept the license agreement and click **Next**.

**Step 4.** Provide the following information for the source vCenter Server:

- Provide the address, HTTPS port, SSO credentials, and root password for the source vCenter Server.
- Provide the address, HTTPS port, and credentials for a user with administrative privileges for the ESXi host (or vCenter Server) that is hosting the source vCenter Server.
- Click **Connect**.

**Step 5.** Follow the wizard prompts to accept the certificate and accept the plan to converge the source vCenter Server and external PSC into a single vCenter Server Appliance.

**Step 6.** Follow the wizard prompts to provide the following information for the target environment that will host the new vCenter Server Appliance.

- If you are connecting to a vCenter Server, provide the address, HTTPS port, SSO credentials, and root password for the vCenter Server. Select a data center and an ESXi host (or cluster).
- If you are connecting to an ESXi host, provide the address, HTTPS port, and credentials for a user with administrative privileges for the ESXi host.

**Step 7.** Follow the wizard to configure the new vCenter Server Appliance with

the following information:

- Virtual machine name
- Root user password
- Deployment size (Tiny to X-Large, as described in [Table 1-10](#) in [Chapter 1, “vSphere Overview, Components, and Requirements”](#))
- Storage size (defaults to X-Large, as described in [Table 1-11](#))
- Datastore
- Temporary network used to transfer data from the source vCenter Server to the new vCenter Server

**Step 8.** Click **Finish**.

**Step 9.** Click **Continue** to proceed to Stage 2.

### Note

The identical Stage 1 procedure can be used when upgrading a vCenter Server Appliance with an embedded PSC, except the wizard does not prompt you to accept the plan to converge an external PSC.

For a vCenter Server with an external PSC, you can use the following procedure for Stage 2:

**Step 1.** Review the Stage 2 details and click **Next**.

**Step 2.** Wait for the pre-upgrade check to finish and respond to any of the following messages:

- **Errors:** Read the message, click **Logs** to obtain a support bundle, and troubleshoot. You cannot proceed with the upgrade until errors are corrected.
- **Warnings:** Read the messages and click **Close**.

**Step 3.** Specify the replication technology by choosing one of the following options:

- **This Is the First vCenter Server in the Topology That I Want to**

## Converge

- **This Is a Subsequent vCenter Server** (and also provide the IP address and HTTPS port of the partner vCenter Server)

**Step 4.** On the Select Upgrade Data page, choose the type of data transfer, as described in the section “vCenter Server Data Transfer,” earlier in this chapter.

**Step 5.** Complete the wizard and wait for the transfer and setup operations to complete.

**Step 6.** Decommission the source external PSC.

### Note

The Stage 2 procedure to upgrade a vCenter Server Appliance with an embedded PSC is similar, but instead of being prompted for the replication technology, you are prompted (again) for information for connecting to the source vCenter Server and host environment (that is, the vCenter Server or ESXi host where the source vCenter Server resides).

## Migrating vCenter Server for Windows to vCenter Server Appliance



Migrating a vCenter Server for Windows to vCenter Server Appliance includes the following unique requirements:

- If the vCenter Server service is running as a user other than the Local System account, ensure that the account is a member of the Administrators group and has the Log On as a Service, Act as Part of the Operating System, and Replace a Process Level Token permissions.
- Verify that the vCenter Server and PSC certificates are valid and reachable.
- Verify that the network connection to the domain controller is functioning.
- Verify that neither the source vCenter Server nor the PSC instance is using a DHCP IP address as the system name.



When you migrate vCenter Server for Windows to vCenter Server Appliance, the installer performs an environment precheck that includes the following items:

- Sufficient storage space in the source server
- Validity and compatibility of SSL certificates and system names
- Network connectivity, ports, and DNS resolution
- Database connectivity
- Proper credentials and privileges for the Single Sign-On and Windows administrator accounts
- NTP server validation

The following limitations apply when you migrate vCenter Server for Windows to vCenter Server Appliance 7.0:

- Local Windows OS users and groups are not migrated to the guest OS (Photon OS) of the new appliance. You should remove any vCenter Server permissions to local Windows users prior to the migration.
- At the end of the migration, the source vCenter Server is turned off, and any solutions that are not migrated become unavailable. You should leave the source vCenter Server powered off to avoid network ID conflicts with the target vCenter Server Appliance.
- The migration of Windows-based vCenter Server instances that use custom ports for services other than Auto Deploy, Update Manager, vSphere ESXi Dump Collector, or HTTP reverse proxy is not supported.
- Only one network adapter setting is migrated to the target vCenter Server Appliance. If the source uses multiple IP addresses, you can select which IP address and network adapter settings to migrate.

To migrate a Windows-based vCenter Server with an embedded PSC to vCenter 7.0, you can use the following procedure:

- Step 1.** Run the VMware Migration Assistant on the Windows server that runs vCenter Server and follow these steps:
  - a. Download and mount the vCenter Server installer.

- b. Log on to Windows as an administrator.
- c. Double-click on **VMware-Migration-Assistant.exe**.
- d. Leave the Migration Assistant running until the migration is complete.

**Step 2.** Migrate the vCenter Server instance to an appliance by following these steps:

- a. Using the Stage 1 procedure for upgrading a vCenter Server Appliance as a guide, launch the vCenter Server GUI installer but choose **Migrate** rather than Upgrade.
- b. Provide connection information for the source Windows-based vCenter Server.
- c. Provide information on the target server (vCenter Server or ESXi host) to deploy the vCenter Server Appliance.
- d. Provide appliance information, such as root password, compute deployment size, storage deployment size, and datastore.
- e. Configure the temporary network used to transfer data from the source vCenter Server to the new vCenter Server and select **Continue to Stage 2**.
- f. In Stage 2, provide credentials for the Single Sign-On administrator and for the Windows system. If the Windows system is connected to an Active Directory domain, provide credentials for an appropriate domain user.
- g. Complete the wizard and click **Finish** on the Ready to Complete page.
- h. Click **OK** to confirm the shutdown of the source Center Server.
- i. Monitor the data transfer and configuration process.

## Note

If the Windows-based vCenter Server uses an external Update Manager, run the Migration Assistant on the Update Manager machine before running it on the vCenter Server.

To migrate a Windows-based vCenter Server to vCenter 7.0 with an external PSC, you can use the previous procedure as a guide, but you should run the VMware Migration Assistant in the Windows-based PSC prior to running it in the vCenter Server. You should also decommission the source external PSC following the migration.

## Upgrading ESXi and Virtual Machines

After upgrading to vCenter Server 7.0, you can use Lifecycle Manager to upgrade ESXi hosts and virtual machines, as described in the section “Using vSphere Lifecycle Manager,” later in this chapter.

## Using Update Planner

You can use the Update Planner to examine available vCenter Server updates and upgrades. You can produce interoperability reports for associated VMware products with your source (current) and target vCenter Server versions. You can generate pre-update reports to help ensure that your system meets the minimum software and hardware requirements. The report identifies potential upgrade issues and provides potential remedy actions. To use Update Planner, you must join the VMware Customer Experience Improvement Program (CEIP).

You can use the following procedure to perform an interoperability check of VMware products in your environment against the current vCenter Server version:

- Step 1.** In the vSphere Client, select a vCenter Server in the inventory pane and navigate to **Monitor > Interoperability**.
- Step 2.** Review the Product Interoperability report, which should contain all the available products in your environment.
- Step 3.** If a VMware product in your environment is not automatically detected, you can use the following steps to manually add the product to the list and regenerate the Product Interoperability report:
  - a.** For each missing product, click **Add Product** and select the VMware product and version.
  - b.** Click **Done**.
  - c.** Regenerate the Product Interoperability report and review the product list.

**Step 4.** Click **Export** to save the report as a comma-separated values (CSV) file.

You can use the following steps to create an interoperability report on the compatibility of your environment's VMware products against a target version of vCenter Server:

**Step 1.** In the vSphere Client, select a vCenter Server in the inventory pane and navigate to **Updates > Update Planner**.

**Step 2.** Select a target vCenter Server version (major upgrade or minor update).

**Step 3.** Click **Generate Report > Interoperability**.

**Step 4.** Review the Product Interoperability report, which should contain all the available products in your environment.

**Step 5.** If a VMware product in your environment is not automatically listed or is undetected, you can use the following steps to manually add the product and regenerate the Product Interoperability report:

- a. For each missing product, click **Add Product** and select the VMware product and version.
- b. Click **Done**.
- c. Regenerate the Product Interoperability report and review the product list.

**Step 6.** Click **Export** to save the report as a comma-separated values (CSV) file.

A report that provides pre-update information identifies potential problems that might prevent the completion of a software upgrade or update. It includes a list with actions that you must address to ensure a successful upgrade of vCenter Server. You can use the following steps to run prechecks and generate reports providing pre-update information:

**Step 1.** In the vSphere Client, select a vCenter Server in the inventory pane and navigate to **Updates > Update Planner**.

**Step 2.** Select a target vCenter Server version (major upgrade or minor update).

**Step 3.** Click **Generate Report > Pre-Update Checks**.

**Step 4.** Click **Export** to save the report as a comma-separated values (CSV) file.

**Step 5.** Optionally, click Open Appliance Management or Download ISO.

After you address issues identified in the report, you can use the VAMI to perform administrative tasks to apply patches and updates.

## Using vSphere Lifecycle Manager

In vSphere 7.0, you have choices for methods and tools to facilitate the deployment and lifecycle management of ESXi host. You can use an ESXi installer image or VMware vSphere Auto Deploy to deploy hosts. Using vSphere Auto Deploy can result in hosts deployed in Stateless Mode. Using an ESXi installer image results in hosts being deployed in Stateful Mode. Depending on the deployment method, you can use a variety of tools and methods for host updates and upgrade. For example, you can use Update Manager baselines, the ESXi Image Builder CLI, ESXCLI, or vSphere Auto Deploy. In any case, the ESXi image may change at runtime due to some solution installing software automatically or a service changing a setting.

*VMware vSphere Lifecycle Manager* provides simple, centralized lifecycle management for ESXi hosts and clusters by using images and baselines. Specifically, lifecycle management of a vSphere cluster refers to tasks such as installing and updating host firmware and ESXi. In vSphere 7.0, vSphere Lifecycle Manager encompasses and enhances the functionality that Update Manager provided for earlier vSphere releases. It is a service that runs in vCenter Server and is automatically enabled in the vSphere Client. It can work in an environment with direct access to the Internet or access via a proxy server. It can also work with Update Manager Download Service (UMDS) in a secured network with no access to the Internet. In such cases, you use the UMDS to download updates to the vSphere Lifecycle Manager depot, or you import them manually.

[Chapter 8](#) provides information on the implementation of Lifecycle Manager.

In vSphere 7.0, vSphere Lifecycle Manager enables you to use images or baselines. An *image* represents a desired software specification to be applied to all hosts in a cluster. An image is a description of which software, drivers, and firmware to run on a host. You can apply a single image to all hosts in a cluster to ensure consistency. Updates to software and firmware occur in a single workflow. The use of images with vSphere Lifecycle Manager enables new

functionalities, including image recommendations, automated firmware updates, and hardware compatibility checks.

You can use baselines in vSphere 7.0, much as you could in previous vSphere versions, to perform the following tasks:

- Upgrade ESXi 6.5 and 6.7 hosts
- Patch ESXi 6.5, 6.7, and 7.0 hosts
- Install and update third-party software on ESXi hosts

Starting with vSphere 7.0, you can use vSphere Lifecycle Manager images to perform the following tasks on a set of hosts at the cluster level:

- Install a desired ESXi version on each host
- Install and update third-party software on each ESXi host
- Update the firmware of each ESXi host
- Update and upgrade each ESXi host in a cluster
- Check the hardware compatibility of each host against hardware compatibility lists, such as the *VMware Compatibility Guide* and the *vSAN Hardware Compatibility List*

You can start using Lifecycle Manager images as you create a cluster. Otherwise, you can switch from using baselines to images later. After switching a cluster to use images, you cannot revert the cluster back to using baselines. However, you can move the hosts to another cluster that uses baselines. If you set up an image for a cluster and remediate all the hosts in the cluster, then all standalone VIB and non-integrated agents are deleted from the hosts.

vSphere Lifecycle Manager has several components, including a service named vSphere Lifecycle Manager that runs in vCenter Server and uses the embedded vCenter Server PostgreSQL database. It communicates with agents running in each ESXi host.

vSphere Lifecycle Manager uses a desired-state model based on images that represent both the target software and target configuration of the host. To use images, all hosts must be ESXi 7.0 or later, they must be stateful, and they must be from the same hardware vendor.

## Note

In vSphere 7.0, vSphere Lifecycle Manager images are not supported for clusters with Kubernetes enabled or NSX installed. You can manage such clusters using baselines and baseline groups. You cannot enable vSphere with Kubernetes or install NSX-T on a cluster that is managed with an image.

You can leverage vSphere Lifecycle Manager for VMware Tools and virtual machine hardware upgrade operations on virtual machines running on ESXi 6.5, ESXi 6.7, and ESXi 7.0 hosts.

To get started using vSphere Lifecycle Manager, in the vSphere Client, you can navigate to **Menu > Lifecycle Manager** (which is called the Lifecycle Manager home view) and select a vCenter Server. Here you can configure Lifecycle Manager by using the **Settings** tab. [Table 13-3](#) describes the available settings for Lifecycle Manager remediation.



**Table 13-3** Lifecycle Manager Remediation Settings

Setting	Description
Quick Boot	You can enable Quick Boot to skip the hardware reboot during remediation.
Cluster Settings	You can disable vSphere Distributed Power Management (DPM), HA admission control, and FT for an entire cluster during remediation.
VM Power State	You can power off, suspend, or leave alone the power state of running VMs during remediation.
VM Migration	You can migrate or leave alone powered-off and suspended VMs during remediation.
Maintenance Mode Failures	You can control the number of retries and the delay between retries in the case that a host fails to enter Maintenance Mode.
PXE	You can allow the installation of software for solutions on the

Booted Hosts PXE-booted ESXi hosts.

---

Removable Media Devices You can automatically disconnect all virtual machine removable media devices prior to remediation and automatically reconnect afterward.

---

When working with images, the following settings are applicable: Quick Boot, VM Power State, VM Migration, Maintenance Mode Failures, HA Admission Control, and DPM.

When working with baselines, the following settings are applicable: Quick Boot, VM Power State, VM Migration, Maintenance Mode Failures, PXE Booted Hosts, and Removable Media Devices.

You can perform the following tasks from the Lifecycle Manager home view:

- Browse the vSphere Lifecycle Manager depot.
- Trigger the synchronization of updates with the configured online depots.
- Trigger the synchronization of hardware compatibility data.
- Import offline depots manually.
- Import ISO images to use for the creation of upgrade baselines.
- Create and manage baselines and baseline groups.
- Configure the default vSphere Lifecycle Manager download source.
- Add a URL to an online depot to the list of download sources.
- Enable or disable downloading from a download source.
- Configure host remediation settings.
- Configure virtual machine rollback settings

## About VMware Update Manager

In vSphere 7.0, VMware Update Manager (VUM) has been rebranded as vSphere Lifecycle Manager, which includes new features, such as the ability to provide ESXi lifecycle management at a cluster level.

## VMware Update Manager Download Service (UMDS)



**VMware Update Manager Download Service (UMDS)** is an optional module of vSphere Lifecycle Manager whose primary function is to download data when Lifecycle Manager does not have Internet connectivity. You can configure a web server on UMDS to automatically export the downloads to make them available to Lifecycle Manager. Alternatively, you can export the data from UMDS to a portable media drive. UMDS and Lifecycle Manager must be of the same version and update release.

UMDS is a 64-bit Linux application that is bundled with the vCenter Server Appliance 7.0. You can use that bundle to install UMDS on a separate Linux system. You cannot upgrade UMDS on a Linux system. Instead, you can uninstall UMDS, perform a fresh installation, and continue using an existing patch store. To install UMDS, you can use the following procedure:

- Step 1.** Log on to a supported Linux system, such as Ubuntu (14.04, 18.04, 18.04 LTS, or 20.04 LTS) or Red Hat Enterprise Linux (74, 7.5, 7.7, or 8.1).
- Step 2.** For Red Hat Enterprise 8.1, install the libnsl package Version 2.28 or later.
- Step 3.** Verify that you have administrative privileges.
- Step 4.** Open a command shell.
- Step 5.** Copy VMware-UMDS-7.0.0-build\_number.tar.gz to the Linux server.
- Step 6.** Run the **vmware-install.pl** script.
- Step 7.** When prompted, accept the EULA, select an installation directory, provide proxy settings, and specify the directory for storing patches.

To connect UMDS to third-party vendor websites, you can use the following command:

[Click here to view code image](#)

```
vmware-umds -S --add-url https://web1.vendor1.com/index.html  
/index.xml --url-type HOST
```

To export data from UMDS to a specific location that serves as a shared repository for vSphere Lifecycle Manager, you can use the following command, where *repositoryPath* represents a valid path to the shared repository:

[Click here to view code image](#)

```
vmware-umds -E -export-store repositoryPath
```

## Baselines and Images

vSphere Lifecycle Manager supports the use of baselines and baseline groups that are available in previous vSphere releases for host patching and upgrade operations. It supports multiple types of baselines, including predefined baselines, recommendation baselines, extension baselines, and custom baselines.

Predefined baselines cannot be edited or deleted. To examine the predefined baselines in the vSphere Client, select Lifecycle Manager > Baselines. The predefined baselines are categorized as host security patches, critical host patches, and non-critical host patches.

Recommendation baselines, which are predefined baselines generated by vSAN, appear by default when for vSAN cluster members. You can use recommendation baselines to update a cluster with recommended critical patches, drivers, updates, and supported ESXi version for vSAN. You can create a baseline group containing multiple recommendation baselines, but only if you do not include other baseline types in the group. The vSAN recommendation baselines are typically refreshed every 24 hours for Lifecycle Manager instances with Internet access.

Custom baselines are baselines that you create. Patch baselines can be dynamic or fixed. With dynamic baselines, you specify the criteria for patches to be automatically included in the baseline. You can manually include or exclude patches from dynamic baselines. With fixed baselines, you manually select the patches to include.

If you have the VMware vSphere Lifecycle Manager.Manage Baselines privilege, you can use the following procedure to create a dynamic baseline:

### Key Topic

- Step 1.** In the vSphere Client, select **Menu > Lifecycle Manager**.
- Step 2.** If necessary, select the vCenter Server system in the **Lifecycle Manager** menu.
- Step 3.** Select **Baselines > New > Baseline**.
- Step 4.** In the wizard, provide the following baseline information:
  - a.** Name and (optionally) description

**b. Type: Upgrade, Patch, or Extension**

Click **Next**.

**Step 5.** On the Select Patches Automatically page, set the criteria for adding patches to the baseline:

- a.** Select the **Automatic Update** checkbox.
- b.** On the Criteria tab, specify the values for the following options to restrict which patches are included in the baseline:

- **Patch Vendor**
- Product (You can use an asterisk at the end to allow any version number.)
- Severity
- Category
- Release Date (You can specify a date range.)

Click **Next**.

- c.** On the Matched tab, optionally deselect patches that matched the criteria but that you want to permanently exclude.
- d.** On the Excluded and Select tabs, view which patches are excluded and which patches are included.

**Step 6.** On the Select Patches Manually page, optionally select specific patches from the set of patches that do not meet the criteria for automatic inclusion to include in the baseline.

**Step 7.** On the Summary page, click **Finish**

Extensions baselines contain additional (VMware or third-party) software modules for hosts, such as device drivers. You can install additional modules by using extension baselines and update the modules using patch baselines.

## Note

In vSphere 7.0, the vendor name of VMware for inbox components has changed from VMware, Inc to VMware. If you filter the components by VMware, the results contain both VMware, Inc for 6.x patches and VMware for 7.0 patches.

If a user has the View Compliance Status privilege, you can use the Updates tab for a selected object to view the object's compliance with baselines or images. You can select a host or cluster that is managed with baselines and click on Updates > Baselines. From there, you can do the following tasks:

- Check the compliance of hosts or clusters against baselines and baseline groups.
- Attach and detach baselines and baseline groups to hosts or clusters.
- Perform a remediation precheck.
- Stage patches or extensions to prepare for remediation.
- Check the compliance of ESXi hosts against an image.
- Remediate hosts against baselines and baseline groups.
- Remediate hosts that are part of a vSAN cluster against system-managed baselines.

You can select a cluster that is managed with an image and click on Updates > Images. From there, you can do the following tasks:

- Export, import, and edit the image used by the cluster.
- Upgrade the firmware of the ESXi hosts in the cluster.
- Check for and examine recommended images for the cluster.
- Check for hardware compatibility for a selected ESXi version against the vSAN HCL.
- Check the compliance of the ESXi hosts against the image.
- Run a remediation precheck.
- Remediate the ESXi hosts against the cluster's image.

You can select a host and then select Updates > Hosts > Hardware Compatibility to check the host hardware against the *VMware Compatibility Guide*. You can select a host and then select Updates > Hosts and then select VMware Tools or VM Hardware to check and upgrade the VMware Tools version and virtual hardware version of the virtual machines.

Table 13-4 provides definitions of vSphere Lifecycle Manager terms.

**Table 13-4** Lifecycle Manager Definitions

<b>Term</b>	<b>Definition</b>
Update	A software release that makes small changes to the current version, such as vSphere 7.0 Update 1, 7.0 Update 2, and so on.
Upgrade	A software release that introduces major changes to the software. For example, you can upgrade from vSphere 6.5 to 6.7 and 7.0.
Patch	A small software update that provides bug fixes or enhancements to the current version of the software, such as 7.0a, 7.0 Update 1a, and so on.
VIB (vSphere Installation Bundle)	The smallest installable software package (metadata and binary payload) for ESXi.
VIB metadata	An XML file that describes the contents of the VIB, including dependency information, textual descriptions, system requirements, and information about bulletins.
Standalone VIB	A VIB that is not included in a component.
Depot	The hosted version of updates provided by VMware, OEMs, and third-party software vendors, containing the metadata and the actual VIBs.
Offline bundle/offline depot	An archive (ZIP file) that contains VIBs and metadata that you use for offline patching and updates. A single offline bundle might contain multiple base images, vendor add-ons, or components.
OEM (original equipment manufacturer)	A VMware partner, such as Dell, HPE, or VMware Cloud on AWS.

Third-party software provider      A provider of I/O filters, device drivers, CIM modules, and so on.

---

## Note

During the synchronization of a depot, vSphere Lifecycle Manager downloads only the VIB metadata.

In earlier vSphere releases, VIBs are packaged into bulletins. In vSphere 7.0, VIBs are packaged into components, which are created by VMware, OEMs, and third-party software providers. A component is much like a bulletin with extra metadata containing the component name and version. VMware bundles components together into fully functional ESXi images. OEMs bundle components into add-ons that are delivered via the VMware online depot as offline bundles. Third-party vendors create and ship drivers packaged as components.

vSphere Lifecycle Manager can consume bulletins and components. It lists the available components as bulletins when baselines are used to manage a host or cluster. When images are used, vSphere Lifecycle Manager works only with components.

The **ESXi base image**, which is the ESXi image that VMware provides with each release of ESXi, is a complete set of components that can boot up a server. A base image has a friendly name, has a unique version that corresponds to the ESXi release, and is hosted in the VMware online depot. Alternatively, you can download an ESXi installer ISO file and an offline bundle (ZIP file) that contains the ESXi version from [my.vmware.com](https://my.vmware.com).

A vendor **add-on** is a collection of components that you can use to customize an ESXi image with OEM content and drivers. You cannot use vendor add-ons on their own. You can use a vendor add-on to add, update, or remove components that are part of the ESXi base image. You can use the vSphere Client to view the list of components that a vendor add-on adds to or removes from an ESXi base image.

Prior to vSphere 7.0, OEMs create custom images by merging their content with the stock VMware-provided image. OEMs release custom images in accordance

with the major and update releases of vSphere. Starting with vSphere 7.0, in addition to releasing custom ISO images and offline bundles, OEMs can release ZIP files that contain only the vendor add-on. This approach decouples the release cycle of OEMs from the release cycle of VMware.

vSphere Lifecycle Manager can consume software updates delivered as an online depot, as an offline depot, or as an installable ISO image. An online depot is a hosted version of the software updates. Starting with vSphere 7.0, the default, the VMware online depot provides vendor add-ons to hosts. The default depot also contains VMware-certified, ESXi-compatible I/O device drivers. You can use the vSphere Client to access third-party online depots containing additional components.

Offline bundles are ZIP files that contain the software metadata and the respective VIBs. Starting with vSphere 7.0, an OEM can distribute an add-on ZIP files that contains the delta between the OEM custom image and the base image that VMware provides.

A **baseline** is a set of bulletins. Patch baselines, extension baselines, and upgrade baselines contain patch bulletins, extension bulletins, and ESXi images, respectively. You can attach baselines to hosts, check compliance of a host with its associated baseline, and remediate (update) hosts by using the baseline.

You can classify baselines based on the following:

- Update type (such as patch baselines, extension baselines, and upgrade baselines)
- Content (such as fixed or dynamic)
- Predefined, recommendation, or custom baselines
- Predefined host patches (such as host security patches, critical host patches, and non-critical host patches)

You cannot modify or delete the predefined baselines. You can use the predefined baselines to create custom patch, extension, and upgrade baselines. Recommendation baselines are baselines generated automatically by vSAN. You can use recommendation baselines only with vSAN clusters.

A baseline group is a set of non-conflicting baselines that you can apply as a single entity. A host baseline group can contain a single upgrade baseline plus patch and extension baselines. For efficiency, you can attach and apply baselines

and baseline groups to container objects (such as folders, vApps, and clusters) rather than to the individual underlying objects (virtual machines and hosts).

To check a cluster's compliance against an image, you can select the cluster, select Updates > Hosts > Image, and click the Check Compliance button. When you check a cluster's compliance with a Lifecycle Manager image, one of the following four compliance states is identified for each member host:

- **Compliant:** The host's image matches the image applied to the cluster.
- **Non-Compliant:** The host's image does not match the image applied to the cluster. Some potential causes are differences in the ESXi version, the firmware version, or the set of components. Another potential cause is that the host contains a standalone VIB.
- **Incompatible:** The cluster's image cannot be applied to the host. Some potential reasons are that the host's ESXi version is later than the version in the image, the host has insufficient RAM, or the host's hardware is incompatible with the cluster's image.
- **Unknown:** No compliance state information is available, perhaps because the host was recently added or the cluster's image was edited.

## ESXi Quick Boot

You can use Quick Boot with vSphere Lifecycle Manager to optimize the host patching and upgrade operations. When Quick Boot is enabled, vSphere Lifecycle Manager skips the hardware reboot (the BIOS or UEFI firmware reboot), which reduces the time a host spends in Maintenance Mode and reduces the risk of failures.

Starting with vSphere 6.7, Quick Boot is supported on a limited set of hardware platforms and drivers. Quick Boot is not supported on ESXi hosts that use TPM or on ESXi 6.7 hosts that use passthrough devices. To determine whether your system is compatible with Quick Boot, you can run the following command in the ESXi Shell:

[Click here to view code image](#)

```
/usr/lib/vmware/loadesx/bin/loadESXCheckCompat.py
```

To enable Quick Boot, you can use the following procedure:





- Step 1.** In the vSphere Client, navigate to **Menu > Lifecycle Manager**.
- Step 2.** From the Lifecycle Manager drop-down menu, select a vCenter Server.
- Step 3.** Select **Settings > Host Remediation > Images**.
- Step 4.** Click **Edit**.
- Step 5.** In the Edit Cluster Settings window, select the **Quick Boot** checkbox.
- Step 6.** Click **Save**.

## ESXi Firmware Updates

You can use Lifecycle Manager images to perform firmware updates on cluster member hosts. Firmware updates are not available for clusters that are managed with baselines. To apply firmware updates, you must use a vendor-provided firmware and driver add-on in the image. Firmware updates are available in a special vendor depot that you access through a vendor-specific hardware support manager plug-in that registers itself as a vCenter Server extension.

The hardware support manager enables you to select a firmware add-on to include in an image and the firmware versions to be installed on the hosts. During remediation, vSphere Lifecycle Manager requests the selected hardware support manager to update the firmware on the hosts in accordance with the firmware add-on specified in the image.

In vSphere 7.0, you can deploy hardware support manager plug-ins from Dell and HPE. Dell's plug-in, which you deploy as an appliance, is part of the Dell OpenManage Integration for VMware vCenter Server (OMIVV). HPE's plug-in, which you deploy as an appliance, is part of the HPE iLO Amplifier management tool. You should follow the vendor's specific deployment and configuration documentation. These are the main steps:

- Step 1.** Deploy and power on the virtual appliance.
- Step 2.** Register the virtual appliance as a vCenter Server extension.
- Step 3.** Use the plug-in's UI in the vSphere Client.

You can use the following procedure to manage the firmware on cluster member hosts that are managed with a single image:

- Step 1.** In the vSphere Client, select the cluster in the inventory pane.
- Step 2.** Examine the cluster member hosts and verify that they are from the same vendor.
- Step 3.** Select the cluster and click **Updates > Hosts > Image**.
- Step 4.** In the Image card, click **Edit**.
- Step 5.** In the Edit Image card, click **Firmware and Drivers Addon > Select**.
- Step 6.** In the dialog box, select a hardware support manager.
- Step 7.** Select a firmware add-on from the provided list and review the right panel, which contains information, such as whether the selected add-on contains the necessary drivers for the ESXi versions in the cluster.
- Step 8.** Click **Select**.
- Step 9.** In the Image card, validate and save the image, which triggers a compliance check against the new image.
- Step 10.** In the Image Compliance card, review the results.
- Step 11.** If any host in the cluster has firmware that is non-compliant with the new image firmware, remediate that host or the cluster, using the following steps:
  - a.** Optionally, in the Image Compliance card, click **Run Pre-Check** for the cluster or for a selected host.
  - b.** In the Image Compliance card, initiate remediation.
  - c.** To remediate all hosts in the cluster, click the **Remediate All** button. If the remediation of a single host fails, the remediation for the cluster is aborted.
  - d.** Alternatively, to remediate a single host, click the vertical ellipsis icon for the host and select **Remediate**.

## Note

The host vendor must match the selected hardware support manager vendor. Otherwise, a compliance check will report that the hosts that are from a different vendor or will report an incompatibility. Firmware remediation fails.

## Hardware Compatibility Checks

You can use vSphere Lifecycle Manager to automate the process of validating hardware compliance against the *VMware Compatibility Guide (VCG)* and *vSAN Hardware Compatibility List (vSAN HCL)*, based on ESXi version. Hardware compatibility checks for vSAN cluster member hosts use the vSAN HCL for the I/O devices used by vSAN and the VCG for all other devices.

To check the hardware compatibility for a cluster, select the cluster, click Updates > Hosts > Hardware Compatibility, and click Run Checks.

## Exporting and Importing Cluster Images

You can export a cluster image from one vCenter Server instance and import it to another vCenter Server instance. You can use the following procedure to export the image:

- Step 1.** In the vSphere Client, select the cluster in the inventory pane and navigate to **Updates > Hosts > Image**.
- Step 2.** Click the horizontal ellipsis icon and select **Export**.
- Step 3.** In the Export Image dialog box, select a file format and click **Export**.

### Note

If you intend to use the image in another vCenter Server, export it as a JSON file and as a ZIP file. You can import both the JSON file and the ZIP file to the target vCenter Server system.

You can use the following procedure to import a cluster image:

- Step 1.** In the vSphere Client, select the cluster in the navigation pane and navigate to **Updates > Hosts > Image**.
- Step 2.** Click the horizontal ellipsis icon and select **Import**.
- Step 3.** In the Import Image dialog box, select a JSON file and click **Next**.
- Step 4.** Optionally, use the Edit Image card to modify the following image

elements:

- ESXi version
- Vendor add-on
- Firmware and drivers add-on
- Components

**Step 5.** Resolve any issues with conflicting components or unresolved dependencies.

**Step 6.** Optionally, click **Validate**.

**Step 7.** Click **Save**.

## **Backup and Restore Scenarios**

After switching a cluster from using baselines to using images, if you restore the vCenter Server from a backup made prior to the switch, the restored cluster reverts to being managed by baselines.

If you restore a vCenter Server back to a point in time prior to when you used Lifecycle Manager to remediate a cluster with a new image containing new components, a compliance check on the cluster reveals that the hosts are incompatible. In such a case, Lifecycle Manager expected the hosts to be using a different image. To fix this, you should upgrade the cluster to the new image and remediate.

## **Upgrading Virtual Machines**

You can use vSphere Lifecycle Manager to upgrade VMware Tools and the virtual machine hardware version of multiple virtual machines in a folder, vApp, host, or cluster. You can upgrade a virtual machine regardless of its power state and let Lifecycle Manager change the power state of the virtual machine as needed and return it to its original state when ready. To upgrade VMware Tools, Lifecycle Manager powers on the virtual machine. To upgrade the virtual machine hardware, Lifecycle Manager powers off the virtual machine.

You can use the following procedure to upgrade the hardware for a set of virtual machines by using vSphere Lifecycle Manager:

**Step 1.** In the vSphere client, select a virtual machine container object, such as

a folder or cluster and navigate to **Updates > Hosts > VM Hardware**.

**Step 2.** Select the specific virtual machines to upgrade.

**Step 3.** Click **Upgrade to Match Host**.

**Step 4.** Optionally, use **Scheduling Options** to configure the upgrade as a scheduled task.

**Step 5.** Optionally, use **Rollback Options** and configure the following:

- a. Select or deselect the **Take Snapshot of VMs** checkbox. (It is enabled by default.)
- b. Select a period for keeping the snapshots, either indefinitely or for a fixed period.
- c. Provide a snapshot name.
- d. Optionally, provide a description.
- e. Optionally, select the checkbox to include the virtual machine memory in the snapshot.

**Step 6.** Click **Upgrade to Match Host**.

You can use the same procedure to upgrade VMware Tools, except you should choose Upgrade VMware Tools to Match Host in step 3. On the VMware Tools page, you can choose Set Auto Update. Optionally, you can use a virtual machine's Updates tab to turn on the Automatically Upgrade on Reboot feature.

For more information concerning VMware Tools and virtual machine hardware, such as procedures for interactive installation, host compatibility, and log levels, see [Chapter 14, "Managing Virtual Machines."](#)

## Managing ESXi Hosts

In some cases, such as when you are performing host maintenance and troubleshooting activities, you may want to restart the management agents on an ESXi host. To do so using the Direct Console User Interface (DCUI), select Troubleshooting > Restart Management Agents. This operation restarts all installed management agents and services that are running in /etc/init.d on the ESXi host. Typically, these agents include **hostd**, **ntpd**, **sfcdb**, **slpd**, **wsman**, and **vobd**. When you restart the management agents, current users of the vSphere Client and vSphere Host Client lose connectivity to the host.

When troubleshooting network connectivity with a host, you can choose Test Management Network in the host's DCUI, which automatically pings the host's configured default gateway, pings the DNS server, and attempts to resolve its hostname. You can specify additional addresses for the ping tests. To restore network connectivity for a host or to renew a DHCP lease, you can choose Restart the Management Networking in the host's DCUI.

In rare situations such as the following, you may want to restore a standard switch to an ESXi host, which you can do with the Restore Standard Switch option in the DCUI:

- The distributed switch used for management is not functioning.
- You need to restore the host connection to vCenter, which uses a distributed switch.
- You do not want vCenter Server to manage the host and want to ensure that the management connection does not use a distributed switch.

You can use the vSphere Host Client interface to manage the host's services. To get started, log on to the vSphere Host Client as the root user or as another user with local administrator privileges and navigate to Manage > Services. Here you can examine the state of each ESXi service. To change the state of a service, right-click on the service and select Start, Stop, or Reset. You can also change a service's startup policy such that it automatically starts with the host or associated firewall ports or is only started manually, as illustrated in [Figure 13-1](#). You can perform similar operations using the vSphere Client by selecting the host and navigating to Configure > System > Services.

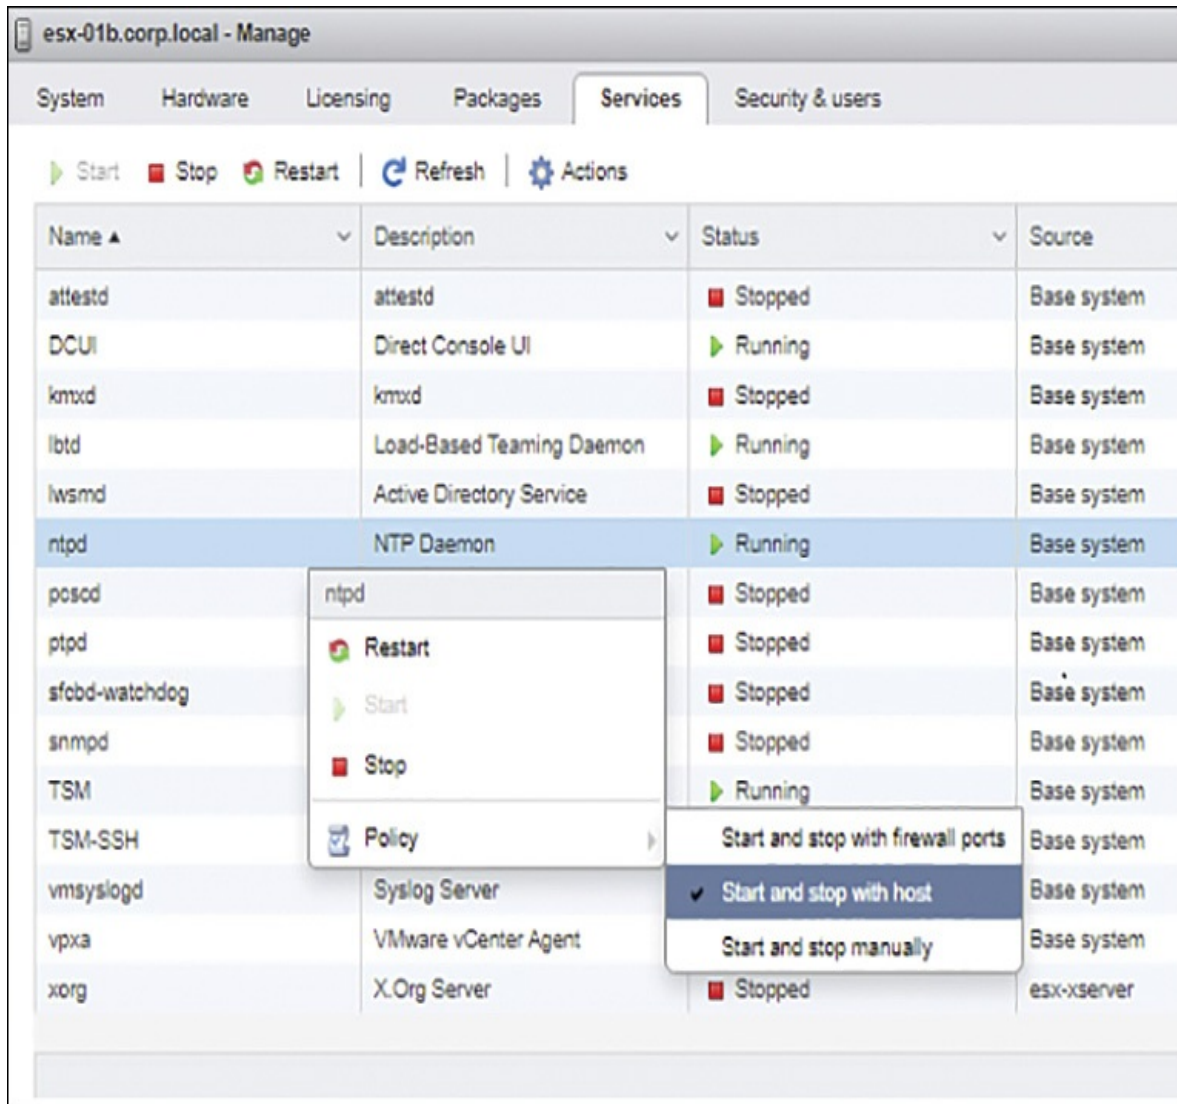


FIGURE 13-1 Managing Host Services in the vSphere Host Client

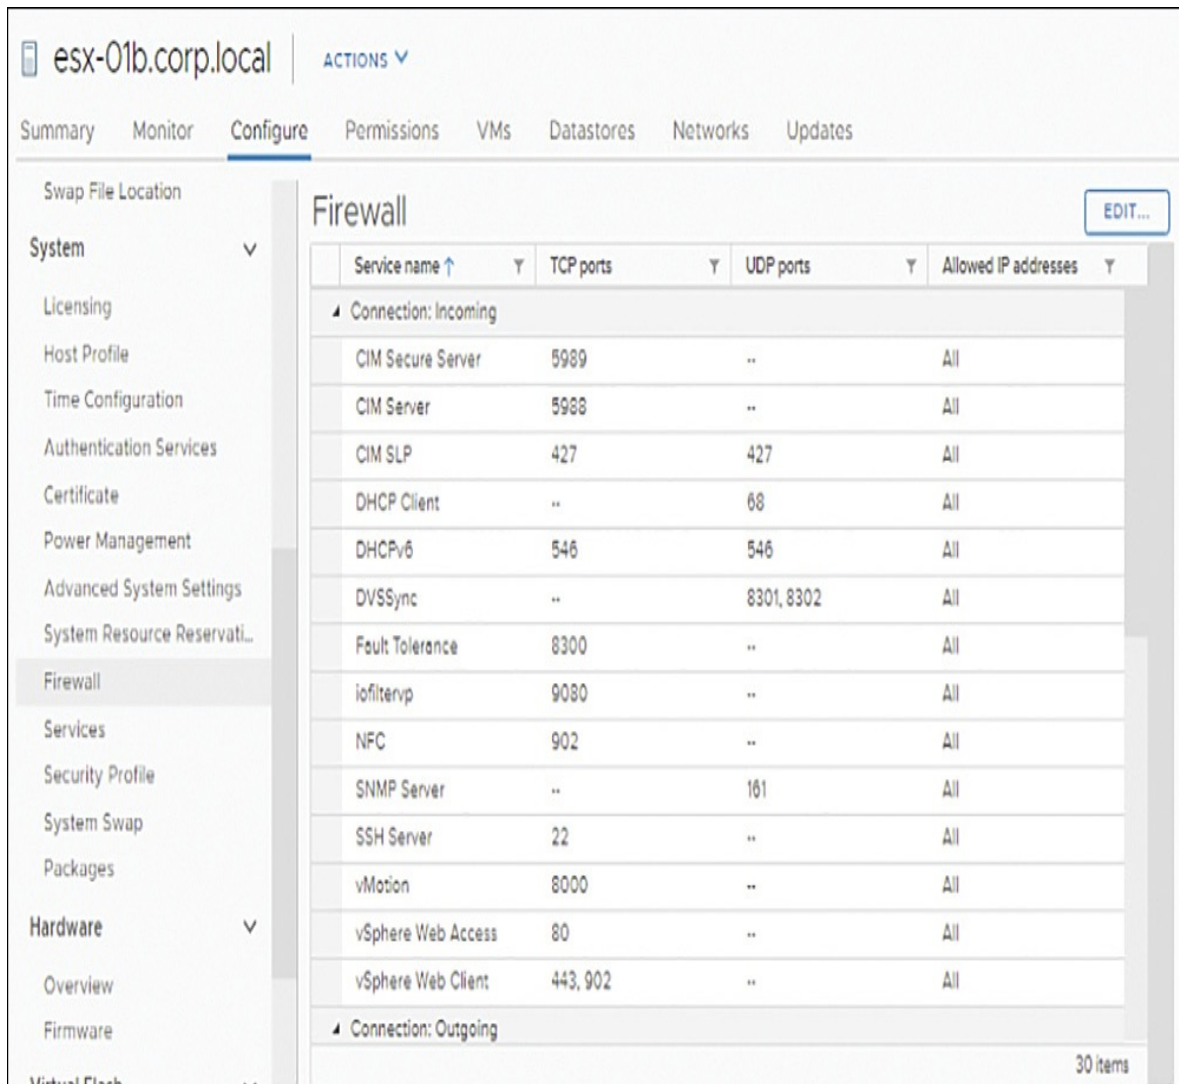
To manage firewall rules on an ESXi host, you can select the host in the vSphere Client and navigate to Configure > System > Firewall, as illustrated in Figure 13-2. Here you can view the currently allowed incoming and outgoing firewall services. The details for each service include the service name, the associated TCP ports, the associated UDP ports, and the allowed IP addresses. To make changes, you can use the Edit button to view all the currently defined services, select the services you want to allow, and optionally restrict the available IP addresses for the service. To perform similar operations in the vSphere Host Client, navigate to Networking > Firewall Rules.

In the vSphere Client, you can right-click on a specific host and choose from a



set of available actions. For example, to address a vCenter Server connection to a host, you can choose Connection > Disconnect, wait for the task to complete, and choose Connection > Connect. To remove a host from the vCenter Server inventory, you can first choose Enter Maintenance mode, then choose Remove from Inventory.

If you want to perform maintenance activities on a host, such as upgrading its hardware, you can choose Maintenance Mode > Enter Maintenance Mode. Following the completion of a maintenance activity, you can select Maintenance Mode > Exit Maintenance Mode. To test a host’s ability to be used with Distributed Power Management (DPM), you can choose Power > Enter Standby Mode.



**FIGURE 13-2** Firewall Rules in the vSphere Client



The following are some of the other options you can select for a host in the vSphere client:

- Certificates (for example, renewing certificates or refreshing CA certificates)
- Host Profiles (including Attach, Detach, Extract, Change Host Profile, or Remediate)
- Export System Logs
- Assign License
- Settings
- Move To
- Add Permission

Details on these tasks are provided elsewhere in this book.

## Monitoring and Managing vCenter Server

You can use the vSphere Client and the vCenter Server Appliance Management Interface (commonly called the VAMI) to configure, monitor, and manage components and resource usage of the vCenter Server Appliance.

The “VIMTOP” and “vCenter Server Appliance Management Interface (VAMI)” sections in [Chapter 10, “Managing and Monitoring Clusters and Resources,”](#) provide information about monitoring the resource usage of the services and database running in the vCenter Server Appliance. If database resources are low, you can consider adjusting the statistics interval, statistics level, task retention, and event retention settings. If you determine that the appliance is low on disk space, you can add more space. To increase the disk space for a vCenter Server Appliance 7.0, you can use the following procedure:

- Step 1.** Use the vSphere Client (or Host Client) to navigate the vSphere environment that is hosting the vCenter Server Appliance.
- Step 2.** Select the vCenter Server Appliance (virtual machine), edit its settings, and increase the virtual disk size.
- Step 3.** Use SSH to connect to the vCenter Server as the root user.
- Step 4.** Run the following command:

[Click here to view code image](#)

```
com.vmware.appliance.system.storage.resize
```

Optionally, to see the impact of the command, you can use the following command to examine the total and used storage (in kilobytes) before and after step 4:

[Click here to view code image](#)

```
com.vmware.appliance.version1.resources.storage.stats.list
```

### Note

If you want to keep a complete history of tasks and events for your vCenter Server, do not use the database retention options.

## Monitoring and Managing vCenter Server with the VAMI

You can use the VAMI to monitor and manage specific components of your vCenter Server.

To access the VAMI, use a web browser to connect to `https://vCenter-FQDN:5480`. By default, you can log on using the root account and the password that was set during the vCenter Server deployment.

### Note

If you are using Internet Explorer, verify that TLS 1.0, TLS 1.1, and TLS 1.2 are enabled in the security settings.

After logging in to the VAMI as root, you can perform any of the tasks described in [Table 13-5](#).

**Table 13-5** Tasks in vCenter Server Appliance Management Interface



Task	Steps/Details
------	---------------

View vCenter  
Server health  
status

1. Click Summary.
2. Examine the color of the Overall Health badge.

---

Reboot  
vCenter  
Server

1. Click Summary
2. Click Actions > Reboot and click Yes.

---

Shut down  
vCenter  
Server

1. Click Summary
2. Click Actions > Shutdown and click Yes.

---

Create a  
support  
bundle

1. Click Summary
2. Click Actions > Create Support Bundle and save the bundle to your desktop.

---

Monitor CPU  
and memory  
use of vCenter  
Server

1. Click Monitor
2. Click the CPU & Memory tab and use the date range drop-down menu to specify a timeframe.

---

Monitor disk  
use of vCenter  
Server

1. Click Monitor
2. Click the Disks tab.

---

Monitor  
network use  
of vCenter  
Server

1. Click Monitor
2. Click the Network tab and use the date range drop-down menu to specify a timeframe.
3. Examine the graph grid and use the provided table to select a packet or transmit byte rate to monitor.

---

Monitor  
database use  
of vCenter  
Server

1. Click Monitor
2. Click the Database tab and use the date range drop-down menu to specify a timeframe.
3. At the base of the graph, select specific database components to include in the graph.

---

Enable or  
disable each  
type of  
vCenter  
Server access

1. Click Access
2. Configure each of the following options:

- Enable SSH Login.
- Enable DCUI
- Enable Console CLI
- Enable Bash Shell

---

Configure  
network  
settings for  
vCenter  
Server

1. Click Networking
2. Click Edit and fill in the following networking details:
  - DNS settings
  - IPv4 settings
  - IPv6 settings
  - Proxy server settings

When configuring a proxy server, you can enable HTTPS, FTP, and HTTP options. You should provide the proxy server's IP address or host name, user credentials, and a port number.

---

Configure the firewall rules for the vCenter Server

1. Click Firewall
2. Examine the existing set of rules and choose from the following commands to change the rule set.
  - Add
  - Edit
  - Delete
  - Reorder
3. For each rule, include the appropriate NIC, IP address, subnet, and action (Accept, Ignore, Reject, or Return).

---

Configure the time settings for the vCenter Server

1. Click Time
2. Click Time Zone > Edit and select the appropriate time zone.

3. Click Time Synchronization > Edit and select Mode to Disable, Host, or NTP.

---

Start, stop, and restart a service in the vCenter Server

1. Click Services
2. Select a service and click Start, Stop, or Restart.

---

Configure settings for updating vCenter Server

1. Click Update
2. Click Settings, set the options for automatic update checking, and set the repository to the default or to a custom (HTTPS or FTPS) URL. Optionally provide credentials if a custom repository is used.
3. Click Check Updates to manually check for updates.

---

Change the root user password in vCenter Server

1. Click Administration
2. Click Password > Change. Set the password and the password expiration details. If you set the password to expire, provide the number of days and an address for the email warning.

---

Configure log forwarding on the vCenter

1. Click Syslog

## Server

2. Click Configure (or Edit, if you previously configured syslog hosts), enter up to three destination hosts in the Create Forward Configuration pane, and set the Protocol and Port.
  
3. Select one of the following protocol options:
  - **TLS:** Transport Layer Security
  
  - **TCP:** Transmission Control Protocol
  
  - **RELP:** Reliable Event Logging Protocol
  
  - **UDP:** User Datagram Protocol

---

### Configure and schedule backups of the vCenter Server

1. Click Backup
  
2. Click Backup > Backup Now to initiate a backup or click Backup > Configure to schedule backups.
  
3. Use the Activity table to monitor backups.

---

Table 13-6 describes the possible colors (icons) for the Health Status badge.



**Table 13-6** Health States

Col or	Mea ning	Description
Green	Good	All components are healthy.
Yellow	Warning	One or more components may become overloaded soon.
Orange	Alert	One or more components may be degraded. Non-security patches may be available.
Red	Critical	One or more components may be in an unusable state, and vCenter Server may become unresponsive soon. Security patches may be available.
Gray	Unknown	No data is available.

### Note

When responding to an alert, you should begin by examining the details in the Health Messages pane.

You can use the following procedure to reconfigure the FQDN, IP address, or the primary network identifier (PID) of the vCenter Server:

**Step 1.** Log in to the VAMI using your administrator SSO credentials.

**Step 2.** Select **Networking** > **Network Settings** > **Edit**.

- Step 3.** Select the appropriate NIC and click **Next**.
- Step 4.** In the Edit Settings pane, change the host name and IP address. Click **Next**.
- Step 5.** In the SSO Credentials pane, provide the administrator SSO credentials. (You must provide the credentials for the administrator account in the SSO domain, such as [administrator@vsphere.local](#).)
- Step 6.** In the Ready to Complete pane, click **Finish**.
- Step 7.** Monitor the progress in the taskbar, and when you are redirected to the new IP address, log in using your administrator SSO credentials.
- Step 8.** On the Networking page, verify the new host name and IP address.
- Step 9.** Re-register all deployed plug-ins.
- Step 10.** Regenerate all custom certificates.
- Step 11.** If vCenter Server HA was enabled, reconfigure vCenter HA.
- Step 12.** If Active Directory was enabled, reconfigure Active Directory.
- Step 13.** If Hybrid Link Mode was enabled, reconfigure Hybrid Link with Cloud vCenter Server.

### **Note**

If you set an IP address as a system name during the deployment of the appliance, you can later change the primary network identifier to a fully qualified domain name. If vCenter High Availability (HA) is enabled, you must disable the vCenter HA setup before reconfiguring the primary network identifier.

## **Monitoring and Managing vCenter Server with the vSphere Client**

You can use the vSphere Client to monitor and manage specific components of your vCenter Server, as described in this section.

### **Common vCenter Server Management Tasks**

Using the vSphere Client, you can log on, select your vCenter Server in the inventory, and then perform any of the tasks listed in [Table 13-7](#).



**Table 13-7** Management Tasks in the vSphere Client

Task	Steps	Notes
Assign a license to your vCenter Server	<ol style="list-style-type: none"> <li>1. Select the Configure tab.</li> <li>2. Select Settings &gt; Licensing &gt; Assign License.</li> <li>3. Select an existing license or enter a new license key.</li> </ol>	This task requires the Global.Licenses privilege.
Configure statistics settings for vCenter Server	See the “ <a href="#">Configuring Statistics Collection Settings</a> ” section, later in this chapter.	—
Configure runtime settings for vCenter Server	<ol style="list-style-type: none"> <li>1. Select the Configure tab.</li> <li>2. Navigate to Settings &gt; General &gt; Edit &gt; Runtime Settings and configure the following settings: <ul style="list-style-type: none"> <li>▪ <b>vCenter Server Unique ID:</b> Specify a valid number from 0</li> </ul> </li> </ol>	This task requires the Global.Settings privilege.

through 63. (The default value is randomly generated.)

- **vCenter Server Managed Address:** Can be IPv4, IPv6, a fully qualified domain name, an IP address, or another address format.
  
- **vCenter Server Name:** Specify the name of the vCenter Server system (which should match the vCenter Server DNS name).

---

Configure user directory settings for vCenter Server

1. Select the Configure tab.
  
2. Navigate to Settings > General > Edit > User Directory and configure the following, which will be applied to vCenter Server Single Sign-On identity sources:

This task requires the Global.Settings privilege.

- **User Directory Timeout:** The timeout, in seconds, for directory server connection
  
- **Query Limit:** Must be enabled to set a query limit size

- **Query Limit Size:** The maximum number of queries.
  
- **Validation:** When enabled, has vCenter Server validate (synchronize) known users against the directory server
  
- **Validation Period:** The number of minutes between validation executions.

---

Configure mail sender settings for vCenter Server

1. Select the Configure tab.
  
2. Navigate to Settings > General > Edit > Mail and configure the following:

- **Mail Server:** SMTP server information
  
- **Mail Sender:** Full email address (including domain name) of the sender

This task requires the Global.Settings privilege.

Configure  
SNMP  
settings for  
vCenter  
Server

This task requires the  
Global.Settings  
privilege.

1. Select the Configure tab.
2. Navigate to Settings > General > Edit > SNMP receivers and configure the following for one or more receivers:
  - **Enable Receiver:** Whether a receiver is enabled or disabled
  - **Primary Receiver URL:** Host name or IP address of the SNMP receiver
  - **Receiver Port:** Port number (between 1 and 65535) of the receiver
  - **Community String:** The community identifier

---

View port  
settings for  
vCenter  
Server

1. Select the Configure tab.
2. Navigate to Settings > General > Edit > Ports and examine the ports used by the web service.

---

Configure  
timeout  
settings for  
vCenter  
Server

1. Select the Configure tab.
2. Navigate to Settings > General > Edit > Timeout Settings and configure the following:
  - **Normal:** Timeout for normal operations
  - **Long:** Timeout for long operations
3. Restart the vCenter Server system.

This task requires the Global.Settings privilege.

---

Configure  
logging  
settings for  
vCenter  
Server

1. Select the Configure tab.
2. Navigate to Settings > General > Edit > Logging settings and select one of the following logging options:
  - None (disable logging)

This task requires the Global.Settings privilege.

- Error (errors only)
- Warning (errors and warnings)
- Info (normal logging)
- Verbose (verbose)
- Trivia (extended verbose)

---

Configure database settings for vCenter Server

1. Select the Configure tab.
2. Navigate to Settings > General > Edit > Database and configure the following:

- Maximum connection
- Task retention (in days)
- Event retention (in days)

To limit the database growth and save storage space, you can configure the database to periodically discard task and event information.

---

Verify SSL

See the section “[Verifying SSL](#)”

—



certificates for legacy hosts

[Certificates for Legacy Hosts](#),” later in this chapter.

---

Configure advanced settings for vCenter Server (which are kept in the vpxd.cfg file)

1. Select the Configure tab.
2. Navigate to Advanced Settings > Edit Settings. Add or modify advanced settings entries, providing a name and value for each entry. (Note that you cannot delete entries.)
3. If necessary, restart the vCenter Server system.

Use Advanced Settings only when instructed to do so by VMware. This task requires the Global.Settings privilege.

---

Send a message to users who are logged in to vCenter Server

1. Select the Configure tab.
2. Navigate to Settings > Message of the Day, click Edit, and enter a message, such as a maintenance announcement or request for users to log out.

---

Start, stop, and restart the services (nodes) in the vCenter Server

1. Navigate to Administration > System Configuration.
2. Select a node and click Reboot

This task requires that your user account be a member of the SystemConfiguration Administrators group.

Node.

---

View the health status of vCenter Server services and nodes

1. Navigate to Administration > Deployment > System Configuration.
2. Examine the Health Status badge for each service and node (refer to [Table 13-6](#)).

This task requires that your user account be a member of the SystemConfiguration Administrators group.

---

Export a support bundle

1. Navigate to Administration > Deployment > System Configuration.
2. Select a node, click Export Support Bundle, select the services that you want to include from the two available categories (cloud infrastructure and virtual appliance), and click Export Support Bundle.

This task requires that your user account be a member of the SystemConfiguration Administrators group.

---

## Configuring Statistics Collection Settings

vCenter Server and ESXi use data counters, organized in metric groups, to collect data statistics. A data counter is a unit of information that is relevant to a specific inventory object or device. For example, the disk metric group includes separate data counters for collecting disk read rate, disk write rate, and disk usage data. Statistics for each counter are rolled up at specific collection intervals. Each data counter consists of several attributes that are used to

determine the statistical value collected.

Collection levels determine the number of counters for which data is gathered during each collection interval. Collection intervals determine the time period during which statistics are aggregated, calculated, rolled up, and archived in the vCenter Server database. Together, the collection interval and collection level determine how much statistical data is collected and stored in a vCenter Server database.

You can configure statistical data collection intervals to set the frequency at which statistical queries occur, the length of time statistical data is stored in the database, and the type of statistical data that is collected. To configure a statistical interval, you can log on to the vSphere Client as a user with Performance.ModifyIntervals privilege and use the following procedure:

- Step 1.** In the vSphere Client, select the vCenter Server instance in the inventory pane and navigate to **Configure > Settings**.
- Step 2.** Select **General > Edit**.
- Step 3.** To enable a statistics interval, select its checkbox.
- Step 4.** Set the following interval attributes using the provided drop-down menus.
  - **Interval Duration:** The time interval during which statistical data is collected.
  - **Save For:** The amount of time to keep statistics in the database.
  - **Statistics Level:** The level of statistical data to be collected (1 to 4).
- Step 5.** Optionally, for Database Size, set the following options and examine the estimated database size and number of rows:
  - Number of Physical Hosts
  - Number of Virtual Machines

**Step 6.** Click **Save**.

[Table 13-8](#) provides details on the default data collection intervals.

**Table 13-8** Collection Intervals

---

Collection	Collecti	Default Behavior
------------	----------	------------------

<b>Interval (Archive Length)</b>	<b>on Frequen cy</b>	
1 day	5 minut es	<p>Real-time (20-second) statistics are rolled up to create one data point every 5 minutes. The result is 288 data points every day.</p> <p>You can change the interval duration and archive length of the 1-day collection interval by configuring the statistics settings.</p>
1 week	30 minut es	<p>1-day statistics are rolled up to create one data point every 30 minutes. The result is 336 data points every week.</p> <p>You cannot change the default settings of the 1-week collection interval.</p>
1 month	2 hours	<p>1-week statistics are rolled up to create one data point every 2 hours. The result is 360 data points every month.</p> <p>You cannot change the default settings of the 1-month collection interval.</p>
1 year	1 day	<p>1-month statistics are rolled up to create one data point every day. The result is 365 data points each year.</p> <p>You can change the archive length of the 1-year collection interval by configuring the statistics settings.</p>

The default statistics level for all statistical intervals is 1. You can set the statistics level to a value between 1 and 4, inclusive. The lower the level, the smaller the number of statistics counters used. Level 4 uses all statistics counters, but it is typically used only for debugging purposes. When setting a statistics level for a specific statistics interval, you must use a value less than or equal to the statistics level for the preceding statistics interval. [Table 13-9](#)

provides a summary of the metrics that are included for each statistics level.

**Table 13-9** Statistics Levels

<b>Level</b>	<b>Metrics</b>	<b>Best Practice</b>
1	<p>Cluster Services: All metrics</p> <p>CPU: cpuentitlement, totalmhz, usage (average), usagemhz</p> <p>Disk: capacity, maxTotalLatency, provisioned, unshared, usage (average), used</p> <p>Memory: consumed, mementitlement, overhead, swapinRate, swapoutRate, swapused, totalmb, usage (average), vmmemctl (balloon)</p> <p>Network: usage (average), IPv6</p> <p>System: heartbeat, uptime</p> <p>Virtual Machine Operations: numChangeDS, numChangeHost, numChangeHostDS</p>	<p>Use this level for long-term performance monitoring when device statistics are not required.</p>
2	<p>Level 1 metrics plus the following:</p> <p>CPU: idle, reservedCapacity</p>	<p>Use this level for long-term performance monitoring when device statistics are not required but you want to monitor more than the basic statistics.</p>

Disk: All metrics, excluding numberRead and numberWrite

Memory: All metrics, excluding memUsed and maximum and minimum rollup values

Virtual Machine Operations: All metrics

---

3	Level 1 and Level 2 metrics plus the following:  Metrics for all counters, excluding minimum and maximum rollup values  Device metrics	Use this level for short-term troubleshooting or when device statistics are required.  Due to the large quantity of data, use this level for the shortest (day or week) collection interval that suits your use case.
4	All metrics supported by the vCenter Server, including minimum and maximum rollup values	Same best practice as Level 3, except that Level 4 should be used only when you require metric rollup values that are not available in Level 3.

---

### Note

If you increase the collection level, you may need to allocate more storage and system resources to avoid a decrease in performance.

## Verifying SSL Certificates for Legacy Hosts

You can configure vCenter Server and the vSphere Client to check for valid SSL certificates before connecting to a host for operations such as adding a host or making a remote console connection to a virtual machine. vCenter Server 5.1 and vCenter Server 5.5 always connect to ESXi hosts using SSL thumbprint

certificates. Starting with vCenter Server 6.0, VMware Certificate Authority signs the SSL certificates by default. You can instead replace certificates with certificates from a third-party CA. Thumbprint Mode is supported only for legacy hosts. To configure SSL certificate validation, you can use the following procedure:

- Step 1.** In the vSphere Client, select the vCenter Server object in the inventory and navigate to **Configure > Settings > General**.
- Step 2.** Click **Edit** and select **SSL Settings**.
- Step 3.** Determine the host thumbprint for each legacy host that requires validation:
  - a.** Log in to the DCUI.
  - b.** Select **System Customization > View Support Information** and examine the thumbprint displayed in the right column.
- Step 4.** Compare each host thumbprint with the thumbprint listed in the vCenter Server SSL Settings dialog box.
- Step 5.** If the thumbprints match, select the checkbox for the host.
- Step 6.** Click **Save**. Hosts that were not selected in step 5 are disconnected.

## Updating the vCenter Server

You can use the VAMI or the vCenter Server Appliance shell to install patches.

### Patching with VAMI

To update a vCenter Server Appliance, you may want to first stage the patches to the appliance, using the following procedure:

- Step 1.** Log on as root to the VAMI.
- Step 2.** Click **Update**.
- Step 3.** Click **Check Updates** and select one of the following:
  - a.** **Check URL**
  - b.** **Check CDROM**
- Step 4.** Optionally, choose **Run Pre-check**.

**Step 5.** In the Staging Options section, click **Stage**.

### Note

Other staging options include Stage and Install, Unstage, and Resume.

If you choose to use the Check URL option, the vCenter Server uses the configured VMware repository URL. The default VMware repository URL requires Internet access. If your vCenter Server is not connected to the Internet or if required by your security policy, you can configure a custom repository URL for your vCenter Server patches by using the following procedure:

- Step 1.** Download the vCenter Server Appliance patch ZIP file from VMware's website (<https://my.vmware.com/web/vmware/downloads>).
- Step 2.** On a local web server, create a repository directory under the root.
- Step 3.** Extract the ZIP file into the repository directory.
- Step 4.** Log on as root to the VAMI.
- Step 5.** Select **Update > Settings**.
- Step 6.** Set the **Repository settings**: Choose **use specified repository**, provide the URL and (optionally) the user credentials.
- Step 7.** Click **OK**.

After staging the patches to the vCenter Server, you can install the patches by using the following procedure:

- Step 1.** Log on as root to the VAMI.
- Step 2.** Ensure that the patches are staged or use the staging procedure, but for the staging options, select **Stage and Install**.
- Step 3.** Click **Update**.
- Step 4.** Select the range of patches to apply and click **Install**.
- Step 5.** Read and accept the end-user license agreement (EULA).
- Step 6.** Wait for the installation to complete and then click **OK**.
- Step 7.** If a reboot is required, click **Summary > Reboot**.



## Note

You should perform the previous procedure only during a maintenance period because the services provided by the vCenter Server become unavailable during the patch installation. As a precaution, you should also back up the vCenter Server prior to patching.

To configure vCenter Server to schedule automatic checks for available patches in the configured repository URL, you can use the following procedure:

**Step 1.** Log on as root to the VAMI.

**Step 2.** Verify that the correct repository URL is set and available.

**Step 3.** Click **Update > Settings**.

**Step 4.** Select **Check for Updates Automatically** and set the day and time (in UTC) to check for available patches.

## Patching with the vCenter Server Appliance Shell

An alternate method for patching vCenter Server is to use its shell. For patching purposes, you should log on to the shell as a user with the super administrator role and use the commands in [Table 13-10](#), as needed.

**Table 13-10** Commands for Patching

<b>Purpose</b>	<b>Command/Utility</b>
View the full list of patches and software packages installed in the vCenter Server Appliance, in chronological order	<b>software-packages list --history</b>
View details about a specific patch	<b>software-packages list --patch</b> <i>patch_name</i>
Configure the vCenter Server to use a custom repository URL	<b>update.set --currentURL</b> <i>http://webserver /repo</i> [-- <b>username</b> <i>username</i> ] [-- <b>password</b> <i>password</i> ]
Configure the vCenter Server to use the default VMware repository URL	<b>update.set --currentURL default</b>

Enable automatic checks for vCenter Server Appliance patches in the current repository	<b>update.set --CheckUpdates enabled [--day <i>day</i>] [--time <i>HH:MM:SS</i>]</b>
Stage the patches from an attached ISO image	<b>software-packages stage --iso</b>
Stage the patches from the current repository URL	<b>software-packages stage --url</b>
Stage the patches from the repository URL that is not currently configured in the vCenter Server	<b>software-packages stage --url <i>URL_of_the_repository</i></b>
Install the staged patches	<b>software-packages install --staged</b>
Install patches directly from an attached ISO image	<b>software-packages install --iso</b>
Install patches directly from the configured repository URL	<b>software-packages install --url</b>
Install patches directly from a repository URL that is not currently configured in the vCenter Server	<b>software-packages install --url <i>URL_of_the_repository</i></b>
Reboot the vCenter Server following a patch installation	<b>shutdown reboot -r "patch reboot"</b>

## Note

To stage only third-party patches, include the **--thirdParty** option with the **software-packages stage** command. To directly accept the EULA, include the **--acceptEulas** option.

## Managing a vCenter HA Cluster

During normal operation, the vCenter HA cluster mode is Enabled, such that it is protecting vCenter Server from hardware and software failures. When the cluster detects the failure of the Active node, the Passive node attempts an automatic failover, and the Passive node becomes the Active node. You can choose to

perform a manual failover by choosing the Initiate Failover option in the vCenter HA settings in the vSphere Client. When performing a manual failover, you can choose to synchronize first or to force the failover without synchronization.

You can change the vCenter HA cluster mode to Maintenance when preparing to perform some maintenance activities. If the Passive or Witness nodes are unavailable (or recovering from a failure), a vCenter HA cluster mode may be disabled, in which case the Active node continues as a standalone vCenter Server. When the cluster is operating in either Maintenance Mode or Disabled Mode, an Active node can continue serving client requests even if the Passive and Witness nodes are lost or unreachable.

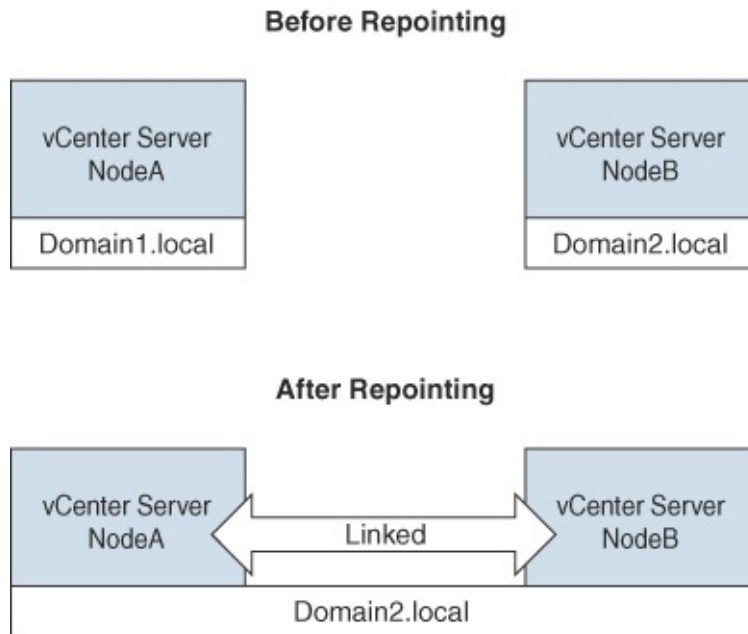
To change the vCenter HA cluster mode, you can select the Active node in the vSphere Client, select Configure > Settings > vCenter HA > Edit, and choose Enable vCenter HA, Maintenance Mode, Disable vCenter HA, or Remove vCenter HA Cluster.

You can use the following procedure to perform backup and restore operations on a vCenter HA cluster:

- Step 1.** Use the Active node's VAMI to obtain a file-based backup of the Active vCenter Server node. (Do not back up the Passive node or Witness node.)
- Step 2.** Before you begin the restore operation, power off and delete all vCenter HA nodes and remove the cluster configuration.
- Step 3.** Restore the Active node from the backup. The Active node is restored as a standalone vCenter Server.
- Step 4.** Reconfigure vCenter HA.

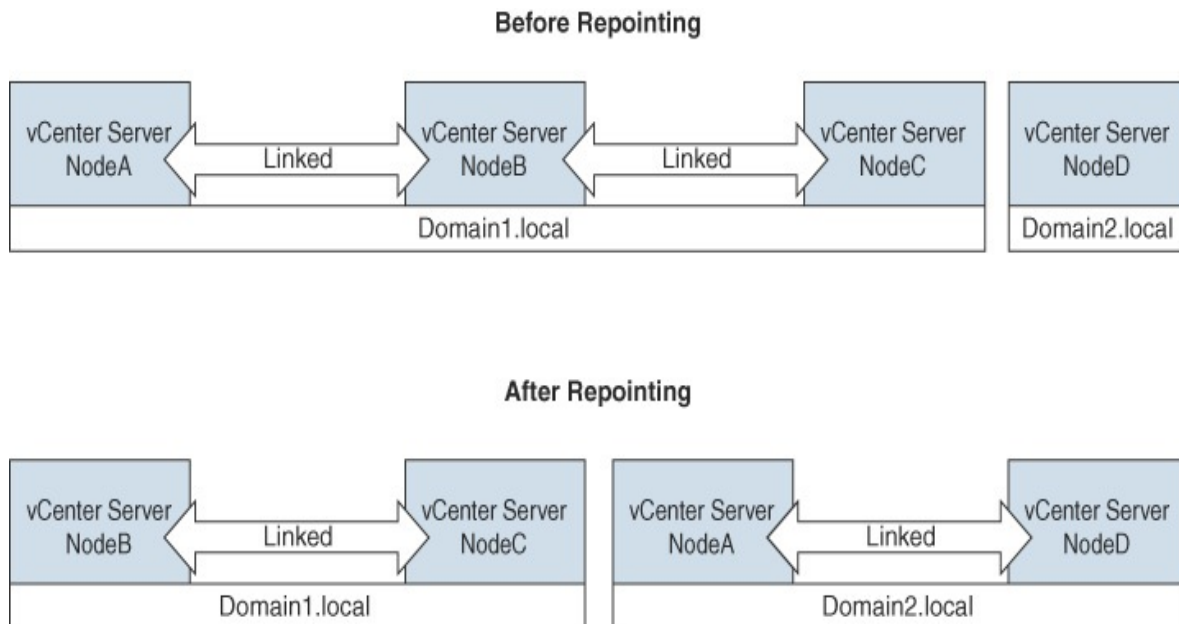
## Repointing a vCenter Server to Another Domain

You can repoint a vCenter Server from one SSO domain to another existing domain. The steps in the procedure depend on whether the vCenter Server is the only node in the source domain and whether it is being repointed to a new domain or to an existing domain. [Figure 13-3](#) illustrates a vCenter Server (Node A) being repointed from a single-node domain to an existing domain.



**FIGURE 13-3** Repointing vCenter to an Existing Domain

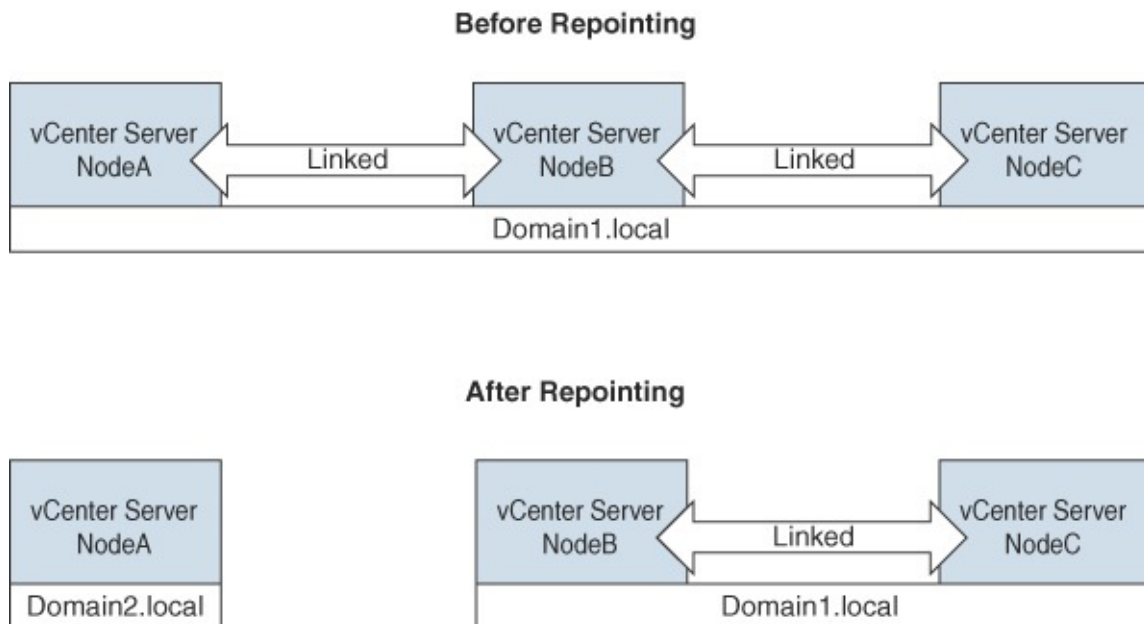
Figure 13-4 illustrates a vCenter Server (Node A) being repointed from a multi-node domain to an existing domain.



**FIGURE 13-4** Repointing vCenter from a Multi-Node Domain

Figure 13-5 illustrates a vCenter Server (Node A) being repointed from a multi-

node domain to a new domain that is created with the **repoint** command.



**FIGURE 13-5** Repointing vCenter to a New Domain

If the source domain contains multiple (linked) vCenter Servers, then the repointing process involves additional steps to shut down the vCenter Server and unregister it from the source domain. If the vCenter Server is repointed to a new domain, you do not need to run a precheck or supply the replication partner parameters. Repointing is only supported with vCenter Server 6.7 Update 1 and later. You should back up each node prior to repointing.

To repoint a vCenter Server to another domain, you can use the following procedure:

**Step 1.** If multiple vCenter Servers exist in the source domain:

- a. Shut down the chosen vCenter Server.
- b. To unregister the chosen vCenter Server from the source domain, log in to one of the other nodes in the source domain and run the following command, where *username* and *password* are credentials for the source SSO domain administrator account:

[Click here to view code image](#)

```
cmsso-util unregister --node-pnid Target_vCenter_FQDN  
--username username --passwd password
```

- c. Power on the chosen vCenter Server.

**Step 2.** Ensure that the chosen vCenter Server is powered on.

**Step 3.** If joining an existing domain:

- a. Ensure that a target replication partner (a vCenter Server in the existing domain) is powered on.
- b. Optionally, run the following precheck mode command from the chosen vCenter Server, which fetches tagging and authorization (roles and privileges) data and checks for conflicts between the source and destination:

[Click here to view code image](#)

```
cmsso-util domain-repoint -m pre-check --src-emb-admin
Administrator --replication-partner-fqdn Replication_
Partner_FQDN --replication-partner-admin Existing_Domain_
Admin --dest-domain-name Existing_Domain
```

The precheck writes the conflicts to the /storage/domain-data directory.

- c. Optionally, check conflicts and apply one of the following resolutions to all conflicts or separately to each conflict:
  - **Copy:** Creates a duplicate copy of the data in the target domain
  - **Skip:** Skips copying the data in the target domain
  - **Merge:** Merges the conflict without creating duplicates

The default resolution mode for Tags and Authorization conflicts is **Copy**.

**Step 4.** Run the following **execute** command, which applies any precheck data and either repoints the chosen vCenter Server to the existing domain or creates a new domain for repointing. If the chosen vCenter Server is being pointed to a new domain, you do not need to supply the replication partner parameters:

[Click here to view code image](#)

```
cmsso-util domain-repoint -m execute --src-emb-admin
Administrator --replication-partner-fqdn Replication_
Partner_FQDN --replication-partner-admin Existing_Domain_
Admin --dest-domain-name Existing_Domain
```

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15](#), “[Final Preparation](#),” and the exam simulation questions on the companion website.

### Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 13-11](#) lists these key topics and the page number on which each is found.



**Table 13-11** Key Topics for [Chapter 13](#)

Key Topic Element	Description	Page Number
List	Restoring vCenter Server	<a href="#">516</a>
<a href="#">Table 13-2</a>	vCenter Server 7.0 compatibility	<a href="#">518</a>
List	Prerequisites for upgrading vCenter Server Appliance	<a href="#">519</a>
Section	Migrating vCenter Server for Windows to vCenter Server	<a href="#">522</a>
<a href="#">Table 13-3</a>	Lifecycle Manager remediation settings	<a href="#">528</a>
List	Creating a dynamic baseline	<a href="#">530</a>
List	Enabling Quick Boot	<a href="#">536</a>
<a href="#">Table 13-5</a>	Management tasks using the VAMI	<a href="#">544</a>
<a href="#">Table 13-7</a>	Management tasks using the vSphere Client	<a href="#">548</a>

### Complete Tables and Lists from Memory

Print a copy of [Appendix B](#), “[Memory Tables](#)” (found on the companion website), or at least the section for this chapter, and complete the tables and lists

from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

[VMware vSphere Lifecycle Manager](#)

[image](#)

[VMware Update Manager Download Service \(UMDS\)](#)

[ESXi base image](#)

[add-on](#)

[baseline](#)

## Review Questions

1. You need to restore the vCenter Server from a file-based backup. Which of the following will not be restored?
  - a. Resource pool hierarchy and setting
  - b. vSphere DRS cluster state
  - c. Cluster-host membership
  - d. vSphere DRS configuration and rules
2. You plan to upgrade a Windows-based vCenter Server to vCenter Server Appliance 7.0 and want to transfer data in the background. Which of the following can be included in the background transfer?
  - a. Configuration data only
  - b. Configuration data and performance data
  - c. Performance data
  - d. Data from the external database
3. You are configuring remediation setting for Lifecycle Manager. Which of the following settings are only available when working with baselines?
  - a. PXE Booted Hosts and Removable Media Devices



- b.** Quick Boot and VM Power State
  - c.** VM Migration and Admission Control
  - d.** VM Migration and Maintenance Mode Failures
- 4.** Your vCenter Server is offline, and the distributed switch for an ESXi host management network is not functioning. Which of the following steps might fix the ESXi management connectivity?
  - a.** Use the vSphere Host Client to restart ESXi networking.
  - b.** Use the vSphere Client to restart ESXi networking.
  - c.** Use SSH to restart ESXi networking.
  - d.** In the DCUI, select Restore Standard Switch.
- 5.** You are repointing a vCenter Server to an existing domain. In which of the following scenarios would you need to run a precheck?
  - a.** Multiple vCenter Servers exist in the target domain.
  - b.** Multiple vCenter Servers exist in the source domain.
  - c.** A single vCenter Server exists in the target domain.
  - d.** A single vCenter Server exists in the source domain.

## Chapter 14

# Managing Virtual Machines

This chapter covers the following topics:

- [Creating and Configuring Virtual Machines](#)
- [Managing Virtual Machines](#)
- [Advanced Virtual Machine Management](#)
- [Content Libraries](#)

This chapter contains information related to Professional VMware vSphere 7.x (2V0-21.20) exam objectives 4.7, 5.6, 7.1, 7.2, 7.3, 7.6, and 7.11.4.

This chapter provides details on managing virtual machines.

## “Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess whether you should study this entire chapter or move quickly to the “Exam Preparation Tasks” section. In any case, the authors recommend that you read the entire chapter at least once. [Table 14-1](#) 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 and Review Questions.](#)”

**Table 14-1** “Do I Know This Already?” Section-to-Question Mapping

<b>Foundation Topics Section</b>	<b>Questions</b>
Creating and Configuring Virtual Machines	1–3
Managing Virtual Machines	4–6
Advanced Virtual Machine Management	7, 8
Content Libraries	9, 10

- 
1. You are creating a virtual machine in your vSphere 7.0 environment, and you want the virtual disk and NVDIMM devices to share the same PMem resources. Which of the following options should you choose?
    - a. In the Memory settings, select PMem.
    - b. In the Memory settings, select Standard.
    - c. In the Storage Type settings, select PMem.
    - d. In the Storage Type settings, select Standard.
  2. You want to change the logging for the VMware Tools installation, such that vminst.log is sent to the host, but vmmsi.log remains in the virtual machine. Which option should you choose?
    - a. vmx.log.guest.level = “warning”
    - b. vmx.log.guest.level = “info”
    - c. vmx.log.guest.level = “verbose”
    - d. vmx.log.guest.level = “trivia”
  3. You want to deploy new virtual machines using linked clones. Which of the following should you use?
    - a. vSphere API
    - b. vSphere Client
    - c. vSphere Host Client
    - d. vCenter Management Interface
  4. You are updating a virtual machine and want to use hardware Version 14. Which of the following compatibility settings should you choose?
    - a. ESXi 7.0 and later
    - b. ESXi 6.7 Update 2 and later
    - c. ESXi 6.7 and later
    - d. ESXi 6.5 and later
  5. You want to control the host compatibility for your virtual machines at various levels of the inventory. On which of the following objects can you set the Default VM Compatibility option?

- a. Cluster
  - b. VM folder
  - c. Virtual machine
  - d. Template
6. Which option should you choose to minimize the time required to create a virtual machine snapshot?
- a. Snapshot the memory and quiesce the file system.
  - b. Snapshot the memory but do not quiesce the file system.
  - c. Quiesce the file system but do not snapshot the memory.
  - d. Do not quiesce the file system or snapshot the memory.
7. You want to enable Microsoft virtualization-based security (VBS) for a Windows virtual machine in a vSphere environment. Which of the following is a requirement?
- a. vSphere 7.0 or later
  - b. Virtual machine hardware Version 17 or later
  - c. IOMMU
  - d. Windows 8 or later
8. You are considering whether to use vGPUs for some of the virtual machines in your vSphere environment. Which of the following is not a common use case for vGPUs?
- a. Fast provisioning
  - b. High-end graphics in VDI
  - c. Machine learning
  - d. Artificial intelligence
9. You are setting permissions for a vCenter Server. You want to ensure that a specific user can manage the vCenter Server's content libraries and content but can only view content libraries belonging to other vCenter Servers. Which settings should you make?
- a. Grant the read-only role as a global permission and the administrator role on the vCenter Server.

- b. Grant the content library administrator role as a global permission and the administrator role on the vCenter Server.
  - c. Grant just the administrator role on the vCenter Server.
  - d. Grant just the content library administrator role on the vCenter Server.
10. You want to add items to the content library. Which of the following is not a valid choice?
- a. You can import a vApp.
  - b. You can select a virtual machine and choose Clone to Template in Library.
  - c. You can import an ISO.
  - d. You can migrate a virtual machine to the library.

## Foundation Topics

### Creating and Configuring Virtual Machines

You can use the vSphere Client to create and manage virtual machines. The associated procedures are intuitive. The following sections summarize the procedures and provide some related details, such as the required privileges.

#### Creating a New Virtual Machine



The privileges are required for creating a virtual machine:

- Virtual Machine.Inventory.Create New on the destination folder or data center
- Virtual Machine.Configuration.Add New Disk on the destination folder or data center (when adding a new disk)
- Virtual Machine.Configuration.Add Existing Disk on the destination folder or data center (when adding an existing disk)
- Virtual Machine.Configuration.Configure Raw Device on the destination

folder or data center (when using a Raw Device Mapping (RDM) or SCSI pass-through device)

- Virtual Machine.Configuration.Configure Host USB Device on the destination folder or data center (when attaching a virtual USB device backed by a host USB device)
- Virtual Machine.Configuration.Advanced Configuration on the destination folder or data center (when configuring advanced virtual machine settings)
- Virtual Machine.Configuration.Change Swapfile Placement on the destination folder or data center (when configuring swap file placement)
- Virtual Machine.Configuration.Toggle Disk Change Tracking on the destination folder or data center (when enabling change tracking on the virtual machine's disks)
- Resource.Assign Virtual Machine to Resource Pool on the destination host, cluster, or resource pool
- Datastore.Allocate Space on the destination datastore or datastore folder
- Network.Assign Network on the network where the virtual machine will be connected

To create a virtual machine, you can use the New Virtual Machine wizard and select Create a New Virtual Machine. In the wizard, you should provide all required information, including compute resource (host, cluster, or resource pool, or vApp), storage type and location, virtual machine compatibility, guest OS, Windows virtualization-based security (for a Windows virtual machine), and hardware customization.

When selecting the storage type on a host that has PMem memory, you can select either the Standard or PMem radio button. If you chose PMem storage for a virtual machine, its default virtual disk, new virtual disk, and NVDIMM devices share the same PMem resources. You should adjust the size of newly added devices. The wizard alerts you if issues exist.

## **Powering on a VM**

To power on a virtual machine, in the vSphere client, you can right-click the virtual machine and choose Power On. The following are some likely causes of power-on failures:

- The evaluation period (or license) has expired.
- Permissions are insufficient.
- There is insufficient storage space to create files, such as the swap file.
- The assigned MAC address conflicts with VMware reserved MAC addresses.
- The operation would violate admission control.

## Opening a Console to a VM

To open a console to a virtual machine, you can use an integrated web-based console or the independent VMware Remote Console (VMRC). To use the integrated web-based console, you should ensure that the virtual machine is powered on, select it in the inventory pane, and either choose Launch Web Console in the vSphere Client or Open Browser Console in the vSphere Host Client.

To use the VMRC to access a virtual machine, you should first ensure that it is installed on your local system and, if necessary, prepare a proxy server. Then you can launch it from the vSphere Client or the vSphere Host Client. In the vSphere Client, select the virtual machine in the inventory pane and select Summary > Launch Remote Console. In the vSphere Host Client, select the virtual machine in the inventory pane and select Console > Launch Remote Console.

To configure a proxy server for VMware Remote Console, you can browse to `vmrc://settings` or use the menu if VMware Remote Console is already open. Choose Preferences in the appropriate menu:

- **Windows:** Select VMRC > Preferences.
- **macOS:** Select VMware Remote Console > Preferences.
- **Linux:** Select File > Remote Console Preferences.

The main steps are to select the Enable Proxy for Remote Virtual Machine option and to set the appropriate settings, such as the proxy server's host name or IP (IPv4 or IPv6) address and port, and optionally provide user credentials. The specific steps depend on the OS type (Windows, Linux, or macOS).

## Note

In VMRC Version 11.0, the `VMWARE_HTTPSPROXY` environment variable, which is used to set a proxy server in previous versions of VMRC, is ignored after the previous procedure is applied. To use authentication with the proxy server, you must use the previous procedure instead of the environment variable.

## Installing and Upgrading VMware Tools

You can install VMware Tools in the guest OS of a virtual machines to enable several features that improve manageability and smooth user interaction. To interactively install VMware Tools, in the vSphere Client, you can right-click on a virtual machine, select Guest OS > Install VMware Tools, and select Mount, which connects the virtual machine's first virtual CD-ROM disk drive to the appropriate VMware Tools ISO file, based on the guest operating system. If Autorun is configured in the guest OS, the VMware Tools installation may begin automatically. Otherwise, you might need to interactively launch the installer in the guest OS. For example, you might need to launch `d:\setup.exe` in a Windows 64-bit guest OS. In many cases, the default installation is adequate. If you need non-default components, such as the Guest Introspection Thin Agent driver, select Custom Setup.

The open source implementation of VMware Tools for Linux is Open VM Tools. Whenever a new VMware Tools version is available, such as following an ESXi upgrade, you should consider upgrading your virtual machines. You should always upgrade VMware Tools prior to upgrading the virtual machine hardware. You can use the same procedure as you used to install VMware Tools, except you choose Upgrade VMware Tools.

Previous versions of vSphere allow you to use Update Manager to upgrade virtual machine hardware and VMware Tools. In vSphere 7.0, you can use the vSphere Client directly to upgrade the hardware and VMware Tools for a set of virtual machines in a container, such as a folder or cluster, as described in the “Upgrading Virtual Machines” section in [Chapter 13](#).

VMware Tools Lifecycle Management provides a simplified and scalable approach for installing and upgrading VMware Tools. You can configure your virtual machine to automatically check for and apply VMware Tools upgrades



each time you power on your virtual machine. Automatic Tools upgrade is not supported for Solaris or Netware guests. The prerequisites for automatic VMware Tools upgrade are that the virtual machines must be hosted by ESX/ESXi 3.5 or later, must be managed by vCenter Server 3.5 or later, must be using VMware Tools shipped with ESX/ESXi 3.5 or later, and must be running a guest OS that is supported for ESX/ESXi 3.5 and vCenter Server 3.5 or later.

You can set the `vmx.log.guest.level` option as described in [Table 14-2](#) to control the use of log files for VMware Tools installation.

**Table 14-2** Installer Log Options

Value	Description
<code>vmx.log.guest.level = "off"</code>	Logging to the host is disabled. (This is the default value.)
<code>vmx.log.guest.level = "error"</code>	<code>vminst.log</code> and <code>vmmsi.log</code> remain in the virtual machine and are not sent to the host.
<code>vmx.log.guest.level = "warning"</code>	<code>vminst.log</code> and <code>vmmsi.log</code> remain in the virtual machine and are not sent to the host.
<code>vmx.log.guest.level = "notice"</code>	<code>vminst.log</code> and <code>vmmsi.log</code> remain in the virtual machine and are not sent to the host.
<code>vmx.log.guest.level = "info"</code>	<code>vminst.log</code> is sent to the host, but <code>vmmsi.log</code> remains in the virtual machine.
<code>vmx.log.guest.level = "verbose"</code>	<code>vminst.log</code> and <code>vmmsi.log</code> are sent to the host.
<code>vmx.log.guest.level = "trivia"</code>	<code>vminst.log</code> and <code>vmmsi.log</code> are sent to the host.

When using the `setup.exe` command to install VMware Tools, you can use the `/mg` or `"LOGMODE=G"` options to control and suppress logging to the host. To suppress logging during automatic upgrades, you can set the `install-vmxGuestLogDisabled` parameter to `true` in the `tools.conf` file. To use the `tools.conf` file in some versions of Windows, you might need to create the file and deal with a hidden application data or program data file. To do so, you could open a text editor (such as Notepad) using Run as Administrator. If you change the `tools.conf` file, you do not need to restart VMware Tools. By default, the Tools service checks the configuration file for changes every 5 seconds.

## Shutting Down a Guest

To stop a virtual machine gracefully, in the vSphere client, you can right-click the virtual machine and choose Power > Shutdown Guest. This operation requires that VMware Tools is running in the guest OS. It safely stops the guest OS and powers down the virtual machine. If Shutdown Guest is not available for a virtual machine, a likely cause is that VMware Tools is not installed or not running.

## Cloning a Virtual Machine

You can clone a virtual machine to a template. The following privileges are required for cloning a virtual machine to a template:

- Virtual Machine.Provisioning.Create Template from Virtual Machine on the source virtual machine
- Virtual Machine.Inventory.Create from Existing on the virtual machine folder where the template is created
- Resource.Assign Virtual Machine to Resource Pool on the destination host, cluster, or resource pool
- Datastore.Allocate Space on all datastores where the template is created

To clone a virtual machine to template, in the vSphere client, right-click the virtual machine, select Clone > Clone as Template, and complete the wizard. In the wizard, provide a template name, folder, compute resource, and datastore.

### Note

You cannot change the storage policy if you clone an encrypted virtual machine.

You can clone a virtual machine to create a new virtual machine. The following privileges are required to clone a virtual machine to create a new virtual machine:



- Virtual Machine.Provisioning.Clone Virtual Machine on the virtual machine you are cloning
- Virtual Machine.Inventory.Create from Existing on the data center or virtual machine folder
- Virtual Machine.Configuration.Add New Disk on the data center or virtual machine folder
- Resource.Assign Virtual Machine to Resource Pool on the destination host, cluster, or resource pool
- Datastore.Allocate Space on the destination datastore or datastore folder
- Network.Assign Network on the network to which you assign the virtual machine
- Virtual Machine.Provisioning.Customize on the virtual machine or virtual machine folder (when customizing the guest operating system)
- Virtual Machine.Provisioning.Read Customization Specifications on the root vCenter Server (when customizing the guest operating system).

You can clone a virtual machine to create a new virtual machine by right-clicking the virtual machine and selecting Clone > Clone to Virtual Machine. In the wizard, you should provide all required information, such as name, compute resource, compatibility, and storage. The procedure is much like the procedure in the “Deploying a Virtual Machine from a Template” section, later in this chapter, including the option to customize the guest OS.

### **Note**

You cannot use the vSphere Client to clone a virtual machine using linked clones or instant clones. You can do so with API calls.

If the source virtual machine has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have an available PMem resource. If the virtual machine has virtual PMem hard disks but does not have an NVDIMM device, the destination host or cluster must have an available PMem resource. Otherwise, all hard disks of the destination virtual machine use the storage policy and datastore selected for the configuration files of the source virtual

machine.

## **Converting Between a VM and a Template**

You can easily convert a virtual machine to a template. In the vSphere client, you can right-click a powered-down virtual machine and choose Template > Convert to Template.

You can convert a template to a virtual machine. This is useful when you want to install new software and guest OS updates in the template. To do so, in the vSphere client, right-click the template and select Convert Template to Virtual Machine. You need the following privileges:

- Virtual Machine.Provisioning.Mark as Virtual Machine on the source template
- Resource.Assign Virtual Machine to Resource Pool on the target resource pool

## **Deploying a Virtual Machine from a Template**

The following privileges are required for deploying a virtual machine from a template:

- Virtual Machine.Inventory.Create from Existing on the data center or virtual machine folder
- Virtual Machine.Configuration.Add New Disk on the data center or virtual machine folder (when adding a new virtual disk)
- Virtual Machine.Provisioning.Deploy Template on the source template
- Resource.Assign Virtual Machine to Resource Pool on the target host, cluster, or resource pool
- Datastore.Allocate Space on the target datastore
- Network.Assign Network on the target network (when adding a new network card)
- Virtual Machine.Provisioning.Customize on the template or template folder (when customizing the guest operating system)
- Virtual Machine.Provisioning.Read Customization Specifications on the root vCenter Server (when customizing the guest operating system)

To deploy a virtual machine from template, in the vSphere client, right-click the template and select Clone Deploy from Template. In the wizard, you should provide all required information, such as name, compute resource, compatibility, storage, and guest customization options. The guest customization choices are Select an Existing Specification, Create a Specification, and Create a Specification from an Existing Application.

## Customizing the Guest OS

When you clone a virtual machine to a template or to a new virtual machine, you have the option to customize the guest OS. In addition, other scenarios may allow you to customize a guest OS. This section describes guest OS customization.

You can customize the guest OS to change the computer name, network settings, and guest OS licensing to prevent conflicts in the environment. During a cloning operation, you can provide the customization settings or select a prebuilt customization specification.

Guest OS customization requires a supported guest OS installed on SCSI node 0:0 and VMware Tools. Windows guest customization requires ESXi Version 3.5 or later. Linux guest customization requires Perl in the guest OS. To customize a Linux guest OS, you need to install VMware Tools 10.10.10 or later and enable the enable-custom-scripts option (which is disabled by default).

Optionally, you can create a custom application for vCenter Server to use to generate computer names and IP addresses during guest customization. To do so, create a custom script based on the sample reference script (sample-generate-name-ip.pl) found at <https://kb.vmware.com/s/article/2007557> and configure the associated vCenter Server advanced settings. For example, set config.guestcust.name-ip-generator.program to c:\perl\bin\perl.exe and set config.guestcust.name-ip-generator.arg1 to c:\sample-generate-name-ip.pl.

You can use the following procedure to create a guest customization specification for Linux:



- Step 1.** In the vSphere Client, navigate to **Menu > Policies and Profiles > VM Customization Specifications**.
- Step 2.** Click the **Create a New Specification** icon.

- Step 3.** On the Name and Target OS page, enter a name and a description for the customization specification, select **Linux** as the target guest OS, and click **Next**.
- Step 4.** On the Computer Name page, configure one of the following options in order to assign the computer name:
- **Use the Virtual Machine Name**
  - **Enter a Name in the Clone/Deploy Wizard**
  - **Enter a Name** (For this option, enter a name in the provided box and optionally select the **Append a Numeric Value** checkbox.)
  - **Generate a Name Using the Custom Application Configured with vCenter Server** (For this option, optionally enter a parameter to pass to the application.)
- Step 5.** Enter the domain name and click **Next**.
- Step 6.** Select the time zone and click **Next**.
- Step 7.** On the Customization Script page, optionally provide a script to run in the guest OS and click **Next**.
- Step 8.** On the Network page, choose one of the following options:
- **Use Standard Network Settings:** Select this option to use DHCP to assign IP configuration.
  - **Manually Select Custom Settings:** Select this option to have vCenter Server prompt the user to provide the IP configuration for each virtual NIC when using the guest customization specification.
- Click **Next**.
- Step 9.** On the DNS Settings page, enter the DNS server and domain settings.
- Step 10.** Complete the wizard and click **Finish**.



To create a guest customization specification for Windows, you can use the previous procedure with the following modifications:

- On the Name and Target OS page, select Windows as a target guest OS and

optionally select Generate a New Security Identity (SID).

- On the Set Registration Information page, enter the virtual machine owner's name and organization and click Next.
- On the Windows License page, provide a Windows product key. For a Windows Server specification, either select the Per Seat option or configure the maximum concurrent connections for the Per Server option. Click Next.
- On the Set Administrator Password page, configure the password, optionally select Automatically Logon as Administrator option, and click Next.
- On the Networking page, if you choose the Manually Select Custom Settings, use the DNS tab to provide DNS server details and click WINS to provide WINS details.
- On the Set Workgroup or Domain page, either provide a workgroup name or provide user credentials and a domain name and click Next.

Whenever you create a new virtual machine by deploying from template or by cloning, you can use the wizard to select the Customize the Operating System checkbox and select the appropriate specification. To customize an existing virtual machine, in the vSphere client, right-click the virtual machine in the inventory pane, select Guest OS > Customize Guest OS, and select the appropriate specification.

As needed, you can manage guest customization specifications by navigating to Menu > Policies and Profiles > VM Customization Specifications. Here, you can import guest customization specifications. You can also select a particular specification and select one of the following actions:

- Edit Customization Spec
- Duplicate Customization Spec
- Export Customization Spec
- Delete Customization Spec

## Deploying OVF/OVA Templates

Another method for deploying virtual machines is to leverage *Open Virtual Format (OVF) templates* or *Open Virtual Appliance (OVA) templates*. You can use the vSphere Client to deploy an OVF or OVA template. You can export a

virtual machine, virtual appliance, or vApp as an OVF or OVA template to create virtual appliances that can be imported by other users. Compared to other methods, using an OVF template to export and import virtual machines provides the following benefits:

- Compressed data (which means faster downloads and uploads)
- Validation of the OVF by vCenter Server prior to importing
- Encapsulation of multiple virtual machines

OVA is essentially a single-file distribution of an OVF package. Prior to vSphere 6.5, the Client Integration Plug-In is required to export and import OVF and OVA templates. Starting in vSphere 6.5, you can only export to OVF. Deploying a virtual machine from an OVF template is commonly referred to as *deploying an OVF*.

To deploy an OVF, you can use the following procedure:

- Step 1.** In the vSphere Client, right-click on a cluster in the inventory pane and select **Deploy OVF Template**.
- Step 2.** On the Select OVF Template page, specify the path to the OVF file as a URL or local file and click **Next**.
- Step 3.** Use the wizard to provide information for the new virtual machine, such as name, folder, and compute resource.
- Step 4.** On the Review Details page, verify the OVF template details, such as publisher, download size, and size on disk. Click **Next**.
- Step 5.** Complete the wizard by providing typical details for a new virtual machine, such as storage policy, storage location, and network configuration.
- Step 6.** Optionally, customize the deployment properties on the Customize Template page.
- Step 7.** Optionally, select a binding service provider on the vService Bindings page.
- Step 8.** On the Ready to Complete page, click **Finish**.

## Managing Virtual Machines



This section covers reoccurring activities that you might perform regarding virtual machines.

## Configuring Virtual Machine Hardware

When creating or upgrading a virtual machine, you can configure the virtual machine compatibility setting, which controls the ESXi versions on which the virtual machine can run. The compatibility setting controls which virtual machine hardware version is used. The main use cases for configuring the compatibility setting to a version earlier than the default for the host are to maintain compatibility with older hosts and to standardize virtual machine deployment in the environment. The main downside of configuring the compatibility setting to a version earlier than the default for the host is that the virtual machine may not be able to use virtual hardware features supported by the host and may not achieve the best performance. [Table 14-3](#) describes virtual machine compatibility options.

**Table 14-3** Virtual Machine Compatibility Options

Compatibility Setting	Hardware Version
ESXi 7.0 and later	Hardware Version 17
ESXi 6.7 Update 2 and later	Hardware Version 15
ESXi 6.7 and later	Hardware Version 14
ESXi 6.5 and later	Hardware Version 13
ESXi 6.0 and later	Hardware Version 11
ESXi 5.5 and later	Hardware Version 10
ESXi 5.1 and later	Hardware Version 9
ESXi 5.0 and later	Hardware Version 8
ESX/ESXi 4.0 and later	Hardware Version 7
ESX/ESXi 3.5 and later	Hardware Version 4

### Note

ESXi 5.0 allows you to run virtual machines with ESX/ESXi 3.5 and later compatibility (hardware Version 4) but does not allow you to create them.

The compatibility setting impacts the supported features for the virtual machine. [Table 14-4](#) lists some of the feature sets available in recent hardware versions.

**Table 14-4** Features by Recent Virtual Machine Hardware Versions

<b>Feature</b>	<b>Versio n 17</b>	<b>Versio n 15</b>	<b>Versio n 14</b>	<b>Versio n 13</b>	<b>Versio n 11</b>
Maximum memory (GB)	6128	6128	6128	6128	4080
Maximum number of logical processors	256	256	128	128	128
Maximum number of cores (virtual CPUs) per socket	64	64	64	64	64
NVMe controllers	4	4	4	4	No
Maximum NICs	10	10	10	10	10
USB 3.1 SuperSpeedPlus	Yes	No	No	No	No
Maximum video memory (GB)	4	2	2	2	2
Dynamic DirectPath	Yes	No	No	No	No
PCI hot adding support	Yes	Yes	Yes	Yes	Yes
Virtual precision clock device	Yes	No	No	No	No
Virtual watchdog timer device	Yes	No	No	No	No
Virtual SGX device	Yes	No	No	No	No
Virtual RDMA	Yes	Yes	Yes	Yes	No
NVDIMM controller	1	1	1	No	No
NVDIMM devices	64	64	64	No	No
Virtual I/O MMU	Yes	Yes	Yes	No	No
Virtual TPM	Yes	Yes	Yes	No	No
Microsoft VBS	Yes	Yes	Yes	No	No

To control the default hardware compatibility for new virtual machines, you can set the Default VM Compatibility setting at the host, cluster, or data center level. The settings on a host override the settings on a cluster, which override the

settings on the data center. To make the settings on a host or cluster, you must have the Host.Inventory.Modify Cluster privilege. To make the settings on a data center, you must have the Datacenter.Reconfigure Datacenter privilege.

You can upgrade the compatibility level of an existing virtual machine but should first upgrade VMware Tools. For example, you can select a virtual machine and use the Compatibility > Schedule VM Compatibility Upgrade option to upgrade the compatibility the next time you restart the virtual machine. Optionally, you can select Only Upgrade After Normal Guest OS Shutdown to upgrade compatibility during regularly scheduled guest maintenance.

You can change the number of virtual CPUs used by a virtual machine. Specifically, you can set the number of cores and the number of cores per socket. In ESXi 7.0, the maximum number of virtual CPU sockets is 128. To configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

By default, you cannot add CPU resources to a virtual machine when it is turned on. To change this behavior, you can enable the virtual machine's CPU hot adding option, but the following conditions apply:

**Key  
Topic**

- For best results, set virtual machine compatibility to ESXi 5.0 or later.
- Hot adding multicore virtual CPUs requires compatibility set to ESXi 5.0 or later.
- You cannot use hot adding to increase the number of virtual CPUs for a virtual machine with 128 or fewer virtual CPUs.
- You can use hot adding to increase the number of virtual CPUs for a virtual machine that already has more than 128 virtual CPUs.
- You can disable hot adding for virtual machines with guest operating systems that do not support CPU hot adding.
- For virtual machines with compatibility set to ESXi 4.x and later, to support CPU hot adding, set Number of Cores per Socket to 1.
- Hot adding CPU resources to a virtual machine disconnects and reconnects all USB passthrough devices.

To enable CPU hot adding, the following prerequisites apply:

- The latest VMware Tools version must be installed.
- The guest operating system must support CPU hot adding.
- Virtual machine compatibility must be set to a minimum of ESX/ESXi 4.x or later.
- The virtual machine must be turned off.
- The Virtual Machine.Configuration.Settings privilege is required.

CPU identification (CPU ID) masks control the visibility of CPU features to the guest OS. Masking CPU features can impact a virtual machine’s availability for migration using vMotion. For example, if you mask the AMD No eXecute (NX) or the Intel eXecute Disable (XD) bits, you prevent the virtual machine from using those features, but you allow the virtual machine to hot migrate to hosts that do not include this capability.

### Note

Changing the CPU compatibility masks can result in an unsupported configuration. Do not manually change the CPU compatibility masks unless instructed to do so by VMware Support or a VMware Knowledge Base article.

During specific management operations, such as creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine, you can set the provisioning policy for the virtual disk file. You can use Storage vMotion or other cross-datastore migrations to transform virtual disks from one format to another. The available virtual disk provisioning policies are described in [Table 14-5](#).

**Table 14-5** Virtual Disk Provisioning Policies

Provisioning Policy	Description	Sample Use Case
Thick provisioned lazy	All the provisioned space is allocated during creation. Data is zeroed on the first demand from the virtual machine.	Minimizes the risk of exhausting free datastore space.

zeroed

Thick provisioned eager zeroed	All the provisioned space is allocated and zeroed (erased) during creation.	Required for some clustering features, such as Fault Tolerance. Provides the best performance.
Thin provisioned	Space is allocated and zeroed on demand, as needed, up to the provisioned space.	Fastest method to provision and migrate virtual disks.

You can change a virtual disk from the thin format to thick format by navigating to Datastore > Files in the vSphere Client and choosing the Inflate action for the virtual disk file. The vSphere Client does not provide a deflate option. To change a virtual disk provisioning type from thick to thin, you can migrate the virtual machine storage and select the appropriate policy.

Creating and growing a virtual disk provisioned for thick provisioned eager zeroed may take significantly longer than with a virtual disk provisioned for thick provisioned lazy zeroed.

You can configure virtual machines with virtual disks greater than 2 TB (large-capacity virtual disks), but you must meet resource and configuration requirements. The maximum size for large-capacity virtual disks is 62 TB. You should avoid using the maximum size because some operations, such as those involving snapshots and linked clones, may not finish when the maximum amount of disk space is allocated to a virtual disk. Operations such as snapshot quiesce, cloning, Storage vMotion, and vMotion in environments without shared storage can take significantly longer to finish. The following conditions and limitations apply to virtual machines with large-capacity disks:

- You must use a guest OS that supports large-capacity virtual hard disks.
- Target hosts for migration and cloning operations must use ESXi 6.0 or later.
- NFS, vSAN, and VMFS Version 5 or later datastores are supported.
- Fault Tolerance is not supported.
- BusLogic Parallel controllers are not supported.

To increase the size of a virtual disk, you need the following privileges:

- Virtual Machine.Configuration.Modify Device Settings on the virtual machine
- Virtual Machine.Configuration.Extend Virtual Disk on the virtual machine
- Datastore.Allocate Space on the datastore

To control how a virtual disk is impacted by snapshots, you can set the disk mode for a virtual disk to the settings described in [Table 14-6](#).

**Table 14-6** Virtual Disk Mode Settings

<b>Disk Mode</b>	<b>Description</b>
Dependent	Included in snapshots.
Independent–Persistent	Not included in snapshots.  All data written is written permanently to disk.
Independent–Nonpersistent	Not included in snapshots.  Changes are discarded when you turn off or reset the virtual machine.

You can set shares for a virtual disk, and they work much like CPU or memory shares for a virtual machine. The disk shares provide a relative priority for accessing the disk during periods of disk I/O contention for the underlying storage. The values Low, Normal, High, and Custom are compared to the sum of all shares of all virtual machines on the host. To control the maximum amount of disk I/O for a virtual disk, you can set the virtual disk’s Limit–IOPS value. By default, the virtual disk is set for normal shares and unlimited IOPS.

You can add virtual disks to virtual machines, including new virtual disks, existing virtual disks, and raw device mappings (RDMs). To add an RDM to a virtual machine, you need to use an account with the Virtual Machine.Configuration.Configure Raw Device privilege, select a target LUN, choose where to store the mapping file, choose a compatibility mode (physical or virtual), and select a disk mode. Disk modes are not available for RDMs using physical compatibility mode.

A storage controller is included by default when you create a virtual machine. You can add additional SCSI controllers (BusLogic Parallel, LSI Logic Parallel,

LSI Logic SAS, and VMware Paravirtual SCSI), AHCI, SATA, and NVMe Express (NVMe) controllers. The following limitations apply to storage controllers:

- ESXi 4.x and later compatibility is required for LSI Logic SAS and VMware Paravirtual SCSI.
- ESXi 5.5 and later compatibility is required for AHCI SATA
- ESXi 6.5 and later compatibility is required for NVMe
- BusLogic Parallel controllers do not support large-capacity disks.
- Disks on VMware Paravirtual SCSI controllers may not provide the expected performance if they have snapshots or if the host's memory is overcommitted.

Before changing the storage controller type, you should ensure that the guest OS has the drivers for the target controller type, or the disks will become inaccessible. Likewise, in the following cases, adding storage controller types to a virtual machine that uses BIOS firmware may cause boot problems and require you to fix the issue by entering the BIOS setup:

- If the virtual machine boots from LSI Logic SAS or VMware Paravirtual SCSI, and you add a disk that uses BusLogic, LSI Logic, or AHCI SATA controllers
- If the virtual machine boots from AHCI SATA, and you add BusLogic Parallel or LSI Logic controllers

### **Note**

Adding additional disks to virtual machines that use EFI firmware does not cause boot problems.

## **Editing Virtual Machine Options**

You can edit a virtual machine and use the VM Options tab for multiple purposes, such as setting VMware Tools scripts, controlling user access to the remote console, configuring startup behavior, and changing the virtual machine name, as summarized in [Table 14-7](#).

**Table 14-7** Settings on the Virtual Machine Options Tab

<b>Options</b>	<b>Description</b>
General Options	You can set the virtual machine name. In addition, you can view the configuration file location, working location, and guest OS type and version.
VMware Remote Console Options	You can set the locking behavior and simultaneous connections.
Encryption	You can adjust the encryption settings.
Power Management	You can configure the suspend virtual machine behavior.
VMware Tools	You can set the VMware Tools scripts behavior, VMware Tools upgrade settings, and guest OS time synchronization settings.
Boot Options	You can set the boot delay and force entry into the BIOS or EUFI setup screen.
Advanced	You can adjust the acceleration and logging settings, debugging and statistics, swap file location, and latency sensitivity.
Fibre Channel NPIV	You can set the virtual node and port worldwide names (WWNs).

A virtual machine name must be unique within the folder where the virtual machine is located. If you move a virtual machine to a different datastore folder or host that already has a virtual machine of the same name, you must change the virtual machine's name to keep it unique. Changing a virtual machine name impacts how the virtual machine is identified by vCenter Server and does not impact file (or folder) names or the guest OS. After changing a virtual machine name, you can leverage Storage vMotion to migrate the virtual machine, which renames the associated files to match the new virtual machine name.

You can encrypt a virtual machine by editing its storage policies or by editing VM Options. Before encrypting a virtual machine, you must meet the following prerequisites:



**Key  
Topic**

- Establish a trusted connection with the KMS and select a default KMS.
- Create an encryption storage policy (or plan to use the sample VM encryption policy).
- Ensure that the virtual machine is powered off.
- Verify that you have the required privileges:
  - Verify that you have the Cryptographic Operations.Encrypt New Privilege.
  - If the host encryption mode is not enabled, verify that you have the Cryptographic Operations.Register Host privilege.

You can use the following procedure to encrypt a virtual machine:

**Step 1.** In the vSphere Client, right-click the virtual machine in the inventory pane and navigate to **VM Policies > Edit VM Storage Policies**.

**Step 2.** Use one of the following methods:

- Select an encryption storage policy to apply to the virtual machine and its virtual disks and click **OK**.
- Click **Configure per Disk**, select the encryption storage policy for VM, select an encryption or other storage policies for each virtual disk, and click **OK**.

**Step 3.** Optionally, encrypt the virtual machine, or both virtual machine and disks, from the Edit Settings menu in the vSphere Client:

- a. Right-click the virtual machine and select **Edit Settings**.
- b. Select the **VM Options** tab and open **Encryption**. Choose an encryption policy. If you deselect all disks, only the VM home is encrypted.

**Step 4.** Click **OK**.

## Configuring Guest User Mappings

You can enable guest OS access for some of your SSO user account to facilitate

some administrative tasks, such as upgrading VMware Tools. You can use the following procedure to enroll SSO users to user accounts in guest operating systems by using SSO certificates. Subsequent guest management requests use an SSO SAML token to log in to the guest OS:

**Step 1.** In the vSphere Client, select a powered-on virtual machine in the inventory pane and navigate to **Configure > Guest User Mappings**.

**Step 2.** Enter your username and password and click **Log In**.

**Step 3.** In the Guest User Mappings pane, click the **Add** button.

**Step 4.** In the dialog box, select the SSO user.

**Step 5.** Specify a guest OS username and click **OK**.

## Editing OVF Details

You can use the following procedure to edit a virtual machine's OVF settings to customize the OVF environment, OVF transport, and boot behavior after OVF deployment. This information is preserved when you export the virtual machine as an OVF template:

**Step 1.** In the vSphere Client, select a virtual machine in the inventory pane and navigate to **Configure > Settings > vApp Options**.

**Step 2.** Click the **Edit** button.

**Step 3.** If vApp options are not enabled, select the **Enable vApp Options** checkbox.

**Step 4.** Click the **OVF Details** tab.

**Step 5.** Set the OVF Environment Transport option to one of the following:

- **ISO Image:** Mounts an ISO image with the OVF template to the CD-ROM drive.
- **VMware Tools:** Initializes the guestInfo.ovfEnv variable with the OVF environment document.

**Step 6.** Optionally, enable the Installation Boot option and delay time in seconds to automatically reboot the virtual machine after OVF deployment.

**Step 7.** Click **OK**.

## Creating and Managing Virtual Machine Snapshots

You can use virtual machine snapshots to capture the state and data of a virtual machine at a specific point in time. A snapshot preserves the following information:

- Virtual machine settings
- Power state
- Disk state
- Memory state (optional)

To take a snapshot, in the vSphere client, you can right-click a virtual machine, select Snapshots > Take Snapshot, and provide a snapshot name. Optionally, you can provide a snapshot description and select Snapshot the Virtual Machine's Memory. Also, you can optionally choose Quiesce Guest File System. Quiescing the file system requires the virtual machine to be powered on, VMware Tools to be running, and Snapshot the Virtual Machine's Memory to be deselected.

### Note

To minimize the impact to a running virtual machine and to reduce the time required to take a snapshot, do not snapshot the memory state or quiesce the guest file system.

After creating a snapshot, you can use the Snapshot Manager to view the snapshot hierarchy of the virtual machine, which appears as a tree with branches, as illustrated in [Figure 5-2 in Chapter 5, “vCenter Server Features and Virtual Machines.”](#) To open the Snapshot Manager, in the vSphere client, you can right-click the virtual machine and choose Snapshots > Manage Snapshots. In the Snapshot Manager, the snapshot that appears above the You Are Here icon is the parent snapshot. If you revert to a snapshot, that snapshot becomes the parent snapshot. If you take a snapshot of a virtual machine that already has at least one snapshot, the new snapshot is a child of the parent snapshot.

To revert a virtual machine to a specific snapshot, select the snapshot in the

Snapshot Manager for the virtual machine and select Revert To. To do this, you must have the Virtual Machine.Snapshot Management.Revert to Snapshot privilege on the virtual machine.

When you revert the virtual machine to a snapshot, you return its virtual disks and settings to the state captured in the snapshot. If the snapshot includes the memory state, reverting to the snapshot returns the virtual machine's memory to that state. You can revert the virtual machine to any available snapshot in the Snapshot Manager. Subsequent snapshots from this point create a new branch of the snapshot tree. When you revert to a snapshot, no snapshots are removed, but you lose the virtual machine's current disk state. In other words, all changes to disk data made since the last snapshot are permanently lost. If you revert to a snapshot without memory state, the virtual machine is in the powered-off state.

You can delete a snapshot for a running virtual machine without disrupting its end users. Deleting a snapshot removes your ability to revert to that snapshot's state in the future. To delete a specific snapshot, select the snapshot in the Snapshot Manager for the virtual machine and select Delete. Optionally, to delete all snapshots, select Delete All.

If the virtual machine is in a state where it has no snapshots but has one or more delta disks contributing to the active state of the virtual machine, then the vSphere Client may provide a "Consolidation Needed" warning. In this state, you can right-click the virtual machine in the vSphere client and select Snapshots > Consolidate. The system merges the data from delta disks into the base disks and delete the delta disks. In normal conditions, your virtual machine is in a state where the Consolidate option is not available.

## Migrating Virtual Machines



To migrate a virtual machine by using the vSphere Client, you can right-click the virtual machine in the inventory pane, choose Migrate, and complete the wizard. The details for completing the wizard depend on the migration type. The required privileges for each migration type are covered in [Chapter 5](#). You can use the Recent Tasks pane to monitor the progress of your migration.

To cold migrate a virtual machine, you can use the following procedure:

**Step 1.** In the vSphere Client, right-click a powered-off virtual machine and select **Migrate**.

**Step 2.** Select one of the following migration types:

- **Change Compute Resource Only**
- **Change Storage Only**
- **Change Both Compute Resource and Storage**
- **Migrate Virtual Machine(s) to a Specific Datacenter**

Click **Next**.

**Step 3.** If you select a migration type that includes a cross-host migration, select the destination compute resource (**Host, Cluster, Resource Pool, or vApp**), verify that no issues exist in the Compatibility panel, and click **Next**.

**Step 4.** If you select a migration type that includes a cross-datastore migration, select the virtual disk format (**Same as Source, Thin Provisioned, Thick Provisioned Lazy Zeroed, or Thick Provisioned Eager Zeroed**), select the appropriate policy in the VM Storage Policy menu, and select the destination, as described here:

- To store all the virtual machines in a datastore, select the datastore and click **Next**.
- To store all the virtual machines in a Storage DRS cluster, select the cluster and click **Next**.
- To store the virtual machine configuration files and virtual disks in separate locations, click **Advanced** and configure the destination for the configuration files and each virtual disk. Click **Next**.

**Step 5.** For cross-host migrations, select the destination network for the virtual machines and click **Next**. Alternatively, you can click **Advanced** to assign separate networks to individual virtual machine network adapters.

**Step 6.** Click **Finish**.

To perform a hot cross-host (vMotion) migration, you can apply the previous cold migration procedure, with the following changes:

- Start with a powered-on virtual machine.

- Select to change the compute resource only.
- You are not prompted to select a destination datastore.
- Select either Schedule vMotion with High Priority or Schedule Regular vMotion.

To perform a hot cross-data store (Storage vMotion) migration, you can apply the previous cold migration procedure, with the following changes:

- Start with a powered-on virtual machine.
- Select to change storage only.
- You are not prompted to select a destination host.

To perform a hot cross-host and cross-data store (vMotion without shared storage) migration, you can apply the previous cold migration procedure, with the following changes:

- Start with a powered-on virtual machine.
- Select to change both the compute resource and storage.
- Select either Schedule vMotion with High Priority or Schedule Regular vMotion.

## **Advanced Virtual Machine Management**

This section covers topics related to the configuration and management of virtual machines that are not covered elsewhere in the book.

### **Managing OVF Templates**

To export a virtual machine into a self-contained OVF template, you can select the virtual machine and select Actions > Template > Export OVF Template. You must have the vApp.Export privilege. In the export wizard, you must provide a name and can optionally provide a description and configure advanced options. In the advanced options, you can include details concerning BIOS, UUID, MAC address, boot order, PCI slots, and other settings.

You can browse the VMware Virtual Appliance Marketplace to discover and download virtual appliances provided by VMware and VMware partners. The cost and licensing for each appliance are controlled by the provider.

## Virtualization-Based Security

Starting with vSphere 6.7, you can enable **Microsoft virtualization-based security (VBS)** on supported Windows guest operating systems. VBS is a Microsoft feature for Windows 10 and Windows Server 2016 operating systems that use hardware and software virtualization to enhance system security by creating an isolated, hypervisor-restricted, specialized subsystem. Windows typically uses hashed credentials stored in memory, including Active Directory credentials, that may be subject to the pass-the-hash exploit. In VBS, you can enable a feature called Credential Guard that keeps account of hash information outside the memory of the Windows instance, mitigating pass-the-hash. If the hardware TPM chip is not available or is not enabled in the BIOS, Windows still uses VBS, and you can still enable Credential Guard, but the credentials are not as secure.

On a traditional (non-virtual) Windows server, to prepare for VBS, you should ensure that its BIOS, firmware, and operating system are set to use UEFI firmware, Secure Boot, hardware virtualization (Intel VT/ADM-V), and IOMMU. You can enable VBS in the Windows operating system. When you reboot Windows, the Microsoft hypervisor loads and leverages virtualization to bring up additional Windows components, including the credential management subsystem, in a separate memory space. All subsequent communications between Windows and Windows components are via RPC calls run through a Microsoft hypervisor-based communications channel.

In vSphere, to use VBS, you must use virtual hardware Version 14 or later. The virtual machine must be set to use UEFI firmware, Secure Boot, hardware virtualization (Intel VT/ADM-V), and IOMMU. In the virtual machine settings, enable the Virtualization Based Security checkbox on the VM Options tab. Finally, you must enable VBS by editing the group policy.

Enabling VBS for a virtual machine does not automatically enable virtual TPM, but you can add a virtual TPM device. A virtual TPM doesn't have a hardware-based vault. Instead, the data that it secures is written to the NVRAM file, which is encrypted using VM encryption, providing strong encryption and virtual machine portability.

## Managing VMs by Using PowerCLI

**VMware PowerCLI** is a command-line and scripting tool built on Windows PowerShell that provides cmdlets for managing and automating VMware

products, including vSphere. You can install PowerCLI on a workstation or server in your vSphere environment and use PowerCLI to automate some aspects of your virtual machine management.

The main prerequisites for installing PowerCLI 12.0 on a Windows system are the presence of .NET Framework 4.7.2 or later and Windows PowerShell 5.1. For Linux and macOS systems, the requirements are .NET CORE 3.1 and PowerShell 7. The main steps to install PowerCLI are to download the product to the system and run the following command in the PowerShell console:

[Click here to view code image](#)

```
Install-Module VMware.PowerCLI -Scope CurrentUser
```

In many cases, you need to change the execution policy, which by default is set to the most secure policy (Restricted). For example, to change the policy to RemoteSigned, you can use the following command:

[Click here to view code image](#)

```
Set-ExecutionPolicy RemoteSigned
```

The Connect-VIServer cmdlet allows you to connect to a vCenter Server. The Get-VM cmdlet allows you to collect information about virtual machines. You can use the following commands to connect to a vCenter Server named server1.vsphere.local (using the [administrator@vsphere.local](#) account and the password VMware1!) and display information for all of its managed virtual machines:

[Click here to view code image](#)

```
Connect-VIServer -Server server1.vsphere.local -Protocol http  
-User 'administrator@vsphere.local' -Password 'VMware1!'
```

```
Get-VM
```

To start a virtual machine named win-01, you can use the following commands:

```
Get-VM win-01 | Start-VM
```

You can use PowerCLI to create virtual machines from specifications provided in an XML file. The XML content could provide detailed specifications for multiple virtual machines. For example, you can use the following sample XML content, which represents the minimum specifications for two virtual machines named MyVM1 and MyVM2, each having a 100 GB virtual disk:



[Click here to view code image](#)

```
<CreateVM>
<VM>
  <Name>MyVM1</Name>
  <HDDCapacity>100</HDDCapacity>
</VM>
<VM>
  <Name>MyVM2</Name>
  <HDDCapacity>100</HDDCapacity>
</VM>
</CreateVM>
```

If you save the sample content to a file named MyVMs.xml, you can use the following commands to read the file, parse the XML content into a variable, and create a virtual machine based on each specification:

[Click here to view code image](#)

```
[xml]$s = Get-Content myVM.xml
$s.CreateVM.VM | foreach {New-VM -VMHost $vmHost1 -Name $_.Name
-DiskGB $_.HDDCapacity}
```

You can use PowerCLI to migrate virtual machines. Consider a scenario where you need to automate frequent, massive migrations of virtual machines between data-stores to prepare for storage array upgrades. At the lowest level, you need a command that migrates a virtual machine to a specified datastore. For example, you can use the following command to migrate a virtual machine named MyVM1 to a datastore named DS2:

[Click here to view code image](#)

```
Get-VM MyVM1 | Move-VM -Datastore DS2
```

## Configuring VMs to Support vGPUs



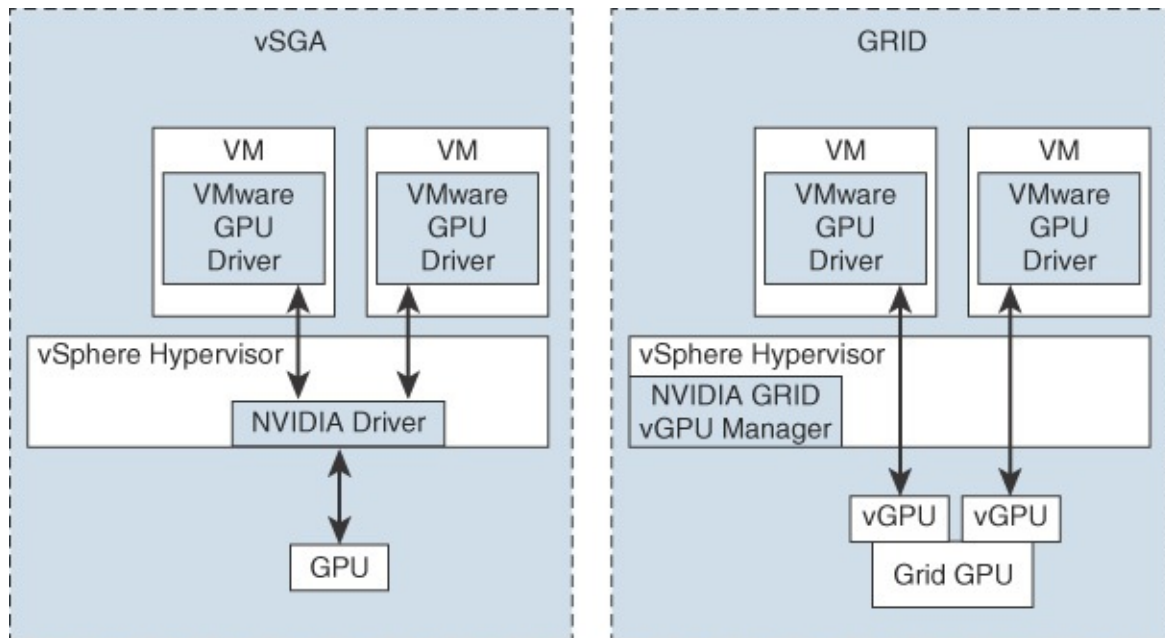
In vSphere 7.0, you can enable virtual machines to use the processing power of available **graphics processing units (GPUs)**. GPUs are specialized processors developed for parallel processing, primarily for rendering graphical images. In vSphere, the main use case for a GPU is to support high-end graphics in virtual desktop infrastructure (VDI). Recently, the need to support artificial intelligence (AI) and machine learning (ML) has also emerged as a major use case.

You can use GPUs in different manners in a vSphere environment. For AI/ML use cases, the GPU configuration choice is mostly impacted by the size and complexity of the problem being solved. For VDI, the GPU configuration choice is impacted by the end user’s graphics needs. The configuration involves either sharing GPUs with multiple virtual machines or dedicating some GPUs to specific virtual machines. [Table 14-8](#) summarizes the potential GPU configuration for specific AI/ML use cases.

**Table 14-8** Use Cases and GPU Configurations

<b>GPU Configuration</b>	<b>Sample Use Cases</b>	<b>Details</b>
GPU sharing	ML development and testing	Good fit for small problems and for the ML inference phase
Dedicated GPU	Data science	Commonly used for development and training in ML models
Dedicated multiple GPUs per VM	Advanced power ML users tackling large problems	Highest-performing GPU model

For the VMware Horizon 7 VDI use case, depending on your hardware, you may have multiple options for sharing GPUs. For example, with NVIDIA hardware, you can choose to share GPUs using the NVIDIA vGPU (GRID) technology or the Virtual Shared Graphics Acceleration (vSGA) technology. In the vSGA model, the vSphere hypervisor presents a virtual VMware SVGA 3D GPU to each virtual machine. In the GRID model, each hardware GPU presents multiple virtual GPUs that the hypervisor passes through to the virtual machines. In the GRID model, you can use a vGPU profile to assign a portion of the GPU hardware to a virtual machine. The vSGA model tends to be flexible and cost-effective for supporting virtual desktops running office, video, and 2D CAD applications. But the performance of the GRID model may be preferred for virtual desktops running 3D modeling software. For a side-by-side comparison of the vSGA and GRID models, see [Figure 14-1](#).



**FIGURE 14-1** Comparison of the vSGA and GRID Models

The procedure to configure the GPU hardware, ESXi host, and virtual machine depends on your choice for GPU configuration. For example, you can use the following procedure to implement the GRID model using a vGPU profile (named grid\_p100-8a) to allow a virtual machine to use up to 8 GB of the GPU's memory:

- Step 1.** Obtain the NVIDIA vGPU software and license.
- Step 2.** In the vSphere Client, select the ESXi host in the inventory pane and navigate to **Configure > Hardware > Graphics > Host Graphics**.
- Step 3.** Click **Edit**.
- Step 4.** Select the **Shared Direct (vendor shared passthrough graphics)** option.
- Step 5.** Reboot the host and enter Maintenance Mode.
- Step 6.** In the ESXi Shell, enter the following command but replace the path with the actual path to the downloaded VIB file:

[Click here to view code image](#)

```
esxcli software vib install -v /vmfs/volumes/ARL-ESX14-DS1/
NVIDIA/NVIDIA-VMware_ESXi_6.7_Host_Driver_390.42-10EM.
670.0.0.7535516.vib
```

- Step 7.** Exit Maintenance Mode.
- Step 8.** Edit the virtual machine settings and select the option to add a new device.
- Step 9.** From the New PCI device drop-down, select **NVIDIA GRID vGPU**.
- Step 10.** From the GPU Profile dropdown, select an appropriate profile, such as **grid\_p100-8a**.
- Step 11.** For the virtual machine guest OS, install the appropriate NVIDIA vGPU driver.

## Content Libraries

This section provides details for implementing and using content libraries to provide templates, ISOs, and other content across multiple vCenter Servers in a vSphere environment.

### Introduction to Content Libraries

A **content library** is a container object for virtual machine templates, vApp templates, ISO images, and other files that you may want to share among multiple vCenter Servers in a vSphere environment. Content libraries allow you to share templates and other files in a manner that provides consistency, compliance, efficiency, and automation when deploying workloads at scale.

A content library contains and manages content in the form of library items. A single library item consists of one file or multiple files. For example, an OVF template is a set of files with the OVF, VMDK, and MF file extensions. When you upload an OVF template to the library, you upload the entire set of files, which the library represents as a single item.

When creating a content library, you can choose to create a local content library or a subscribed content library. With a local library, you store and manage content in a single vCenter Server instance. After creating a local library, you can publish it to make it available for subscription. From another vCenter Server instance, you can create a subscribed content library with a subscription to the published library. With a subscribed library, you can control when to download the subscribed content, either immediately or as needed.

Historically, content libraries supported OVF templates but not standard virtual machine templates. Starting with vSphere 6.7 Update 1, content libraries support

virtual machine templates in addition to OVF templates.

## Creating a Content Library

To create a content library, you must have one of the following privileges on the vCenter Server instance:

- Content Library.Create Local Library
- Content Library.Create Subscribed Library

In addition, you must have the Datastore.Allocate Space privilege on the target datastore.

You can use the following procedure to create a content library:

**Step 1.** In the vSphere Client, select **Menu > Content Libraries** and click the **Create a new content library** icon.

**Step 2.** On the Name and Location page, enter a name and select a vCenter Server instance for the content library. Click **Next**.

**Step 3.** On the Configure Content Library page, select the type of content library that you want to create:

- **Local Content Library:** By default, a local content library is accessible only in the vCenter Server instance where you create it. Optionally, you can select **Enable Publishing** to make the content of the library available to other vCenter Server instances.
- **Subscribed Content Library:** This option creates a content library that subscribes to a published content library.

Click **Next**.

**Step 4.** On the Add Storage page, select a storage location for the content library contents and click **Next**.

**Step 5.** On the Ready to Complete page, review the details and click **Finish**.

## Publishing a Content Library

You can publish an existing content library. For example, to publish an existing local, non-subscribed library, you can use the following procedure:

- Step 1.** Use the vSphere Client to navigate to **Content Libraries**.
- Step 2.** Right-click on an existing content library, select **Edit Settings**.
- Step 3.** Select the Enable Publishing checkbox.
- Step 4.** Click the **Copy Link** button to copy the URL of your library that you can paste into the settings of a subscribed library.
- Step 5.** Select **Enable User Authentication** for access to this content library and set a password for the library.
- Step 6.** Click **OK**.

### Note

When you enable authentication for the content library, you effectively set a password on the static username vcspx, which you cannot change. This is a user account that is not associated with vCenter Single Sign-On or Active Directory.

## Subscribing to a Content Library

When using the previous procedure to create a subscribed content library, you must provide the following information:



- Step 1.** In the **Subscription URL** text box, enter the URL address of the published library.
- Step 2.** If authentication is enabled on the published library, select **Enable Authentication** and enter the publisher password.
- Step 3.** Select a download method for the contents of the subscribed library: **Immediately** or **When Needed**.
- Step 4.** If prompted, accept the SSL certificate thumbprint. The SSL certificate thumbprint is stored on your system until you delete the subscribed content library from the inventory.

## Note

The transfer service on the vCenter Server is responsible for importing and exporting content between the subscriber and the publisher, using HTTP NFC.

## Content Library Permissions

Content libraries are not direct children of the vCenter Server object in the vSphere inventory. Instead, content libraries are direct children of the global root. This means that permissions set on a vCenter Server do not apply to content libraries, even if they are set to propagate to child objects. To assign a permission on a content library, an administrator must grant the permission to the user as a global permission. Global permissions support assigning privileges across solutions from a global root object.

Consider the following scenarios:

- If a user is granted the read-only role as a global permission and the administrator role at a vCenter Server level, the user can manage the vCenter Server's content libraries and content but can only view content libraries belonging to other vCenter Servers.
- If a user is granted the content library administrator role as a global permission, the user can manage all content libraries and content in all vCenter Server instances.
- If a user is not granted any global permission but is granted the administrator role at a vCenter Server level, the user cannot view or manage any libraries or content, including the vCenter Server's local content libraries.

vCenter Server provides a predefined sample role, content library administrator, that allows you to give users or groups the necessary privileges to manage selected content libraries. You can modify the role or use it as an example to create custom roles. If a user is assigned the content library administrator role on a library, that user can perform the following tasks on that library:

- Create, edit, and delete local or subscribed libraries.

- Synchronize a subscribed library and synchronize items in a subscribed library.
- View the item types supported by the library.
- Configure the global settings for the library.
- Import items to a library.
- Export library items.

### **Note**

You cannot set permissions on a content library directly.

## **Content Library Synchronization Options**

When configuring the subscribing library, you can choose either to download all libraries' content immediately or download library content only when needed. The first option starts the full synchronization process immediately. It includes the full content, including the metadata and actual data. The latter option starts the synchronization process for just the metadata immediately. The metadata contains information about the actual content data, allowing users to view and select the associated templates and ISOs. In this case, the actual data is synchronized only as needed when subscribed library objects are demanded. The impact of the on-demand synchronization is that storage space may be saved for the subscribing library, but a delay may exist each time a library item is selected.

To enable automatic synchronization, select the option Enable Automatic Synchronization with the External Library in the subscribed library settings. Consider the fact that the automatic synchronization requires a lot of storage space because you download full copies of all the items in the published library.

The content library synchronization method has an impact on VM provisioning time and datastore space usage. If an object is not already downloaded when you go to use it, you may have to wait while the subscribed content library downloads it from the published library. To optimize VM provisioning time, consider setting the download method to Immediately. To optimize datastore space usage, consider setting the download method to When Needed.

## **Adding Items to a Content Library**



You can import items such as OVA/OVF templates and vApps to a content library from your local machine or from a web server. You can also import ISO images, certificates, and other files. You can add to a content library an item that resides on a web server, or you can add items to a content library by importing files from your local file system.

You can import an OVF package to use as a template for deploying virtual machines and vApps. You can also import other types of files, such as scripts or ISO files. To import a file, in the vSphere client, right-click a content library, choose Import Item, select a file, and assign the item name.

You can also add content to a library by cloning VMs or templates to the library, as described in the following steps:

**Step 1.** In the vSphere Client, navigate to the virtual machine or template that you want to clone and select one of the following cloning tasks:

- Right-click a virtual machine and select **Clone > Clone to Template in Library**.
- Right-click a VM template and select **Clone to Library**.

**Step 2.** Depending on the selection in the previous step, complete the cloning wizard. For example, if you selected a VM template and chose **Clone to Library**, then you can use the following steps to create a new template in the content library:

- a. Select the **Clone As** option and choose to create a new template.
- b. From the Content Libraries list, select the library in which you want to add the template.
- c. Enter a name and description for the template.
- d. Optionally, select the configuration data that you want to include in the template. You can select to preserve the MAC addresses on the network adapters and include extra configuration.
- e. Click **OK**.

## Deploying VMs by Using a Content Library

You can deploy virtual machines from the VM templates in your content library by using this procedure:

- Step 1.** In the vSphere client, select **Home > Content Libraries**.
- Step 2.** Select a content library and click the **Templates** tab.
- Step 3.** Right-click a VM template and select **New VM from This Template**.
- Step 4.** On the Select Name and Location page, enter a name and select a location for the virtual machine.
- Step 5.** Optionally, to apply a customization specification to your virtual machine, select the **Customize the Operating System** checkbox and click **Next**.
- Step 6.** On the Customize Guest OS page, select a customization specification or create a new one and click **Next**.
- Step 7.** On the Select a Resource page, select a host, a cluster, a resource pool, or a vApp to run the deployed VM template and click **Next**.
- Step 8.** On the Review Details page, verify the template details and click **Next**.

## Exam Preparation Tasks

As mentioned in the section “How to Use This Book” in the Introduction, you have some choices for exam preparation: the exercises here, [Chapter 15, “Final Preparation,”](#) and the exam simulation questions on the companion website.

## Review All the Key Topics

Review the most important topics in this chapter, noted with the Key Topics icon in the outer margin of the page. [Table 14-9](#) lists these key topics and the page number on which each is found.



**Table 14-9** Key Topics for [Chapter 14](#)

<b>Key Topic Element</b>	<b>Description</b>	<b>Page Number</b>
List	Permissions to create a virtual machine	<a href="#">568</a>

List	Permissions to clone a virtual machine	572
List	Creating a Linux guest customization specification	575
List	Creating a Windows guest customization specification	576
List	Conditions for CPU hot adding	580
List	Requirements for virtual machine encryption	584
Section	Migrating virtual machines	587
Section	Configuring VMs to support vGPUs	592
List	Subscribing to a content library	596

## Complete Tables and Lists from Memory

Print a copy of [Appendix B, “Memory Tables”](#) (found on the companion website), or at least the section for this chapter, and complete the tables and lists from memory. [Appendix C, “Memory Tables Answer Key”](#) (also on the companion website), includes completed tables and lists to check your work.

## Define Key Terms

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

[Microsoft virtualization-based security \(VBS\)](#)

[VMware PowerCLI](#)

[content library](#)

[Open Virtual Format \(OVF\) template](#)

[Open Virtual Appliance \(OVA\) template](#)

[graphics processing unit \(GPU\)](#)

## Review Questions

1. Which of the following is a requirement for guest OS customization?
  - a. ESXi 5.0 or later
  - b. VMware Tools 11.0 or later

- c. A supported guest OS installed on SCSI node 0:0
    - d. A supported guest OS installed on any SCSI node
- 2. You want to create a virtual machine that can use up to 4096 MB video memory. Which compatibility option should you choose?
  - a. ESXi 7.0 and later
  - b. ESXi 6.7 Update 2 and later
  - c. ESXi 6.7 and later
  - d. ESXi 6.5 and later
- 3. You are snapshotting production virtual machines and want to minimize the impact to users of the guest OS and its applications. Which option should you choose?
  - a. Snapshot the memory and quiesce the file system.
  - b. Snapshot the memory but do not quiesce the file system.
  - c. Quiesce the file system but do not snapshot the memory.
  - d. Do not quiesce the file system or snapshot the memory.
- 4. In your vSphere 7.0 environment, you want to export a virtual machine for portability to other systems. Which approach should you use?
  - a. Export to OVF
  - b. Export to OVA
  - c. Export as a VM template
  - d. Export as a VMDK
- 5. You want to add items to the content library. Which of the following is not a valid choice for importing?
  - a. ISO file
  - b. OVA
  - c. A running virtual machine
  - d. OVF

## Chapter 15

# Final Preparation

Congratulations on making it through all the technical chapters in this book. Now you are ready for your final preparation for taking the Professional VMware vSphere 7.x (2V0-21.20) exam. This chapter contains two sections: “Getting Ready” and “Taking the Exam.”

## Getting Ready

Here is a list of actions and considerations that you should address prior to taking the exam:

- Review the VCP-DCV 2020 or current certification requirements, as described in the Introduction in this book.
- If you are taking an updated version of the VMware Professional vSphere 7.x exam, rather than 2V0-21.20, download the corresponding online appendix from the Pearson companion website. Use the appendix, written by the authors of this book for each exam update, as your guide for preparing for the exam. The appendix covers product and exam changes. It provides an updated table mapping exam objective to chapters and sections in the original book and in the appendix.
- Gain hands-on experience with vSphere 7. If you have not done so already, you should access a vSphere 7 environment and use it to practice performing the procedures described in this book. If you do not have a suitable vSphere 7 environment, consider using VMware Hands-on Labs (<http://labs.hol.vmware.com>) to search for vSphere and vSAN offerings. You can follow the lab exercises as designed or use them as a playground to gain hands-on familiarity with specific vSphere features that may be difficult for you to implement in your own lab, such as vSAN or vVols. The following are examples of labs that might be helpful:
  - HOL-2111-01-SDC-VMware vSphere: What’s New

- HOL-2111-01-SDC-VMware vSphere: Advanced Topics
- HOL-2111-01-SDC-VMware vSphere: Security Getting Started
- HOL-2108-01-HCI: What's New in vSAN: Getting Started
- Practice performing most of the procedures that are described in this book.
- Review the following items in each chapter in this book until you have them committed to memory:
  - The “Do I Know This Already” questions at the beginning of each chapter
  - The key topics table at the end of each chapter
  - The key terms at the end of each chapter
  - The “Review Questions” at the end of each chapter
- Take the practice exams included with the book. The standard edition includes two exams, and the premium edition includes two more exams. If you miss any questions, read the provided explanation and the related section in this book. Continue taking the practice exams until you feel you know and can explain each answer. If you are preparing for an updated version of the exam, be sure to look for updated practice exam questions.
- Prior to registering for the exam, create an account with VMware Certification, at [https://mylearn.vmware.com/MgrReg/login.cfm?ui=www\\_cert](https://mylearn.vmware.com/MgrReg/login.cfm?ui=www_cert).
- Create an account at Pearson Vue ([www.vue.com](http://www.vue.com)), which delivers all VMware career certification exams. To register for the 2V0-21.20 exam, use the information in the Introduction of this book to navigate to 2V0-21.20 exam details and click the Schedule Exam link. Follow the link to log in with your VMware Certification account. After a successful login, you will be redirected to the Pearson Vue website to complete the registration. In the wizard, you select the exam location, date, and time and provide payment.

## Note

Currently, you can choose to take the exam at home or in a Pearson Vue testing center. To take the exam at home, you must meet strict requirements, such as compatibility for audio, camera, and bandwidth. Pay careful attention to all the requirements and precheck information before

choosing this option.

## Taking the Exam

Here is a list of recommendations for the day of the exam:

- Bring two forms of identification that include your photo and signature. You cannot bring personal items such as laptops, tablets, phones, watches, pagers, wallets, or notes into the examination room. You may be able to place some of these items into a locker, but you should avoid bringing larger items into the training facility.
- Arrive at the exam center 30 minutes prior to the scheduled exam start time so you have ample time to complete the sign-in procedure and address personal needs. During the sign-in procedure, you should expect to place personal belongings in a locker, provide credentials, review the test regulations, and sign the agreement.
- Be sure to pay attention to the rules and regulations concerning the exam. For example, follow the venue’s protocol for requesting help during the exam and for signaling your completion of the exam. Each venue’s rules may be unique.
- Pay close attention to the wording of each question and each choice. (The exam format is multiple choice, provided via a web-based user interface.) The following are some examples of what to expect:
  - Some questions may ask you to select “which statement is correct,” and some questions may ask you to select “which statement is incorrect.”
  - Most questions call for you to select a single choice from a list of multiple choices. Whenever a question calls for you to select more than one choice, it does so explicitly by including a phrase such as “Choose two.”
  - Read each question carefully enough to ensure that you successfully interpret feature names and terminology. For example, when a question contains the word *heartbeat*, you need to carefully determine if it is referring to an HA network heartbeat, an HA datastore heartbeat, a VMFS heartbeat, a VMware Tools heartbeat, or some other heartbeat.
  - Questions tend to be written in a concise manner, and at first glance you

might think that insufficient details are provided. For example, a question could provide a symptom and ask you to select three actions that you should take to troubleshoot the issue. Your first thought might be that you would take analytical steps or remediation steps that are not provided as choices for the question. You might even consider the provided choices to be unpractical or insufficient. Do not get frustrated. Just select the appropriate choices that fit the question.

- Questions that ask you to identify more than one choice to accomplish a specific task may not clearly state whether the steps in the selected choices must be done or whether the steps in any one of the selected choices can be done. Although you might wish the question contained better clarity, you should see that only one solution fits the question.
- Strive for good time management during the exam. For the 2V0-21.20 exam, you need to answer 70 questions in 130 minutes. You could, for example, allow yourself 1.5 minutes per question, leaving 25 minutes at the end of the exam to review your marked questions. A timer is provided in the top-right corner of the exam user interface, along with the number of remaining questions. You might want to watch for the following milestones and pick up your pace when necessary:
  - With 50 questions remaining, at least 100 minutes left
  - With 30 questions remaining, at least 70 minutes left
  - With 10 questions remaining, at least 40 minutes left
- Don't allow yourself to spend too much time on a question. For example, if a question is tricky or if more than one choice seems to fit for a question that calls for a single choice, determine which choice is most likely the choice that VMware wants. In some cases, you might find it easier to answer the question by focusing on which choices to eliminate.
- Be sure to answer all questions. You will not be penalized for incorrect answers. Guessing provides an opportunity for a higher score.
- Whenever you are unsure of an answer or feel rushed to make a decision, select the box to mark the question for review.
- After answering all the questions, use the Review Page, which identifies all questions that you marked for review and all questions that are incomplete. If sufficient time remains, use the links provided on the review page to return to any questions that you marked or any questions that are identified



as incomplete.

## Appendix A

# Answers to the “Do I Know This Already?” Quizzes and Review Questions

## “Do I Know This Already?” Answers

### Chapter 1

1. b. Explanation: vCenter Server Essentials Plus, which is included in vSphere Essentials Plus, supports vMotion.
2. b. Explanation: Proactive HA minimizes VM downtime by proactively detecting hardware failures and placing the host in Quarantine Mode or Maintenance Mode.
3. d. Explanation: Security Token Service (STS) provides Security Assertion Markup Language (SAML) tokens, which are used to authenticate users to other vCenter components instead of requiring users to authenticate to each component.
4. a. Explanation: During vCenter Server Appliance deployment, you must create an SSO domain or join an existing SSO domain.
5. a. Explanation: You should deploy vCenter Server Appliance in the Large Environment configuration, which requires 37 GB, because the Medium Environment configuration supports only 4000 virtual machines.
6. d. Explanation: Changing the host boot type between legacy BIOS and UEFI is not supported after you install ESXi 7.0.
7. a. Explanation: If you plan to schedule file-based backups using the vCenter Server Appliance Management Interface (VAMI), you must

prepare an FTP, FTPS, HTTP, HTTPS, or SCP server with sufficient disk space to store the backups. Although the NFS and SMB protocols are supported for backup, the vCenter Server Appliance GUI installer does not support restoration from a backup with these protocols. To perform a restore from the NFS or SMB protocol, use the vCenter Server Management API.

8. c. Explanation: For Windows users, VMware supports Microsoft Edge 38 and later, Microsoft Internet Explorer 11.0.96 and later, Mozilla Firefox 45 and later, Google Chrome 50 and later, and Safari 5.1 and later. For Mac users, VMware supports Safari 5.1 and later, Mozilla Firefox 45 and later, and Google Chrome 50 and later.
9. d. Explanation: vRealize Automation is cloud automation software that speeds up the delivery of infrastructure and application resources on-premises and in the public cloud.
10. b. Explanation: VMware Cloud Foundation (VCF) delivers a simple path to the hybrid cloud by leveraging a common infrastructure and consistent operational model for on-premises and off-premises data centers.

## Chapter 2

1. c and f. Explanation: NPIV requires the use of virtual machines with RDMs. An RDM is a mapping file containing metadata that resides in a VMFS datastore.
2. a. Explanation: A vSphere pod requires ephemeral storage to store Kubernetes objects, such as logs, emptyDir volumes, and ConfigMaps.
3. c. Explanation: Symmetric Multiprocessing Fault Tolerance (SMP-FT) is supported when PFFT is set to 0 and Data Locality is set to Preferred or Secondary.
4. b. Explanation: If you choose RAID 5/6 and PFTT=2 for a 100 GB virtual disk, the required capacity is 150 GB, and the usable capacity is 67%.
5. a. Explanation: The following are the VAAI primitives for NAS: Full File Clone, Fast File Clone/Native Snapshot Support, Extended Statistics, and Reserve Space.

6. d. Explanation: Protocol endpoints (PEs) are logical I/O proxies, used for communication with virtual volumes and the virtual disk files.
7. a. Explanation: When the VMware NMP receives an I/O request, it calls the appropriate PSP, the PSP selects an appropriate physical path, and the NMP issues the I/O request.
8. c. Explanation: A VM storage policy for tag-based placement is helpful for storage arrays that do not support VASA and their storage characteristics are not visible to the vSphere client.
9. a. Explanation: The available vSAN storage policies include PFTT, SFTT, Data Locality, Failure Tolerance Method, Number of Disk Stripes per Object, Flash Read Cached Reservation, Force Provisioning, Object Space Reservation, Disable, Object Checksum, and IOPS Limit for Object.
10. b. Explanation: If the space used on datastore A is 82% and on datastore B is 79%, the difference is 3. If the threshold is 5, Storage DRS will not make migration recommendations from datastore A to datastore B.

## Chapter 3

1. d. Explanation: On a vSS, you can set the following network policies: Teaming and Failover, Security, Traffic Shaping, and VLAN.
2. b. Explanation: The following NIC teaming options are available on vSS and vDS: Route Based on Originating Virtual Port, Route Based on IP Hash, Route Based on Source MAC Hash, and Use Explicit Failover Order.
3. a. Explanation: Distributed virtual switches can do both inbound and outbound traffic shaping, whereas standard virtual switches handle just outbound traffic shaping.
4. c. Explanation: If you reserved 1.0 Gbps for virtual machine system traffic on a distributed switch with 8 uplinks, then the total aggregated bandwidth available for virtual machine reservation on the switch is 8.0 Gbps. Each network resource pool can reserve a portion of the 8 Gbps capacity.
5. c. Explanation: When marking traffic, you can create a rule to

configure qualifiers to identify the data to be tagged and set Action to Tag.

6. c. Explanation: A vDS supports up to 64 LAGs.
7. a. Explanation: The required vDS configuration for the virtual switch teaming policy health check is at least two active physical NICs and two hosts.
8. c. Explanation: Cisco Discovery Protocol (CDP) support was introduced with ESX 3.x. CDP is available for standard switches and distributed switches that are connected to Cisco physical switches. Link Layer Discovery Protocol (LLDP) is supported in vSphere 5.0 and later for vDS (5.0.0 and later), but not for vSS.
9. a. Explanation: With DirectPath I/O in a vSphere 7.0 environment, a virtual machine can be part of a cluster, but it cannot migrate across hosts.
10. a. Explanation: The available services for a custom stack are Management, vMotion, IP-based storage, Provisioning, Fault Tolerance logging, vSphere Replication, vSphere Replication NFC, and vSAN.

## Chapter 4

1. d. Explanation: Intel EVC Mode Nehalem (Level L2) includes the Intel Penryn feature set and exposes additional CPU features, including SSE4.2 and POPCOUNT.
2. c. Explanation: When the DRS migration threshold is set to Level 3, the default level, DRS expands on Level 2 by making recommendations to improve VM happiness and cluster load distribution.
3. b. Explanation: Resource pools are container objects in the vSphere inventory that are used to compartmentalize the CPU and memory resources of a host, a cluster, or a parent resource pool. You can delegate control over each resource pool to specific individuals and groups.
4. a. Explanation: Starting in vSphere 6.7, DRS uses a new two-pass algorithm to allocate resource reservations to its children. In the

second pass, excess pool reservation is allocated proportionally, limited by the virtual machine's configured size.

5. d. Explanation: Set Define Host Failover Capacity By to Dedicated Failover Hosts to designate hosts to use for failover actions.
6. b. Explanation: The medium virtual machine monitoring level sets Failure Interval to 60 seconds and Reset Period to 24 hours.
7. c. Explanation: For PDL and APD failures, you can set VMCP to either issue event alerts or to power off and restart virtual machines. For APD failures only, you can additionally control the restart policy for virtual machines by setting it to Conservative or Aggressive.
8. b. Explanation: Predictive DRS requires vCenter Server 6.5 or later, it must be enabled on both vCenter Server and vROps, and the vCenter Server and vROps clocks must be synchronized.
9. d. Explanation: vSphere Fault Tolerance can accommodate symmetric multiprocessor (SMP) virtual machines with up to eight vCPUs.
10. a. Explanation: If a vCenter service fails, VMware Service Lifecycle Manager restarts it. VMware Service Lifecycle Manager is a service running in vCenter server.

## Chapter 5

1. d and f. Explanation: Other than an optional folder to contain data centers, you cannot create any object in the inventory until you create a data center
2. d. Explanation: A subscribed library is a library whose content you cannot change or publish. It receives its content from a published library.
3. b. Explanation: The file named `<vmName>-flat.vmdk` is the virtual machine data disk file. It is commonly called the flat file. `<vmName>..vmdk` is the virtual machine data disk file.
4. a. Explanation: The parent (current) snapshot is always the snapshot that appears immediately above the You Are Here icon in the Snapshot Manager. The parent snapshot is not always the snapshot that you took most recently.

5. a. Explanation: The SIO controller provides serial and parallel ports and floppy devices, and it performs system management activities. One SIO controller is available to the virtual machine, but it cannot be configured or removed.
6. c. Explanation: The VM Options table includes General Options, Encryption Options, Power Management, VMware Tools, VBS, Boot Options, Advanced Options, Fibre Channel NPIV, and vApp Options.
7. a. Explanation: If the vCenter Server instances exist in separate vCenter Single Sign-On domains, you can use vSphere APIs or the vSphere SDK to migrate virtual machines. The cross-vCenter Server and long-distance vMotion features require an Enterprise Plus license.
8. c. Explanation: You can simultaneously perform one Storage vMotion and four vMotion operations involving a specific host.
9. b. Explanation: To avoid network saturation, you can use traffic shaping to limit the average bandwidth and peak bandwidth available to vMotion traffic.
10. c. Explanation: One of the most popular use cases for instant clones is a just-in-time virtual desktop infrastructure (VDI) deployment. Instant clones enable you to perform large-scale deployments by allowing you to create virtual machines from a controlled point in time. For example, VMware Horizon uses instant clones to improve the provisioning process for virtual desktops.

## Chapter 6

1. d. Explanation: Use cases for vSphere with Tanzu include providing a familiar single stack for containers and virtual machines and streamlining the development of modern applications.
2. b. Explanation: vRealize Operations (vROps) is commonly used for continuous performance optimization and intelligent remediation.
3. a. Explanation: vRealize Log Insight (vRLI) is commonly used to decrease time and effort spent on root cause analysis and centralized log management and analysis.
4. a and e. Explanation: Horizon includes instant clones, which together with VMware Dynamic Environment Manager and VMware App

Volumes dynamically provides just-in-time (JIT) delivery of user profile data and applications to stateless desktops.

5. d and e. Explanation: VMware App Volumes is a set of application and user management solutions for VMware Horizon, Citrix Virtual Apps and Desktops, and Remote Desktop Services Host (RDSH) virtual environments.
6. d. Explanation: vSphere Replication does not require separate licensing. Instead, it is included as a feature of specific vSphere license editions, including vSphere Standard.
7. b and c. Explanation: Site Recovery Manager (SRM) use cases include disaster recovery and data center migrations. For data replication, SRM integrates with vSphere Replication and supported storage-based replication products.
8. c. Explanation: VMware Cloud Foundation (VCF) is a hybrid cloud platform built on full-stack hyperconverged infrastructure (HCI) technology. Its main components include Cloud Builder and SDDC Manager.
9. a. Explanation: VMware HCX is a workload mobility platform that simplifies application migration, workload rebalancing, and business continuity across on-premises data centers, private clouds, and hybrid clouds.
10. a. Explanation: NSX is commonly used with zero-trust security and multi-cloud networking.

## Chapter 7

1. d. Explanation: The key size requirement is 2048 to 16,384 bits, not 1024 to 16,384 bits
2. a. Explanation: vCenter Server supports these certificate modes for ESXi: VMware Certificate Authority, Custom Certificate Authority, and Thumbprint Mode.
3. a. Explanation: In vCenter Server 7.0, the system roles include read-only, administrator, no access, no cryptography administrator, trusted infrastructure administrator, and no trusted infrastructure administrator.



4. a. Explanation: To migrate a virtual machine with Storage vMotion, the user must have the Resource.Migrate Powered On Virtual Machine on the virtual machine or folder and Datastore.Allocate Space on the destination datastore.
5. d. Explanation: In normal lockdown mode, user accounts that are in the Exception Users list and that have administrator privileges on the host can access the DCUI. Also, users identified in the host's DCUI.Access advanced option can access the DCUI.
6. a. Explanation: By default, this password must have at least eight characters, one lowercase character, one numeric character, and one special character.
7. b. Explanation: To list the available security associations, you can use the command **esxcli network ip ipsec sa list** in ESXi.
8. b. Explanation: Concerning encrypted vMotion across vCenter Server instances, you must use the vSphere APIs, encrypted vMotion migration of unencrypted virtual machines is supported, and vMotion migration of encrypted virtual machines is not supported.
9. d. Explanation: Some settings can be disabled in a vSphere environment. To reduce potential risk, consider setting `isolation.tools.gui.launchmenu.change` to TRUE.
10. b. Explanation: You can implement VMware NSX Data Center for vSphere (NSX) to add a distributed logical firewall, microsegmentation, and additional security measures to your vSphere environment.

## Chapter 8

1. a. Explanation: The ESXi installation prerequisites include downloading the ESXi installer ISO and preparing the hardware system to boot from it.
2. b. Explanation: There is a default installation script included with the ESXi installer. The default `ks.cfg` installation script is in the initial RAM disk at `/etc/vmware/weasel/ks.cfg`.
3. c. Explanation: Prior to running the deployment command, you can run a pre-deployment check by using the command **vcsa-deploy**

**install --verify-only** *path-to-JSON-file*.

4. a. Explanation: VECS does not store ESXi certificates. ESXi certificates are stored locally on the ESXi hosts in the `/etc/vmware/ssl` directory.
5. d. Explanation: All users must have the object of class *inetOrgPerson*. All groups must have the object of class *groupOfUniqueNames*. All groups must have the group membership attribute *uniqueMember*.
6. c. Explanation: A user must be a member of the CAAdmins group to perform most certificate management operations, such as using the **certool** command.
7. a. Explanation: You should not select the Use Machine Account option if you plan to rename the machine. Instead, you need to select Use Service Principle Name (SPN) and provide the SPN, UPN, and password.
8. a. Explanation: In scenarios where vCenter Server is installed in a secured network with no Internet access, you can install the Update Manager Download Service (UMDS) and use it to download updates.
9. d. Explanation: For the Passive node, connect NIC 0 to the Management network and NIC 1 to the vCenter HA network. For the Witness node, connect NIC 1 to the vCenter HA network.
10. b. Explanation: The following kernel boot options have been deprecated and are no longer supported in ESXi 7.0: **--no-auto-partition**, **autoPartitionCreateUSBCoreDumpPartition**, and **autoPartitionDiskDumpPartitionSize**.

## Chapter 9

1. a. Explanation: To add physical adapters to a vSS, you can select the host, navigate to Configure > Networking > Virtual Switches, select the switch, and select Manage Physical Adapters. In the wizard, click the Add Adapters (green plus sign) button.
2. b. Explanation: You can set the VLAN ID to 0 (external switch tagging), 1 to 4094 (virtual switch tagging), or 4095 (virtual guest tagging)

3. d. Explanation: You can change the general setting of a vDS, including Name, Number of Uplinks, Network I/O Control (enable or disable), and Description.
4. a. Explanation: You can change the Advanced settings for a vDS, including MTU (in bytes), Multicast Filtering Mode (Basic or IGMP/MLD Snooping), Discovery Protocol, and Administrator Contact.
5. a. Explanation: When creating a VMkernel adapter, you should configure the VMkernel Adapter IP, MTU, Stack, and Available Services settings.
6. c. Explanation: NIOC applies shares to each of the following network traffic types: management traffic, Fault Tolerance (FT) traffic, NFS traffic, vSAN traffic, vMotion traffic, vSphere Replication (VR) traffic, vSphere Data Protection backup traffic, and virtual machine traffic.
7. b. Explanation: For port mirroring, you can select one of the following session types: distributed port monitoring, remote mirroring source, remote mirroring destination, or encapsulated remote mirroring (L3) source.
8. a. Explanation: The number of ports in each port channel must match the number of physical NICs that will be aggregated on the host (the minimum is two). The same hashing algorithm must be used for the port channel and the associated LAG on the vDS. All the NICs in a LAG must be configured with the same speed and duplexing.
9. b. Explanation: You can enable the VLAN and MTU and the Teaming and Failover health checks.
10. d. Explanation: You can use the following to qualify data: IP address (to identify packets by source and destination addresses and ports), MAC address (to identify packets by source and destination addresses and by VLAN), system traffic (to identify data by system traffic type).

## Chapter 10

1. c. Explanation: If you initially selected the Skip Quickstart option, you should add hosts manually. If you previously used Quickstart but

selected Configure Networking Settings Later, you can add hosts by using Quickstart but must manually configure the host networking.

2. d. Explanation: The default value for both CPU and Memory Limit is Unlimited.
3. a. Explanation: Define Host Failover Capacity can be set to Cluster Resource Percentage, Slot Policy (powered-on VMs), Dedicated Host Failures, or Disabled.
4. b. Explanation: To configure Virtual Machine Component Protection (VMCP) in a vSphere HA cluster, you can select Failures and Responses > Datastore with PDL and choose Issue Events or Power Off and Restart VMs.
5. c. Explanation: You need to install a supported vendor-supplied vSphere Client plug-in and register the proactive HA provider. Ensure that vSphere HA and DRS are enabled. To allow remediation actions, you need the Host.Config.Quarantine and Host.Config.Maintenance privileges.
6. b. Explanation: If the guest OS reports that memory paging is occurring, you should increase the memory size of the virtual machine.
7. a. Explanation: If a virtual disk is thin provisioned, you may be maximizing the use of your storage space while decreasing the virtual machine's performance and increasing its risk of denial of service.
8. b. Explanation: %DRPTX is the percentage of transmit packets dropped. It indicates that the physical network adapter cannot meet the demand, perhaps due to load from other virtual machines.
9. a. Explanation: Warning, Information, and Audit are valid event types. Another valid type is Alert, which indicates that a fatal problem has occurred in the system.
10. c. Explanation: The ESXi host agent log contains data related to the agent that manages and configures the ESXi host and its virtual machines.

## Chapter 11

1. c. Explanation: In hybrid clusters, magnetic disks are used for capacity,

and flash devices serve as a read cache and a write buffer. In a hybrid cluster, 70% of the flash space is used for the read cache, and 30% is used for the write buffer.

2. a. Explanation: Distributed switches with Network I/O Control (NIOC) Version 2 cannot be used for vSAN Quickstart. To use Quickstart to configure a vSAN cluster, the hosts must have similar network configuration.
3. d. Explanation: A datastore is expandable when the backing storage device has free space immediately after the datastore extent.
4. b. Explanation: Multiple IP addresses or DNS names can be used with NFS Version 4.1 multipathing.
5. b. Explanation: To configure SIOC, select a datastore, enable Configuration > Properties > Storage I/O Control, click Advanced, and modify the threshold for contention. (This value must be between 10 ms and 100 ms.)
6. d. Explanation: If you want to perform a maintenance activity on an SDRS cluster member datastore or its underlying storage devices, you can place it in Maintenance Mode. To place a datastore in Maintenance Mode using the vSphere Client, right-click the datastore in the inventory pane, select Enter SDRS Maintenance Mode, and optionally apply any recommendations.
7. a, d, and e. Explanation: To configure an ESXi 7.0 host to access shared NVMe devices using RoCE Version 2, you create a new VMkernel virtual network adapter on a virtual switch and configure its uplink to use the RDMA paired uplink (for example, vmnic9). Navigate to Configure > Networking > RDMA Adapters to review the bindings. Select Add Software NVMe over RDMA Adapter.
8. b. Explanation: In vPMem mode, a virtual machine can directly access PMem resources and use the resources as regular memory.
9. c. Explanation: FIXED and LB\_RR are path selection schemes (PSS) used by the HPP module. VMW\_PSP\_FIXED is the NMP path selection module that uses a preferred path.
10. a and b. Explanation: To configure vVols in a vSphere 7.0 environment, you need to configure the following components on the

storage side: protocol endpoints, storage containers, storage profiles, and replication configurations (if you plan to use vVols with replication).

## Chapter 12

1. d. Explanation: To join vCenter Server systems in Enhanced Linked Mode, you need to connect them to the same SSO domain.
2. a. Explanation: In a vSphere environment, you cannot change the object, user, or user group in a permission, but you can change the role.
3. b and c. Explanation: The Certificate Management page shows the certificate types in the VMware Endpoint Certificate Service (VECS). By default, the types are machine SSL certificates and trusted root certificates.
4. a. Explanation: To perform certificate management for ESXi, you must have the Certificates.Manage Certificates privilege.
5. d. Explanation: You can change the required length, change the character class requirement, and allow passphrases by using the Security.PasswordQualityControl advanced option.
6. a. Explanation: An orange icon indicates that the status of the device has changed, and you must reboot the host before you can use the device.
7. b. Explanation: If the host acceptance level is VMwareAccepted, you cannot install VIBs at the PartnerSupported level.
8. a. Explanation: You can use the following command to run the Secure Boot validation script on an upgraded ESXi host:  
**`/usr/lib/vmware/secureboot/bin/secureBoot.py -c`**
9. b. Explanation: To configure a trust relationship between a KMS and vCenter, in the vSphere Client, select the vCenter Server, navigate to Configuration > Key Management Servers, and click Add.
10. c. Explanation: To configure vSphere Trust Authority, you need to configure the trusted key provider for the trusted hosts on the trusted cluster (using Register-KeyProvider and Set-KeyProvider).

## Chapter 13

1. c. Explanation: The vCenter Server Appliance Management Interface (VAMI) provides a file-based backup feature for the vCenter Server. Alternatively, you can perform image-based backups by using the vSphere API.
2. a. Explanation: You can configure a backup to stream the data to a target by using FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB.
3. d. Explanation: To upgrade a vSphere 6.5 or 6.7 environment to vSphere 7.0, you should upgrade the major components in the following order: vCenter Server, ESXi hosts, VMware Tools, virtual machine hardware.
4. c. Explanation: If you choose to transfer configuration, historical, and performance data, you can transfer the configuration data during the upgrade and transfer the remaining data in the background following the upgrade. The option to transfer data in the background following an upgrade is applicable only to scenarios where the source vCenter Server uses an external database.
5. d. Explanation: A vSphere Installation Bundle (VIB) is the smallest installable software package (metadata and binary payload) for ESXi.
6. a. Explanation: To enable Quick Boot, navigate to Menu > Lifecycle Manager, select a vCenter Server, and enable the Quick Boot checkbox in the cluster settings at Settings > Host Remediation > Images.
7. c. Explanation: Firmware updates are available in a special vendor depot that you access through a vendor-specific hardware support manager plug-in that registers itself as a vCenter Server extension. Firmware updates are not available for clusters that are managed with baselines.
8. c. Explanation: To change the state of a service, you can right-click on the service and select Start, Stop, or Reset. You can also change a service's startup policy such that it automatically starts with the host or associated firewall ports or is started only manually.
9. b. Explanation: When the Health Status badge color is orange, it indicates an alert, and one or more components may be degraded. Non-security patches may be available.

10. a. Explanation: Optionally, you can check conflicts and apply one of the following resolutions to all conflicts or separately to each conflict: Copy, Skip, or Merge.

## Chapter 14

1. c. Explanation: When selecting the storage type on a host that has PMem memory, you can select either the Standard or PMem radio button. If you chose PMem storage for a virtual machine, its default virtual disk, new virtual disk, and NVDIMM devices share the same PMem resources.
2. b. Explanation: With the `vmx.log.guest.level = "info"` setting, `vminst.log` is sent to the host, but `vmmsi.log` remains in the virtual machine.
3. a. Explanation: You cannot use the vSphere Client to clone a virtual machine using linked clones or instant clones. You can do so with API calls.
4. c. Explanation: The compatibility setting controls which virtual machine hardware version is used. Setting the compatibility to ESXi 6.7 and later uses hardware Version 14.
5. a. Explanation: To control the default hardware compatibility for new virtual machines, you can set the Default VM Compatibility setting at the host, cluster, or data center level.
6. d. Explanation: To minimize the impact to a running virtual machine and to reduce the time required to take a snapshot, do not snapshot the memory state or quiesce the guest file system.
7. c. Explanation: In vSphere, to use VBS, you must use virtual hardware Version 14 or later. The virtual machine must be set to use UEFI firmware, Secure Boot, hardware virtualization (Intel VT/ADM-V), and IOMMU.
8. a. Explanation: The main use case for vGPUs is to support high-end graphics in virtual desktop infrastructure (VDI). Recently, the need to support artificial intelligence (AI) and machine learning (ML) has also emerged as a major use case.
9. a. Explanation: If a user is granted the read-only role as a global



permission and the administrator role at a vCenter Server level, then the user can manage the vCenter Server's content libraries and content but can only view content libraries belonging to other vCenter Servers.

10. d. Explanation: You can import items such as OVA/OVF templates and vApps to a content library. You can also import ISO images. You can also add content to the library by cloning VMs or templates to the library.

## Review Question Answers

### Chapter 1

1. c. Explanation: Fault Tolerance is supported for up to two vCPUs in vSphere 7.0 Standard.
2. d. Explanation: VMware Directory Service for the vCenter Single Sign-On (SSO) domain (vsphere.local) is a service in vCenter Server Appliance.
3. b. Explanation: To install ESXi 7.0, ensure that the hardware system has 4 GB or more of physical RAM. VMware recommends 8 GB or more for production environments.
4. c and d. Explanation: You must run the GUI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the appliance.
5. b. Explanation: VMware Cloud Foundation (VCF) is the industry's most advanced hybrid cloud platform. It provides a complete set of software-defined services for compute, storage, networking, security, and cloud management to run enterprise apps in private or public environments.

### Chapter 2

1. d. Explanation: VMFS Version 6 provides access for ESXi 6.5 and later hosts.
2. c. Explanation: Ruby vSphere Console (RVC) is a command-line interface used for managing and troubleshooting vSAN. RVC provides a cluster-wide view and is included with the vCenter Server

deployment.

3. a. Explanation: VASA storage providers are software components that integrate with vSphere to provide information about the physical storage capabilities.
4. a. Explanation: MRU is the default path selection policy for most active/passive storage devices.
5. b. Explanation: Thick eager zeroed is the slowest method for virtual disk creation, but it is the best for guest performance.

## Chapter 3

1. b. Explanation: When you enable traffic shaping for a standard switch or port group, you can configure the options Average Bandwidth, Peak Bandwidth, and Burst Size.
2. c. Explanation: At the distributed port level, you can override policies applied to the distributed port group and apply unique policies to a distributed port.
3. a. Explanation: In vSphere 7.0, the default settings for a distributed port group are static binding, elastic port allocation, and eight ports.
4. d. Explanation: Single Root I/O Virtualization (SR-IOV) is a feature that allows a single Peripheral Component Interconnect Express (PCIe) device to appear as multiple devices. It is useful for supporting an application in a guest OS that is sensitive to network latency. SR-IOV-enabled devices provide virtual functions (VFs) to the hypervisor or guest operating system.
5. b. Explanation: After configuring NetFlow on a vDS, you can configure monitoring policies on vDS port groups and ports.

## Chapter 4

1. b. Explanation: Enhanced vMotion Compatibility (EVC) is a cluster feature. The source and target processors must come from the same vendor class (AMD or Intel) to be vMotion compatible. Clock speed, cache size, and number of cores can differ between source and target processors.

2. b. Explanation: When the DRS Migration Threshold is set to Level 2, DRS expands on Level 1 by making recommends in situations that are at or close to resource contention. It does not make recommendations just to improve virtual machine happiness or cluster load distribution.
3. b. Explanation: If the primary host detects datastore heartbeats for a secondary host but no network heartbeats or ping responses, it assumes that the secondary host is isolated or in a network partition.
4. d. Explanation: Hosts must use static IP addresses or guarantee that IP addresses assigned by DHCP persist across host reboots.
5. a. Explanation: To use Wake-on-LAN (WoL) with DPM, you must ensure that vMotion is configured, the vMotion NIC must supports WoL, and the physical switch port must be set to automatically negotiate the link speed.

## Chapter 5

1. b. Explanation: Although making snapshots may be a useful step for a backup utility, a snapshot is not by itself a backup. A snapshot does not provide a redundant copy of data. If the base flat file is lost or corrupt, you cannot restore the virtual machine by reverting to a snapshot.
2. a. Explanation: You can enable or disable hardware acceleration. You can set debugging and statistics to run normally, record debugging information, record statistics, or record statistics and debugging.
3. a. Explanation: To migrate virtual machines with disks larger than 2 TB, the source and destination ESXi hosts must be Version 6.0 or later.
4. c. Explanation: During a vMotion migration without shared storage the virtual disk data is transferred over the vMotion network.
5. a. Explanation: During an instant clone (vmFork) operation, the system quiesces and stuns the source virtual machine, creates and transfers a checkpoint, customizes the destination MAC address and UUID, and forks the memory and disk.

## Chapter 6

1. b. Explanation: vRealize Orchestrator (vRO) is a key component of vRA that provides custom workflows to support anything as a service

(XaaS).

2. d. Explanation: VMware Horizon is commonly used for remote users, kiosk and task users, and call centers.
3. b. Explanation: To configure replication, in the vSphere Client, navigate to Home > Site Recovery > Open Site Recovery.
4. d. Explanation: VMware HCX is a workload mobility platform that simplifies application migration, workload rebalancing, and business continuity across on-premises data centers, private clouds, and hybrid clouds.
5. a. Explanation: VMware AppDefense provides data center endpoint security that protects applications running in a virtualized environment.

## Chapter 7

1. c. Explanation: Do not use CRL distribution points, authority information access, or certificate template information in any custom certificates.
2. b. Explanation: You cannot modify permissions on the following entities that derive permissions from the root vCenter Server system: custom fields, licenses, roles, statistics intervals, or sessions.
3. d. Explanation: By default, the services that are running in ESXi include DCUI, Load-Based Teaming, CIM Server, and VMware vCenter Agent.
4. b. Explanation: You should remove any expired or revoked certificates from the vCenter Server to avoid MITM attacks.
5. d. Explanation: Prerequisites for enabling UEFI boot for your virtual machines are UEFI firmware, virtual hardware Version 13 or later, VMware Tools Version 10.1 or later, and an operating system that supports UEFI Secure Boot.

## Chapter 8

1. d. Explanation: In the first stage when using the GUI installer, you navigate through the installation wizard, choose the deployment type,

provide the appliance settings, and deploy the OVA. In the second stage you use a wizard to configure the appliance time synchronization, configure vCenter Single Sign-On (SSO), and start the services in the newly deployed appliance.

2. d. Explanation: When adding an Active Directory over LDAP identity source, you need to provide required information such as the name, the base DN for users, and the base DN for groups.
3. d. Explanation: The vSphere Lifecycle Manager service is available via the vSphere Client immediately after vCenter Server deployment. No special steps are required to install vSphere Lifecycle Manager.
4. c. Explanation: A rule can identify target hosts by boot MAC address, SMBIOS information, BIOS UUID, vendor, model, or fixed DHCP IP address.
5. b. Explanation: After selecting Remediate and selecting the hosts, you need to click Pre-check Remediation to determine whether the selected hosts are ready for remediation.

## Chapter 9

1. b. Explanation: You can set VLAN ID to 0 (external switch tagging), 1 to 4094 (virtual switch tagging), or 4095 (virtual guest tagging).
2. c. Explanation: As a rollback plan, you should export the distributed switch configuration prior to upgrading. In the export wizard, choose the option to include the distributed port groups.
3. c. Explanation: Edit the distributed port group setting. In the settings, click General and then, from the Network Resource Pool drop-down menu, select the network resource pool and click OK.
4. d. Explanation: The provisioning stack supports traffic for virtual machine cold migration, cloning, and snapshot migration. It also supports the Network File Copy (NFC) traffic used for cloning virtual disks during long-distance vMotion. You can use this stack to isolate provisioning traffic by placing it on a separate gateway. The default stack provides networking support for management traffic and for all VMkernel traffic types.
5. a. Explanation: To enable NetFlow in a distributed port group, select

the distributed port group, select Configure > Policies, click Edit, and then, on the Monitoring page, select Enable NetFlow or Disable NetFlow.

## Chapter 10

1. c. Explanation: Optionally, you can set Memory Reservation to a numeric value (the default is 0) and a unit of measure (MB, GB, MHz, or GHz).
2. c. Explanation: In the vRealize Operations (vROps) GUI, locate the appropriate vCenter Server adapter instance. Select the adapter, choose Advanced Settings, and set Provide Data to vSphere Predictive DRS to True.
3. d. Explanation: To disable admission control, set Define Host Failover Capacity to Disabled.
4. a. Explanation: VIMTOP is a tool you can run in vCenter Server Appliance to see resource usage for services that are running.
5. d. Explanation: vmware.log is in the same folder as the virtual machine configuration file.

## Chapter 11

1. c. Explanation: To provide the encryption keys for a vSAN datastore, you must implement a key management server (KMS) cluster server that is KMIP 1.1 compliant and is in the vSphere compatibility matrices.
2. c. Explanation: To enable deduplication and compression for an existing vSAN cluster, edit the vSAN services in the cluster and enable Deduplication and Compression. (Separate options to enable just deduplication or just decompression are not provided.)
3. c. Explanation: The options in the Datastore Brower include Upload Files, Upload Folder, Download, New Folder, Copy to, Move to, Rename to, Delete, and Inflate.
4. d. Explanation: The default storage module that claims NVMe-oF is HPP. NVMe-oF targets cannot be claimed by NMP.

5. a. Explanation: When using the vSphere Client to manage the storage providers, you can select a storage provider and choose Synchronize Storage Providers, Rescan, Remove, or Refresh Certificate.

## Chapter 12

1. b. Explanation: To assign a global permission, you should use the vSphere Client with a user account that has the Permissions.Modify Permission privilege on the root object of all inventory hierarchies.
2. d. Explanation: A yellow alarm is raised if a certificate's status is Expiring Shortly (that is, if it expires in less than eight months).
3. a. Explanation: With this setting, passwords of one and two character classes are disabled, and so are passphrases. An uppercase character that begins a password and a number that ends a password do not count toward the number of character classes used.
4. d. Explanation: If the "Host secure boot was disabled" message appears in the vSphere Client, you must re-enable Secure Boot to resolve the problem. If the "No cached identity key, loading from DB" message appears, you must disconnect and reconnect the host.
5. a. Explanation: You can leverage the files that you exported during the configuration of vTA vSphere as your backup. If you need to restore vTA, you can use the exported files to reconfigure vTA.

## Chapter 13

1. b. Explanation: After a restore completes, the following configurations are restored: virtual machine resource settings, resource pool hierarchy and setting, cluster-host membership, DRS configuration, and rules.
2. d. Explanation: If you choose to transfer configuration, historical, and performance data, you can transfer the configuration data during the upgrade and transfer the remaining data in the background following the upgrade. The option to transfer data in the background following an upgrade is only applicable to scenarios where the source vCenter Server uses an external database.
3. b. Explanation: When working with baselines, the following settings are applicable: Quick Boot, VM Power State, VM Migration,

Maintenance Mode Failures, PXE Booted Hosts, and Removable Media Devices. When working with images, the following settings are applicable: Quick Boot, VM Power State, VM Migration, Maintenance Mode Failures, HA Admission Control, and DPM.

4. d. Explanation: In rare situations, such as when the distributed switch used for management is not functioning, you may want to restore a standard switch to an ESX host, which you can do with the Restore Standard Switch option in the DCUI.
5. b. Explanation: If the source domain contains multiple (linked) vCenter Servers, the repointing process involves additional steps to shut down the vCenter Server and unregister it from the source domain.

## Chapter 14

1. c. Explanation: Guest OS customization requires a supported guest OS installed on SCSI node 0:0 and VMware Tools. Windows guest customization requires ESXi Version 3.5 or later. Linux guest customization requires Perl in the guest OS.
2. a. Explanation: The compatibility setting controls which virtual machine hardware version is used. Setting the compatibility to ESXi 7.0 and later uses hardware Version 17, which is the minimum version that support 4 GB video memory.
3. d. Explanation: To minimize the impact to a running virtual machine and to reduce the time required to take a snapshot, do not snapshot the memory state or quiesce the guest file system.
4. a. Explanation: You can export a virtual machine, virtual appliance, or vApp as an OVF or OVA template to create virtual appliances that can be imported by other users. Starting in vSphere 6.5, you can only export to OVF.
5. c. Explanation: The transfer service on the vCenter Server is responsible for importing and exporting content between the subscriber and the publisher, using HTTP NFC.



# Index

## Numbers

- 7.0 features, vSphere, [53–54](#)
- 802.1ax, [93](#)
- 802.1q tagging attacks, [97](#)
- 802.3ad, [93](#)

## A

- absent component state, vSAN, [52](#)
- acceptance levels, ESXi hosts, [496](#)
- accessing
  - CIM, controlling access, [491–492](#)
  - datastore browsers, controlling, [261](#)
  - vCenter Server, restricting access, [261](#)
- accounts
  - lockouts, [485–487](#)
  - vCenter Cloud account permissions, [210–213](#)
  - VMware Certification accounts, [604](#)
- acknowledging triggered alarms, [399–400](#)
- active nodes, vCenter HA, [14](#)
- AD (Active Directory), [21](#)
  - ESXi user management, [497–498](#)
  - Federation Services, [313–314](#)
  - identity sources, [307–309](#)
- adapters

- host physical network adapters, managing with vDS, [351](#)
- network adapters, [181](#)
- VMkernel network adapters, migrating to vDS, [352](#)

add-ons, [7–8](#)

- security, [275](#)
- Skyline, [206](#)
  - integration, [206](#)
  - use cases, [206](#)
  - vSphere Health and vSAN Health, [53](#)

vCenter Converter, [205](#)

- integration, [205–206](#)
- use cases, [205](#)

vendor add-ons, [534](#)

vSphere Replication, [206](#), [219–220](#)

- integration, [220–221](#)
- use cases, [220](#)

vSphere with Tanzu, [173](#), [204](#)

- integration, [205](#)
- use cases, [204](#)

administration

- Administration server, vCenter Server, [11](#)
- vCenter Server
  - Administration server, [11](#)
  - restricting access, [261](#)
  - restricting administrative privileges, [260–261](#)

admission control

- HA, [371](#)
- VM resources, monitoring/managing, [390–391](#)
- vSphere HA, [146–148](#)

advanced performance charts, [377–379](#)

affinity rules, DRS, [369–370](#)

agents

- Host Agent, [11](#)

- vCenter Server Agent, [11](#)
- alarms
  - actions, [401](#)
  - advanced use cases, [401](#)
  - defining, [399](#), [400](#)
  - triggered alarms, viewing/acknowledging, [399–400](#)
- alerts, [397](#)
- allocating ports, [112–113](#)
- AMD, EVC modes, [133](#)
- anti-affinity rules, [83](#)
  - DRS, [369–370](#)
  - SDRS, [448–449](#)
- Appliance Shell, patching, [554–556](#)
- appliances, vCenter Server
  - migrating vCenter Server for Windows to vCenter Server appliance, [522–524](#)
  - upgrading, [519–522](#)
- applications
  - App Volumes, [217–218](#)
    - integration, [218–219](#)
    - use cases, [218](#)
  - AppDefense, [227](#), [277–278](#)
    - integration, [227–228](#)
    - use cases, [227](#)
  - monitoring
    - VAMI, [396](#)
    - VM, [150](#), [372](#)
  - vApps, [170](#)
- array-based failovers, iSCSI, [74](#)
- ATS (Atomic Test and Set), [70](#)
- ATS Only Flags, [70](#)
- attacks
  - 802.1q tagging attacks, [97](#)
  - denial-of-service attacks, [269](#)

- double-encapsulation attacks, [97](#)
- multicast brute-force attacks, [97](#)
- spanning tree attacks, [97](#)
- audit events, [397](#)
- authentication. *See also* [authorization](#); [security](#)
  - account lockouts, ESXi, [485–487](#)
  - ESXi passwords, [485–487](#)
  - groups, [474–475](#)
  - LDAP authentication sources, [309–310](#)
  - permissions, [242](#)
    - changing validation settings, [502](#)
    - editing, [476–477](#)
    - global permissions, [476](#)
    - setting, [475–476](#)
  - privileges, [475](#), [496–497](#)
  - proxies, configuring, [498–499](#)
  - roles, [475](#)
  - smart cards, ESXi, [499](#)
  - SSO, [242](#), [472](#)
    - configuring, [305](#)
    - enabling with Windows Session Authentication, [472–473](#)
    - enabling/disabling users, [310–311](#)
    - Enhanced Linked Mode, [474](#)
    - group authentication, [474–475](#)
    - identity sources, [305–307](#)
    - policy configuration, [311–312](#)
    - STS management, [473](#)
    - user authentication, [474–475](#)
  - users, [474–475](#)
  - VMAFD, [236](#)
  - vSphere Authentication Proxy, [257](#)
  - Windows Session Authentication, enabling SSO, [472–473](#)
- authorization. *See also* [authentication](#); [security](#)

- groups, [474–475](#)
- permissions, [242](#)
  - changing validation settings, [502](#)
  - editing, [476–477](#)
  - global permissions, [476](#)
  - setting, [475–476](#)
- privileges, [475](#), [496–497](#)
- roles, [475](#)
- SSO, [242](#), [472](#)
  - configuring, [305](#)
  - enabling with Windows Session Authentication, [472–473](#)
  - enabling/disabling users, [310–311](#)
  - Enhanced Linked Mode, [474](#)
  - group authentication, [474–475](#)
  - identity sources, [305–307](#)
  - policy configuration, [311–312](#)
  - STS management, [473](#)
  - user authorization, [474–475](#)
- users, [474–475](#)

Auto Deploy

- ESXi host installations, [292–297](#)
- security, [491](#)
- stateless caching, [292](#)

automation

- Automation Mode, DRS, [134](#)
- cloud automation, [28](#)
- SDRS
  - automation levels, [82](#)
  - overriding datastore cluster automation, [448](#)
- vRA, [209–213](#)
- vRealize Automation, [26–27](#)

average bandwidth, [101](#)

AWS, VMC on, [28](#), [226](#)

Azure VMware Solution, [226–227](#)

## **B**

backups

    vCenter Appliance File-Based Backup and Restore, [7](#)

    vCenter Server, [23](#), [514–517](#), [538–539](#)

bandwidth

    average bandwidth, [101](#)

    burst size, [101](#)

    inbound traffic shaping, [105](#)

    network resource pools, [106–108](#)

    peak bandwidth, [101](#)

    traffic shaping policies, [101](#)

base images, [534–535](#)

baselines, [527](#), [530–535](#)

basic multicasting filtering, [116–117](#)

behaviors, SDRS, [82–83](#)

binding ports, [112–113](#)

block primitives

    ATS, [70](#)

    ATS Only Flag, [70](#)

    VAAI, [70–71](#)

    Write Same (Zero), [71](#)

    XCOPY, [70](#)

blocking ports, [105](#)

boot devices, vSAN, [68](#)

bootups

    ESXi Quick Boot, [535–536](#)

    Secure Boot, ESXi, [258](#)

branches, snapshot trees, [177](#)

brute-force attacks (multicast), [97](#)

burst size, [101](#)

## C

- CA, VMCA as intermediate, [238–239](#)
- caching, stateless, [292](#)
- capacity reservation settings, vSphere HA, [420](#)
- CDP (Cisco Discovery Protocol), [117–118](#)
- certificates
  - client certificates
    - managing, [477–478](#)
    - vCenter Server, [261](#)
  - CSR, [238–239](#)
  - custom certificates, [241](#)
    - managing, [478–479](#)
    - VMCA, [237](#)
  - ESXi, [240](#)
    - changing certificate mode, [479–480](#)
    - custom certificates, [480](#)
    - expiration, [481](#)
    - host certificate modes, [241](#), [242](#)
    - management, [479–481](#)
    - switching to VMCA Mode, [480–481](#)
  - identity services, [236–237](#)
  - machine SSL certificates, [240](#), [241](#)
  - management, [477–482](#)
  - overview, [236–238](#)
  - requirements, [238–241](#)
  - solution user certificates, [240–241](#)
  - SSL certificates, verifying legacy ESXi hosts, [554](#)
  - unsupported certificates, VMCA, [238](#)
  - vCenter single sign-on SSL signing certificates, [240](#)
  - VECS, [236–237](#)
    - solution user certificate stores, [240–241](#)
    - stores, [303–304](#)

- VMAFD, [236](#)
- VMCA, [236–237](#), [239](#)
  - configuring, [303–305](#)
  - custom certificates, [237](#)
  - as intermediate CA, [237](#), [239](#)
  - management, [303–305](#)
  - management modes (recommended), [237–238](#)
  - unsupported certificates, [238](#)
- vmdir, [236](#)
- vmdir SSL certificates, [240](#)
- vSphere Virtual Machine Encryption certificates, [240](#)
- certification
  - exam preparation
    - “getting ready,” [603–604](#)
    - Pearson Vue, [604](#)
  - VMware Certification accounts, [604](#)
- charts, performance, [375](#)
  - advanced performance charts, [377–379](#)
  - overview performance charts, [375–377](#)
  - troubleshooting, [383](#)
- chipsets, [181](#)
- CIM, controlling access, [491–492](#)
- claim rules, multipathing management, [458](#)
- CLI (Command-Line Interfaces)
  - ESXCLI commands, [483–484](#)
  - installers, [298–299](#), [301–302](#)
  - PowerCLI
    - commands, [484–485](#)
    - VM management, [590–592](#)
- client certificates
  - management, [477–478](#)
  - vCenter Server, [261](#)
- cloning VM, [194](#), [572–573](#)



- cold clones, [194](#)
- hot clones, [194](#)
- instant clones, [195–196](#)
- linked clones, [194](#)
- rapid provisioning VM with templates, [195](#)
- cloud computing
  - automation, [28](#)
  - Azure VMware Solution, [226–227](#)
  - HCX, [224–226](#)
  - hybrid clouds, [28](#)
  - private clouds, [28](#)
  - vCenter Cloud account permissions, [210–213](#)
  - VCF, [28](#), [223–224](#)
  - VMC on AWS, [28](#), [226](#)
  - VMware vCloud Director, [28](#)
  - VMware vCloud Suite, [28](#)
- clusters, [167–168](#)
  - datastore clusters
    - requirements, [83–84](#)
    - SDRS, [81–84](#)
    - vSphere clusters versus, [131](#)
  - hosts, moving into clusters, [251](#)
  - images, importing/exporting, [538](#)
  - SDRS datastore cluster automation, overriding, [448](#)
  - user-defined vSAN clusters, [53](#)
  - vCenter HA clusters, managing, [557–558](#)
  - vSAN clusters
    - creating with Quickstart, [415](#)
    - encryption, [432–435](#)
    - expanding, [422–424](#)
    - increasing space efficiency, [430–432](#)
    - managing devices in clusters, [429–430](#)
    - persistent logging in vSAN clusters, [68](#)

- requirements, [66](#)
- vSphere clusters
  - configuring, [130](#)
  - configuring with Quickstart, [365–367](#)
  - creating, [364](#)
  - datastore clusters versus, [131](#)
  - DPM, [152–153](#)
  - DRS, [130–131](#), [134–139](#), [152](#), [368](#), [369–370](#), [384–385](#)
  - EVC, [130](#), [131–133](#), [367–368](#)
  - HA, [143–149](#), [370–371](#)
  - overview, [130–131](#)
  - Predictive DRS, [152](#), [370](#)
  - resource pools, [139–143](#), [368–369](#), [385–386](#)
- cold clones, [194](#)
- cold migrations, [186](#), [250](#)
- comments, RFC, [92](#)
- community nodes, PVLAN, [110](#)
- compatibility
  - hardware
    - compatibility checks, [537](#)
    - VM hardware, [180–182](#)
  - VM
    - compatibility options, [578](#), [579–580](#)
    - hardware, [180–182](#)
- compliance
  - compliance status (VM), vSAN, [52](#)
  - vROps, [275](#)
- components
  - vSAN component states
    - absent component state, [52](#)
    - degraded component state, [52](#)
  - vSphere
    - core components, [6](#)

- optional components, [6](#)
- compression, vSAN, [59](#), [60](#)
- compute requirements
  - vCenter Server, [14–15](#)
  - vSphere, [14–15](#)
- configuration files, [174–175](#)
- Config-vVol, [73](#)
- connecting devices, VM security, [267](#)
- consoles, opening to VM, [569–570](#)
- consumed capacity, vSAN, [51](#)
- content libraries, [7](#), [171–173](#), [594–595](#)
  - adding items, [598–599](#)
  - creating, [595](#)
  - permissions, [597](#)
  - publishing, [596](#)
  - subscriptions, [596](#)
  - synchronization options, [598](#)
  - VM deployments, [599](#)
- converting VM to templates, [573](#)
- copying/pasting VM security, [267](#)
- CPU, [181](#)
  - performance analysis, [379–383](#)
  - resources, adding to VM, [580–581](#)
  - troubleshooting
    - usage, [380](#)
    - utilization, [381](#)
- cross-datastore migrations, [186](#)
- cross-host migrations, [186](#)
- cross-vCenter Server migrations, [186–187](#)
- CSR (Certificate Signing Requests), [238–239](#)
- customizing
  - certificates, [241](#)
  - ESXi certificates, [480](#)

- management, [478–479](#)
- VMCA, [237](#)
- ESXi services, [493–494](#)
- guest OS on VM, [574–576](#)
- TCP/IP stacks, [122](#)

## D

- data centers, [166–167](#)
  - NSX, requirements, [26](#)
  - vRealize Suite, requirements, [26–27](#)
  - vSAN, requirements, [25–26](#)
  - vSphere Client data center-level management, [111](#)
- Data Locality, vSAN, [57](#)
- data transfers, vCenter Server, [519](#)
- database files, [179](#)
- databases, vCenter Server, [11](#), [297](#)
- datastores, [41](#), [169](#)
  - browser access, controlling, [261](#)
  - clusters
    - requirements, [83–84](#)
    - SDRS, [81–84](#)
    - vSphere clusters versus, [131](#)
  - cross-datastore migrations, [186](#)
  - NFS datastores, [43–45](#), [444–446](#)
  - PMem datastores, [455](#)
  - SDRS datastores
    - configuring, [447–449](#)
    - managing, [447–449](#)
    - overriding cluster automation, [448](#)
    - recommendations, [448](#)
  - utilization, troubleshooting, [381](#)
  - VMFS datastores, [41–43](#), [438–443](#)

- vSAN datastores, [45](#), [51](#)
  - extending across two sites, [427–428](#)
  - viewing, [418–419](#)
- vVols datastores, [45](#)
- Data-vVol, [73](#)
- deduplication, vSAN, [59](#), [60](#)
- default TCP/IP stacks, [121](#)
- degraded component state, vSAN, [52](#)
- DEK (Data Encryption Keys), [61–62](#), [270](#)
- delta disk files, [179](#)
- denial-of-service attacks, VM security, [269](#)
- deploying
  - Auto Deploy
    - ESXi host installations, [292–297](#)
    - security, [491](#)
    - stateless caching, [292](#)
  - vCenter Server, [298–303](#)
    - databases, [297](#)
    - PSC, [297–298](#)
  - VCSA deployments
    - with CLI installers, [301–302](#)
    - with GUI installers, [299–301](#)
- VM
  - from templates, [249](#), [574](#), [577](#), [585–586](#)
  - using content libraries, [599](#)
- vSAN, [54–59](#), [422](#)
- desktop virtualization
  - App Volumes, [217–219](#)
  - VMware Horizon, [215–217](#)
- device connections, VM security, [267](#), [269–270](#)
- device latency, troubleshooting, [382](#)
- directories
  - AD, [21](#), [307–309](#)

- ESXi hosts, joining to directory services, [257](#)
  - vmdir, [236](#)
  - VMware Directory Service, [11](#)
- DirectPath I/O, [118–119](#), [343](#)
- disabling
  - MOB, [490](#)
  - SSO users, [310–311](#)
  - unexposed features, network security, [266–267](#)
  - vSAN, [421](#)
- discovery protocols, [117–118](#)
- disk groups, vSAN, [51](#)
- disk shrinking, [267](#)
- distributed file locking, RDM, [39](#)
- distributed port groups, [103](#), [337–338](#), [353–354](#)
- Distributed Power Management. *See* [DPM](#)
- Distributed Resource Schedulers. *See* [DRS](#)
- DNS (Domain Name System), [21–22](#)
- domains
  - DNS, [21–22](#)
  - FQDN, [21–22](#)
  - vCenter Server
    - domain discovery, [21](#)
    - repointing to other domains, [558–560](#)
  - vCenter Single Sign-On domains, [11–12](#)
- double-encapsulation attacks, [97](#)
- DPM (Distributed Power Management), [7](#), [23–24](#), [152–153](#)
- drives (large-capacity), vSAN support, [54](#)
- dropped packets, troubleshooting, [383](#)
- DRS (Distributed Resource Scheduler), [7](#), [130–131](#), [134](#)
  - advanced options, [369–370](#)
  - affinity rules, [369–370](#)
  - anti-affinity rules, [369–370](#)
  - Automation Mode, [134](#)

- clusters, creating, [368](#)
- evacuation workflows, [136](#)
- Memory metric for load balancing, [135](#)
- migration sensitivity, [138–139](#)
- monitoring/managing resource usage, [384–385](#)
- network-aware DRS, [135](#)
- NVM support, [136](#)
- Predictive DRS, [152](#), [370](#)
- recent enhancements, [134–137](#)
- rules, [137–138](#)
- VM
  - distribution, [135](#)
  - initial placements, [135–136](#)
  - scoring, [136–137](#)
- DVD/CD-ROM drives, [181](#)
- dynamic name resolution, RDM, [39](#)

## E

- eager zeroed thick virtual disks, [79](#)
- editing
  - ESXi host profiles, [319](#)
  - OVF templates, [585–586](#)
  - permissions, [476–477](#)
  - VM, [583–585](#)
  - vSAN settings, [417–418](#)
- editions, vSphere, [8–10](#)
- elastic port allocation, [113](#)
- encrypted vMotion, [192](#)
- encryption
  - DEK, [61–62](#), [270](#)
  - Encrypted vSphere vMotion, [272–273](#)
  - KEK, [61–62](#), [270](#), [271](#)

- VM, [270–272](#), [505–506](#)
- vSAN, [61–62](#)
- vSAN clusters, [432–435](#)
- vSphere Virtual Machine Encryption certificates, [240](#)
- Enhanced Linked Mode, [12–13](#), [474](#)
- enhanced reservations, resource pools, [142](#)
- ephemeral binding, [113](#)
- erasure coding
  - RAID 5, [60–61](#)
  - RAID 6, [60–61](#)
  - vSAN, [59](#), [60–61](#)
- esxcli commands, [483–484](#)
  - HPP, [457](#)
  - multipathing, [456–457](#)
  - NMP, [456](#)
- ESXi (ESX Integrated)
  - account lockouts, [485–487](#)
  - AD and user management, [497–498](#)
  - base images, [534–535](#)
  - certificates, [240](#), [241–242](#)
    - changing certificate mode, [479–480](#)
    - custom certificates, [480](#)
    - expiration, [481](#)
    - management, [479–481](#)
    - switching to VMCA Mode, [480–481](#)
  - CIM, controlling access, [491–492](#)
  - configuring with host profiles, [317–318](#)
  - customizing services, [493–494](#)
  - ESXCLI commands, [483–484](#)
  - firewalls
    - configuring, [492–493](#)
    - ports, [255–256](#)
  - firmware updates, [536–537](#)



## hosts

- acceptance levels, [496](#)
- accessing, [257](#)
- advanced system settings, [321](#)
- assigning privileges, [496–497](#)
- certificate modes, [241](#), [242](#)
- configuring security, [482–483](#)
- installing, [286](#)
- installing, Auto Deploy installations, [292–297](#)
- installing, interactive installations, [286–288](#)
- installing, scripted installations, [288–292](#)
- joining to directory services, [257](#)
- kernel options, [321–322](#)
- managing, [540–542](#)
- profiles, applying, [318–319](#)
- profiles, applying permissions, [319–320](#)
- profiles, configuring ESXi with host profiles, [317–318](#)
- profiles, editing, [319](#)
- scripts and host configuration management, [483–485](#)
- syslog configurations, [405–407](#)
- TPM, [500–501](#)
- UEFI Secure Boot, [499–500](#)
- verifying legacy hosts with SSL certificates, [554](#)
- VIB, [496](#)
- VMware Tools, [320–321](#)
- vSAN encryption, [61–62](#)
- kernel options, [321–322](#)
- Lockdown Mode, [494–495](#)
  - logs, [401–404](#), [501](#)
- MOB
  - controlling access, [257–258](#)
  - disabling, [490](#)
- networking security recommendations, [490](#)

- passwords, [256](#), [485–487](#)
- PCI, [489](#)
- PCIe devices, [489](#)
- PowerCLI commands, [484–485](#)
- Quick Boot, [535–536](#)
- RDMA support, [453](#)
- Secure Boot, [258](#)
- security
  - configuring with host profiles, [482–483](#)
  - profiles, [254–255](#)
  - recommendations, [481–482](#)
- Shell security and SSH, [487–489](#)
- smart card authentication, [499](#)
- TPM chips, [258](#)
- upgrading, [524](#)
- vCenter Server security, [253](#)
  - controlling MOB access, [257–258](#)
  - ESXi firewall ports, [255–256](#)
  - ESXi host access, [257](#)
  - ESXi password hardening, [256](#)
  - ESXi Secure Boot, [258](#)
  - ESXi security profiles, [254–255](#)
  - TPM chips, [258](#)
  - vSphere Authentication Proxy, [257](#)
  - vSphere built-in features, [254](#)
- vSphere Authentication Proxy, [257](#)
  - web proxy settings, [490–491](#)

ESXi Server, [6](#)

- installing, [15–16](#), [17](#)
- network requirements, [20–21](#)
- ports, required, [20–21](#)
- storage requirements, [17](#)
- system requirements, [15–16](#)

## ESXTOP

- metrics, [393–395](#)

- monitoring/managing resources, [393–395](#)

EtherChannel, [93](#)

## Ethernet

- FCoE, [38](#)

  - packets. *See* [frames](#)

  - switches (physical), [92–93](#)

evacuation workflows, DRS, [136](#)

EVC (Enhanced vMotion Compatibility), [130](#), [131–132](#)

- Intel modes, [132–133](#)

- vSphere cluster configuration, [367–368](#)

events, [396](#)

- alerts, [397](#)

- audit events, [397](#)

- information events, [397](#)

- streaming to remote syslog servers, [398–399](#)

- types of, [397](#)

- viewing

  - System Event Log, [397](#)

  - in vSphere Client, [397](#)

- warning events, [397](#)

exam preparation

- exam day recommendations, [604–606](#)

- “getting ready,” [603–604](#)

- Pearson Vue, [604](#)

- taking the exam, [604–606](#)

expanding

- reservations, resource pools, [141](#)

- vSAN clusters, [422–424](#)

expiration, ESXi certificates, [481](#)

exporting/importing cluster images, [538](#)

EZT for shared disks, vSAN, [53](#)

## F

### Fabric

NVMe over Fabric, [46](#), [452–453](#)

SCSI over Fabric, [452](#)

### failovers

array-based failovers with iSCSI, [74](#)

FC failovers, [74](#)

HA, [143](#), [144](#)

host-based failovers with iSCSI, [74](#)

path failovers and VM, [74](#)

storage multipathing/failover, [74](#)

array-based failovers with iSCSI, [74](#)

FC failovers, [74](#)

host-based failovers with iSCSI, [74](#)

path failovers and VM, [74](#)

### failures

HA response to, [145–146](#)

tolerance, vSAN

PFTT, [57](#)

SFTT, [57](#)

fault domains, vSAN, [64–65](#), [426–427](#)

Fault Tolerance. *See* [FT](#)

FC (Fibre Channel), [37](#)

failovers, [74](#)

FC-NVMe, [451](#)

NVMe over FC requirements, [47](#)

FCD (First Class Discs), [45–46](#)

FC-NVMe (Fiber Channel over NVMe), [451](#)

FCoE (Fiber Channel over Ethernet), [38](#)

features, vSphere, [7](#)

Federation Services, AD, [313–314](#)

Fibre Channel. *See* [FC](#)

- File Services (vSAN), [54](#), [62–63](#), [436–438](#)
- file system operations, RDM, [39](#)
- file-based persistent volumes, vSAN, [54](#)
- files
  - RDM
    - distributed file locking, [39](#)
    - file permissions, [39](#)
  - VM
    - configuration files, [174–175](#)
    - snapshot files, [175](#)
    - structure of, [173–174](#)
    - virtual disk files, [175](#)
- filtering
  - I/O filters, encryption, [41](#), [271](#)
  - multicast filtering, [116](#)
    - basic multicasting filtering, [116–117](#)
    - multicast snooping, [117](#)
  - network traffic, [109–110](#)
- firewalls
  - ESXi firewalls
    - configuring, [492–493](#)
    - ports, [255–256](#)
  - networks security, [262](#)
- firmware updates, ESXi, [536–537](#)
- First Class Discs. *See* [FCD](#)
- fixed port allocation, [113](#)
- flat files, [178](#)
- folders, [167](#)
- Forged Transmits, [101](#)
- FQDN (Fully Qualified Domain Names), [21–22](#)
- frames, jumbo, [97–98](#)
- FT (Fault Tolerance), [7](#), [153–157](#), [373](#)

## G

- “getting ready,” exam preparation, [603–604](#)
- global permissions, [247](#), [476](#)
- GPU, VM support, [592–594](#)
- GRID models, VM, [593](#)
- groups, authentication, [474–475](#)
- guest OS installations on VM, [250](#), [574–576](#)
- guest user mappings, VM, [585](#)
- guests, shutting down, [572](#)
- GUI (Graphical User Interface)
  - installers, [298–301](#)
  - installing, [23](#)

## H

- HA (High Availability)
  - Admission Control, [146–148](#)
  - advanced options, [148–149](#)
  - benefits of, [144](#)
  - best practices, [151](#)
  - configuring
    - admission control, [371](#)
    - advanced options, [370](#)
    - HA clusters, [370–371](#)
  - detecting host issues, [144](#)
  - failovers, [144](#)
  - heartbeats, [146](#)
  - Proactive HA, [7](#), [151](#), [372](#)
  - requirements, [145](#)
  - response to failures, [145–146](#)
  - vCenter HA, [6](#), [14](#)
    - active nodes, [14](#)

- cluster management, [557–558](#)
- implementing, [316–317](#)
- passive nodes, [14](#)
- requirements, [24–25](#)
- witness nodes, [14](#)
- vCenter Server HA, [145](#), [157](#)
- vSphere clusters, [143](#)
  - configuring HA clusters, [370–371](#)
  - failovers, [143](#)
- vSphere HA, [7](#)
  - capacity reservation settings, [420](#)
  - configuring, [419–420](#)
  - vSAN, [419–420](#)
- hard disks, [181](#)
- hardware
  - compatibility checks, [537](#)
  - host hardware, monitoring/managing resources/health, [386–387](#)
- VM
  - configuring hardware, [578–583](#)
  - feature sets, [578–579](#)
  - hardware compatibility, [180–182](#)
  - vSAN requirements, [65–66](#)
- HCX (Hybrid Cloud Extension), [224](#)
  - integration, [225–226](#)
  - services, [224–225](#)
  - use cases, [225](#)
- Health Check, vDS, [115–116](#), [354](#)
- healthy object state, vSAN, [52](#)
- heartbeats, [146](#)
- High Availability. *See* [HA](#)
- home namespace (VM), vSAN, [52](#)
- Host Agent, vCenter Server, [11](#)
- hosts, [168–169](#)

## ESXi hosts

- acceptance levels, [496](#)
- advanced system settings, [321](#)
- assigning privileges, [496–497](#)
- configuring security with host profiles, [482–483](#)
- installing, [286](#)
- installing, Auto Deploy installations, [292–297](#)
- installing, interactive installations, [286–288](#)
- installing, scripted installations, [288–292](#)
- kernel options, [321–322](#)
- managing, [540–542](#)
- profiles, applying, [318–319](#)
- profiles, applying permissions, [319–320](#)
- profiles, configuring ESXi with host profiles, [317–318](#)
- profiles, editing, [319](#)
- scripts and host configuration management, [483–485](#)
- TPM, [500–501](#)
- UEFI Secure Boot, [499–500](#)
- verifying legacy hosts with SSL certificates, [554](#)
- VIB, [496](#)
- VMware Tools, [320–321](#)

failovers, iSCSI, [74](#)

hardware, monitoring/managing resources/health, [386–387](#)

host physical network adapters, managing with vDS, [351](#)

issues, detecting with HA, [144](#)

moving into clusters, [251](#)

profiles, [7](#), [170–171](#), [482–483](#)

### vDS

- adding hosts, [350–351](#)

- managing host physical network adapters with vDS, [351](#)

- removing hosts, [352](#)

hot clones, [194](#)

hot cross-host migrations. *See* [vMotion](#)



- hot migrations, [186](#)
- Hot-Plug plug-in (NVMe), vSAN, [53](#)
- HPP (High Performance Plug-Ins)
  - esxcli commands, [457](#)
  - NVMe, [454](#)
  - VMware HPP, [47](#)
    - best practices, [48](#)
    - path selection schemes, [47–48](#)
    - vSphere support, [47](#)
- HTML5-based vSphere Client, [8](#)
- hybrid clouds. *See* [cloud computing](#); [HCX](#)

## I

- IDE 0, [181](#)
- IDE 1, [181](#)
- identification
  - NPIV, [40](#)
  - VLAN ID, standard port groups, [333](#)
- Identity Federation, [313–314](#)
- identity services, [236](#)
  - VECS, [236–237](#), [240–241](#)
  - VMAFD, [236](#)
  - VMCA, [236–237](#), [239](#)
    - custom certificates, [237](#)
    - as intermediate CA, [237](#), [239](#)
    - management modes (recommended), [237–238](#)
    - unsupported certificates, [238](#)
  - vmdir, [236](#)
- identity sources
  - AD, [307–309](#)
  - SSO, [305–307](#)
- IEEE 802.1ax, [93](#)

IEEE 802.3ad, [93](#)

images

- cluster images, importing/exporting, [538](#)

- ESXi base images, [534–535](#)

Improved Virtual Disks. *See* [FCD](#)

inbound traffic shaping, vDS, [105](#)

information events, [397](#)

infrastructure services, vSphere, [21–23](#)

installing

- CLI installers, [298–299](#), [301–302](#)

- ESXi, [258](#)

- ESXi hosts, [286](#)

  - Auto Deploy installations, [292–297](#)

  - interactive installations, [286–288](#)

  - scripted installations, [288–292](#)

- ESXi Server, [15–16](#), [17](#)

- guest OS on VM, [250](#)

- GUI, [23](#)

- GUI installers, [298–301](#)

- VIB, ESXi hosts, [496](#)

- VMware Enhanced Authentication plug-ins, [303](#)

- VMware Tools, [320–321](#), [570–571](#)

vSphere

- deploying vCenter Server components, [297–305](#)

- ESXi hosts, [286–297](#)

- initial vSphere configuration, [315–322](#)

- SSO configurations, [305–314](#)

instant clones, [195–196](#)

integrated file services, vSAN, [53](#)

Intel

- EVC modes, [132–133](#)

- SGX, VM security, [505](#)

interactive ESXi host installations, [286–288](#)

- intermediate CA, VMCA as, [238–239](#)
- inter-VM anti-affinity rules, [448–449](#)
- inventories
  - hierarchies, [243–244](#)
    - vCenter Server inventory configuration, [315–316](#)
- inventory objects, [166](#)
  - clusters, [167–168](#)
  - data centers, [166–167](#)
  - datastores, [169](#)
  - folders, [167](#)
  - hosts, [168–169](#)
  - networks, [169](#)
  - resource pools, [168](#)
  - templates, [170](#)
  - vApps, [170](#)
  - VM, [169](#)
- I/O (Input/Output)
  - DirectPath I/O, [118–119](#), [343](#)
  - filters, [41](#), [271](#)
  - latency load balancing, SDRS, [81–82](#)
  - NIOC, [84](#), [105–106](#)
    - configuring, [340–341](#)
    - resource pools, [106–108](#)
  - redirects, vSAN, [53](#)
  - requests, PSA, [78](#)
  - SIOC, [84](#)
    - configuring, [449–451](#)
    - management, [449–451](#)
    - monitoring shares, [450](#)
    - setting shares/limits, [450](#)
    - thresholds, [450–451](#)
  - SR-IOV, [119–121](#), [343–345](#)
- IP addresses, [92](#)

IPsec (Internet Protocol Security), [262–263](#)

iSCSI, [37](#)

- array-based failovers, [74](#)

- host-based failovers, [74](#)

isolation

- isolated nodes, PVLAN, [110](#)

- networks security, [262](#)

IVD (Improved Virtual Disks). *See* [FCD](#)

## J

JSON templates, VCSA deployments with CLI installers, [302](#)

jumbo frames, [97–98](#)

## K

KEK (Key Encryption Keys), [61–62](#), [270](#), [271](#)

kernels, ESXi, [321–322](#)

Key Management Servers, security, [502](#)

keyboards, [181](#)

KMS, vSAN encryption, [61–62](#)

Kubernetes, [45–46](#), [54](#)

## L

LACP (Link Aggregation Control Protocol), [93](#), [113–115](#)

LAG (Link Aggregation Groups), [346–349](#)

LAN (Local Area Networks). *See* [PVLAN](#); [VLAN](#)

large-capacity drives, vSAN support, [54](#)

latency

- sensitivity, [392](#)

- troubleshooting

  - device latency, [382](#)

- dropped packets, [383](#)
  - VMkernel, [382](#)
- lazy zeroed thick virtual disks, [79](#)
- LDAP (Lightweight Data Access Protocol)
  - authentication sources, [309–310](#)
  - OpenLDAP, [309–310](#)
- libraries, content, [7](#), [171–173](#), [594–595](#)
  - adding items, [598–599](#)
  - creating, [595](#)
  - permissions, [597](#)
  - publishing, [596](#)
  - subscriptions, [596](#)
  - synchronization options, [598](#)
  - VM deployments, [599](#)
- licenses
  - License Service, vCenter Server, [11](#)
  - vSAN, [67](#), [418](#)
  - vSphere, [9](#)
- Lifecycle Manager (vSphere), [157](#)
  - baselines, [530–535](#)
  - definitions, [532–533](#)
  - ESXi
    - firmware updates, [536–537](#)
    - hosts, [526–529](#)
    - Quick Boot, [535–536](#)
  - hardware compatibility checks, [537](#)
  - remediation settings, [528](#)
  - UMDS, [529–530](#)
  - vSAN, [53](#), [54](#)
- limits
  - resource pools, setting limits with, [141](#)
  - VM resources, monitoring/managing, [389–390](#)
- linear snapshots, [176](#)

linked clones, [194](#)

links

- Enhanced Linked Mode, [12–13](#), [474](#)

- LACP, [93](#), [113–115](#)

- LAG, [346–349](#)

- LLDP, [117–118](#)

LLDP (Link Layer Discovery Protocol), [117–118](#)

load balancing

- Memory metric for load balancing, DRS, [135](#)

- SDRS

  - I/O latency load balancing, [81–82](#)

  - ongoing balancing, [81](#)

  - space utilization load balancing, [81](#)

load-based NIC teaming, [105](#)

local storage, [37](#)

Lockdown Mode, ESXi, [494–495](#)

lockouts (account), [485–487](#)

logs

- ESXi logs, [401–404](#), [501](#)

- levels, [404–405](#)

- syslog, ESXi host configurations, [405–407](#)

- System Event Log, viewing, [397](#)

- system logs, uploading to VMware, [404](#)

- vCenter Server logs, [404](#)

- VM security, [267–268](#)

- vRLI, [27](#), [208–209](#), [407](#)

- vSAN, persistent logging in vSAN clusters, [68](#)

lookups, vCenter Lookup Service, [11](#)

LSO (Large Segmentation Offload). *See* [TSO](#)

LUN (storage devices), [37](#)

## M

- MAC Address Changes, [100–101](#)
- MAC addresses, [92](#)
- machine SSL certificates, [240](#), [241](#)
- Maintenance Mode
  - DRS and evacuation workflows, [136](#)
  - vSAN, [424–426](#)
- managing, [463–464](#)
  - client certificates, [477–478](#)
  - custom certificates, [478–479](#)
  - data center-level management, vSphere Client, [111](#)
  - DPM, [7](#), [23–24](#), [152–153](#)
  - DRS resource usage, [384–385](#)
  - ESXi
    - AD and user management, [497–498](#)
    - certificates, [479–481](#)
    - hosts, [540–542](#)
      - scripts and host configuration management, [483–485](#)
  - host hardware resources/health, [386–387](#)
  - host physical network adapters, managing with vDS, [351](#)
  - inventory objects, [166–170](#)
  - Key Management Servers, security, [502](#)
  - multipathing, [456](#)
    - claim rules, [458](#)
    - esxcli commands, [456–457](#)
    - with vSphere Client, [457–458](#)
  - NFS datastores, [444–446](#)
  - NVMe, [451–454](#)
  - OVF templates, [589](#)
  - RDM, [439–446](#)
  - resource pool resource usage, [385–386](#)
  - SAN management agents, RDM, [40](#)
  - SDRS, [447–449](#)
  - security certificates, [477–482](#)

- client certificates, [477–478](#)
- custom certificates, [478–479](#)
- ESXi certificates, [479–481](#)
- VMCA, recommended management modes, [237–238](#)
- SPBM, [40–41](#), [79](#), [459–461](#)
  - applying storage policies to VM, [462–463](#)
  - VASA, managing storage providers, [462](#)
  - VASA, registering storage providers, [461](#)
- SRM, [221–222](#)
- SSO, [472](#)
  - enabling with Windows Session Authentication, [472–473](#)
  - Enhanced Linked Mode, [474](#)
  - STS management, [473](#)
- STS, [473](#)
- VAMI, [396](#)
- vCenter HA clusters, [557–558](#)
- vCenter Server
  - backups, [514–517](#), [538–539](#)
  - data transfers, [519](#)
  - importing/exporting cluster images, [538](#)
  - managing, [542–543](#)
  - migrating vCenter Server for Windows to vCenter Server appliance, [522–524](#)
  - patching with Appliance Shell, [554–556](#)
  - patching with VAMI, [554–556](#)
  - repointing to other domains, [558–560](#)
  - restores, [538–539](#)
  - Update Planner, [524–526](#)
  - updating, [554–557](#)
  - upgrading appliances, [519–522](#)
  - upgrading ESXi, [524](#)
  - upgrading VM, [524](#)
  - with VAMI, [543–547](#)



with vSphere Client, [547–554](#)

## VM

configuring hardware, [578–583](#)

editing options, [583–585](#)

editing OVF templates, [585–586](#)

GRID models, [593](#)

guest user mappings, [585](#)

migrating, [587–589](#)

PowerCLI, [590–592](#)

snapshots, [586–587](#)

VBS, [590](#)

vGPU support, [592–594](#)

virtual disks, [581–583](#)

vSGA models, [593](#)

## VM resources, [393](#)

admission control, [390–391](#)

alarms, [399–401](#)

ESXTOP, [393–395](#)

events, [396–399](#)

impact of VM configurations, [392–393](#)

latency sensitivity, [392](#)

limits, [389–390](#)

metrics, [388](#)

reservations, [389–390](#)

shares, [389–390](#)

VAMI, [396](#)

VIMTOP, [396](#)

VMware tools, [391–392](#)

Windows Perfmon, [391–392](#)

## VMCA, [303–305](#)

## VMFS datastores, [438–443](#)

## vSAN cluster devices, [429–430](#)

## vSphere 7.0 upgrades, [517–518](#)

- vSphere Lifecycle Manager, [157](#), [315](#)
  - baselines, [530–535](#)
  - definitions, [532–533](#)
  - ESXi firmware updates, [536–537](#)
  - ESXi hosts, [526–529](#)
  - ESXi Quick Boot, [535–536](#)
  - hardware compatibility checks, [537](#)
  - remediation settings, [528](#)
  - UMDS, [529–530](#)
- vSphere resources, [373](#)
- vTA, [502–504](#)
- mappings, guest user, [585](#)
- marking policies, [109–110](#)
- memory, [181](#)
  - NVM, DRS support, [136](#)
  - NVMe
    - FC-NVMe, [451](#)
    - HPP, [454](#)
    - managing, [451–454](#)
    - over Fabric, [452–453](#)
    - over PCIe, [451](#)
    - over RDMA, [451](#), [453](#)
  - PMem
    - datastores, [455](#)
    - devices, [454](#)
    - vPMem, [454](#)
    - vPMemDisk, [455](#)
  - RDMA
    - ESXi and RDMA support, [453](#)
    - NVMe over RDMA, [451](#), [453](#)
  - usage, troubleshooting, [380–381](#)
  - vSAN performance, [53](#)
- memory files, [179](#)

- Memory metric for load balancing, DRS, [135](#)
- memory objects, vSAN, [52](#)
- Mem-vVol, [73](#)
- metrics
  - ESXTOP, [393–395](#)
  - VM resources, monitoring/managing, [388](#)
  - vSphere performance metrics, [374](#)
- microsegmentation, [276–277](#)
- migrating
  - DRS migration sensitivity, [138–139](#)
  - vCenter Server for Windows to vCenter Server appliance, [522–524](#)
  - VM, [185–186](#), [587–589](#)
    - cold migrations (relocation), [186](#), [250](#)
    - cross-datastore migrations, [186](#)
    - cross-host migrations, [186](#)
    - cross-vCenter Server migrations, [186–187](#)
    - hot migrations, [186](#)
    - limitations, [187–188](#)
    - with Storage vMotion, [251](#)
    - to vDS, [353](#)
    - with vMotion, [250](#)
  - VMkernel network adapters to vDS, [352](#)
  - vMotion, [189–191](#)
    - data flow, [191](#)
    - encrypted vMotion, [192](#)
    - multi-NIC vMotion, [190](#)
    - storage vMotion, [192–193](#)
- mirroring ports, [111–112](#), [345–346](#)
- MOB (Managed Object Browsers)
  - controlling access, [257–258](#)
  - disabling, [490](#)
- modifying vDS, [336](#)
- monitoring

- applications with VM, [150](#), [372](#)
- DRS resource usage, [384–385](#)
- host hardware resources/health, [386–387](#)
- network monitoring policies, [108](#)
- port states, [111](#), [353–354](#)
- resource pool resource usage, [385–386](#)
- SIOC shares, [450](#)
- vCenter Server, [542–543](#)
  - VAMI, [543–547](#)
  - vSphere Client, [547–554](#)
- VM, [372](#)
- VM resources, [393](#)
  - admission control, [390–391](#)
  - alarms, [399–401](#)
  - ESXTOP, [393–395](#)
  - events, [396–399](#)
  - impact of VM configurations, [392–393](#)
  - latency sensitivity, [392](#)
  - limits, [389–390](#)
  - metrics, [388](#)
  - reservations, [389–390](#)
  - shares, [389–390](#)
  - VAMI, [396](#)
  - VIMTOP, [396](#)
  - VMware tools, [391–392](#)
  - Windows Perfmon, [391–392](#)
- vSphere resources, [373](#)
- moving
  - hosts into clusters, [251](#)
  - VM into resource pools, [250](#)
- MPP (third-party), VMware native modules and PSA, [78](#)
- MTU (Maximum Transmission Units), [97–98](#)
- multicast brute-force attacks, [97](#)

- multicast filtering, [116](#)
  - basic multicasting filtering, [116–117](#)
  - multicast snooping, [117](#)
- multi-NIC vMotion, [190](#)
- multipathing, [456](#)
  - esxcli commands, [456–457](#)
  - managing
    - claim rules, [458](#)
    - with vSphere Client, [457–458](#)
  - NMP, esxcli commands, [457](#)
  - storage multipathing/failover, [74](#)
    - array-based failovers with iSCSI, [74](#)
    - FC failovers, [74](#)
    - host-based failovers with iSCSI, [74](#)
    - path failovers and VM, [74](#)

## N

- NAI primitives, VAAI, [71](#)
- naming conventions, RDM
  - dynamic name resolution, [39](#)
  - user-friendly persistent names, [39](#)
- NAS/NFS, [38](#)
- NetFlow, [108](#), [336–337](#)
- Network File Systems. *See* [NFS](#)
- Network Time Protocol. *See* [NTP](#)
- network-aware DRS, [135](#)
- networks, [169](#)
  - adapters, [181](#)
    - host physical network adapters, managing with vDS, [351](#)
    - VMkernel network adapters, migrating to vDS, [352](#)
  - bandwidth
    - average bandwidth, [101](#)

- burst size, [101](#)
- inbound traffic shaping, [105](#)
- peak bandwidth, [101](#)
- resource pools, [106–108](#)
- CDP, [117–118](#)
- data center-level management, vSphere Client, [111](#)
- DirectPath I/O, [118–119](#)
- discovery protocols, [117–118](#)
- ESXi
  - security recommendations, [490](#)
  - server requirements, [20–21](#)
- EtherChannel, [93](#)
- IEEE 802.1ax, [93](#)
- IEEE 802.3ad, [93](#)
- IP addresses, [92](#)
- LACP, [93](#)
- LLDP, [117–118](#)
- MAC addresses, [92](#)
- marking policies, [109–110](#)
- monitoring policies, [108](#)
- MTU, [97–98](#)
- multicast filtering, [116](#)
  - basic multicasting filtering, [116–117](#)
  - multicast snooping, [117](#)
- NetFlow, [108](#)
- NIC
  - load-based NIC teaming, [105](#)
  - teaming policies, [98–100](#)
  - vNIC, [93–94](#)
- NIOC, [84](#), [105–108](#)
- NSX Data Center, [228–229](#)
- opaque networks, [18](#), [93](#)
- physical Ethernet switches, [92–93](#)

- physical networks, [17](#), [92](#), [351](#)
- ports
  - allocating, [112–113](#)
  - binding, [112–113](#)
  - mirroring, [111–112](#)
  - state monitoring, [111](#)
  - vMotion, [111](#)
- PVLAN, [110](#), [342](#)
- resource allocation policies, NIOC, [105–106](#)
- resource pools, [106–108](#), [341–342](#)
- RFC, [92](#)
- security, [262](#)
  - firewalls, [262](#)
  - IPsec, [262–263](#)
  - isolation, [262](#)
  - policies, [100–101](#), [264–265](#)
  - recommendations, [263–264](#)
  - segmentation, [262](#)
- SR-IOV, [119–121](#)
- TCP/IP, [92](#)
- terminology, [92–93](#)
- traffic
  - filtering, [109–110](#)
  - shaping policies, [101](#)
- TSO, [118](#)
- vCenter Server requirements, [18–20](#)
- vDS, [94](#)
  - architecture, [102](#)
  - distributed port groups, [103](#)
  - Health Check, [115–116](#)
  - inbound traffic shaping, [105](#)
  - LACP, [113–115](#)
  - marking policies, [109–110](#)

- multicast filtering, [116–117](#)
- multicast snooping, [117](#)
- NetFlow, [108](#)
- network policies, [104–105](#)
- port blocking policies, [105](#)
- port state monitoring, [111](#)
- teaming policies, [99](#)
- traffic filtering, [109–110](#)
- uplink port groups, [103](#)
- vSphere Client data center-level management, [111](#)
- vSS comparison, [103–104](#)
- virtual networks, [17](#)
  - advanced features, [355–356](#)
  - DirectPath I/O, [343](#)
  - distributed port groups, [337–338](#), [353–354](#)
  - LAG, [346–349](#)
  - network resource pools, [341–342](#)
  - NIOC, [340–341](#)
  - NSX Data Center, [228–229](#)
  - policies, [355–356](#)
  - port mirroring, [345–346](#)
  - PVLAN, [342](#)
  - SR-IOV, [343–345](#)
  - standard port groups, [332–334](#)
  - TCP/IP stacks, [339–340](#)
  - vDS, [334–337](#), [351–354](#)
  - VMkernel adapters, [338–339](#)
  - vSS, [330–332](#)
- virtual ports, [94](#)
- virtual switches, [94](#)
- VLAN, [94–95](#), [101–102](#)
  - PVLAN, [110](#), [342](#)
  - VLAN ID, standard port groups, [333](#)



## VMkernel

- adapter settings, [122](#)

- TCP/IP stacks, [121–122](#)

VMkernel TCP/IP networking layer, [18](#)

vNIC, [93–94](#)

vRNI, [27](#), [214–215](#)

## vSAN

- best practices, [67–68](#)

- characteristics, [414](#)

- configuring, [419–420](#)

- creating clusters with Quickstart, [415](#)

- deploying with vCenter Server, [422](#)

- disabling, [421](#)

- editing settings, [417–418](#)

- encryption in clusters, [432–435](#)

- expanding clusters, [422–424](#)

- extending datastores across two sites, [427–428](#)

- fault domains, [426–427](#)

- File Services, [436–438](#)

- increasing space efficiency in clusters, [430–432](#)

- licensing, [418](#)

- Maintenance Mode, [424–426](#)

- managing devices in clusters, [429–430](#)

- manually enabling, [416–417](#)

- preparing, [414](#)

- requirements, [67](#)

- restarting, [421–422](#)

- shutting down, [421–422](#)

- storage policies, [435–436](#)

- viewing datastores, [418–419](#)

- viewing storage providers, [436](#)

- vSphere HA, [419–420](#)

## vSphere

- network requirements, [17–21](#)
- segmenting, [18](#)
- standard switches, [18](#)
- vSphere Client data center-level management, [111](#)
- vSS, [94](#), [95–97](#)
  - network policies, [98](#)
  - vDS comparison, [103–104](#)
- NFS (Network File Systems)
  - datastores, [43–45](#), [444–446](#)
  - NAS/NFS, [38](#)
- NIC (Network Interface Cards)
  - DirectPath I/O, [118–119](#)
  - load-based NIC teaming, [105](#)
  - multi-NIC vMotion, [190](#)
  - teaming policies, [98–100](#)
  - TSO, [118](#)
  - vNIC, [93–94](#)
- NIOC (Network I/O Control), [84](#), [105–106](#)
  - configuring, [340–341](#)
  - resource pools, [106–108](#)
- NMP (Native Multipathing Plug-ins), [75–76](#), [78](#), [456](#)
- notifications
  - alarms
    - actions, [401](#)
    - advanced use cases, [401](#)
    - creating definitions, [400](#)
    - definition elements, [399](#)
    - viewing/acknowledging triggered alarms, [399–400](#)
  - VCG notification service, vSAN, [54](#)
- NPIV (N-Port ID Virtualization), [40](#)
- NSX, [8](#)
  - requirements, [26](#)
  - VMware, [276–277](#)

NSX Data Center, [228–229](#)

integration, [229](#)

use cases, [229](#)

NTP (Network Time Protocol), [22–23](#)

NVDIMM

controllers, [182](#)

devices, [182](#)

NVM (Non-Volatile Memory), DRS support, [136](#)

NVMe (Non-Volatile Memory Express), [46](#)

controllers, [182](#)

FC-NVMe, [451](#)

Hot-Plug plug-in, [53](#)

HPP, [454](#)

managing, [451–454](#)

over Fabric, [46](#), [452–453](#)

over FC requirements, [47](#)

over PCIe, [451](#)

over PCIe requirements, [46](#)

over RDMA, [451](#), [453](#)

over RDMA (RoCE Version 2) requirements, [46](#)

VMware HPP, [47](#)

best practices, [48](#)

path selection schemes, [47–48](#)

vSphere support, [47](#)

## O

objects

inventory hierarchies, [243–244](#)

states, vSAN

healthy object state, [52](#)

unhealthy object state, [52](#)

storage, vSAN, [51](#)

- Observer (vSAN), [53](#)
- opaque networks, [18](#), [93](#)
- OpenLDAP, [309–310](#)
- optimizing performance, [379–383](#)
- OS, guest installations on VM, [250](#), [574–576](#)
- Other-vVol, [73](#)
- OVA templates, deploying VM, [577](#)
- overview performance charts, [375–377](#)
- OVF templates
  - deploying VM, [577](#)
  - editing details, [585–586](#)
  - managing, [589](#)

## P

- packets
  - dropped packets, troubleshooting, [383](#)
  - Ethernet packets. *See* [frames](#)
- parallel ports, [181](#)
- passive nodes, vCenter HA, [14](#)
- passwords
  - ESXi, [485–487](#)
  - ESXi password hardening, [256](#)
  - SSO password policies, vCenter Server security, [260](#)
- patching, vCenter Server
  - Appliance Shell, [556–557](#)
  - VAMI, [554–556](#)
- path failovers, VM, [74](#)
- PCI (Peripheral Component Interconnect)
  - controllers, [181](#)
  - devices, [182](#)
  - ESXi, [489](#)
- PCIe (PCI express)

- devices, ESXi, [489](#)
  - NVMe over PCIe, [46](#), [451](#)
- peak bandwidth, [101](#)
- Pearson Vue
  - exam preparation, [604](#)
  - requirements, [604](#)
- Perfmon, Windows, [391–392](#)
- performance
  - charts
    - advanced performance charts, [377–379](#)
    - overview performance charts, [375–377](#)
    - troubleshooting, [383](#)
  - counters, VM security, [268](#)
  - CPU performance analysis, [379–383](#)
  - optimizing, [379–383](#)
  - troubleshooting, [379–383](#)
  - vSAN, memory as performance service metric, [53](#)
  - vSphere
    - charts, [375–379](#)
    - metrics, [374](#)
- permissions
  - applying with vCenter Server, [251–253](#)
  - authentication, [242](#)
  - authorization, [242](#)
  - best practices, [248](#)
  - content libraries, [597](#)
  - diagrams, [246–247](#)
  - editing, [476–477](#)
  - ESXi hosts, [319–320](#)
  - file permissions, RDM, [39](#)
  - global permissions, [247](#), [476](#)
  - hosts, moving into clusters, [251](#)
  - inventory hierarchies, [243–244](#)

- privileges, [244–245](#), [246](#), [248–251](#)
- roles, [245–246](#), [248](#)
- setting, [475–476](#)
- validation settings, changing, [502](#)
- vCenter Cloud account permissions, [210–213](#)
- VM
  - cold migration (relocation), [250](#)
  - creating, [249](#)
  - deploying from templates, [249](#)
  - guest OS installations, [250](#)
  - migrating with Storage vMotion, [251](#)
  - migrating with vMotion, [250](#)
  - moving into resource pools, [250](#)
  - taking snapshots, [249](#)
- persistent logging in vSAN clusters, [68](#)
- persistent volumes (file-based), vSAN, [54](#)
- PFTT (Primary Level of Failures to Tolerate), [57](#)
- physical compatibility mode, RDM, [39](#)
- physical Ethernet switches, [92–93](#)
- physical networks, [17](#), [92](#), [351](#)
- planning
  - fault domain planning, vSAN, [64–65](#)
  - Update Planner, [524–526](#)
  - VSAN, [63–64](#)
- Platform Service Controllers. *See* [PSC](#)
- plug-ins
  - HPP
    - esxcli commands, [457](#)
    - NVMe, [454](#)
  - NMP, esxcli commands, [457](#)
  - NVMe Hot-Plug plug-in, [53](#)
  - PSP, PSA, [76–77](#)
  - SATP, PSA, [76](#)

- vCenter Server plug-ins, [11](#)
- VMware Enhanced Authentication plug-ins, [303](#)
- VMware HPP, [47](#)
  - best practices, [48](#)
  - path selection schemes, [47–48](#)
  - vSphere support, [47](#)

## PMem

- datastores, [455](#)
- devices, [454](#)
- vPMem, [454](#)
- vPMemDisk, [455](#)

Pods, vSphere with Kubernetes, [45–46](#)

pointing devices, [182](#)

## ports

- allocating, [112–113](#)
- binding, [112–113](#)
- blocking policies, [105](#)
- distributed port groups, [103](#), [337–338](#), [353–354](#)
- ESXi firewall ports, [255–256](#)
- ESXi Server, required ports, [20–21](#)
- mirroring, [111–112](#), [345–346](#)
- network resource pools, [106–108](#)
- NPIV, [40](#)
- parallel ports, [181](#)
- resource allocation policies, NIOC, [105–106](#)
- serial ports, [182](#)
- standard port groups
  - configuring, [332–334](#)
  - creating, [332–334](#)
- states
  - monitoring, [111](#), [353–354](#)
  - vMotion, [111](#)
- uplink port groups, [103](#)

- vCenter Server, restricting access, [18–20](#)
- virtual ports, [94](#)
- post-revert snapshot trees, [176](#)
- power management, DPM, [7](#), [23–24](#), [152–153](#)
- PowerCLI (VMware)
  - commands, [484–485](#)
  - VM management, [590–592](#)
  - vSAN, [53](#)
- powering on VM, [569](#)
- Predictive DRS, [152](#), [370](#)
- preparing for exams
  - exam day recommendations, [604–606](#)
  - “getting ready,” [603–604](#)
  - Pearson Vue, [604](#)
  - taking the exam, [604–606](#)
- private clouds. *See* [cloud computing](#)
- privileges, [244–245](#), [246](#)
  - administrative privileges (vCenter Server), restricting, [260–261](#)
  - ESXi host assignments, [496–497](#)
  - required privileges for common tasks, [248–251](#)
  - vCenter Server, [475](#)
- Proactive HA, [7](#), [151](#), [372](#)
- Promiscuous Mode, [100](#)
- promiscuous nodes, PVLAN, [110](#)
- provisioning
  - rapid provisioning VM with templates, [195](#)
  - TCP/IP stacks, [122](#)
  - thin provisioning, vSAN, [59](#)
  - virtual disks, [183](#), [581](#)
- proxies (authentication), configuring, [498–499](#)
- PSA (Pluggable Storage Architectures), [74–75](#)
  - I/O requests, [78](#)
  - PSP, [76–77](#)



SATP, [76](#)

tasks, [77–78](#)

VMware native modules and third-party MPP, [78](#)

VMware NMP, [75–76](#), [78](#)

PSC (Platform Service Controllers), vCenter Server, [10](#), [11](#), [297–298](#)

PSP (Path Selection Plug-ins), [76–77](#)

public clouds. *See* [cloud computing](#)

publishing content libraries, [596](#)

PVLAN (Private VLAN), [110](#), [342](#)

community nodes, [110](#)

isolated nodes, [110](#)

promiscuous nodes, [110](#)

## Q

Quick Boot (ESXi), [535–536](#)

Quickstart

vSAN clusters, creating, [415](#)

vSphere cluster configuration, [365–367](#)

## R

RAID 5 erasure coding, [60–61](#)

RAID 6 erasure coding, [60–61](#)

rapid provisioning VM with templates, [195](#)

RDM (Raw Device Mappings)

benefits of, [39–40](#)

diagrams, [38](#)

distributed file locking, [39](#)

dynamic name resolution, [39](#)

file permissions, [39](#)

file system operations, [39](#)

management, [439–446](#)

- NPIV, [40](#)
- physical compatibility mode, [39](#)
- SAN management agents, [40](#)
- snapshots, [39](#)
- use cases, [39](#)
- user-friendly persistent names, [39](#)
- virtual compatibility mode, [38–39](#)
- VM, [582–583](#)
- vMotion, [40](#)
- RDMA (Remote Direct Memory Access)
  - ESXi and RDMA support, [453](#)
  - NVMe over RDMA, [451](#), [453](#)
  - NVMe over RDMA (RoCE Version 2) requirements, [46](#)
- Ready Node (vSAN), [53](#)
- recovery, SRM, [221–222](#)
  - relocation (cold migrations), [186](#), [250](#)
- remote syslog servers, streaming events to, [398–399](#)
- repair objects, vSAN witness deployments, [54](#)
- replication
  - VRMS, [24](#)
  - VRS, [24](#)
  - vSphere Replication, [206](#), [219–221](#)
  - vSphere requirements, [24](#)
- Replication objects (vSphere), vSAN, [54](#)
- repointing vCenter Server to other domains, [558–560](#)
- requirements
  - ESXi Server
    - network requirements, [20–21](#)
    - storage requirements, [17](#)
    - system requirements, [15–16](#)
  - HA, [145](#)
  - NSX, [26](#)
  - Pearson Vue, [604](#)

- SDDC, [25–27](#)
- security certificates, [238–241](#)
- vCenter HA, [24–25](#)
- vCenter Server
  - compute requirements, [14–15](#)
  - network requirements, [18–20](#)
  - storage requirements, [16](#)
- vRealize Suite, [26–27](#)
- vSAN, [25–26](#)
- vSphere
  - compute requirements, [14–15](#)
  - DPM, [23–24](#)
  - GUI installer, [23](#)
  - network requirements, [17–21](#)
  - replication requirements, [24](#)
  - SDDC, [25–27](#)
  - storage requirements, [16–17](#)
  - system requirements, [15–16](#)
  - user interfaces, [23](#)
  - vCenter HA requirements, [24–25](#)
  - vCenter Server file-based backup and restore, [23](#)
- reservations
  - resource pools, [140–141](#)
  - VM resources, monitoring/managing, [389–390](#)
- resource allocation policies, NIOC, [105–106](#)
- resource pools, [106–108](#), [139](#), [168](#)
  - creating, [368–369](#)
  - enhanced reservations, [142](#)
  - expandable reservations, [141](#)
  - limits, [141](#)
  - monitoring/managing resource usage, [385–386](#)
  - network resource pools, [341–342](#)
  - reservations, [140–141](#)

- shares, [140](#), [141](#), [142–143](#)
- use cases, [139](#)
- VM, moving into resource pools, [250](#)
- resource usage, monitoring/managing
  - alarms
    - actions, [401](#)
    - advanced use cases, [401](#)
    - creating definitions, [400](#)
    - definition elements, [399](#)
    - viewing/acknowledging triggered alarms, [399–400](#)
  - events, [396](#)
    - alerts, [397](#)
    - audit events, [397](#)
    - information events, [397](#)
    - streaming to remote syslog servers, [398–399](#)
    - types of, [397](#)
    - viewing in vSphere Client, [397](#)
    - viewing System Event Log, [397](#)
    - warning events, [397](#)
- VM resources
  - admission control, [390–391](#)
  - DRS, [384–385](#)
  - ESXTOP, [393–395](#)
  - host hardware, [386–387](#)
  - latency sensitivity, [392](#)
  - limits, [389–390](#)
  - reservations, [389–390](#)
  - resource pools, [385–386](#)
  - shares, [389–390](#)
  - VAMI, [396](#)
  - VIMTOP, [396](#)
  - VMware tools, [391–392](#)
  - Windows Perfmon, [391–392](#)

- restarting vSAN, [421–422](#)
- restores
  - vCenter Appliance File-Based Backup and Restore, [7](#)
  - vCenter Server, [23](#), [538–539](#)
- restricting vCenter Server
  - access, [261](#)
  - administrative privileges, [260–261](#)
- RFC (Request For Comments), [92](#)
- risk profiles, VM, [268–269](#)
- roles
  - security roles, [245–246](#), [248](#)
  - vCenter Server, [475](#)
- Route Based on IP Hash teaming policy, [99–100](#)
- Route Based on Originating Virtual Port teaming policy, [99](#)
- Route Based on Source MAC Hash teaming policy, [99](#)
- RVC, vSAN, [52](#)

## S

- SAN
  - SAN (Storage Area Networks)
    - management agents, RDM, [40](#)
    - vSAN. *See* individual entry
- SATA controllers, [182](#)
- SATP (Storage Array Type Plug-ins), [76](#)
- scalable shares, resource pools, [142–143](#)
- schedulers, DRS, [7](#)
- scoring VM, DRS, [136–137](#)
- scripted ESXi host installations, [288–292](#)
- SCSI controllers, [182](#)
- SCSI UNIMAP, [59](#)
- SDDC (Software-Defined Data Centers)
  - NSX, requirements, [26](#)

- requirements, [25–27](#)
- VMware SDCC, [27](#)
- vRealize Suite, requirements, [26–27](#)
- vSAN, requirements, [25–26](#)
- SDRS (Storage DRS), [81](#), [84](#)
  - anti-affinity rules, [448–449](#)
  - automation levels, [82](#)
  - behaviors, [82–83](#)
  - datastores
    - configuring, [447–449](#)
    - managing, [447–449](#)
    - overriding cluster automation, [448](#)
    - recommendations, [448](#)
  - I/O latency load balancing, [81–82](#)
  - load balancing, [81–82](#)
  - placement (initial), [81](#)
  - recommendations, [83](#)
  - space utilization load balancing, [81](#)
  - thresholds, [82–83](#)
- Secure Boot
  - ESXi, [258](#)
  - UEFI, [266](#), [499–500](#)
- security. *See also* [authentication](#); [authorization](#)
  - account lockouts, ESXi, [485–487](#)
  - add-ons, [275](#)
  - administrative privileges (vCenter Server), restricting, [260–261](#)
  - AppDefense, [227–228](#), [277–278](#)
  - Auto Deploy, [491](#)
  - certificates
    - CSR, [238–239](#)
    - custom certificates, VMCA, [237](#)
    - identity services, [236–237](#)
    - management, [477–482](#)

- overview, [236–238](#)
- requirements, [238–241](#)
- unsupported certificates, VMCA, [238](#)
- VECS, [236–237](#)
- VMAFD, [236](#)
- VMCA, [236–238](#)
- VMCA as intermediate CA, [237](#), [239](#)
- vmdir, [236](#)
- client certificates, [477–478](#)
- custom certificates, [241](#), [478–479](#)
- distributed file locking, RDM, [39](#)
- encryption
  - DEK, [270](#)
  - Encrypted vSphere vMotion, [272–273](#)
  - KEK, [270](#), [271](#)
  - VM, [270–272](#)
  - vSAN clusters, [432–435](#)
- ESXi, [241–242](#), [494–495](#)
  - account lockouts, [485–487](#)
  - AD and user management, [497–498](#)
  - assigning privileges, [496–497](#)
  - configuring with host profiles, [482–483](#)
  - controlling CIM access, [491–492](#)
  - controlling MOB access, [257–258](#)
  - customizing services, [493–494](#)
  - disabling MOB, [490](#)
  - firewall configuration, [492–493](#)
  - firewall ports, [255–256](#)
  - host access, [257](#)
  - hosts, acceptance levels, [496](#)
  - hosts, joining to directory services, [257](#)
  - hosts, TPM, [500–501](#)
  - hosts, UEFI Secure Boot, [499–500](#)

- hosts, VIB, [496](#)
- log files, [501](#)
- networking security recommendations, [490](#)
- password hardening, [256](#)
- passwords, [485–487](#)
- PCI, [489](#)
- PCIe devices, [489](#)
- recommendations, [481–482](#)
- scripts and host configuration management, [483–485](#)
- Secure Boot, [258](#)
- security profiles, [254–255](#)
- Shell security and SSH, [487–489](#)
- smart card authentication, [499](#)
- TPM chips, [258](#)
- vCenter Server security, [253–258](#)
- vSphere Authentication Proxy, [257](#)
- web proxy settings, [490–491](#)

ESXi certificates

- changing certificate mode, [479–480](#)
- custom certificates, [480](#)
- expiration, [481](#)
- host certificate modes, [241](#), [242](#)
- management, [479–481](#)
- switching to VMCA Mode, [480–481](#)

firewalls

- ESXi firewall configuration, [492–493](#)
- ESXi firewall ports, [255–256](#)
- networks security, [262](#)

identity services, [236](#)

- VECS, [236–237](#), [240–241](#)
- VMAFD, [236](#)
- VMCA, [236–238](#)
- vmdir, [236](#)



- IPsec, [262–263](#)
- Key Management Servers, [502](#)
- machine SSL certificates, [240](#), [241](#)
- networks, [262](#)
  - firewalls, [262](#)
  - IPsec, [262–263](#)
  - isolation, [262](#)
  - policies, [264–265](#)
  - recommendations, [263–264](#)
  - security policies, [100–101](#)
  - segmentation, [262](#)
- NSX Data Center, [228–229](#)
- passwords
  - ESXi, [256](#), [485–487](#)
  - SSO password policies, [260](#)
- permissions
  - applying with vCenter Server, [251–253](#)
  - authentication, [242](#)
  - authorization, [242](#)
  - best practices, [248](#)
  - cold migration (relocation) of VM, [250](#)
  - creating VM, [249](#)
  - deploying from templates, [249](#)
  - diagrams, [246–247](#)
  - global permissions, [247](#)
  - guest OS installations on VM, [250](#)
  - inventory hierarchies, [243–244](#)
  - migrating VM with Storage vMotion, [251](#)
  - migrating VM with vMotion, [250](#)
  - moving hosts into clusters, [251](#)
  - moving VM into resource pools, [250](#)
  - privileges, [244–245](#), [246](#), [248–251](#)
  - roles, [245–246](#), [248](#)

- taking VM snapshots, [249](#)
- privileges, [244–245](#), [248–251](#)
- roles, [245–246](#), [248](#)
- smart cards, [499](#)
- solution user certificates, [240–241](#)
- storage providers, viewing, [436](#)
- STS, [11](#), [473](#)
- VBS, [590](#)
- vCenter Server
  - client certificates, [261](#)
  - controlling datastore browser access, [261](#)
  - ESXi and vCenter Server security, [253–258](#)
  - restricting access, [261](#)
  - restricting administrative privileges, [260–261](#)
  - SSO password policies, [260](#)
  - time synchronization, [261](#)
  - user access, [259–260](#)
- vCenter single sign-on SSL signing certificates, [240](#)
- VECS, [240–241](#)
- VM, [265](#)
  - add-ons, [275](#)
  - AppDefense, [277–278](#)
  - compliance with vROps, [275](#)
  - copying/pasting, [267](#)
  - denial-of-service attacks, [269](#)
  - device connections, [267](#), [269–270](#)
  - disabling unexposed features, [266–267](#)
  - disk shrinking, [267](#)
  - Encrypted vSphere vMotion, [272–273](#)
  - encryption, [270–272](#), [505–506](#)
  - hardening, best practices, [265–266](#)
  - logs, [267–268](#)
  - performance counters, [268](#)

- risk profiles, [268–269](#)
- SGX, [505](#)
- UEFI Secure Boot, [266](#)
- VMware NSX, [276–277](#)
- VMX file size, [268](#)
- vSGX, [274–275](#)
- vTPM, [273–274](#)
- VMCA, [239](#)
- vmdir SSL certificates, [240](#)
- vSGX, [274–275](#)
- vSphere built-in features, ESXi and vCenter Server security, [254](#)
- vSphere Virtual Machine Encryption certificates, [240](#)
- vTA, [258–259](#)
  - configuring, [502–504](#)
  - management, [502–504](#)
- segmenting
  - networks security, [262](#)
  - vSphere networks, [18](#)
- serial ports, [182](#)
- servers
  - Administration server, vCenter Server, [11](#)
  - ESXi Server, [6](#)
    - installing, [15–16, 17](#)
    - network requirements, [20–21](#)
    - required ports, [20–21](#)
    - storage requirements, [17](#)
    - system requirements, [15–16](#)
  - Key Management Servers, security, [502](#)
  - KMS, vSAN encryption, [61–62](#)
  - remote syslog servers, streaming events to, [398–399](#)
  - tcServer, [11](#)
  - vCenter Server, [6](#)
    - Administration server, [11](#)

- appliance deployments, [298–303](#)
- applying permissions, [251–253](#)
- backups, [514–517](#), [538–539](#)
- compute requirements, [14–15](#)
- controlling MOB access, [257–258](#)
- data transfers, [519](#)
- database, [11](#)
- database deployments, [297](#)
- domain discovery, [21](#)
- Enhanced Linked Mode, [12–13](#), [474](#)
- ephemeral binding, [113](#)
- ESXi and vCenter Server security, [253–258](#)
- file-based backup and restore, [23](#)
- HA, [145](#), [157](#)
- Host Agent, [11](#)
- importing/exporting cluster images, [538](#)
- inventory configuration, [315–316](#)
- License Service, [11](#)
- managing, [542–543](#)
- managing, VAMI, [543–547](#)
- managing, vSphere Client, [547–554](#)
- migrating, [522–524](#)
- monitoring, [542–543](#)
- monitoring, VAMI, [543–547](#)
- monitoring, vSphere Client, [547–554](#)
- network requirements, [18–20](#)
- NTP, [23](#)
- patching with Appliance Shell, [554–556](#)
- patching with VAMI, [554–556](#)
- plug-ins, [11](#)
- post-installation, [302–303](#)
- privileges, [475](#)
- PSC, [10](#), [11](#)

- PSC deployments, [297–298](#)
- repointing to other domains, [558–560](#)
- required ports, [18–20](#)
- restores, [538–539](#)
- security, [259–261](#)
- storage requirements, [16](#)
- STS, [11](#)
- tcServer, [11](#)
- topology, [10–14](#)
- Update Planner, [524–526](#)
- updating, [554–557](#)
- upgrading appliances, [519–522](#)
- upgrading ESXi, [524](#)
- upgrading VM, [524](#)
- vCenter Lookup Service, [11](#)
- vCenter Server Agent, [11](#)
- VCSA deployments, [299–302](#)
- VMware Directory Service, [11](#)
- vSAN deployments, [422](#)
- vTA, [258–259](#)
- VCSA, [113](#)
- VMware servers, virtualization, [27](#)
- SFTT (Secondary Level of Failures to Tolerate), vSAN, [57](#)
- SGX, VM security, [505](#)
- shared disks, EZT for, [53](#)
- shares
  - resource pools, [140](#), [141](#), [142–143](#)
  - scalable shares, [142–143](#)
- SIOC
  - monitoring, [450](#)
  - setting shares/limits, [450](#)
- virtual disks, [582](#)
- VM resources, monitoring/managing, [389–390](#)

- shrinking disks, [267](#)
- Shutdown Guest, VM, [572](#)
- shutting down vSAN, [421–422](#)
- SIO controllers, [182](#)
- SIOC (Storage I/O Control), [84](#)
  - configuring, [449–451](#)
  - management, [449–451](#)
  - shares
    - monitoring, [450](#)
    - setting shares/limits, [450](#)
  - thresholds, [450–451](#)
- sizing
  - virtual disks, [582](#)
  - vSAN, [63–64](#)
- Skyline, [206](#)
  - integration, [206](#)
  - use cases, [206](#)
  - vSphere Health and vSAN Health, [53](#)
- smart card authentication, ESXi, [499](#)
- snapshot delta VMDK, [52](#)
- snapshots, [175](#)
  - RDM, [39](#)
  - taking, [249](#)
  - virtual disks, [582](#)
  - VM snapshots, [175, 586–587](#)
    - behaviors, [178–179](#)
    - branches, [177](#)
    - database files, [179](#)
    - delta disk files, [179](#)
    - flat files, [178](#)
    - limitations, [179–180](#)
    - linear snapshots, [176](#)
    - memory files, [179](#)

- parent snapshots, [178](#)
- post-revert snapshot trees, [176](#)
- preserving information, [177–178](#)
- use cases, [177](#)
- virtual disks, [582](#)

snooping, multicast, [117](#)

software

- SGX, VM security, [505](#)
- vSAN requirements, [66](#)
- vSGX, [274–275](#)

Software-Defined Data Centers. *See* [SDDC](#)

software-defined storage models, [40](#)

solution user certificates, [240–241](#)

space efficiency, vSAN, [59–61](#), [430–432](#)

space utilization load balancing, SDRS, [81–82](#)

spanning tree attacks, [97](#)

SPBM (Storage Policy Based Management), [40–41](#), [79](#)

- managing, [459–461](#)
  - applying storage policies to VM, [462–463](#)
  - VASA, managing storage providers, [462](#)
  - VASA, registering storage providers, [461](#)
- vSAN, [52](#)

SR-IOV (Single Root-I/O Virtualization), [119–121](#), [343–345](#)

SRM (Site Recovery Manager), [221–222](#)

- integration, [222](#)
- use cases, [222](#)

SSH, ESXi Shell security, [487–489](#)

SSL (Secure Sockets Layer)

- certificates, verifying legacy ESXi hosts, [554](#)
- machine SSL certificates, [240](#), [241](#)
- vCenter single sign-on SSL signing certificates, [240](#)
- vmdir SSL certificates, [240](#)

SSO (Single Sign-On), [242](#), [472](#)

- configuring, [305](#)
- enabling with Windows Session Authentication, [472–473](#)
- Enhanced Linked Mode, [474](#)
- group authentication, [474–475](#)
- identity sources, [305–307](#)
- password policies, vCenter Server security, [260](#)
- policy configuration, [311–312](#)
- STS management, [473](#)
- users
  - authentication, [474–475](#)
  - enabling/disabling, [310–311](#)
- vCenter Single Sign-On, [10](#), [11–12](#)
- standard port groups
  - configuring, [332–334](#)
  - creating, [332–334](#)
- standard vSAN cluster deployments, [54–55](#)
- stateless caching, [292](#)
- static binding, [112](#)
- storage
  - anti-affinity rules, [83](#)
  - controllers, VM, [583](#)
  - storage
    - datastores, [41](#)
      - NFS datastores, [43–45](#)
      - VMFS datastores, [41–43](#)
      - vSAN datastores, [45](#)
      - vVols datastores, [45](#)
    - devices (LUN), [37](#)
    - ESXi Server requirements, [17](#)
    - FC, [37](#)
    - FCoE, [38](#)
    - iSCSI, [37](#)
    - Kubernetes, [45–46](#)



- local storage, [37](#)
- multipathing/failover, [74](#)
  - array-based failovers with iSCSI, [74](#)
  - FC failovers, [74](#)
  - host-based failovers with iSCSI, [74](#)
  - path failovers and VM, [74](#)
- NAS/NFS, [38](#)
- NFS datastores, [43–45](#), [444–446](#)
- NVMe, [46](#)
  - over FC requirements, [47](#)
  - over PCIe requirements, [46](#)
  - over RDMA (RoCE Version 2) requirements, [46](#)
  - VMware HPP, [47–48](#)
- object-based storage, vSAN, [51](#)
- PSA, [74–75](#)
  - I/O requests, [78](#)
  - PSP, [76–77](#)
  - SATP, [76](#)
  - tasks, [77–78](#)
  - VMware native modules and third-party MPP, [78](#)
  - VMware NMP, [75](#), [78](#)
- RDM, [38–40](#), [439–446](#)
- SDRS, [81](#), [84](#)
  - automation levels, [82](#)
  - behaviors, [82–83](#)
  - configuring, [447–449](#)
  - datastores, [447–449](#)
  - I/O latency load balancing, [81–82](#)
  - load balancing, [81–82](#)
  - managing, [447–449](#)
  - placement (initial), [81](#)
  - recommendations, [83](#)
  - space utilization load balancing, [81](#)

- thresholds, [82–83](#)
- SIOC, [84](#)
  - configuring, [449–451](#)
  - management, [449–451](#)
  - monitoring shares, [450](#)
  - setting shares/limits, [450](#)
  - thresholds, [450–451](#)
- SPBM, [40–41](#), [79](#)
  - applying storage policies to VM, [462–463](#)
  - VASA, managing storage providers, [462](#)
  - VASA, registering storage providers, [461](#)
  - vSAN, [52](#)
- Storage vMotion, [7](#), [251](#)
- VAAI
  - block primitives, [70–71](#)
  - NAI primitives, [71](#)
  - thin provisioning primitives, [71](#)
  - vSphere storage integration, [70–71](#)
- VASA
  - managing storage providers, SPBM, [462](#)
  - registering storage providers, SPBM, [461](#)
  - vSphere storage integration, [69–70](#)
- vCenter Server requirements, [16](#)
- virtual disks, [37](#)
  - eager zeroed thick virtual disks, [79](#)
  - lazy zeroed thick virtual disks, [79](#)
  - thin provisioned virtual disks, [79](#)
  - zeroing out files, [79](#)
- virtualization, [36](#)
  - FC, [37](#)
  - FCoE, [38](#)
  - I/O filters, [41](#)
  - iSCSI, [37](#)

- local storage, [37](#)
- NAS/NFS, [38](#)
- RDM, [38–40](#)
- software-defined storage models, [40–41](#)
- SPBM, [40–41](#)
- storage devices (LUN), [37](#)
- virtual disks, [37](#)
- VMFS, [38](#)
- vSAN, [40](#)
- vVols, [40–41](#), [72–73](#)

VM

- access, [36](#)
- storage policies, [78–79](#)

VMFS, [38](#), [41–43](#), [438–443](#)

vMotion, [192](#)

- data flow, [193](#)
- limitations, [193](#)
- requirements, [193](#)
- Storage vMotion, [7](#), [251](#)

vSAN

- 7.0 features, [53–54](#)
- absent component state, [52](#)
- boot devices, [68](#)
- characteristics, [50–51](#), [414](#)
- cluster requirements, [66](#)
- compression, [59](#), [60](#)
- concepts, [49](#)
- configuring, [419–420](#)
- consumed capacity, [51](#)
- creating clusters with Quickstart, [415](#)
- Data Locality, [57](#)
- datastores, [45](#), [51](#)
- deduplication, [59](#), [60](#)

- degraded component state, [52](#)
- deploying with vCenter Server, [422](#)
- disabling, [421](#)
- disk groups, [51](#)
- editing settings, [417–418](#)
- encryption, [61–62](#)
- encryption in clusters, [432–435](#)
- erasure coding, [59](#), [60–61](#)
- expanding clusters, [422–424](#)
- extending datastores across two sites, [427–428](#)
- EZT for shared disks, [53](#)
- fault domain planning, [64–65](#)
- fault domains, [426–427](#)
- File Services, [54](#), [62–63](#), [436–438](#)
- file-based persistent volumes, [54](#)
- hardware requirements, [65–66](#)
- healthy object state, [52](#)
- increasing space efficiency in clusters, [430–432](#)
- integrated file services, [53](#)
- I/O redirects, [53](#)
- large-capacity drive support, [54](#)
- license requirements, [67](#)
- licensing, [418](#)
- limitations, [59](#)
- Maintenance Mode, [424–426](#)
- managing devices in clusters, [429–430](#)
- manually enabling, [416–417](#)
- memory as performance service metric, [53](#)
- memory objects, [52](#)
- network best practices, [67–68](#)
- network requirements, [67](#)
- NVMe Hot-Plug plug-in, [53](#)
- object-based storage, [51](#)

- Observer, [53](#)
- persistent logging in vSAN clusters, [68](#)
- PFTT, [57](#)
- planning, [63–64](#)
- preparing, [414](#)
- Ready Node, [53](#)
- repair objects after witness deployments, [54](#)
- requirements, [63–67](#)
- restarting, [421–422](#)
- RVC, [52](#)
- SCSI UNIMAP, [59](#)
- SFTT, [57](#)
- shutting down, [421–422](#)
- sizing, [63–64](#)
- Skyline and vSphere Health integration, [53](#)
- snapshot delta VMDK, [52](#)
- software requirements, [66](#)
- space efficiency, [59–61](#)
- SPBM, [52](#)
- standard cluster deployments, [54–55](#)
- storage policies, [79–81](#), [435–436](#)
- stretched cluster deployments, [56–59](#)
- terminology, [51–53](#)
- thin provisioning, [59](#)
- two-host cluster deployments, [55](#)
- unhealthy object state, [52](#)
- user-defined vSAN clusters, [53](#)
- VCG notification service, [54](#)
- viewing datastores, [418–419](#)
- viewing storage providers, [436](#)
- VM compliance status, [52](#)
- VM home namespace, [52](#)
- VM swap objects, [52](#)

- VMDK, [52](#)
- VMware PowerCLI, [53](#)
- vSphere HA, [419–420](#)
- vSphere Health, [53](#)
- vSphere Lifecycle Manager, [53](#), [54](#)
- vSphere Replication objects, [54](#)
- vSphere with Kubernetes integration, [54](#)
- vVols support, [54](#)
- witnesses, [52](#)
- vSAN datastores, [45](#)
- vSphere storage, [16–17](#), [68](#)
  - VAAI, [70–71](#)
  - VASA, [69–70](#)
- vVols
  - configuring, [463–464](#)
  - management, [463–464](#)
- vVols datastores, [45](#)
- streaming events to remote syslog servers, [398–399](#)
- stretched vSAN cluster deployments, [56–59](#)
- STS (Security Token Service), [11](#), [473](#)
- subscribing to content libraries, [596](#)
- swapping VM, [380–381](#)
- Swap-vVol, [73](#)
- switches
  - CDP, [117–118](#)
  - discovery protocols, [117–118](#)
  - distributed port groups, [337–338](#), [353–354](#)
  - physical Ethernet switches, [92–93](#)
  - standard port groups
    - configuring, [332–334](#)
    - creating, [332–334](#)
- vDS, [94](#)
  - adding hosts, [350–351](#)

- advanced settings, [337](#)
- architecture, [102](#)
- configuring, [334–335](#)
- creating, [334–335](#)
- distributed port groups, [103](#)
- Health Check, [115–116](#), [354](#)
- inbound traffic shaping, [105](#)
- LACP, [113–115](#)
- managing host physical network adapters with vDS, [351](#)
- marking policies, [109–110](#)
- modifying, [336](#)
- multicast filtering, [116–117](#)
- multicast snooping, [117](#)
- NetFlow, [108](#), [336–337](#)
- network policies, [104–105](#)
- port blocking policies, [105](#)
- port state monitoring, [111](#)
- removing hosts, [352](#)
- teaming policies, [99](#)
- traffic filtering, [109–110](#)
- upgrading, [335–336](#)
- uplink port groups, [103](#)
- VM, migrating to vDS, [353](#)
- VMkernel network adapters, migrating to vDS, [352](#)
- vSphere Client data center-level management, [111](#)
- vSS comparison, [103–104](#)
- virtual switches, [94](#)
- vSphere standard switches, [18](#)
- vSS, [94](#), [95–97](#)
  - configuring, [330–332](#)
  - creating, [330–332](#)
  - network policies, [98](#)
  - vDS comparison, [103–104](#)

- synchronizing
  - content libraries, [598](#)
  - time, vCenter Server, [261](#)
- syslog
  - ESXi host configurations, [405–407](#)
  - remote syslog servers, streaming events to, [398–399](#)
- System Event Log, [397](#)
- system logs, uploading to VMware, [404](#)
- system requirements, ESXi Server, [15–16](#)

## T

- taking exams, [604–606](#)
- Tanzu, vSphere with, [173](#), [204](#)
  - integration, [205](#)
  - use cases, [204](#)
- TCP/IP (Transmission Control Protocol/Internet Protocol), [92](#)
  - stacks, [121–122](#), [188](#)
  - VMkernel
    - TCP/IP networking layer, [18](#)
    - TCP/IP stacks, [121–122](#), [339–340](#)
- tcServer, [11](#)
- teaming policies, NIC, [98–100](#), [105](#)
- templates, [170](#)
  - JSON templates, VCSA deployments with CLI installers, [302](#)
  - OVA templates, deploying VM, [577](#)
  - OVF templates
    - deploying VM, [577](#)
    - editing details, [585–586](#)
    - managing, [589](#)
  - rapid provisioning VM with templates, [195](#)
  - VM
    - converting to templates, [573](#)



- deploying from templates, [574](#)
- deployments, [249](#)
- tests. See [exam preparation](#)
- thin provisioning
  - VAAI, [71](#)
  - virtual disks, [79](#)
  - vSAN, [59](#)
- thresholds
  - SDRS, [82–83](#)
  - SIOC, [450–451](#)
- time
  - NTP, [22–23](#)
- synchronization, vCenter Server, [261](#)
- tokens, STS, [11](#)
- topologies, vCenter Server, [10–14](#)
- TPM (Trusted Platform Modules)
  - devices, [182](#)
  - ESXi, [258](#)
  - ESXi hosts, [500–501](#)
  - vTPM, [273–274](#)
- traffic filtering (network), [109–110](#)
- traffic shaping policies, [101](#), [105](#)
- transferring data, vCenter Server, [519](#)
- Transmission Control Protocol/Internet Protocol. See [TCP/IP](#)
- triggered alarms, viewing/acknowledging, [399–400](#)
- troubleshooting
  - CPU
    - usage, [380](#)
    - utilization, [381](#)
  - datastores, utilization, [381](#)
  - device latency, [382](#)
  - dropped packets, [383](#)
  - latency

- device latency, [382](#)
- dropped packets, [383](#)
- VMkernel latency, [382](#)
- performance, [379–383](#)
- VM
  - memory usage, [380–381](#)
  - swapping, [380–381](#)
  - VMkernel latency, [382](#)
- TSO (TCP Segmentation Offload), [118](#)
- two-host vSAN cluster deployments, [55](#)

## U

- UEFI Secure Boot, [266](#), [499–500](#)
- UMDS (Update Manager Download Service), [529–530](#)
- unexposed features (network security), disabling, [266–267](#)
- unhealthy object state, vSAN, [52](#)
- updating
  - ESXi firmware updates, [536–537](#)
  - UMDS, [529–530](#)
  - Update Planner, [524–526](#)
  - vCenter Server, [554–557](#)
  - vSphere Lifecycle Manager, [157](#)
    - baselines, [530–535](#)
    - definitions, [532–533](#)
    - ESXi firmware updates, [536–537](#)
    - ESXi hosts, [526–529](#)
    - ESXi Quick Boot, [535–536](#)
    - hardware compatibility checks, [537](#)
    - remediation settings, [528](#)
    - UMDS, [529–530](#)
  - VUM. *See* [vSphere Lifecycle Manager](#)
- upgrading

- ESXi, [524](#)
- vCenter Server
  - appliances, [519–522](#)
  - ESXi, [524](#)
  - Update Planner, [524–526](#)
  - VM, [524](#)
- vDS, [335–336](#)
- VM, [524](#), [539–540](#)
- VMware Tools, [570–571](#)
- to vSphere 7.0, [517–518](#)
- uplink port groups, [103](#)
- uploading system logs to VMware, [404](#)
- USB (Universal Serial Bus)
  - controllers, [182](#)
  - devices, [182](#)
- Use Explicit Failover Order teaming policy, [99](#)
- user interfaces, vSphere requirements, [23](#)
- user-defined vSAN clusters, [53](#)
- user-friendly persistent names, RDM, [39](#)
- users
  - authentication, [474–475](#)
  - ESXi, AD and user management, [497–498](#)
  - guest user mappings, VM, [585](#)
  - SSO users, enabling/disabling, [310–311](#)

## V

- VAAI (vStorage API for Array Integration)
  - block primitives, [70–71](#)
  - NAI primitives, [71](#)
  - thin provisioning primitives, [71](#)
  - vSphere storage integration, [70–71](#)
- VAIO (vSphere API for I/O Filtering), [271](#)

- validation settings (permissions), changing, [502](#)
- VAMI (vCenter Server Application Management Interface)
  - monitoring/managing resources, [396](#)
  - vCenter Server
    - monitoring/managing, [543–547](#)
    - patching, [554–556](#)
    - updating, [554–557](#)
- vApps, [170](#)
- VASA (vStorage API for Storage Awareness)
  - SPBM
    - managing storage providers, [462](#)
    - registering storage providers, [461](#)
  - vSphere storage integration, [69–70](#)
- VBS (Virtualization-Based Security), [590](#)
- vCenter Appliance File-Based Backup and Restore, [7](#)
- vCenter Cloud account permissions, [210–213](#)
- vCenter Converter, [205](#)
  - integration, [205–206](#)
  - use cases, [205](#)
- vCenter HA, [6](#), [14](#)
  - active nodes, [14](#)
  - clusters, managing, [557–558](#)
  - implementing, [316–317](#)
  - passive nodes, [14](#)
  - requirements, [24–25](#)
  - witness nodes, [14](#)
- vCenter Lookup Service, [11](#)
- vCenter Server, [6](#)
  - Administration server, [11](#)
  - Appliance Shell, patching, [554–556](#)
  - backups, [514–517](#), [538–539](#)
  - client certificates, [261](#)
  - cluster images, importing/exporting, [538](#)

- compute requirements, [14–15](#)
- cross-vCenter Server migrations, [186–187](#)
- data transfers, [519](#)
- database, [11](#)
- deployments
  - appliances, [298–303](#)
  - databases, [297](#)
  - post-installation, [302–303](#)
  - PSC, [297–298](#)
- domain discovery, [21](#)
- Enhanced Linked Mode, [12](#), [13](#), [474](#)
- ephemeral binding, [113](#)
- ESXi and vCenter Server security, [253](#)
  - controlling MOB access, [257–258](#)
  - ESXi firewall ports, [255–256](#)
  - ESXi host access, [257](#)
  - ESXi password hardening, [256](#)
  - ESXi Secure Boot, [258](#)
  - ESXi security profiles, [254–255](#)
  - TPM chips, [258](#)
  - vSphere Authentication Proxy, [257](#)
  - vSphere built-in features, [254](#)
- file-based backup and restore, [23](#)
- HA, [145](#), [157](#)
- Host Agent, [11](#)
- inventory configuration, [315–316](#)
- License Service, [11](#)
- logs, [404](#)
- managing, [542–543](#)
  - VAMI, [543–547](#)
  - vSphere Client, [547–554](#)
- migrating, [522–524](#)
- MOB, controlling access, [257–258](#)

- monitoring, [542–543](#)
  - VAMI, [543–547](#)
  - vSphere Client, [547–554](#)
- network requirements, [18–20](#)
- NTP, [23](#)
- patching
  - with Appliance Shell, [554–556](#)
  - with VAMI, [554–556](#)
- permissions, applying, [251–253](#)
- plug-ins, [11](#)
- ports, required ports, [18–20](#)
- privileges, [475](#)
- PSC, [10](#), [11](#)
- repointing to other domains, [558–560](#)
- restores, [538–539](#)
- security
  - client certificates, [261](#)
  - controlling datastore browser access, [261](#)
  - restricting access, [261](#)
  - restricting administrative privileges, [260–261](#)
  - SSO password policies, [260](#)
  - time synchronization, [261](#)
  - user access, [259–260](#)
- storage requirements, [16](#)
- STS, [11](#)
- tcServer, [11](#)
- time synchronization, [261](#)
- topology, [10–14](#)
- updating, [554–557](#)
- upgrading
  - appliances, [519–522](#)
  - ESXi, [524](#)
  - Update Planner, [524–526](#)

- VM, [524](#)
- VAMI, monitoring/managing resources, [396](#)
- vCenter Lookup Service, [11](#)
- vCenter Server Agent, [11](#)
- VCSA deployments, [113](#)
  - with CLI installers, [301–302](#)
  - with GUI installers, [299–301](#)
- VMware Directory Service, [11](#)
- vSAN deployments, [422](#)
- vTA, [258–259](#)
- vCenter Single Sign-On, [10](#), [11–12](#), [240](#), [242](#)
- VCF (VMware Cloud Foundation), [28](#), [223](#)
  - integration, [224](#)
  - use cases, [223](#)
- VCG notification service, vSAN, [54](#)
- vCloud Suite (VMware), [28](#)
- VCSA (vCenter Server Appliance)
  - deployments
    - with CLI installers, [301–302](#)
    - with GUI installers, [299–301](#)
  - ephemeral binding, [113](#)
- vDS (vSphere Distributed Switches), [94](#)
  - advanced settings, [337](#)
  - architecture, [102](#)
  - configuring, [334–335](#)
  - creating, [334–335](#)
  - distributed port groups, [103](#)
  - Health Check, [115–116](#), [354](#)
  - hosts
    - adding, [350–351](#)
    - managing host physical network adapters with vDS, [351](#)
    - removing, [352](#)
  - inbound traffic shaping, [105](#)

- LACP, [113–115](#)
- marking policies, [109–110](#)
- modifying, [336](#)
- multicast filtering, [116–117](#)
- multicast snooping, [117](#)
- NetFlow, [108](#), [336–337](#)
- network policies, [104–105](#)
- ports
  - blocking policies, [105](#)
  - state monitoring, [111](#)
- teaming policies, [99](#)
- traffic filtering, [109–110](#)
- upgrading, [335–336](#)
- uplink port groups, [103](#)
- VM, migrating to vDS, [353](#)
- VMkernel network adapters, migrating to vDS, [352](#)
- vSphere Client data center-level management, [111](#)
- vSS comparison, [103–104](#)
- VECS (VMware Endpoint Certificate Store), [236–237](#), [304](#)
  - solution user certificate stores, [240–241](#)
  - stores, [303–304](#)
- vendor add-ons, [534](#)
- vGPU (Virtual Graphical Processing Units), VM support, [592–594](#)
- VIB (vSphere Installation Bundles), [258](#), [496](#)
- viewing
  - events
    - System Event Log, [397](#)
    - in vSphere Client, [397](#)
  - System Event Log, [397](#)
  - triggered alarms, [399–400](#)
- vSAN
  - datastores, [418–419](#)
  - storage providers, [436](#)



VIMTOP, monitoring/managing resources, [396](#)

virtual compatibility mode, RDM, [38–39](#)

virtual disks, [37](#)

- anti-affinity rules, [83](#)

- configuring, [581–582](#)

- database files, [179](#)

- delta disk files, [179](#)

- eager zeroed thick virtual disks, [79](#)

- files, [175](#)

- flat files, [178](#)

- increasing size, [582](#)

- lazy zeroed thick virtual disks, [79](#)

- memory files, [179](#)

- provisioning, [183](#), [581](#)

- shares, [582](#)

- snapshots, [582](#)

- thin provisioned virtual disks, [79](#)

- zeroing out files, [79](#)

Virtual Machine File Systems. *See* [VMFS](#)

virtual machines. *See* [VM](#)

virtual networks, [17](#)

- advanced features, [355–356](#)

- DirectPath I/O, [343](#)

- distributed port groups, [337–338](#), [353–354](#)

- LAG, [346–349](#)

- network resource pools, [341–342](#)

- NIOC, [340–341](#)

- NSX Data Center, [228–229](#)

- policies, [355–356](#)

- port mirroring, [345–346](#)

- PVLAN, [342](#)

- SR-IOV, [343–345](#)

- standard port groups, [332–334](#)

TCP/IP stacks, [339–340](#)

vDS

adding hosts, [350–351](#)

advanced settings, [337](#)

configuring, [334–335](#)

creating, [334–335](#)

Health Check, [354](#)

managing host physical network adapters with vDS, [351](#)

modifying, [336](#)

NetFlow, [336–337](#)

removing hosts, [352](#)

upgrading, [335–336](#)

VM, migrating to vDS, [353](#)

VMkernel network adapters, migrating to vDS, [352](#)

VMkernel adapters, [338–339](#)

vSS, creating, [330–332](#)

virtual ports, [94](#)

virtual switches, [94](#)

Virtual Volumes. *See* [vVols](#)

virtualization

App Volumes, [217–219](#)

desktops

App Volumes, [217–219](#)

VMware Horizon, [215–217](#)

NPIV, [40](#)

NSX Data Center, [228–229](#)

SR-IOV, [119–121](#), [343–345](#)

storage virtualization, [36](#)

FC, [37](#)

FCoE, [38](#)

I/O filters, [41](#)

iSCSI, [37](#)

local storage, [37](#)

- NAS/NFS, [38](#)
- RDM, [38–40](#)
- software-defined storage models, [40–41](#)
- SPBM, [40–41](#)
- storage devices (LUN), [37](#)
- virtual disks, [37](#)
- VMFS, [38](#)
- vSAN, [40](#)
- vVols, [40–41](#), [72–73](#)
- VBS, [590](#)
- VMware Horizon, [215–217](#)
- VMware servers, [27](#)
- VLAN (Virtual Local Area Networks), [94–95](#), [101–102](#)
  - PVLAN, [110](#), [342](#)
  - VLAN ID, standard port groups, [333](#)
- VM (Virtual Machines), [169](#)
  - adding CPU resources, [580–581](#)
  - advanced settings, [184–185](#)
  - anti-affinity rules, [83](#), [448–449](#)
  - application monitoring, [150](#), [372](#)
  - chipsets, [181](#)
  - cloning, [194](#), [572–573](#)
    - cold clones, [194](#)
    - hot clones, [194](#)
    - instant clones, [195–196](#)
    - linked clones, [194](#)
    - rapid provisioning VM with templates, [195](#)
  - compatibility options, [578](#), [579–580](#)
  - compliance status, vSAN, [52](#)
  - configuring, [372](#)
    - files, [174–175](#)
    - impact of, [392–393](#)
  - content libraries, [594–595](#)

- adding items, [598–599](#)
- creating, [595](#)
- permissions, [597](#)
- publishing, [596](#)
- subscriptions, [596](#)
- synchronization options, [598](#)
- VM deployments, [599](#)

CPU, [181](#), [380](#)

creating, [249](#), [568–569](#)

deploying

- from templates, [249](#), [574](#), [577](#), [585–586](#)
- using content libraries, [599](#)

DRS

- initial VM placements, [135–136](#)
- scoring VM, [136–137](#)
- VM distribution, [135](#)

DVD/CD-ROM drives, [181](#)

editing options, [583–585](#)

encryption, [270–272](#), [505–506](#)

file structures, [173–174](#)

FT, [153–157](#), [373](#)

GRID models, [593](#)

guest OS

- customizing, [574–576](#)
- installations, [250](#)

guest user mappings, [585](#)

hard disks, [181](#)

hardening, best practices, [265–266](#)

hardware

- compatibility, [180–182](#)
- configuring, [578–583](#)
- feature sets, [578–579](#)

IDE 0, [181](#)

- IDE 1, [181](#)
- inter-VM anti-affinity rules, [448–449](#)
- keyboards, [181](#)
- memory, [181](#)
- migrating, [185–186](#), [587–589](#)
  - cold migrations (relocation), [186](#), [250](#)
  - cross-datastore migrations, [186](#)
  - cross-host migrations, [186](#)
  - cross-vCenter Server migrations, [186–187](#)
  - hot migrations, [186](#)
  - limitations, [187–188](#)
  - with Storage vMotion, [251](#)
  - to vDS, [353](#)
  - with vMotion, [250](#)
  - vMotion, [189–193](#)
- Mode settings, [582](#)
- monitoring/managing resources, [372](#), [393](#)
  - admission control, [390–391](#)
  - ESXTOP, [393–395](#)
  - impact of VM configurations, [392–393](#)
  - latency sensitivity, [392](#)
  - limits, [389–390](#)
  - metrics, [388](#)
  - reservations, [389–390](#)
  - shares, [389–390](#)
  - VAMI, [396](#)
  - VIMTOP, [396](#)
  - VMware tools, [391–392](#)
  - Windows Perfmon, [391–392](#)
- moving into resource pools, [250](#)
- network adapters, [181](#)
- NPIV, [40](#)
- NVDIMM

- controllers, [182](#)
- devices, [182](#)
- NVMe controllers, [182](#)
- opening consoles to VM, [569–570](#)
- options, [183–184](#)
- OVF templates
  - editing details, [585–586](#)
  - managing, [589](#)
  - VM deployments, [577](#)
- parallel ports, [181](#)
- path failovers, [74](#)
- PCI
  - controllers, [181](#)
  - devices, [182](#)
- pointing devices, [182](#)
- PowerCLI, VM management, [590–592](#)
- powering on, [569](#)
- rapid provisioning with templates, [195](#)
- RDM, [439–446](#), [582–583](#)
- SATA controllers, [182](#)
- SCSI controllers, [182](#)
- SDRS, inter-VM anti-affinity rules, [448–449](#)
- security, [265](#)
  - add-ons, [275](#)
  - AppDefense, [277–278](#)
  - compliance with vROps, [275](#)
  - copying/pasting, [267](#)
  - denial-of-service attacks, [269](#)
  - device connections, [267](#), [269–270](#)
  - disabling unexposed features, [266–267](#)
  - disk shrinking, [267](#)
  - Encrypted vSphere vMotion, [272–273](#)
  - encryption, [270–272](#)

- hardening, best practices, [265–266](#)
- logs, [267–268](#)
- performance counters, [268](#)
- risk profiles, [268–269](#)
- SGX, [505](#)
- UEFI Secure Boot, [266](#)
- VMware NSX, [276–277](#)
- VMX file size, [268](#)
- vSGX, [274–275](#)
- vTPM, [273–274](#)
- serial ports, [182](#)
- settings, [149](#)
- Shutdown Guest, [572](#)
- SIO controllers, [182](#)
- snapshots, [175](#), [586–587](#)
  - behaviors, [178–179](#)
  - branches, [177](#)
  - database files, [179](#)
  - delta disk files, [179](#)
  - flat files, [178](#)
  - limitations, [179–180](#)
  - linear snapshots, [176](#)
  - memory files, [179](#)
  - parent snapshots, [178](#)
  - post-revert snapshot trees, [176](#)
  - preserving information, [177–178](#)
  - RDM, [39](#)
  - taking, [249](#)
  - use cases, [177](#)
  - virtual disks, [582](#)
- SPBM, applying storage policies to VM, [462–463](#)
- storage
  - accessing, [36](#)

- controllers, [583](#)
- policies, [78–79](#)
- TCP/IP stacks, [188](#)
- templates
  - converting VM to templates, [573](#)
  - deploying VM from templates, [574](#)
- TPM devices, [182](#)
- traffic shaping policies, [101](#)
- troubleshooting
  - memory usage, [380–381](#)
  - swapping, [380–381](#)
- upgrading, [524](#), [539–540](#)
- USB
  - controllers, [182](#)
  - devices, [182](#)
- VBS, [590](#)
- vGPU support, [592–594](#)
- virtual disks, [581](#)
  - configuring, [581–582](#)
  - files, [175](#)
  - increasing size, [582](#)
  - provisioning, [183](#)
  - shares, [582](#)
  - snapshots, [582](#)
- VMCI, [182](#)
- VMCP, [150](#), [371](#)
- VM-host affinity rule, [137](#)
- VM-VM affinity rule, [137–138](#)
- VMware Tools, [183](#), [570–571](#)
- vSGA models, [593](#)
- vSphere Virtual Machine Encryption certificates, [240](#)
- VM home namespace, vSAN, [52](#)
- VM swap objects, vSAN, [52](#)



- VMAFD (VMware Authentication Framework Daemon), [236](#)
- VMC on AWS, [28](#), [226](#)
- VMCA (VMware Certificate Authority), [236–237](#), [239](#)
  - configuring, [303–305](#)
  - custom certificates, [237](#)
  - ESXi certificates, VMCA Mode, [480–481](#)
  - as intermediate CA, [237](#), [239](#)
  - management, [303–305](#)
  - management modes (recommended), [237–238](#)
  - unsupported certificates, [238](#)
- VMCI (Virtual Machine Communication Interface), [182](#)
- VMCP (VM Component Protection), [150](#), [371](#)
- vmdir (VMware Directory Service), [236](#), [240](#)
- VMDK (Virtual Machine Disks)
  - snapshot delta VMDK, [52](#)
  - vSAN, [52](#)
- VMFS (Virtual Machine File Systems), [38](#), [41–43](#), [438–443](#)
- VMkernel
  - adapter settings, [122](#), [338–339](#)
  - latency, troubleshooting, [382](#)
  - network adapters, migrating to vDS, [352](#)
  - TCP/IP
    - networking layer, [18](#)
    - stacks, [121–122](#), [339–340](#)
- vMotion, [7](#), [189](#)
  - data flow, [191](#)
  - encrypted vMotion, [192](#), [272–273](#)
- EVC
  - AMD modes, [133](#)
  - Intel modes, [132–133](#)
  - vSphere clusters, [130](#), [131–133](#), [367–368](#)
- multi-NIC vMotion, [190](#)
- port states, [111](#)

- RDM, [40](#)
- requirements, [189–191](#)
- storage vMotion, [7](#), [192](#)
  - data flow, [193](#)
  - limitations, [193](#)
  - requirements, [193](#)
  - VM migration, [251](#)
- TCP/IP stacks, [122](#)
- VM migration, [250](#)

## VMware

- AppDefense, [227–228](#), [277–278](#)
- Azure VMware Solution, [226–227](#)
- Enhanced Authentication plug-ins, [303](#)
- HCX, [224–226](#)
- HPP, [47](#)
  - best practices, [48](#)
  - path selection schemes, [47–48](#)
  - vSphere support, [47](#)
- NMP, [75–76](#), [78](#)
- NSX, [276–277](#)
- NSX Data Center, [228–229](#)
- NVMe, [46](#)
  - over Fabric, [46](#), [452–453](#)
  - over FC requirements, [47](#)
  - over PCIe requirements, [46](#)
  - over RDMA (RoCE Version 2) requirements, [46](#)
  - VMware HPP, [47–48](#)
- PowerCLI
  - commands, [484–485](#)
  - vSAN, [53](#)
- private clouds, [28](#)
- PSA, VMware native modules and third-party MPP, [78](#)
- SDDC, [27](#)

- server virtualization, [27](#)
- Skyline, [206](#)
  - integration, [206](#)
  - use cases, [206](#)
  - vSphere Health and vSAN Health, [53](#)
- SRM, [221–222](#)
- system logs, uploading, [404](#)
- vCenter Converter, [205–206](#)
- VCF, [28](#), [223–224](#)
- vCloud Suite, [28](#)
- VECS, [236–237](#), [304](#)
  - solution user certificate stores, [240–241](#)
  - stores, [303–304](#)
- VM resources, monitoring/managing, [391–392](#)
- VMAFD, [236](#)
- VMC on AWS, [28](#), [226](#)
- VMCA, [236–237](#), [239](#)
  - custom certificates, [237](#)
  - as intermediate CA, [237](#), [239](#)
  - management modes (recommended), [237–238](#)
  - unsupported certificates, [238](#)
- vmdir, [236](#)
- VMware Tools, [320–321](#)
- vRA, [209](#)
  - integration, [210–213](#)
  - use cases, [210](#)
- vRealize Suite, [8](#), [207](#)
  - requirements, [26–27](#)
  - vRA, [209–213](#)
  - vRLI, [208–209](#)
  - vRNI, [214–215](#)
  - vRO, [213–214](#)
  - vROps, [207–208](#)

- vRLI, [208](#)
  - integration, [208–209](#)
  - use cases, [208](#)
- vRNI, [214](#)
  - integration, [215](#)
  - use cases, [214–215](#)
- vRO, [213–214](#)
  - integration, [214](#)
  - use cases, [214](#)
- vROps, [8](#), [26](#), [207](#)
  - compliance, [275](#)
  - integration, [208](#)
  - Predictive DRS, [152](#)
  - use cases, [207](#)
- vSphere Lifecycle Manager, [315](#)
- vSphere Replication, [206](#), [219–221](#)
- vSphere with Tanzu, [204–205](#)
- VMware Certification accounts, [604](#)
- VMware Cloud Foundation. *See* [VCF](#)
- VMware Directory Service, [11](#)
- VMware Horizon, [215–216](#)
  - integration, [216–217](#)
  - use cases, [216](#)
- VMware Service Lifecycle Manager, [157](#)
- VMware Tools, [183](#)
  - installing, [570–571](#)
  - upgrading, [570–571](#)
- VMware vCloud Director, [28](#)
- VMX file size, VM security, [268](#)
- vNIC (Virtual Network Interface Cards), [93–94](#)
- vPMem (Virtual PMem), [454](#)
- vPMemDisk (Virtual PMem Disks), [455](#)
- vRA (vRealize Automation), [26–27](#), [209](#)

- integration, [210–213](#)
  - use cases, [210](#)
- vRealize Log Insight. *See* [vRLI](#)
- vRealize Network Insight. *See* [vRNI](#)
- vRealize Operations. *See* [vROps](#)
- vRealize Suite, [8](#), [207](#)
  - requirements, [26–27](#)
  - vRA, [209](#)
    - integration, [210–213](#)
    - use cases, [210](#)
  - vRLI, [208](#)
    - integration, [208–209](#)
    - use cases, [208](#)
  - vRNI, [214](#)
    - integration, [215](#)
    - use cases, [214–215](#)
  - vRO, [213–214](#)
    - integration, [214](#)
    - use cases, [214](#)
  - vROps, [8](#), [26](#), [207](#)
    - compliance, [275](#)
    - integration, [208](#)
    - Predictive DRS, [152](#)
    - use cases, [207](#)
- vRLI (vRealize Log Insight), [27](#), [208](#), [407](#)
  - integration, [208–209](#)
  - use cases, [208](#)
- VRMS (vSphere Replication Management Service), [24](#)
- vRNI (vRealize Network Insight), [27](#), [214](#)
  - integration, [215](#)
  - use cases, [214–215](#)
- vRO (vRealize Orchestrator), [213–214](#)
  - integration, [214](#)

- use cases, [214](#)
- vROps (vRealize Operations), [26](#), [207](#)
  - compliance, [275](#)
  - integration, [208](#)
  - Predictive DRS, [152](#)
  - use cases, [207](#)
- VRS (vSphere Replication Service), [24](#)
- vSAN (virtual SAN), [8](#), [40](#)
  - 7.0 features, [53–54](#)
  - absent component state, [52](#)
  - boot devices, [68](#)
  - characteristics, [50–51](#), [414](#)
  - cluster requirements, [66](#)
  - clusters
    - creating with Quickstart, [415](#)
    - encryption, [432–435](#)
    - expanding, [422–424](#)
    - increasing space efficiency, [430–432](#)
    - managing devices in clusters, [429–430](#)
  - compression, [59](#), [60](#)
  - concepts, [49](#)
  - configuring, [419–420](#)
  - consumed capacity, [51](#)
  - Data Locality, [57](#)
  - datastores, [45](#), [51](#)
    - extending across two sites, [427–428](#)
    - viewing, [418–419](#)
  - deduplication, [59](#), [60](#)
  - degraded component state, [52](#)
  - deployments, [54–59](#), [422](#)
  - disabling, [421](#)
  - disk groups, [51](#)
  - editing settings, [417–418](#)

- encryption, [61–62](#)
- erasure coding, [59](#), [60–61](#)
- EZT for shared disks, [53](#)
- fault domains, [64–65](#), [426–427](#)
- File Services, [54](#), [62–63](#), [436–438](#)
- file-based persistent volumes, [54](#)
- hardware requirements, [65–66](#)
- healthy object state, [52](#)
- integrated file services, [53](#)
- I/O redirects, [53](#)
- large-capacity drive support, [54](#)
- licensing, [67](#), [418](#)
- limitations, [59](#)
- Maintenance Mode, [424–426](#)
- manually enabling, [416–417](#)
- memory as performance service metric, [53](#)
- memory objects, [52](#)
- networks
  - best practices, [67–68](#)
  - requirements, [67](#)
- NVMe Hot-Plug plug-in, [53](#)
- object-based storage, [51](#)
- Observer, [53](#)
- persistent logging in vSAN clusters, [68](#)
- PFTT, [57](#)
- planning, [63–64](#)
- preparing, [414](#)
- Ready Node, [53](#)
- repair objects after witness deployments, [54](#)
- requirements, [25–26](#), [63–67](#)
- restarting, [421–422](#)
- RVC, [52](#)
- SCSI UNIMAP, [59](#)

- SFTT, [57](#)
- shutting down, [421–422](#)
- sizing, [63–64](#)
- Skyline and vSphere Health integration, [53](#)
- snapshot delta VMDK, [52](#)
- software requirements, [66](#)
- space efficiency, [59–61](#)
- SPBM, [52](#)
- standard cluster deployments, [54–55](#)
- storage policies, [79–81](#), [435–436](#)
- stretched cluster deployments, [56–59](#)
- terminology, [51–53](#)
- thin provisioning, [59](#)
- two-host cluster deployments, [55](#)
- unhealthy object state, [52](#)
- user-defined vSAN clusters, [53](#)
- VCG notification service, [54](#)
- VM compliance status, [52](#)
- VM home namespace, [52](#)
- VM swap objects, [52](#)
- VMDK, [52](#)
- VMware PowerCLI, [53](#)
- vSphere HA, [419–420](#)
- vSphere Health, [53](#)
- vSphere Lifecycle Manager, [53](#), [54](#)
- vSphere Replication objects, [54](#)
- vSphere with Kubernetes integration, [54](#)
- vVols support, [54](#)
- witnesses, [52](#), [54](#)
- vSGA models, VM, [593](#)
- vSGX (Virtual Intel Software Guard Extension), [274–275](#)
- vSphere
  - add-on products, [7–8](#)



- Auto Deploy, security, [491](#)
- components
  - core components, [6](#)
  - optional components, [6](#)
- configuring, [315](#)
  - ESXi host profiles, [317–322](#)
  - vCenter HA implementation, [316–317](#)
  - vCenter Server inventory configuration, [315–316](#)
  - VMware Tools, [320–321](#)
  - VMware vSphere Lifecycle Manager, [315](#)
  - vSphere Client, [315](#)
- editions, [8–10](#)
- features, [7](#)
- infrastructure services, [21–23](#)
- installing
  - deploying vCenter Server components, [297–305](#)
  - ESXi hosts, [286–297](#)
  - initial vSphere configuration, [315–322](#)
  - SSO configurations, [305–314](#)
- inventory objects, [166–170](#)
- Kubernetes, [45–46](#), [54](#)
- licenses, [9](#)
- Lifecycle Manager, [157](#)
  - baselines, [530–535](#)
  - definitions, [532–533](#)
  - ESXi firmware updates, [536–537](#)
  - ESXi hosts, [526–529](#)
  - ESXi Quick Boot, [535–536](#)
  - hardware compatibility checks, [537](#)
  - remediation settings, [528](#)
  - UMDS, [529–530](#)
  - vSAN, [53](#), [54](#)
- managing resources, [373](#)

- monitoring resources, [373](#)
- networks
  - requirements, [17–21](#)
  - segmenting, [18](#)
- performance
  - charts, [375–379](#)
  - metrics, [374](#)
- Pods, vSphere with Kubernetes, [45–46](#)
- Replication objects, vSAN, [54](#)
- requirements
  - compute requirements, [14–15](#)
  - DPM, [23–24](#)
  - GUI installer, [23](#)
  - network requirements, [17–21](#)
  - NSX, [26](#)
  - replication requirements, [24](#)
  - SDDC, [25–27](#)
  - storage requirements, [16–17](#)
  - system requirements, [15–16](#)
  - user interfaces, [23](#)
  - vCenter HA requirements, [24–25](#)
  - vCenter Server file-based backup and restore, [23](#)
  - vRealize Suite, [26–27](#)
  - vSAN, [25–26](#)
- storage integration, [68](#)
  - VAAI, [70–71](#)
  - VASA, [69–70](#)
- upgrading to vSphere 7.0, [517–518](#)
- vSphere Host Client, [8](#)
- vSphere Authentication Proxy, [257](#)
- vSphere Client
  - data center-level management, [111](#)
  - events, viewing, [397](#)

- HTML5-based, [8](#)
- multipathing management, [457–458](#)
- port state monitoring, [111](#)
- vCenter Server, monitoring/managing, [547–554](#)
- vSphere configurations, [315](#)
- vSphere clusters
  - configuring, [130](#), [365–367](#)
  - creating, [364](#)
  - datastore clusters versus, [131](#)
  - DPM, [152–153](#)
  - DRS, [130–131](#), [134](#)
    - advanced options, [369–370](#)
    - affinity rules, [369–370](#)
    - anti-affinity rules, [369–370](#)
    - Automation Mode, [134](#)
    - creating DRS clusters, [368](#)
    - evacuation workflows, [136](#)
    - Memory metric for load balancing, [135](#)
    - migration sensitivity, [138–139](#)
    - monitoring/managing resource usage, [384–385](#)
    - network-aware DRS, [135](#)
    - NVM support, [136](#)
    - Predictive DRS, [152](#), [370](#)
    - recent enhancements, [134–137](#)
    - rules, [137–138](#)
    - scoring VM, [136–137](#)
    - VM distribution, [135](#)
    - VM initial placements, [135–136](#)
- EVC, [130](#), [131–132](#)
  - AMD modes, [133](#)
  - configuring, [367–368](#)
  - Intel modes, [132–133](#)
- HA, [143](#)

- Admission Control, [146–148](#)
- admission control, [371](#)
- advanced options, [148–149](#), [370](#)
- benefits of, [144](#)
- best practices, [151](#)
- configuring HA clusters, [370](#)
- detecting host issues, [144](#)
- failovers, [143](#), [144](#)
- heartbeats, [146](#)
- requirements, [145](#)
- response to failures, [145–146](#)
- vCenter Server, [145](#)
- overview, [130–131](#)
- resource pools, [139](#)
  - creating, [368–369](#)
  - expandable reservations, [141](#)
  - limits, [141](#)
  - monitoring/managing resource usage, [385–386](#)
  - reservations, [140–141](#)
  - shares, [140](#), [141](#), [142–143](#)
  - use cases, [139](#)
- vSphere HA, [7](#)
  - Admission Control, [146–148](#)
  - advanced options, [148–149](#)
  - benefits of, [144](#)
  - best practices, [151](#)
  - capacity reservation settings, [420](#)
  - configuring, [419–420](#)
  - detecting host issues, [144](#)
  - failovers, [143](#)
  - heartbeats, [146](#)
  - requirements, [145](#)
  - response to failures, [145–146](#)

- vCenter Server, [145](#)
- vSAN, [419–420](#)
- vSphere clusters, failovers, [144](#)
- vSphere Health
  - Skyline and vSAN Health integration, [53](#)
  - vSAN Health, [53](#)
- vSphere Host Client, [8](#)
- vSphere Lifecycle Manager, [315](#)
- vSphere Replication, [6](#), [206](#), [219–220](#)
  - integration, [220–221](#)
  - use cases, [220](#)
- vSphere Replication Management Service. *See* [VRMS](#)
- vSphere Replication Service. *See* [VRS](#)
- vSphere standard switches, [18](#)
- vSphere Virtual Machine Encryption certificates, [240](#)
- vSphere with Tanzu, [173](#), [204](#)
  - integration, [205](#)
  - use cases, [204](#)
- vSS (vSphere Standard Switches), [94](#), [95–97](#)
  - configuring, [330–332](#)
  - creating, [330–332](#)
  - network policies, [98](#)
  - vDS comparison, [103–104](#)
- vTA (vSphere Trust Authority), [258–259](#)
  - configuring, [502–504](#)
  - management, [502–504](#)
  - operations, [504](#)
- vTPM (Virtual Trusted Platform Module), [273–274](#)
- Vue (Pearson)
  - exam preparation, [604](#)
  - requirements, [604](#)
- VUM (VMware Update Manager). *See* [vSphere Lifecycle Manager](#)
- vVols (Virtual Volumes), [40–41](#), [72](#)

- architecture, [72](#)
- characteristics, [72–73](#)
- configuring, [463–464](#)
- Config-vVol, [73](#)
- datastores, [45](#)
- Data-vVol, [73](#)
- limitations, [73](#)
- management, [463–464](#)
- Mem-vVol, [73](#)
- Other-vVol, [73](#)
- Swap-vVol, [73](#)
- vSAN, [54](#)

## W

- warning events, [397](#)
- web proxies, ESXi security settings, [490–491](#)
- Windows Perfmon, [391–392](#)
- Windows Session Authentication, enabling SSO, [472–473](#)
- witness nodes, vCenter HA, [14](#)
- witnesses, vSAN, [52](#), [54](#)
- workflows, evacuation, [136](#)
- Write Same (Zero), [71](#)

## X - Y - Z

- XCOPY (Extended Copy), [70](#)
- zeroing out files, [79](#)

Exclusive Offer – 40% OFF

## Pearson IT Certification Video Training

livelessons

[pearsonitcertification.com/video](https://pearsonitcertification.com/video)

Use coupon code **PITCVIDEO40** during checkout.



## Video Instruction from Technology Experts



### Advance Your Skills

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



### Train Anywhere

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



### Learn

Learn from trusted author trainers published by Pearson IT Certification.

## Try Our Popular Video Training for FREE!

[pearsonitcertification.com/video](https://pearsonitcertification.com/video)

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

PEARSON  
IT CERTIFICATION

[pearsonitcertification.com/video](https://pearsonitcertification.com/video)

ALWAYS LEARNING

PEARSON







**REGISTER YOUR PRODUCT** at [PearsonITcertification.com/register](https://PearsonITcertification.com/register)  
**Access Additional Benefits and SAVE 35% on Your Next Purchase**

- Download available product updates.
- Access bonus material when applicable.
- Receive exclusive offers on new editions and related products.  
(Just check the box to hear from us when setting up your account.)
- Get a coupon for 35% for your next purchase, valid for 30 days. Your code will be available in your PITC cart. (You will also find it in the Manage Codes section of your account page.)

Registration benefits vary by product. Benefits will be listed on your account page under Registered Products.

---

**PearsonITcertification.com—Learning Solutions for Self-Paced Study, Enterprise, and the Classroom**

Pearson is the official publisher of Cisco Press, IBM Press, VMware Press, Microsoft Press, and is a Platinum CompTIA Publishing Partner—CompTIA's highest partnership accreditation.

At [PearsonITcertification.com](https://PearsonITcertification.com) you can

- Shop our books, eBooks, software, and video training.
- Take advantage of our special offers and promotions ([pearsonitcertification.com/promotions](https://pearsonitcertification.com/promotions)).
- Sign up for special offers and content newsletters ([pearsonitcertification.com/newsletters](https://pearsonitcertification.com/newsletters)).
- Read free articles, exam profiles, and blogs by information technology experts.
- Access thousands of free chapters and video lessons.

**Connect with PITC – Visit [PearsonITcertification.com/community](https://PearsonITcertification.com/community)**

Learn about PITC community events and programs.



## PEARSON IT CERTIFICATION

Addison-Wesley • Cisco Press • IBM Press • Microsoft Press • Pearson IT Certification • Prentice Hall • Que • Sams • VMware Press

ALWAYS LEARNING

PEARSON



## Appendix B

# Memory Tables

## Chapter 1

**Table 1-4** Available vSphere Features

---

Available vSphere Features	Description
----------------------------	-------------

---

A feature introduced in vSphere 7.0 that enables you to back up and restore the vCenter Server Appliance instances.

---

A feature that provides live virtual machine migrations with negligible disruption from a source ESXi host to a target ESXi host.

---

A feature that provides automated failover protection for VMs against host, hardware, network, and guest OS issues. In the event of host system failure, it performs cold migrations and restarts failed VMs on surviving hosts.

---

A feature that places and starts VMs on appropriate ESXi hosts and hot-migrates VMs using vMotion when there is contention for compute resources.

---

A feature that performs live migrations with negligible disruption of VMs from a source datastore to a target datastore.

---

A feature that provides automated live failover protection for VMs against host, hardware, network, and guest OS issues.

---

A feature that optimizes power consumption in an ESXi cluster.

---

A feature that minimizes VM downtime by proactively detecting hardware failures and placing the host in Quarantine Mode or

Maintenance Mode.

A centralized repository used to manage and distribute templates, ISO files, scripts, vApps, and other files associated with VMs.

A feature that provides a means to apply a standard configuration to a set of ESXi hosts.

**Table 1-6** vCenter Server Editions

<b>Feature</b>	<b>Essentials</b>	<b>Essentials Plus</b>	<b>For Standard Edition</b>
Number of ESXi hosts	3 (2 CPU max)	3 (2 CPU max)	2000
vCenter License	Packaged with vSphere license in Essentials	Packaged with vSphere license in Essentials Plus	Sold separately from vSphere license
Basic vCenter features, like single pane of glass management, Lifecycle Manager, and VMware Converter	Supported	Supported	Supported
Common vCenter features like vMotion, vSphere HA, and vSphere Replication	Not supported	Supported	Supported
Advanced features like vCenter Server High Availability (VCHA) and vCenter Server Backup and Restore	N/A	N/A	Supported

**Table 1-10** Compute Specifications for vCenter Server Appliance

<b>Component</b>	<b>Number of CPUs</b>	<b>Memory</b>
------------------	-----------------------	---------------

---

Tiny Environment

Up to 10 hosts or 100 virtual machines

---

Small Environment

Up to 100 hosts or 1000 virtual machines

---

Medium Environment

Up to 400 hosts or 4000 virtual machines

---

Large Environment

Up to 1000 hosts or 10,000 virtual machines

---

X-Large Environment

Up to 2000 hosts or 35,000 virtual machines

---

**Table 1-11** Storage Sizes for vCenter Server Appliance

<b>Deployment Size</b>	<b>Default Storage Size</b>	<b>Large Storage Size</b>	<b>X-Large Storage Size</b>
Tiny		1490 GB	
Small		1535 GB	
Medium		1700 GB	
Large		1765 GB	
X-Large		1905 GB	

**Table 1-12** Required Ports for vCenter Sever

<b>Port</b>	<b>Description</b>	<b>Required for</b>
22	System port for SSHD	vCenter Server

	(Must be open for upgrade of the appliance.)
Port for direct HTTP connections; redirects requests to HTTPS port 443	vCenter Server
Required to be open to join Active Directory	vCenter Server
LDAP port for directory services for the vCenter Server group	vCenter Server to vCenter Server
Default port used by vCenter Server to listen for connections from the vSphere Web Client and SDK clients	vCenter Server to vCenter Server
vSphere Syslog Collector port for vCenter Server and vSphere Syslog Service port for vCenter Server Appliance	vCenter Server
Default port that the vCenter Server system uses to send data to managed hosts	vCenter Server
vSphere Syslog Collector TLS port for vCenter Server	vCenter Server
Control interface RPC for Single Sign-On	vCenter Server
RPC port for VMware Certificate Authority (VMCA) APIs	VMCA
Authentication framework management	vCenter Server
vCenter Server Appliance Management Interface (VAMI)	vCenter Server
ESXi Dump Collector port	vCenter Server
Secure Token Service (internal ports)	vCenter Server
vSphere Client (internal ports)	vCenter Server
VMware vSphere Authentication Proxy	vCenter Server
vSphere Lifecycle Manager SOAP port used by vSphere Lifecycle	vSphere Lifecycle Manager
Manager client plug-in	

vSphere Lifecycle Manager Web Server	vSphere Lifecycle Manager
Port used by ESXi hosts to access host patch files from vSphere Lifecycle Manager server	
vSphere Lifecycle Manager Web SSL port used by vSphere Lifecycle Manager client plug-in for uploading host upgrade files to vSphere Lifecycle Manager server	vSphere Lifecycle Manager
vSphere Web Client HTTPS	vCenter Server

## Chapter 2

**Table 2-2** Comparison of VMFS Version 5 and Version 6

<b>VMFS Features and Functionalities</b>	<b>Version 5</b>	<b>Version 6</b>
Access for ESXi hosts Version 6.5 and later	Yes	Yes
Access for ESXi hosts Version 6.0 and earlier	Yes	No
Datastores per host	512	512
512n storage devices	Yes	Yes (default)
512e storage devices	Yes (Not supported on local 512e devices.)	
4Kn storage devices	No	
Automatic space reclamation	No	
Manual space reclamation through the <b>esxcli</b> command.	Yes	
Space reclamation from guest OS	Limited	
GPT storage device partitioning	Yes	

MBR storage device partitioning	Yes	
		For a VMFS5 datastore that has been previously upgraded from VMFS3.
Storage devices greater than 2 TB for each VMFS extent	Yes	
Support for virtual machines with large-capacity virtual disks, or disks greater than 2 TB	Yes	
Support of small files (1 KB)	Yes	
Default use of ATS-only locking mechanisms on storage devices that support ATS	Yes	
Block size	Standard 1 MB	Standard 1 MB
Default snapshots	VMFSsparse for virtual disks smaller than 2 TB	SEsparse
	SEsparse for virtual disks larger than 2 TB	
Virtual disk emulation type		512n
vMotion		Yes
Storage vMotion across different datastore types		Yes
High Availability and Fault Tolerance		Yes
DRS and Storage DRS		Yes
RDM		Yes

**Table 2-4** Comparison of NFS Version 3 and Version 4.1 Support for vSphere



## Features and Solutions

<b>NFS Features and Functionalities</b>	<b>Ver</b>	<b>Version 4.1 sion 3</b>
vMotion and Storage vMotion	Y e s	Yes
High Availability (HA)	Y e s	Yes
Fault Tolerance (FT)	Y e s	Yes (Supports the new FT mechanism introduced in vSphere 6.0 that supports up to four vCPUs, not the legacy FT mechanism.)
Distributed Resource Scheduler (DRS)	Y e s	
Host Profiles	Y e s	
Storage DRS	Y e s	
Storage I/O Control	Y e s	
Site Recovery Manager	Y e s	
Virtual Volumes	Y e s	
vSphere Replication	Y e	

	s
vRealize Operations Manager	Yes

**Table 2-7 RAID Configuration Comparison**

RAID Configuration	PF TT	Data Size	Required Capacity	Usable Capacity
RAID 1 (mirroring)	1	100 GB		
RAID 5 or RAID 6 (erasure coding) with four fault domains	1	100 GB		
RAID 1 (mirroring)	2	100 GB		
RAID 5 or RAID 6 (erasure coding) with six fault domains	2	100 GB		
RAID 1 (mirroring)	3	100 GB		
RAID 5 or RAID 6 (erasure coding) with six fault domains	3	N/A		

**Table 2-12 vSAN Storage Policies**

Policy	Description
1	This policy defines how many host and device failures a VM object can withstand. For $n$ failures tolerated, data is stored in $n+1$ location. (This includes parity copies with RAID 5 or 6.) If no storage policy is selected at the time of provisioning a VM, this policy is assigned by default. Where fault domains are used, $2n+1$ fault domains, each with hosts adding to the capacity, are required. If an ESXi host isn't in a fault domain, it is considered to be in a single-host fault domain. The default setting for this policy is 1, and the maximum is 3.

In stretched clusters, this policy defines how many additional host failures

can be tolerated after a site failure's PFTT has been reached. If PFTT = 1, SFTT = 2, and one site is inaccessible, two more host failures can be tolerated. The default setting for this policy is 1, and the maximum is 3.

---

If PFTT = 0, this option is available. The options for this policy are None, Preferred, and Secondary. This allows objects to be limited to one site or one host in stretched clusters. The default setting for this policy is None.

---

This policy defines whether the data replication mechanism is optimized for performance or capacity. If RAID-1 (Mirroring)—Performance is selected, there will be more space consumed in the object placement but better performance for accessing the space. If RAID-5/6 (Erasure Coding)—Capacity is selected, there will be less disk utilization, but performance will be reduced.

---

This policy determines the number of capacity devices where each VM object replica is striped. Setting this above 1 can improve performance but consumes more resources. The default setting for this policy is 1, and the maximum is 12.

---

This policy defines the amount of flash capacity that is reserved for read caching of VM objects. This is defined as a percentage of the size of the VMDK. This is supported only in hybrid vSAN clusters. The default setting for this policy is 0%, and the maximum is 100%.

---

If set to yes, this policy forces provisioning of objects, even when policies cannot be met. The default setting for this policy is no.

---

This policy defines the percentage of VMDK objects that must be thick provisioned on deployment. The options are as follows:

- Thin provisioning (default value)
  
- 25% reservation
  
- 50% reservation

- 75% reservation
- Thick provisioning

---

A checksum is used end-to-end in validating the integrity of the data to ensure that data copies are the same as the original. In the event of a mismatch, incorrect data is overwritten. If this policy is set to yes, a checksum is not calculated. The default setting for this policy is no.

---

This policy sets a limit for IOPS of an object. If set to 0, there is no limit.

---

## Chapter 3

**Table 3-2** Advantages and Disadvantages of IP Hash NIC Teaming

<b>Advantages</b>	<b>Disadvantages</b>
A more even distribution of the load compared to Route Based on Originating Virtual Port and Route Based on Source MAC Hash	
A potentially higher throughput for virtual machines that communicate with multiple IP addresses	

**Table 3-3** Comparison of vSS and vDS Features

<b>Feature</b>	<b>vSS</b>	<b>vDS</b>
Layer 2 switch		
VLAN segmentation (802.1q tagging)		
IPv6 support		
NIC teaming		
Outbound traffic shaping		
Cisco Discovery Protocol (CDP)		
Inbound traffic shaping		

VM network port block

---

Private VLANs

---

Load-based NIC teaming

---

Data center-level management

---

Network vMotion

---

Per-port policy settings

---

Port state monitoring

---

NetFlow

---

Port mirroring

---

**Table 3-4** vDS Health Checks

<b>Health Check</b>	<b>Required vDS Configuration</b>
Checks whether the VLAN trunk ranges on the distributed switch match the trunk port configuration on the connected physical switch ports.	
Checks for matching MTU settings on the distributed switch, the physical network adapter, and the physical switch ports.	
Checks whether the virtual switch teaming policy matches the physical switch port-channel settings.	

**Table 3-5** SR-IOV Requirements

<b>Component</b>	<b>Requirements</b>
Physical host	
Physical network adapter	

---

Physical function (PF) driver in ESXi

---

Guest OS

---

Virtual function (VF) driver in guest OS

---

---

## Chapter 4

**Table 4-4** Resource Pool Use Cases

<b>Use Case</b>	<b>Details</b>
Flexible hierarchical organization	Add, remove, modify, and reorganize resource pools, as needed.
Resource isolation	Use resource pools to allocate resources to separate departments, in such a manner that changes in a pool do not unfairly impact other departments.

---

Managing multitier applications.	Manage the resources for a group of virtual machines (in a specific resource pool), which is easier than managing resources per virtual machine.
----------------------------------	--

**Table 4-6** Virtual Machine Shares

Setting	CPU Share Value	Memory Share Value
High		
Normal	1000 per vCPU	10 per MB
Low		

**Table 4-9** Advanced vSphere HA Options

Option	Description
das.isolationaddressX	Provides the addresses to use to test for host isolation when no heartbeats are received from other hosts in the cluster. If this option is not specified (which is the default setting), the management network default gateway is used to test for isolation. To specify multiple addresses, you can set das.isolationaddressX, where X is a number between 0 and 9.
	Specifies whether to use the default gateway IP address for isolation tests.
das.isolationshutdowntimeout	For scenarios where the host's isolation response is to shut down, specifies the period of time that the virtual machine is permitted to shut down before the system powers it off.
	Defines the maximum bound on the memory slot size.
	Defines the maximum bound on the CPU slot size.
	Defines the default memory resource value assigned to a virtual machine whose memory reservation is not specified or is zero. This is used for the Host Failures Cluster Tolerates admission control policy.
	Defines the default CPU resource value assigned to a virtual machine

whose CPU reservation is not specified or is zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default of 32 MHz is used.

das.h earb eatds perho st	Specifies the number of heartbeat datastores required per host. The default is 2. The acceptable values are 2 to 5.
das.c onfig .fdm. isolat ionPo licyD elayS ec	Specifies the number of seconds the system delays before executing the isolation policy after determining that a host is isolated. The minimum is 30. A lower value results in a 30-second delay.
	Determines whether vSphere HA should enforce VM–VM anti-affinity rules even when DRS is not enabled.

**Table 4-10** VM Monitoring Settings

Setting	Failure Interval	Reset Period
High		1 hour
Medium		24 hours
Low		7 days

## Chapter 5

**Table 5-2** Virtual Machine Files

Description
Virtual machine configuration file
Additional virtual machine configuration file
Virtual disk characteristics (metadata) file
Virtual disk data file (commonly called a flat file)



Virtual machine BIOS or UEFI configuration file
Virtual machine snapshot file
Virtual machine snapshot data file
Virtual machine swap file
Virtual machine suspend file
Current virtual machine log file
Old virtual machine log file, where # is a number starting with 1

**Table 5-4** Virtual Machine Options

Category	Description
General Options	
Encryption Options	
Power Management	
VMware Tools	Settings allow you to choose how to respond to specific power operations. For example, you can choose whether to power off the virtual machine or shut down the guest when the red power-off button is clicked.
Virtualization Based Security (VBS)	
Boot Options	

Advanced  
Options

---

Fibre Channel NPIV	Settings allow the virtual machine to use N_Port ID Virtualization (NPIV), including whether to generate new worldwide names (WWNs).
vApp Options	Settings allow you to control vApp functionality for the virtual machine, such as enable/disable and IP allocation policy. vApp settings that are made directly to a virtual machine override settings made on the vApp.

---

## Chapter 6

**Table 6-2** Required Permissions for the vCenter Cloud Account

---

<b>Object</b>	<b>Permissions</b>
Datastore	
Datastore cluster	
Folder	
Global	

---

Network

---

Permissions

---

Resource

---

Content library

- Add library item
- Create local library
- Create subscribed library
- Delete library item
- Delete local library
- Delete subscribed library
- Download files
- Evict library item
- Evict subscribed library
- Probe subscription information

- Read storage
- Sync library item
- Sync subscribed library
- Type introspection
  
- Update configuration settings
- Update files
- Update library
- Update library item
- Update local library
- Update subscribed library
- View configuration settings

---

## Tags

- Assign or unassign vSphere tag

- Create a vSphere tag
- Create a vSphere tag category
- Delete vSphere tag
- Delete vSphere tag category
- Edit vSphere tag
- Edit vSphere tag category
- Modify UsedBy field for category
- Modify UsedBy field for tag

---

## vApp

- Import
- vApp application configuration.

---

## Virtual machine inventory

---

## Virtual machine interaction

---

## Virtual machine configuration

- Add existing disk
- Add new disk
- Add or remove
- Remove disk
- Advanced
- Change CPU count
- Change resource
- Extend virtual disk
- Disk change tracking

- Memory
- Modify device settings
- Rename
- Set annotation
- Settings
- Swapfile placement

---

## Virtual machine provisioning

---

## Virtual machine state

- Create snapshot
  - Remove snapshot
  - Revert to snapshot
-

**Table 6-3** Required vCenter Server Privileges for Horizon (without instant clones)

<b>Privilege Group</b>	<b>Privileges to Enable</b>
Folder	
Datastore	
Virtual Machine	
Resource	
Global	
Host (for Storage Accelerator)	



---

Profile Driven Storage (for vSAN or Virtual Volumes)

---

**Table 6-5** VMware HCX Services

---

**Service Description**

Advanced	Creates secured connections between HCX instances, supporting migration, replication, disaster recovery, and management operations.  Deployed as a virtual appliance.
Advanced	Optimizes the performance of the connection provided by HCX Interconnect through a combination of deduplication, compression, and line conditioning techniques.  Deployed as a virtual appliance.
Advanced	Extends (that is, provides Layer 2 adjacency) the virtual machine networks between source and remote HCX-enabled environments.  Deployed as a virtual appliance.
Advanced	Migrates a set of virtual machines using VMware vSphere Replication in parallel between HCX-enabled sites.
Advanced	Migrates a single virtual machine between HCX-enabled sites with no service interruption, using vMotion.
Advanced	Protects virtual machines from disaster by using replication and

dv  
an  
ce  
d

---

E  
nt  
er  
pr  
is  
e

Allows you to group virtual machines by application, network, or other aspects for migration and monitoring.

---

E  
nt  
er  
pr  
is  
e

Leverages HCX Sentinel software in the guest OS to migrate Windows and Linux virtual machines to a vSphere-enabled data center.

Uses a gateway appliance at the source and a receiver appliance at the destination.

---

E  
nt  
er  
pr  
is  
e

Migrates a set of virtual machines in parallel, using VMware vSphere Replication and vMotion between HCX-enabled sites with no service interruption.

---

E  
nt  
er  
pr  
is  
e

Integrates HCX functionality with the VMware SRM for protection and orchestrated recovery operations.

---

E  
nt  
er  
pr  
is  
e

Optimizes network traffic for HCX Interconnect and Network Extension services.

The Application Path Resiliency service creates multiple tunnel flows for both Interconnect and Network Extension traffic.

The TCP Flow Conditioning service adjusts and optimizes the

segment size to reduce fragmentation and reduce the overall packet rate.

---

E nt er pr is e	Integrates HCX Network Extension with NSX Dynamic Routing to enable optimal networking between migrated virtual machines and other virtual machines.
	Works with new or existing network extensions to NSX-T 3.0 Data Center.

---

## Chapter 7

**Table 7-2** Core Identity Services in vSphere

---

<b>Service</b>	<b>Description</b>
----------------	--------------------

---

	Serves as an identity source that handles SAML certificate management for authentication with vCenter Single Sign-On.
--	---

---

	Issues certificates for VMware solution users, machine certificates for machines on which services are running, and ESXi host certificates. VMCA can be used as is, or it can be used as an intermediary certificate authority.
--	---

---

	Includes VMware Endpoint Certificate Store (VECS) and several internal authentication services.
--	---

---

**Table 7-6** Certificates in vSphere

---

<b>Certificate</b>	<b>Pro vider Details</b>
--------------------	----------------------------------

---

ESXi certificate	Stored locally on an ESXi host in the /etc/vmware/ssl directory when the host is first added to vCenter Server and when it reconnects.
------------------	--

---

Machine SSL certificate	Stored in VECS.
-------------------------	-----------------

---

Used to create SSL sockets for SSL client connections, for server verification, and for secure communication such as HTTPS and LDAPS.

Used by the reverse proxy service, the vCenter Server service (vpxd), and the VMware Directory service (vmdir).

Uses X.509 Version 3 certificates to encrypt session information.

---

Solution user certificate

Stored in VECS.

Used by solution users to authenticate to vCenter Single Sign-On through SAML token exchange.

---

vCenter Single Sign-On SSL signing certificate

Used throughout vSphere for authentication, where a SAML token represents the user's identity and contains group membership information.

You can manage this certificate from the command line. Changing this certificate in the file system leads to unpredictable behavior.

---

VMware Directory Service (vmdir) SSL certificate

Starting with vSphere 6.5, the machine SSL certificate is used as the vmdir certificate.

---

vSphere Virtual Machine Encryption Certificates

Used for virtual machine encryption, which relies on an external key management server (KMS).

Depending on how the solution authenticates to the KMS, it might generate certificates and store them in VECS.

---

**Table 7-9** System Roles in vCenter Server 7.0

---

## System Description

m

### Role

---

Read-only	Allows the user to view the state of an object and details about the object. For example, users with this role can view virtual machine attributes but cannot open the VM console.
Administrator	Includes all privileges of the read-only role and allows the user to view and perform all actions on the object. If you have the administrator role on an object, you can assign privileges to individual users and groups. If you have the administrator role in vCenter Server, you can assign privileges to users and groups in the default SSO identity source. By default, the <a href="#">administrator@vsphere.local</a> user has the administrator role on both vCenter Single Sign-On and vCenter Server.
No access	Prevents users from viewing or interacting with the object. New users and groups are effectively assigned this role by default.
No cryptographic administrator	Includes all privileges of the administrator role, except for cryptographic operations privileges. This role allows administrators to designate users who can perform all administrative tasks except encrypting or decrypting virtual machines or accessing encrypted data.
Trustee infrastructure administrator	Allows users to perform VMware vSphere Trust Authority operations on some objects. Membership in the TrustedAdmins group is required for full vSphere Trust Authority capabilities.

nist  
rat  
or  
rol  
e

---

**Table 7-10** Required Permissions for Common Tasks

---

<b>Task</b>	<b>Required Privileges</b>
-------------	----------------------------

---

Create a virtual machine	
--------------------------	--

---

Deploy a virtual machine from a template	On the destination folder or in the data center:
--	--

- Virtual Machine.Inventory.Create from Existing

- Virtual Machine.Configuration.Add New Disk

On a template or in a template folder:

- Virtual Machine.Provisioning.Deploy Template

On the destination host or cluster or in the resource pool:

- Resource.Assign Virtual Machine to Resource Pool

On the destination datastore or in a datastore folder:

- Datastore.Allocate Space

On the network that the virtual machine will be assigned to:

- Network.Assign Network

---

Take a virtual machine snapshot

On the virtual machine or in a virtual machine folder:

- Virtual Machine.Snapshot Management.Create Snapshot

On the destination datastore or in a datastore folder:

- Datastore.Allocate Space

---

Move a virtual machine into a resource pool

On the virtual machine or in a virtual machine folder:

- Resource.Assign Virtual Machine to Resource Pool
- Virtual Machine.Inventory.Move

In the destination resource pool:

- Resource.Assign Virtual Machine to Resource Pool
-



Install a guest operating system on a virtual machine

On the virtual machine or in a virtual machine folder:

- Virtual Machine.Interaction.Answer Question
- Virtual Machine.Interaction.Console Interaction
- Virtual Machine.Interaction.Device Connection
- Virtual Machine.Interaction.Power Off
- Virtual Machine.Interaction.Power On
- Virtual Machine.Interaction.Reset
- Virtual Machine.Interaction.Configure CD Media
- Virtual Machine.Interaction.Configure Floppy Media
- Virtual Machine.Interaction.Tools Install

---

On a datastore containing the installation media ISO image:

- Datastore.Browse Datastore

On the datastore to which you upload the installation media ISO image:

- Datastore.Browse Datastore
- Datastore.Low Level File Operations

---

Migrate a virtual machine with vMotion

---

Cold migrate (relocate) a virtual machine

---

Migrate a virtual machine with Storage vMotion

---

On the virtual machine or in a virtual machine folder:

- Resource.Migrate Powered On Virtual Machine

On the destination datastore:

- Datastore.Allocate Space

---

Move a host into a cluster

On the host:

- Host.Inventory.Add Host to Cluster

On the destination cluster:

- Host.Inventory.Add Host to Cluster
- Host.Inventory.Modify. cluster

---

**Table 7-11** ESXi Security Profile Services

Service	Default State Description
Direct Console User Interface (DCUI)	Allows you to interact with an ESXi host from the local console host using text-based menus

ESXi Shell	Is available from the DCUI or from SSH
SSH	Allows remote connections through Secure Shell
Load-Based Teaming Daemon	Enables load-based teaming
attestd	Enables the vSphere Trust Authority Attestation Service
kmxd	Enables the vSphere Trust Authority Key Provider Service
Active Directory Service	Is started on hosts after you configure ESXi for Active Directory
NTP Daemon	Enables the Network Time Protocol daemon
PC/SC Smart Card Daemon	Is started on hosts after you enable the host for smart card authentication
CIM Server	Can be used by Common Information Model (CIM) applications
SNMP Server	Enables the SNMP daemon
Syslog Server	Enables the syslog daemon
VMware vCenter Agent (vpxa)	Connects the host to vCenter Server
X.Org Server	Internally used for virtual machine 3D graphics

**Table 7-12** Incoming and Outgoing Firewall Ports

<b>Firewall Service</b>	<b>Incoming Port(s)</b>	<b>Outgoing Port(s)</b>
CIM Server	5988 (TCP)	
CIM Secure Server	5989 (TCP)	
CIM SLP	427 (TCP,UDP)	427 (TCP,UDP)
DHCPv6	546 (TCP,UDP)	547 (TCP,UDP)
DVSSync	8301, 8302 (UDP)	8301, 8302 (UDP)
HBR		44046, 31031

		(TCP)
NFC	902 (TCP)	902 (TCP)
WOL		9 (UDP)
vSAN Clustering	12345, 23451 (UDP)	12345, 23451 (UDP)
DCHP Client	68 (UDP)	68 (UDP)
DNS Client	53 (UDP)	53 (TCP,UDP)
Fault Tolerance		
NSX Distributed Logical Router Service	6999 (UDP)	6999 (UDP)
Software iSCSI Client		3260 (TCP)
rabbitmqproxy		5671 (TCP)
vSAN Transport		
SNMP Server	161 (UDP)	
SSH Server	22 (TCP)	
vMotion		
VMware vCenter Agent		902 (UDP)
vSphere Web Access		
vsanvp	8080 (TCP)	8080 (TCP)
RFB Protocol	5900–5964 (TCP)	
vSphere Life Cycle Manager		
I/O Filter	9080 (TCP)	

**Table 7-14** Network Security Policies

**O Se Description**  
**pti tti**  
**on ng**

The virtual switch forwards all frames to the virtual network adapter.

The virtual switch forwards only the frames that are addressed to the virtual network adapter.

If the guest operating system changes the effective MAC address of the virtual adapter to a value that differs from the MAC address assigned to the adapter in the VMX file, the virtual switch allows the inbound frame to pass.

---

If the guest operating system changes the effective MAC address of the virtual adapter to a value that differs from the MAC address assigned to the adapter in the VMX file, the virtual switch drops all inbound frames to the adapter. If the guest OS changes the MAC address back to its original value, the virtual switch stops dropping the frames and allows inbound traffic to the adapter.

---

The virtual switch does not filter outbound frames. It permits all outbound frames, regardless of the source MAC address.

---

The virtual switch drops any outbound frame from a virtual machine virtual adapter that uses a source MAC address that differs from the MAC address assigned to the virtual adapter in the VMX file.

---

## Chapter 8

**Table 8-2** Information Required for ESXi Installation

<b>Information</b>	<b>Required or Optional</b>	<b>Details</b>
Keyboard layout	Required	
VLAN ID	Optional	
IP address	Optional	
Subnet mask	Optional	
Gateway	Optional	
Primary DNS	Optional	

Secondary DNS	Optional
Host name	Required for static IP settings
Install location	Required
Migrate existing ESXi settings; preserve VMFS datastore	Required if you are installing ESXi on a drive with an existing ESXi installation
Root password	Required

**Table 8-5** Auto Deploy Components

<b>Component</b>	<b>Description/Purpose</b>
Auto Deploy	Uses a rules engine, a set of images, a set of host profiles, and required infrastructure to manage ESXi deployments.
Image Profiles	Assigns image profiles and host profiles to each host.
Host Profiles	Defines host-specific configurations, such as networking, NTP, and host permissions. You can use host customization in conjunction with host profiles to provide details that are unique to each host, such as IP address.
Auto Deploy Server	Serves as a command-line engine for driving Auto Deploy.
Image Builder	Serves as a command-line engine for building images.
Inventory	Manages the vSphere inventory and provides host profiles.
IP Address Manager	Provides IP configuration to the host and redirects the host to the PXE

server.

---

Boots the host and directs it to the TFTP server.

---

Provides the appropriate boot image.

---

Holds a collection of VIBs either online (accessible via HTTP) or offline (accessible via a USB drive or CD/DVD).

---

Holds a collection of VIBs used to install the ESXi server and saved as ZIP files or ISO images. You can obtain image profiles from VMware and VMware partners, and you can create custom image profiles by using ESXi Image Builder.

---

Packages a collection of files (such as drivers) into an archive similar to a ZIP file. Each VIB is released with an acceptance level that cannot be changed. The host acceptance level assigned to each host determines which VIBs can be installed to the host. These are the acceptance levels, from highest to lowest:

- VMwareCertified
  
- VMwareAccepted
  
- PartnerSupported
  
- CommunitySupported

---

**Table 8-9** VECS Stores

Store	Description
	Used by the reverse proxy service on each ESXi host and by the vmdir service.
	Contains all trusted root certificates.



---

Solution user stores: VECS includes one store for each solution user.

- | Machine
  
- | vpxd
  
- | vpxd-  
extension
  
- | vsphere-  
webclient

---

Used by VMCA to support certificate reversion.

---

Other stores: Other stores might be added by solutions. For example, the Virtual Volumes solution adds an SMS store.

---

The SSO domain contains many predefined groups, including the following:

- **Users:** This group contains all users in the SSO domain.
- \_\_\_\_\_: Members of this group can perform domain controller administrator actions on VMware Directory Service.
- \_\_\_\_\_: Each solution user authenticates individually to vCenter Single Sign-On with a certificate. By default, VMCA provisions solution users with certificates. Do not add members to this group explicitly.
- \_\_\_\_\_: Members have administrator privileges for VMCA. Adding members to these groups is not usually recommended, but a user must be a

member of this group to perform most certificate management operations, such as using the **certool** command.

- **SystemConfiguration.BashShellAdministrators**: Members can enable and disable access to the BASH Shell.
- \_\_\_\_\_: Members can view and manage the system configuration and perform tasks such as restarting services.
- \_\_\_\_\_: Members have full write access to all licensing-related data and can add, remove, assign, and un-assign serial keys for all product assets registered in licensing service.
- \_\_\_\_\_: Members can perform SSO administration tasks for VMware Directory Service (vmdir).

**Table 8-10** SSO Policies and Parameters

SSO Policy Setting Parameter	Details	
Password Policy	Description	Password policy description.
		Maximum number of days a password can exist before the user must change it.
		Number of the user's previous passwords that cannot be selected.
		Maximum number of characters that are allowed in the password.
		Minimum number of characters that are allowed in the password, which must be no fewer than the combined minimum of alphabetic, numeric, and special character requirements.
		Minimum number of different character types that are required in the password. The types include special, alphabetic, uppercase, lowercase, and numeric.
Identical adjacent	The number of identical adjacent characters that are supported in a password. The value must be greater than	

	characters	0.
Lockout Policy	Description	Description of the lockout policy.
		Maximum number of failed login attempts that are allowed before the account is locked.
		Time period in which failed login attempts must occur to trigger a lockout.
	Unlock time	The amount of time the account stays locked. The value 0 specifies that an administrator must explicitly unlock the account.
Token Policy		Time difference, in milliseconds, that SSO tolerates between a client clock and a domain controller clock. If the time difference is greater than the specified value, SSO declares the token to be invalid.
		Maximum number of times a token may be renewed before a new security token is required.
		Maximum number of times a single holder-of-key token can be delegated.
		The lifetime value of a bearer token before the token must be reissued.
	Maximum holder-of-key token lifetime	The lifetime value of a holder-of-key token before the token is marked invalid.

**Table 8-11** ESXi 7.0 Kernel Options

**Kernel Option**

**autoPartition** This option, if set to TRUE, defines automatic partitioning of the unused local storage devices at boot time. The boot disk gets

**rtition=TRUE/FA** partitioned with boot bands, ESXi-OSData, and, if the disk is larger than 128 GB, a VMFS partition. Any new empty device discovered will be auto-partitioned as well. Auto-partitioning can be set for only the first unused device with the setting **autoPartitionOnlyOnceAndSkipSsd=TRUE**. On hosts with USB boot and VMFS-L, ESX-OSData does not exist on other local disks.

**LS E (default FA LS E)** If a storage device has both a scratch partition and a coredump partition, the scratch partition is converted to ESX-OSData; otherwise, the first unused disk identified is partitioned with ESX-OSData as well.

---

If this option is set to TRUE, local SSDs are excluded from automatic partitioning.

---

If this option is set to TRUE, SSD/NVMe devices are excluded, and the ESXi host automatically partitions the first unused local disk if there is no VMFS-L ESX-OSData volume.

---

If this option is set to TRUE, ESXi can write kernel crash coredumps to the VMFS-L Locker volume on a USB boot device.

---

This option sets the size of the coredump file (in megabytes) created on the system VMFS-L volume. This is limited to one-half of the space available on the VMFS-L volume.

---

This option, when set to TRUE, automatically creates a coredump file. This is attempted in the following order:

- VMFS-L ESX-OSData
  
- USB VMFS-L
  
- Local VMFS

# Chapter 9

**Table 9-2** VLAN ID Details

<b>VLAN ID</b>	<b>VLAN Tagging Mode</b>	<b>Description</b>
		The virtual switch does not pass traffic associated with a VLAN.
		The virtual switch tags traffic with the entered tag.
		Virtual machines handle VLANs. The virtual switch passes traffic from any VLAN.

Enhanced LACP support for vDS supports the following load-balancing modes (hashing algorithms):

- Destination IP address
- Destination IP address and TCP/UDP port
- Destination IP address and VLAN
- Destination IP address, TCP/UDP port, and VLAN
- Destination MAC address
- Destination TCP/UDP port
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_

- \_\_\_\_\_
- \_\_\_\_\_
- \_\_\_\_\_
- VLAN

## Chapter 10

**Table 10-4** Performance Chart Types

Chart Type	Description	Example
Line chart	Displays metrics for a single object, where data for each metric is represented by a separate line.	For example, A network chart for a host can contain one line showing the number of packets received and another line showing the number of packets transmitted.
Bar chart	Displays metrics for objects, where each bar represents metrics for an object.	A bar chart can display metrics for datastores, where each datastore is represented as a bar. Each bar displays metrics based on the file type, such as virtual disk or snapshot.
Pie chart	Displays metrics for a single object, where each slice represents a category or child object.	A datastore pie chart can display the amount of storage space occupied by each virtual machine or by each file type.
Stacked chart	Displays metrics for child objects.	A host's stacked CPU usage chart displays metrics for the 10 virtual machines on the host that are consuming the most CPU. The Other amount displays the total CPU usage of the remaining virtual machines.

**Table 10-6** CPU Performance Analysis

<b>Symptoms</b>	<b>Likely Causes</b>	<b>Potential Solutions</b>
Host: CPU usage is consistently high.	The host has insufficient CPU resources to meet the demand.	Add the host to a DRS cluster.
	Too many virtual CPUs are running on the host.	Increase the number of hosts in the DRS cluster.
Virtual machine: CPU usage is above 90%. CPU ready is above 20%. Application performance is poor.	Storage or network operations are placing the CPU in a wait state.	Migrate one or more virtual machines to other hosts.
	The guest OS generates too much load for the CPU.	Upgrade the physical CPUs of the host.  Upgrade ESXi to the latest version.  Enable CPU-saving features such as TCP segmentation offload, large memory pages, and jumbo frames.  Increase the amount of memory allocated to the virtual machines, which

may improve cached I/O and reduce CPU utilization.

Reduce the number of virtual CPUs assigned to virtual machines.

Ensure that VMware Tools is installed.

Compare the CPU usage of troubled virtual machines with that of other virtual machines on the host or in the resource pool. (Hint: Use a stacked graph.)

Increase the CPU limit, shares, or reservation on the troubled virtual machine.

---

Host: Memory usage is consistently 94% or higher. Free memory is

The host has insufficient memory resources to meet the demand.

Ensure that VMware Tools is installed and that the balloon driver is enabled for all virtual machines.



6% or less.

Virtual machine: Swapping is occurring. (Memory usage may be high or low.)

Reduce the memory size on oversized virtual machines.

Reduce the memory reservation of virtual machines where it is set higher than needed.

Add the host to a DRS cluster.

Increase the number of hosts in the DRS cluster.

Migrate one or more virtual machines to other hosts.

Add physical memory to the host.

---

Virtual machine: Memory usage is high.

The guest OS is not provided sufficient memory by the virtual machine.

Increase the memory size of the virtual machine.

Guest OS:  
Memory  
usage is  
high.  
Paging is  
occurring.

---

Virtual  
machine:  
CPU ready  
is low.

Guest OS:  
CPU  
utilization  
is high.

---

Datastore:  
Space  
utilization  
is high.

---

Disk:  
Device  
latency is  
greater than  
15 ms.

---

Disk:  
VMkernel  
latency is

The maximum throughput of a storage  
device is not sufficient to meet the  
demand of the current workload.

Migrate the virtual  
machines to  
datastores backed

greater than  
4 ms.  
Queue  
latency is  
greater than  
zero.

by storage devices  
(LUNs) with more  
spindles.

Balance virtual  
machines and their  
disk I/O across the  
available physical  
resources. Use  
Storage DRS I/O  
balancing.

Add more disks  
(spindles) to the  
storage device  
backing the  
datastore.

Configure the  
queue depth and  
cache settings on  
the RAID  
controllers. Adjust  
the  
Disk.SchedNumR  
eqOutstanding  
parameter.

Configure  
multipathing.

Increase the  
memory size of  
the virtual  
machine to  
eliminate any

guest OS paging.  
Increase the guest OS caching of disk I/O.

Ensure that no virtual machine swapping or ballooning is occurring.

Defragment guest file systems.

Use eager zeroed thick provisioned virtual disks.

---

Network:  
The number of packets dropped is greater than zero.  
Latency is high. The transfer rate is low.

The maximum throughput of a physical network adapter is not sufficient to meet the demand of the current workload.

Virtual machine network resource shares are too few.

Network packet size is too large, which results in high network latency. Use the VMware AppSpeed performance monitoring application or a third-party application to check network latency.

Network packet size is too small, which increases the demand for the CPU resources needed for processing each packet. Host CPU, or possibly virtual machine CPU, resources are not enough

Install VMware Tools on each virtual machine and configure the guest OS to use the best-performing network adapter driver (such as vmxnet3).

Migrate virtual machines to other hosts or to other physical network adapters.

Verify that all

to handle the load.

NICs are running in full duplex mode.

Implement TCP Segmentation Offload (TSO) and jumbo frames.

Assign additional physical adapters as uplinks for the associated port groups.

Replace physical network adapters with high-bandwidth adapters.

Place sets of virtual machines that communicate with each other regularly on the same ESXi host.

---

Performance charts are empty.

Some metrics are not available for pre-ESXi 5.0 hosts.

Upgrade hosts to a later version of ESXi.

Data is deleted when you remove objects to vCenter Server or remove them.

Allow time for data collection on objects that were recently added,

Performance chart data for inventory objects that were moved to a new site by

VMware vCenter Site Recovery Manager is deleted from the old site and not copied to the new site.

migrated, or recovered to the vCenter Server.

Performance chart data is deleted when you use VMware vMotion across vCenter Server instances.

Power on all hosts and allow time for real-time statistics to collect.

Real-time statistics are not available for disconnected hosts or powered-off virtual machines.

Allow time for the required roll-ups for non-real-time statistics.

Non-real-time statics are rolled up at specific intervals. For example, 1-day statistics might not be available for 30 minutes after the current time, depending on when the sample period began.

The 1-day statistics are rolled up to create one data point every 30 minutes. If a delay occurs in the roll-up operation, the 1-week statistics might not be available for 1 hour after the current time. It takes 30 minutes for the 1-week collection interval, plus 30 minutes for the 1-day collection interval.

The 1-week statistics are rolled up to create one data point every two hours. If a delay occurs in the roll-up operations, the 1-month statistics might not be available for 3 hours. It takes 2 hours for the 1-month collection interval, plus 1 hour for the 1-week collection interval.

The 1-month statistics are rolled up to create one data point every day. If a delay occurs in the roll-up operations, the statistics might not be available for 1 day and 3 hours. It takes 1 day for the past year collection interval, plus 3 hours for the past month collection interval. During this time, the charts are empty.

**Table 10-9** Key ESXTOP Panels and Metrics

Panel	Statistic	Description
CPU		Percentage of physical CPU core cycles used by the virtual machine.
CPU		Percentage of total time scheduled for the virtual machine without accounting for hyperthreading, system time, co-stopping, and waiting:  $\%RUN = 100\% - \%RDY - \%CSTP - \%WAIT$
CPU		Percentage of time the virtual machine was ready to run but was not provided CPU resources on which to execute. Indicator of CPU contention on the host.
CPU		Percentage of time the virtual machine spent in the blocked or busy wait state, including idle time. %WAIT includes %SWPWT.
CPU		Percentage of time a virtual machine spends in a ready, co-schedule state. A high value indicates that the virtual machine's multiple CPUs are in contention.
CPU		Percentage of time a virtual machine spends waiting for the host to swap memory.
Memory		Amount of physical memory allocated to a virtual machine:

$$\text{MEMSZ} = \text{GRANT} + \text{MCTLSZ} + \text{SWCUR} + \text{“Never Touched”}$$

Memory	Amount of guest physical memory mapped to a virtual machine
Memory	Amount of the memory consumed by the virtual machine:  $\text{CNSM} = \text{GRANT} - \text{Shared Memory}$
Memory	Amount of memory swapped by the virtual machine.
Memory	Rate at which the host swaps in memory from disk for the virtual machine.
Memory	Amount of memory used for virtual machine overhead, which is memory charged to the virtual machine that is not used by the guest OS.
Virtual Machine Storage	Number of read commands issued per second.
Virtual Machine Storage	Number of write commands issued per second.
Virtual Machine Storage	Megabytes read per second.
Virtual Machine Storage	Average latency (in milliseconds) per read.
Network	Number of packets received per second.
Network	Megabits transmitted per second.



Network	Percentage of transmit packets dropped. Indicates that the physical network adapter cannot meet the demand, perhaps due to load from other virtual machines.
Network	Percentage of receive packets dropped. Indicates that insufficient CPU resources are available for network processing.

**Table 10-13** ESXi Log Files

Component	Description
-----------	-------------

VMkernel	Data related to virtual machines and ESXi
VMkernel warnings	Data related to virtual machines
VMkernel summary	Data related to uptime and availability statistics for ESXi
ESXi host agent	Data related to the agent that manages and configures the ESXi host and its virtual machines
vCenter agent	Data related to the agent that communicates with vCenter Server
ESXi Shell	Data related to each command typed into the ESXi Shell as well as shell events
Authentication	Data related to event authentication for the local system
System messages	General log messages that can be used for troubleshooting

ges

Virtual machines		Data related to virtual machine power events, system failure information, tool status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and more
Trusted infrastructure agent	/var/run/log/kmxa.log	Data related to the client service on the ESXi trusted host
Key provider service	/var/run/log/kmxd.log	Data related to the vSphere Trust Authority key provider service
Attestation service	/var/run/log/atstd.log	Data related to the vSphere Trust Authority attestation service
ESX token service	/var/run/log/esxtokend.log	Data related to the vSphere Trust Authority ESXi token service
ESX API forwarder	/var/run/log/esxapiadapter.log	Data related to the vSphere Trust Authority API forwarder
Quick Boot	/var/log/loadESX.log	Data related to restarting an ESXi host through Quick Boot

**Table 10-14** vCenter Server Logging Options

Logging Option	Description
----------------	-------------

No vCenter Server logging occurs.

---

The vCenter Server collects only error entries in its log files.

---

The vCenter Server collects warning and error entries in its log files.

---

The vCenter Server collects information, warning, and error entries in its log files.

---

The vCenter Server collects verbose, information, warning, and error entries in its log files.

---

The vCenter Server collects trivia, verbose, information, warning, and error entries in its log files.

---

## Chapter 11

**Table 11-2** Network Differences in vSAN and non-vSAN Clusters

<b>Factor</b>	<b>vSAN Is Enabled</b>	<b>vSAN Is Not Enabled</b>
Network used by vSphere HA		Management network
Heartbeat datastores		Any datastore that is mounted to multiple hosts in the cluster
Host isolation criteria		Isolation addresses not pingable and management network inaccessible

**Table 11-4** Datastore Browser Options

<b>Option</b>	<b>Description</b>
---------------	--------------------

Upload a local file to the datastore.
Upload a local folder to the datastore.
Download a file from the datastore to the local machine.
Create a folder on the datastore.
Copy selected folders or files to a new location on the datastore or on another datastore.
Move selected folders or files to a new location on the datastore or on another datastore.
Rename selected files.
Delete selected folders or files.
Convert a selected thin virtual disk to thick.

**Table 11-5** Storage Filters

**Filter Description**

Hides storage devices (LUNs) that are used by a VMFS datastore on any host managed by vCenter Server.
Hides storage devices (LUNs) that are used by an RDM on any host managed by vCenter Server.
Hides storage devices (LUNs) that are ineligible for use as VMFS datastore extents because of incompatibility with the selected datastore. Hides LUNs that are not exposed to all hosts that share the original datastore. Hides LUNs that use a storage type (such as Fibre Channel, iSCSI, or local) that is different from the original datastore.
Automatically rescans and updates VMFS datastores following datastore management operations. If you present a new LUN to a host or a cluster, the hosts automatically perform a rescan, regardless of this setting.

**Table 11-7** SCSI over Fabric and NVMe over Fabric Comparison

<b>Shared Storage Capability</b>	<b>SCSI over Fabric</b>	<b>NVMe over Fabric</b>
RDM	Supported	
CoreDump	Supported	
SCSI-2 reservations	Supported	
Shared VMDK	Supported	
vVols	Supported	
Hardware acceleration with VAAI plug-ins	Supported	
Default MPP	NMP	
Limits	LUNs=1024, paths=4096	

## Chapter 12

**Table 12-2** Sample ESXCLI Commands

<b>Command</b>	<b>Description</b>
<b>esxcli system account add</b>	Creates an ESXi host local user account
	Configures an ESXi host local user account
<b>esxcli system account list</b>	Lists ESXi host local user accounts
<b>esxcli system account remove</b>	Deletes an ESXi host local user accounts
	Lists the host's DNS servers
	Lists the ESXi host's physical network adapters
	Displays the shell interactive timeout for the host

**Table 12-4** ESXi Lockdown Mode Behavior

<b>Service</b>	<b>Normal Mode</b>	<b>Normal Lockdown Mode</b>	<b>Strict Lockdown Mode</b>
vSphere Web Services API	All users, based on permissions	vCenter (vpxuser)  Exception users, based on permissions  vCloud Director (vslauser, if available)	vCenter (vpxuser)  Exception users, based on permissions  vCloud Director (vslauser, if available)
CIM providers	Users with administrator privileges on the host	vCenter (vpxuser)  Exception users, based on permissions  vCloud Director (vslauser, if available)	vCenter (vpxuser)  Exception users, based on permissions  vCloud Director (vslauser, if available)
DCUI			

ESXi Shell  
(if enabled)

---

SSH (if  
enabled)

---

## Chapter 13

**Table 13-4** Lifecycle Manager Definitions

---

**Term Definition**

**r**  
**m**

---

A software release that makes small changes to the current version, such as vSphere 7.0 Update 1, 7.0 Update 2, and so on.

---

A software release that introduces major changes to the software. For example, you can upgrade from vSphere 6.5 to 6.7 and 7.0.

---

A small software update that provides bug fixes or enhancements to the current version of the software, such as 7.0a, 7.0 Update 1a, and so on.

---

The smallest installable software package (metadata and binary payload) for ESXi.

---

An XML file that describes the contents of the VIB, including dependency information, textual descriptions, system requirements, and information

---

about bulletins.

---

A VIB that is not included in a component.

---

The hosted version of updates provided by VMware, OEMs, and third-party software vendors, containing the metadata and the actual VIBs.

---

An archive (ZIP file) that contains VIBs and metadata that you use for offline patching and updates. A single offline bundle might contain multiple base images, vendor add-ons, or components.

---

A VMware partner, such as Dell, HPE, or VMware Cloud on AWS.

---

---

A provider of I/O filters, device drivers, CIM modules, and so on.

---

---

**Table 13-8** Collection Intervals

<b>Collection Interval (Archive Length)</b>	<b>Collection Frequency</b>	<b>Default Behavior</b>
1 day		
1 week		



---

1 month

---

1 year

---

## Appendix C

# Memory Tables Answer Key

## Chapter 1

**Table 1-4** Available vSphere Features

<b>Available vSphere Features</b>	<b>Description</b>
vCenter Appliance File-Based Backup and Restore	A feature introduced in vSphere 7.0 that enables you to back up and restore the vCenter Server Appliance instances.
vMotion	A feature that provides live virtual machine migrations with negligible disruption from a source ESXi host to a target ESXi host.
<i>vSphere HA</i>	A feature that provides automated failover protection for VMs against host, hardware, network, and guest OS issues. In the event of host system failure, it performs cold migrations and restarts failed VMs on surviving hosts.
<i>Distributed Resource Scheduler (DRS)</i>	A feature that places and starts VMs on appropriate ESXi hosts and hot-migrates VMs using vMotion when there is contention for compute resources.
Storage vMotion	A feature that performs live migrations with negligible disruption of VMs from a source datastore to a target datastore.
Fault Tolerance	A feature that provides automated live failover protection for VMs against host, hardware, network, and guest OS issues.

(FT)

Distributed Power Management (DPM)	A feature that optimizes power consumption in an ESXi cluster.
<b>Proactive HA</b>	A feature that minimizes VM downtime by proactively detecting hardware failures and placing the host in Quarantine Mode or Maintenance Mode.
Content library	A centralized repository used to manage and distribute templates, ISO files, scripts, vApps, and other files associated with VMs.
Host profiles	A feature that provides a means to apply a standard configuration to a set of ESXi hosts.

**Table 1-6** vCenter Server Editions

Feature	Essentials	Essentials Plus	Foundation	Standard
Number of ESXi hosts	3 (2 CPU max)	3 (2 CPU max)	4	2000
vCenter License	Packaged with vSphere license in Essentials	Packaged with vSphere license in Essentials Plus	Sold separately from vSphere license	Sold separately from vSphere license
Basic vCenter features, like single pane of glass management, Lifecycle Manager, and VMware Converter	Supported	Supported	Supported	Supported
Common vCenter features like vMotion, vSphere HA, and vSphere Replication	Not supported	Supported	Supported	Supported
Advanced features like	N/A	N/A	N/A	Support

vCenter Server High Availability (VCHA) and vCenter Server Backup and Restore

ed

**Table 1-10** Compute Specifications for vCenter Server Appliance

<b>Component</b>	<b>Number of CPUs</b>	<b>Memory</b>
Tiny Environment	2	12 GB
Up to 10 hosts or 100 virtual machines		
Small Environment	4	19 GB
Up to 100 hosts or 1000 virtual machines		
Medium Environment	8	28 GB
Up to 400 hosts or 4000 virtual machines		
Large Environment	16	37 GB
Up to 1000 hosts or 10,000 virtual machines		
X-Large Environment	24	56 GB
Up to 2000 hosts or 35,000 virtual machines		

**Table 1-11** Storage Sizes for vCenter Server Appliance

<b>Deployment Size</b>	<b>Default Storage Size</b>	<b>Large Storage Size</b>	<b>X-Large Storage Size</b>
Tiny	415 GB	1490 GB	3245 GB
Small	480 GB	1535 GB	3295 GB
Medium	700 GB	1700 GB	3460 GB
Large	1065 GB	1765 GB	3525 GB

X-Large                      1805 GB                      1905 GB                      3665 GB

---

**Table 1-12** Required Ports for vCenter Sever

<b>Proto</b>	<b>Description</b>	<b>Required for</b>
TC P 22	System port for SSHD	vCenter Server (Must be open for upgrade of the appliance.)
TC P 80	Port for direct HTTP connections; redirects requests to HTTPS port 443	vCenter Server
TC P 88	Required to be open to join Active Directory	vCenter Server
TC P/ UD P 389	LDAP port for directory services for the vCenter Server group	vCenter Server to vCenter Server
TC P 443	Default port used by vCenter Server to listen for connections from the vSphere Web Client and SDK clients	vCenter Server to vCenter Server
TC P/ UD P 514	vSphere Syslog Collector port for vCenter Server and vSphere Syslog Service port for vCenter Server Appliance	vCenter Server
TC P/ UD P 902	Default port that the vCenter Server system uses to send data to managed hosts	vCenter Server

TC P 151 4	vSphere Syslog Collector TLS port for vCenter Server	vCenter Server
TC P 201 2	Control interface RPC for Single Sign-On	vCenter Server
TC P 201 4	RPC port for VMware Certificate Authority (VMCA) APIs	VMCA
TC P/ UD P 202 0	Authentication framework management	vCenter Server
TC P 548 0	vCenter Server Appliance Management Interface (VAMI)	vCenter Server
TC P/ UD P 650 0	ESXi Dump Collector port	vCenter Server
TC P 708 0, 127 21	Secure Token Service (internal ports)	vCenter Server
TC P	vSphere Client (internal ports)	vCenter Server

708 1	TC P	VMware vSphere Authentication Proxy	vCenter Server
747 5, 747 6	TC P	vSphere Lifecycle Manager SOAP port used by vSphere Lifecycle Manager client plug-in	vSphere Lifecycle Manager
908 4	TC P	vSphere Lifecycle Manager Web Server Port used by ESXi hosts to access host patch files from vSphere Lifecycle Manager server	vSphere Lifecycle Manager
908 7	TC P	vSphere Lifecycle Manager Web SSL port used by vSphere Lifecycle Manager client plug-in for uploading host upgrade files to vSphere Lifecycle Manager server	vSphere Lifecycle Manager
944 3	TC P	vSphere Web Client HTTPS	vCenter Server

## Chapter 2

**Table 2-2** Comparison of VMFS Version 5 and Version 6

<b>VMFS Features and Functionalities</b>	<b>Version 5</b>	<b>Version 6</b>
Access for ESXi hosts Version 6.5 and later	Yes	Yes
Access for ESXi hosts Version 6.0 and earlier	Yes	No

Datastores per host	512	512
512n storage devices	Yes	Yes (default)
512e storage devices	Yes (Not supported on local 512e devices.)	Yes (default)
4Kn storage devices	No	Yes
Automatic space reclamation	No	Yes
Manual space reclamation through the <b>esxcli</b> command.	Yes	Yes
Space reclamation from guest OS	Limited	Yes
GPT storage device partitioning	Yes	Yes
MBR storage device partitioning	Yes	No
	For a VMFS5 datastore that has been previously upgraded from VMFS3.	
Storage devices greater than 2 TB for each VMFS extent	Yes	Yes
Support for virtual machines with large-capacity virtual disks, or disks greater than 2 TB	Yes	Yes
Support of small files (1 KB)	Yes	Yes
Default use of ATS-only locking mechanisms on storage devices that support ATS	Yes	Yes
Block size	Standard 1 MB	Standard 1 MB
Default snapshots	VMFSsparse for virtual disks	SEs



	smaller than 2 TB	parse
	SEsparse for virtual disks larger than 2 TB	
Virtual disk emulation type	512n	512n
vMotion	Yes	Yes
Storage vMotion across different datastore types	Yes	Yes
High Availability and Fault Tolerance	Yes	Yes
DRS and Storage DRS	Yes	Yes
RDM	Yes	Yes

**Table 2-4** Comparison of NFS Version 3 and Version 4.1 Support for vSphere Features and Solutions

<b>NFS Features and Functionalities</b>	<b>Ver 3</b>	<b>Version 4.1</b>
vMotion and Storage vMotion	Yes	Yes
High Availability (HA)	Yes	Yes
Fault Tolerance (FT)	Yes	Yes (Supports the new FT mechanism introduced in vSphere 6.0 that supports up to four vCPUs, not the legacy FT mechanism.)
Distributed Resource Scheduler (DRS)	Yes	Yes
Host Profiles	Yes	Yes

	s	
Storage DRS	Y	No
	e	
	s	
Storage I/O Control	Y	No
	e	
	s	
Site Recovery Manager	Y	No
	e	
	s	
Virtual Volumes	Y	Yes
	e	
	s	
vSphere Replication	Y	Yes
	e	
	s	
vRealize Operations Manager	Y	Yes
	e	
	s	

**Table 2-7 RAID Configuration Comparison**

<b>RAID Configuration</b>	<b>PF TT</b>	<b>Data Size</b>	<b>Required Capacity</b>	<b>Usable Capacity</b>
RAID 1 (mirroring)	1	100 GB	200 GB	50%
RAID 5 or RAID 6 (erasure coding) with four fault domains	1	100 GB	133 GB	75%
RAID 1 (mirroring)	2	100 GB	300 GB	33%
RAID 5 or RAID 6 (erasure coding) with six fault domains	2	100 GB	150 GB	67%
RAID 1 (mirroring)	3	100 GB	400 GB	25%

RAID 5 or RAID 6 (erasure coding) with six fault domains	3	N/A	N/A	N/A
--	---	-----	-----	-----

**Table 2-12 vSAN Storage Policies**

**Policy Description**

y

**Primary Level of Failures to Tolerate (PFTT)**  
 This policy defines how many host and device failures a VM object can withstand. For  $n$  failures tolerated, data is stored in  $n+1$  location. (This includes parity copies with RAID 5 or 6.) If no storage policy is selected at the time of provisioning a VM, this policy is assigned by default. Where fault domains are used,  $2n+1$  fault domains, each with hosts adding to the capacity, are required. If an ESXi host isn't in a fault domain, it is considered to be in a single-host fault domain. The default setting for this policy is 1, and the maximum is 3.

**Secondary Level of Failures to Tolerate (SFSTT)**  
 In stretched clusters, this policy defines how many additional host failures can be tolerated after a site failure's PFTT has been reached. If PFTT = 1, SFSTT = 2, and one site is inaccessible, two more host failures can be tolerated. The default setting for this policy is 1, and the maximum is 3.

**Data Locality**  
 If PFTT = 0, this option is available. The options for this policy are None, Preferred, and Secondary. This allows objects to be limited to one site or one host in stretched clusters. The default setting for this

ality policy is None.

---

Failure Tolerance Method  
This policy defines whether the data replication mechanism is optimized for performance or capacity. If RAID-1 (Mirroring)—Performance is selected, there will be more space consumed in the object placement but better performance for accessing the space. If RAID-5/6 (Erasure Coding)—Capacity is selected, there will be less disk utilization, but performance will be reduced.

---

Number of Disk Stripes per Object  
This policy determines the number of capacity devices where each VM object replica is striped. Setting this above 1 can improve performance but consumes more resources. The default setting for this policy is 1, and the maximum is 12.

---

Flash Read Cache Reservation  
This policy defines the amount of flash capacity that is reserved for read caching of VM objects. This is defined as a percentage of the size of the VMDK. This is supported only in hybrid vSAN clusters. The default setting for this policy is 0%, and the maximum is 100%.

---

Force Provisioning  
If set to yes, this policy forces provisioning of objects, even when policies cannot be met. The default setting for this policy is no.

---

Object Space Reservation

This policy defines the percentage of VMDK objects that must be thick provisioned on deployment. The options are as follows:

Object Space Reservation

- Thin provisioning (default value)
- 25% reservation
- 50% reservation
- 75% reservation
- Thick provisioning

---

Disable Object Checksum

A checksum is used end-to-end in validating the integrity of the data to ensure that data copies are the same as the original. In the event of a mismatch, incorrect data is overwritten. If this policy is set to yes, a checksum is not calculated. The default setting for this policy is no.

Disable Object Checksum

---

IOPS Limit for Object

This policy sets a limit for IOPS of an object. If set to 0, there is no limit.

IOPS Limit for Object

## Chapter 3

**Table 3-2** Advantages and Disadvantages of IP Hash NIC Teaming

Advantages	Disadvantages
A more even distribution of the load compared to Route Based on Originating Virtual Port and Route Based on Source MAC Hash	Highest resource consumption compared to the other load-balancing algorithms
A potentially higher throughput for virtual machines that communicate with multiple IP addresses	Requires changes on the physical network.
	Complex to troubleshoot

**Table 3-3** Comparison of vSS and vDS Features

Feature	vSS	vDS
Layer 2 switch	X	X
VLAN segmentation (802.1q tagging)	X	X
IPv6 support	X	X
NIC teaming	X	X
Outbound traffic shaping	X	X
Cisco Discovery Protocol (CDP)	X	X
Inbound traffic shaping		X
VM network port block		X
Private VLANs		X
Load-based NIC teaming		X
Data center-level management		X
Network vMotion		X
Per-port policy settings		X
Port state monitoring		X

NetFlow	X
Port mirroring	X

**Table 3-4** vDS Health Checks

Health Check	Required vDS Configuration
Checks whether the VLAN trunk ranges on the distributed switch match the trunk port configuration on the connected physical switch ports.	At least two active physical NICs
Checks for matching MTU settings on the distributed switch, the physical network adapter, and the physical switch ports.	At least two active physical NICs
Checks whether the virtual switch teaming policy matches the physical switch port-channel settings.	At least two active physical NICs and two hosts

**Table 3-5** SR-IOV Requirements

Component	Requirements
Physical host	<p>Must use an Intel or AMD processor.</p> <p>Must support IOMMU and SR-IOV.</p> <p>IOMMU and SR-IOV must be enabled in the BIOS.</p>
Physical network adapter	<p>Must be supported by the server vendor for use with the host system and SR-IOV for the specific ESXi release.</p> <p>SR-IOV must be enabled in the firmware.</p> <p>Must use MSI-X interrupts.</p>
Physical function (PF) driver in	Must be certified by VMware.

ESXi	Must be installed on the ESXi host, which may require custom installation.
Guest OS	Must be supported by the NIC vendor for the specific ESXi release.
Virtual function (VF) driver in guest OS	Must be compatible with the NIC and supported on the guest OS release.  Must be Microsoft WLK or WHCK certified for Windows virtual machines.  Must be installed on the operating system and may require custom installation.

## Chapter 4

**Table 4-4** Resource Pool Use Cases

Use Case	Details
Flexible hierarchical organization	Add, remove, modify, and reorganize resource pools, as needed.
Resource isolation	Use resource pools to allocate resources to separate departments, in such a manner that changes in a pool do not unfairly impact other departments.
Access control and delegation	Use permissions to delegate activities, such as virtual machine creation and management, to other administrators.
Separation of resources from hardware	In a DRS cluster, perform resource management independently of the actual hosts.
Managing multitier applications.	Manage the resources for a group of virtual machines (in a specific resource pool), which is easier than managing resources per virtual machine.



**Table 4-6** Virtual Machine Shares

Setting	CPU Share Value	Memory Share Value
High	2000 per vCPU	20 per MB
Normal	1000 per vCPU	10 per MB
Low	500 per vCPU	5 per MB

**Table 4-9** Advanced vSphere HA Options

Option	Description
das.isolationaddressX	Provides the addresses to use to test for host isolation when no heartbeats are received from other hosts in the cluster. If this option is not specified (which is the default setting), the management network default gateway is used to test for isolation. To specify multiple addresses, you can set das.isolationaddressX, where X is a number between 0 and 9.
das.usedefaultisolationaddress	Specifies whether to use the default gateway IP address for isolation tests.
das.isolationshutdowntimeout	For scenarios where the host's isolation response is to shut down, specifies the period of time that the virtual machine is permitted to shut down before the system powers it off.
das.slotmemoryminimum	Defines the maximum bound on the memory slot size.
das.slotcpuinhibit	Defines the maximum bound on the CPU slot size.

das.v mme mory minm b	Defines the default memory resource value assigned to a virtual machine whose memory reservation is not specified or is zero. This is used for the Host Failures Cluster Tolerates admission control policy.
das.v mcpu minm hz	Defines the default CPU resource value assigned to a virtual machine whose CPU reservation is not specified or is zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default of 32 MHz is used.
das.h earb eatds perho st	Specifies the number of heartbeat datastores required per host. The default is 2. The acceptable values are 2 to 5.
das.c onfig .fdm. isolat ionPo licyD elayS ec	Specifies the number of seconds the system delays before executing the isolation policy after determining that a host is isolated. The minimum is 30. A lower value results in a 30-second delay.
das.r espec tvmv manti affini tyrule s	Determines whether vSphere HA should enforce VM-VM anti-affinity rules even when DRS is not enabled.

**Table 4-10** VM Monitoring Settings

Setting	Failure Interval	Reset Period
High	30 seconds	1 hour
Medium	60 seconds	24 hours
Low	120 seconds	7 days

## Chapter 5

**Table 5-2** Virtual Machine Files

<b>File</b>	<b>Description</b>
vmname.vmx	Virtual machine configuration file
vmname.vmxv	Additional virtual machine configuration file
vmname.vmdk	Virtual disk characteristics (metadata) file
vmname-flat.vmdk	Virtual disk data file (commonly called a flat file)
vmname.nvram or nvram	Virtual machine BIOS or UEFI configuration file
vmname.vmsd	Virtual machine snapshot file
vmname.vmsn	Virtual machine snapshot data file
vmname.vswp	Virtual machine swap file
vmname.vms	Virtual machine suspend file
vmware.log	Current virtual machine log file
vmware-#.log	Old virtual machine log file, where # is a number starting with 1

**Table 5-4** Virtual Machine Options

<b>Category</b>	<b>Description</b>
General Options	Settings include virtual machine name, configuration file location, and the working directory location.
Encryption Options	Settings allow you to enable or disable virtual machine encryption or vMotion encryption.
Power Management	Settings allow you to choose how to respond when the guest OS is placed on standby. The choices are to suspend the virtual machine or put the guest OS into standby mode.
VMware Tools	Settings allow you to choose how to respond to specific power operations. For example, you can choose whether to power off the virtual machine or shut down the guest when the red power-

off button is clicked.

---

Virtualization Based Security (VBS)	For virtual machines running the modern Windows OS versions, you can enable VBS to add an extra level of protection.
Boot Options	Settings include firmware, boot delay, and failed boot recovery parameters.
Advanced Options	Settings include logging, debugging, swap file location, and configuration parameters.
Fibre Channel NPIV	Settings allow the virtual machine to use N_Port ID Virtualization (NPIV), including whether to generate new worldwide names (WWNs).
vApp Options	Settings allow you to control vApp functionality for the virtual machine, such as enable/disable and IP allocation policy. vApp settings that are made directly to a virtual machine override settings made on the vApp.

---

## Chapter 6

**Table 6-2** Required Permissions for the vCenter Cloud Account

---

<b>Object</b>	<b>Permissions</b>
Datastore	<ul style="list-style-type: none"><li>■ Allocate space</li><li>■ Browse datastore</li><li>■ Low level file operations</li></ul>
Datastore cluster	<ul style="list-style-type: none"><li>■ Configure a datastore cluster</li></ul>

---

---

## Folder

- Create folder
- Delete folder

---

## Global

- Manage custom attributes
- Set custom attribute

---

## Network

- Assign network

---

## Permissions

- Modify permission

---

## Resource

- Assign VM to resource pool
- Migrate powered-off virtual machine
- Migrate powered-on virtual machine

---

## Content library

- Add library item

- Create local library
- Create subscribed library
- Delete library item
- Delete local library
- Delete subscribed library
- Download files
- Evict library item
- Evict subscribed library
- Probe subscription information
- Read storage
- Sync library item
- Sync subscribed library
- Type introspection

- Update configuration settings
- Update files
- Update library
- Update library item
- Update local library
- Update subscribed library
- View configuration settings

---

## Tags

- Assign or unassign vSphere tag
- Create a vSphere tag
- Create a vSphere tag category
- Delete vSphere tag
- Delete vSphere tag category
- Edit vSphere tag

- Edit vSphere tag category
- Modify UsedBy field for category
- Modify UsedBy field for tag

---

## vApp

- Import
- vApp application configuration.

---

## Virtual machine inventory

- Create from existing
- Create new
- Move
- Remove

---

## Virtual machine interaction

- Configure CD media
- Console interaction
- Device connection



- Power off
- Power on
- Reset
- Suspend
- Tools install

---

## Virtual machine configuration

- Add existing disk
- Add new disk
- Add or remove
- Remove disk
- Advanced
- Change CPU count
- Change resource
- Extend virtual disk

- Disk change tracking
- Memory
- Modify device settings
- Rename
- Set annotation
- Settings
- Swapfile placement

---

## Virtual machine provisioning

- Customize
- Clone template
- Clone virtual machine
- Deploy template
- Read customization specs

Virtual machine state

- Create snapshot
- Remove snapshot
- Revert to snapshot

---

**Table 6-3** Required vCenter Server Privileges for Horizon (without instant clones)

<b>Privilege Group</b>	<b>Privileges to Enable</b>
Folder	<ul style="list-style-type: none"><li>■ Create Folder</li><li>■ Delete Folder</li></ul>
Datastore	<ul style="list-style-type: none"><li>■ Allocate space</li></ul>
Virtual Machine	<p>In Configuration:</p> <ul style="list-style-type: none"><li>■ Add or remove device</li><li>■ Advanced</li><li>■ Modify device settings</li></ul>

### In Interaction:

- Power off
- Power on
- Reset
- Suspend
- Perform wipe or shrink operations

### In Inventory:

- Create new
- Create from existing
- Remove

### In Provisioning:

- Customize
- Deploy template
- Read customization specifications
- Clone template
- Clone virtual machine

Resource	Assign virtual machine to resource pool
Global	Act as vCenter Server
Host (for Storage Accelerator)	:
	Advanced settings (in Configuration)
Profile Driven Storage (for vSAN or Virtual Volumes)	All privileges

**Table 6-5** VMware HCX Services

Service	Licenses	Description
Interconnect	Admin	Creates secured connections between HCX instances, supporting migration, replication, disaster recovery, and management operations.

	c e d	Deployed as a virtual appliance.
WAN Optimization	A d v a n c e d	Optimizes the performance of the connection provided by HCX Interconnect through a combination of deduplication, compression, and line conditioning techniques.  Deployed as a virtual appliance.
Network Extension	A d v a n c e d	Extends (that is, provides Layer 2 adjacency) the virtual machine networks between source and remote HCX-enabled environments.  Deployed as a virtual appliance.
Bulk Migration	A d v a n c e d	Migrates a set of virtual machines using VMware vSphere Replication in parallel between HCX-enabled sites.
vMotion Migration	A d v a n c e d	Migrates a single virtual machine between HCX-enabled sites with no service interruption, using vMotion.
Disaster Recovery	A d	Protects virtual machines from disaster by using replication and recovery.

v  
a  
n  
c  
e  
d

---

Mobility Groups	E n t e r p r i s e	Allows you to group virtual machines by application, network, or other aspects for migration and monitoring.
OS Assisted Migration	E n t e r p r i s e	Leverages HCX Sentinel software in the guest OS to migrate Windows and Linux virtual machines to a vSphere-enabled data center.  Uses a gateway appliance at the source and a receiver appliance at the destination.
Replication Assisted vMotion (RAV)	E n t e r p r i s e	Migrates a set of virtual machines in parallel, using VMware vSphere Replication and vMotion between HCX-enabled sites with no service interruption.
Site Recovery Manager (SRM) Integration	E n t e r p r i s	Integrates HCX functionality with the VMware SRM for protection and orchestrated recovery operations.

	e	
Traffic Engineering: Application Path Resiliency and TCP Flow Conditioning	Enterprise	<p>Optimizes network traffic for HCX Interconnect and Network Extension services.</p> <p>The Application Path Resiliency service creates multiple tunnel flows for both Interconnect and Network Extension traffic.</p> <p>The TCP Flow Conditioning service adjusts and optimizes the segment size to reduce fragmentation and reduce the overall packet rate.</p>
Mobility Optimized Networking (MON)	Enterprise	<p>Integrates HCX Network Extension with NSX Dynamic Routing to enable optimal networking between migrated virtual machines and other virtual machines.</p> <p>Works with new or existing network extensions to NSX-T 3.0 Data Center.</p>

## Chapter 7

**Table 7-2** Core Identity Services in vSphere

Service	Description
VMware Directory Service (vmdir)	Serves as an identity source that handles SAML certificate management for authentication with vCenter Single Sign-On.
VMware Certificate Authority (VMCA)	Issues certificates for VMware solution users, machine certificates for machines on which services are running, and ESXi host certificates. VMCA can be used as is, or it can be used as an intermediary certificate authority.
VMware	Includes VMware Endpoint Certificate Store (VECS) and



Authentication Framework Daemon (VMAFD) several internal authentication services.

**Table 7-6** Certificates in vSphere

Certificate	Provisioned Details	
ESXi certificate	VM CA (default)	Stored locally on an ESXi host in the /etc/vmware/ssl directory when the host is first added to vCenter Server and when it reconnects.
Machine SSL certificate	VM CA (default)	<p>Stored in VECS.</p> <p>Used to create SSL sockets for SSL client connections, for server verification, and for secure communication such as HTTPS and LDAPS.</p> <p>Used by the reverse proxy service, the vCenter Server service (vpxd), and the VMware Directory service (vmdir).</p> <p>Uses X.509 Version 3 certificates to encrypt session information.</p>
Solution user certificate	VM CA (default)	<p>Stored in VECS.</p> <p>Used by solution users to authenticate to vCenter Single Sign-On through SAML token exchange.</p>
vCenter Single Sign-On SSL signing certificate	During installation	Used throughout vSphere for authentication, where a SAML token represents the user's identity and contains group membership information.

You can manage this certificate from the command line. Changing this certificate in the file system leads to unpredictable behavior.

VMware Directory Service (vmdir) SSL certificate	During installation	Starting with vSphere 6.5, the machine SSL certificate is used as the vmdir certificate.
vSphere Virtual Machine Encryption Certificates	Depends	Used for virtual machine encryption, which relies on an external key management server (KMS).  Depending on how the solution authenticates to the KMS, it might generate certificates and store them in VECS.

**Table 7-9** System Roles in vCenter Server 7.0

System Role	Description
Read-only	Allows the user to view the state of an object and details about the object. For example, users with this role can view virtual machine attributes but cannot open the VM console.
Administrator	Includes all privileges of the read-only role and allows the user to view and perform all actions on the object. If you have the administrator role on an object, you can assign privileges to individual users and groups. If you have the administrator role in vCenter Server, you can assign privileges to users and groups in the default SSO identity source. By default, the <a href="mailto:administrator@vsphere.local">administrator@vsphere.local</a> user has the administrator role on both vCenter Single Sign-On and vCenter Server.
No access	Prevents users from viewing or interacting with the object. New users and groups are effectively assigned this role by default.
No cryptography	Includes all privileges of the administrator role, except for cryptographic operations privileges. This role allows administrators to

ptography administrator designate users who can perform all administrative tasks except encrypting or decrypting virtual machines or accessing encrypted data.

---

Trusted Infrastructure Administrator Allows users to perform VMware vSphere Trust Authority operations on some objects. Membership in the TrustedAdmins group is required for full vSphere Trust Authority capabilities.

---

**Table 7-10** Required Permissions for Common Tasks

Task	Required Privileges
Create a virtual machine	On the destination folder or in the data center: <ul style="list-style-type: none"> <li data-bbox="706 1549 1347 1591">■ Virtual Machine.Inventory.Create New</li> <li data-bbox="706 1675 1380 1759">■ Virtual Machine.Configuration.Add New Disk</li> <li data-bbox="706 1843 1307 1883">■ Virtual Machine .Configuration.Add</li> </ul>

Existing Disk

- Virtual Machine.Configuration.Raw Device

On the destination host or cluster or in the resource pool:

- Resource.Assign Virtual Machine to Resource Pool

On the destination datastore or in the datastore folder:

- Datastore.Allocate Space

On the network:

- Network.Assign Network

---

Deploy a virtual machine from a template

On the destination folder or in the data center:

- Virtual Machine.Inventory.Create from Existing

- Virtual Machine.Configuration.Add New Disk

On a template or in a template folder:

- Virtual Machine.Provisioning.Deploy Template

On the destination host or cluster or in the resource pool:

- Resource.Assign Virtual Machine to Resource Pool

On the destination datastore or in a datastore folder:

- Datastore.Allocate Space

On the network that the virtual machine will be assigned to:

- Network.Assign Network

---

Take a virtual machine snapshot

On the virtual machine or in a virtual machine folder:

- Virtual Machine.Snapshot Management.Create Snapshot

On the destination datastore or in a datastore folder:

- Datastore.Allocate Space

---

Move a virtual machine into a resource pool

On the virtual machine or in a virtual machine folder:

- Resource.Assign Virtual Machine to Resource Pool
- Virtual Machine.Inventory.Move

In the destination resource pool:

- Resource.Assign Virtual Machine to

## Resource Pool

---

Install a guest operating system on a virtual machine

On the virtual machine or in a virtual machine folder:

- Virtual Machine.Interaction.Answer Question
- Virtual Machine.Interaction.Console Interaction
- Virtual Machine.Interaction.Device Connection
- Virtual Machine.Interaction.Power Off
- Virtual Machine.Interaction.Power On
- Virtual Machine.Interaction.Reset
- Virtual Machine.Interaction.Configure CD Media
- Virtual Machine.Interaction.Configure Floppy Media
- Virtual Machine.Interaction.Tools Install

On a datastore containing the installation media ISO image:

- Datastore.Browse Datastore

On the datastore to which you upload the installation media ISO image:

- Datastore.Browse Datastore
- Datastore.Low Level File Operations

---

Migrate a virtual machine with vMotion

On the virtual machine or in a virtual machine folder:

- Resource.Migrate Powered on Virtual Machine
- Resource.Assign Virtual Machine to Resource Pool

On the destination host or cluster or in a resource pool:



- Resource.Assign Virtual Machine to Resource Pool

---

Cold migrate (relocate) a virtual machine

On the virtual machine or in a virtual machine folder:

- Resource.Migrate Powered Off Virtual Machine
- Resource.Assign Virtual Machine to Resource Pool

On the destination host or cluster or in a resource pool:

- Resource.Assign Virtual Machine to Resource Pool

On the destination datastore:

- Datastore.Allocate Space

---

Migrate a virtual machine with Storage vMotion

On the virtual machine or in a virtual machine folder:

- Resource.Migrate Powered On Virtual Machine

On the destination datastore:

- Datastore.Allocate Space

Move a host into a cluster

On the host:

- Host.Inventory.Add Host to Cluster

On the destination cluster:

- Host.Inventory.Add Host to Cluster
- Host.Inventory.Modify. cluster

**Table 7-11** ESXi Security Profile Services

Service	Default State	Description
Direct Console User Interface (DCUI)	Running	Allows you to interact with an ESXi host from the local console host using text-based menus
ESXi Shell	Stopped	Is available from the DCUI or from SSH

SSH	Stopped	Allows remote connections through Secure Shell
Load-Based Teaming Daemon	Running	Enables load-based teaming
attestd	Stopped	Enables the vSphere Trust Authority Attestation Service
kmxsd	Stopped	Enables the vSphere Trust Authority Key Provider Service
Active Directory Service	Stopped	Is started on hosts after you configure ESXi for Active Directory
NTP Daemon	Stopped	Enables the Network Time Protocol daemon
PC/SC Smart Card Daemon	Stopped	Is started on hosts after you enable the host for smart card authentication
CIM Server	Running	Can be used by Common Information Model (CIM) applications
SNMP Server	Stopped	Enables the SNMP daemon
Syslog Server	Stopped	Enables the syslog daemon
VMware vCenter Agent (vpxa)	Running	Connects the host to vCenter Server
X.Org Server	Stopped	Internally used for virtual machine 3D graphics

**Table 7-12** Incoming and Outgoing Firewall Ports

<b>Firewall Service</b>	<b>Incoming Port(s)</b>	<b>Outgoing Port(s)</b>
CIM Server	5988 (TCP)	
CIM Secure Server	5989 (TCP)	
CIM SLP	427 (TCP,UDP)	427 (TCP,UDP)
DHCPv6	546 (TCP,UDP)	547 (TCP,UDP)

DVSSync	8301, 8302 (UDP)	8301, 8302 (UDP)
HBR		44046, 31031 (TCP)
NFC	902 (TCP)	902 (TCP)
WOL		9 (UDP)
vSAN Clustering	12345, 23451 (UDP)	12345, 23451 (UDP)
DCHP Client	68 (UDP)	68 (UDP)
DNS Client	53 (UDP)	53 (TCP,UDP)
Fault Tolerance	8100, 8200, 8300 (TCP,UDP)	80, 8100, 8200, 8300 (TCP,UDP)
NSX Distributed Logical Router Service	6999 (UDP)	6999 (UDP)
Software iSCSI Client		3260 (TCP)
rabbitmqproxy		5671 (TCP)
vSAN Transport	2233 (TCP)	2233 (TCP)
SNMP Server	161 (UDP)	
SSH Server	22 (TCP)	
vMotion	8000 (TCP)	8000 (TCP)
VMware vCenter Agent		902 (UDP)
vSphere Web Access	80 (TCP)	
vsanvp	8080 (TCP)	8080 (TCP)
RFB Protocol	5900–5964 (TCP)	
vSphere Life Cycle Manager	80, 9000 (TCP)	80, 9000 (TCP)
I/O Filter	9080 (TCP)	

**Table 7-14** Network Security Policies

**Opt Sett Description  
ion ing**

Pr A The virtual switch forwards all frames to the virtual network

o c adapter.

m c  
is e  
c p  
u t

---

o R The virtual switch forwards only the frames that are addressed to  
us ej the virtual network adapter.  
M e  
o ct  
d  
e

---

M A If the guest operating system changes the effective MAC address of  
A c the virtual adapter to a value that differs from the MAC address  
C c assigned to the adapter in the VMX file, the virtual switch allows  
A e the inbound frame to pass.  
d p  
dr t

---

es R If the guest operating system changes the effective MAC address of  
s ej the virtual adapter to a value that differs from the MAC address  
C e assigned to the adapter in the VMX file, the virtual switch drops all  
h ct inbound frames to the adapter. If the guest OS changes the MAC  
a address back to its original value, the virtual switch stops dropping  
n the frames and allows inbound traffic to the adapter.  
g  
es

---

F A The virtual switch does not filter outbound frames. It permits all  
or c outbound frames, regardless of the source MAC address.  
g c  
e e  
d p  
T t

---

ra R The virtual switch drops any outbound frame from a virtual  
ns ej machine virtual adapter that uses a source MAC address that differs  
m e from the MAC address assigned to the virtual adapter in the VMX  
it ct file.  
s

---

## Chapter 8

**Table 8-2** Information Required for ESXi Installation

<b>Information</b>	<b>Required or Optional</b>	<b>Details</b>
Keyboard layout	Required	Default: US English
VLAN ID	Optional	Range: 0–4094
		Default: None
IP address	Optional	Default: DHCP
Subnet mask	Optional	Default: Based on the configured IP address
Gateway	Optional	Default: Based on the configured IP address and subnet mask
Primary DNS	Optional	Default: Based on the configured IP address and subnet mask
Secondary DNS	Optional	Default: None
Host name	Required for static IP settings	Default: None
Install location	Required	At least 5 GB if you install on a single disk
		Default: None
Migrate existing ESXi settings; preserve VMFS datastore	Required if you are installing ESXi on a drive with an existing ESXi installation	Default: None
Root password	Required	Must contain at least 8 to 40 characters and meet other requirements

Default: None

**Table 8-5** Auto Deploy Components

Component	Description/Purpose
Auto Deploy server	Uses a rules engine, a set of images, a set of host profiles, and required infrastructure to manage ESXi deployments.
Rules engine	Assigns image profiles and host profiles to each host.
Host profile	Defines host-specific configurations, such as networking, NTP, and host permissions. You can use host customization in conjunction with host profiles to provide details that are unique to each host, such as IP address.
Auto Deploy PowerCLI	Servers as a command-line engine for driving Auto Deploy.
Image Builder PowerCLI	Servers as a command-line engine for building images.
vCenter Server	Manages the vSphere inventory and provides host profiles.
DHCP	Provides IP configuration to the host and redirects the host to the PXE server.

server

---

PXE server	Boots the host and directs it to the TFTP server.
------------	---

server

---

TFTP server	Provides the appropriate boot image.
-------------	--------------------------------------

server

---

Software depot	Holds a collection of VIBs either online (accessible via HTTP) or offline (accessible via a USB drive or CD/DVD).
----------------	---

---

Image profile	Holds a collection of VIBs used to install the ESXi server and saved as ZIP files or ISO images. You can obtain image profiles from VMware and VMware partners, and you can create custom image profiles by using ESXi Image Builder.
---------------	---

---

vSphere Installation Bundle (VIB)	Packages a collection of files (such as drivers) into an archive similar to a ZIP file. Each VIB is released with an acceptance level that cannot be changed. The host acceptance level assigned to each host determines which VIBs can be installed to the host. These are the acceptance levels, from highest to lowest:
-----------------------------------	--

(VIB)

- VMwareCertified
- VMwareAccepted
- PartnerSupported
- CommunitySupported

---

**Table 8-9** VECS Stores

---



Store	Description
Machine SSL store (MACHINE_SSL_CERT)	Used by the reverse proxy service on each ESXi host and by the vmdir service.
Trusted root store (TRUSTED_ROOTS)	Contains all trusted root certificates.
Solution user stores:	VECS includes one store for each solution user.
<ul style="list-style-type: none"> <li>  Machine</li> <li>  vpxd</li> <li>  vpxd-extension</li> <li>  vsphere-webclient</li> </ul>	
vSphere Certificate Manager utility backup store (BACKUP_STORE)	Used by VMCA to support certificate reversion.
Other stores	Other stores might be added by solutions. For example, the Virtual Volumes solution adds an SMS store.

The SSO domain contains many predefined groups, including the following:

- **Users:** This group contains all users in the SSO domain.
- **DCAdmins:** Members of this group can perform domain controller administrator actions on VMware Directory Service.
- **SolutionUsers:** Each solution user authenticates individually to vCenter Single Sign-On with a certificate. By default, VMCA provisions solution users with certificates. Do not add members to this group explicitly.
- **CAAdmins:** Members have administrator privileges for VMCA. Adding

members to these groups is not usually recommended, but a user must be a member of this group to perform most certificate management operations, such as using the **certool** command.

- **SystemConfiguration.BashShellAdministrators:** Members can enable and disable access to the BASH Shell.
- **SystemConfiguration.Administrators:** Members can view and manage the system configuration and perform tasks such as restarting services.
- **LicenseSevice.Administrators:** Members have full write access to all licensing-related data and can add, remove, assign, and un-assign serial keys for all product assets registered in licensing service.
- **Administrators:** Members can perform SSO administration tasks for VMware Directory Service (vmdir).

**Table 8-10** SSO Policies and Parameters

SSO Policy Parameter	Policy Setting	Details
Password Policy	Description	Password policy description.
	Maximum lifetime	Maximum number of days a password can exist before the user must change it.
	Restrict reuse	Number of the user's previous passwords that cannot be selected.
	Maximum length	Maximum number of characters that are allowed in the password.
	Minimum length	Minimum number of characters that are allowed in the password, which must be no fewer than the combined minimum of alphabetic, numeric, and special character requirements.
	Character requirements	Minimum number of different character types that are required in the password. The types include special, alphabetic, uppercase, lowercase, and numeric.

	Identical adjacent characters	The number of identical adjacent characters that are supported in a password. The value must be greater than 0.
Lockout Policy	Description	Description of the lockout policy.
	Maximum number of failed login attempts	Maximum number of failed login attempts that are allowed before the account is locked.
	Time interval between failures	Time period in which failed login attempts must occur to trigger a lockout.
	Unlock time	The amount of time the account stays locked. The value 0 specifies that an administrator must explicitly unlock the account.
Token Policy	Clock tolerance	Time difference, in milliseconds, that SSO tolerates between a client clock and a domain controller clock. If the time difference is greater than the specified value, SSO declares the token to be invalid.
	Maximum token renewal count	Maximum number of times a token may be renewed before a new security token is required.
	Maximum token delegation count	Maximum number of times a single holder-of-key token can be delegated.
	Maximum bearer token lifetime	The lifetime value of a bearer token before the token must be reissued.
	Maximum	The lifetime value of a holder-of-key token before the

m holder- token is marked invalid.  
of-key  
token  
lifetime

---

**Table 8-11** ESXi 7.0 Kernel Options

---

<b>Kernel Option</b>	<b>Description</b>
<b>autoPartition=TRUE/FALSE (default FALSE)</b>	<p>This option, if set to TRUE, defines automatic partitioning of the unused local storage devices at boot time. The boot disk gets partitioned with boot bands, ESXi-OSData, and, if the disk is larger than 128 GB, a VMFS partition. Any new empty device discovered will be auto-partitioned as well. Auto-partitioning can be set for only the first unused device with the setting <b>autoPartitionOnlyOnceAndSkipSsd=TRUE</b>. On hosts with USB boot and VMFS-L, ESX-OSData does not exist on other local disks.</p> <p>If a storage device has both a scratch partition and a coredump partition, the scratch partition is converted to ESX-OSData; otherwise, the first unused disk identified is partitioned with ESX-OSData as well.</p>
<b>skipPartitioningSsds=TRUE/FALSE (default FALSE)</b>	<p>If this option is set to TRUE, local SSDs are excluded from automatic partitioning.</p>
<b>autoPartitionOnlyOnceAndSkipSsd=TRUE</b>	<p>If this option is set to TRUE, SSD/NVMe devices are excluded, and the ESXi host automatically partitions the first unused local disk if there is no VMFS-L ESX-OSData volume.</p>

**dSkipSsd=TRUE/FALSE (default)**  
**t**  
**FALSE**  
**E)**

---

**allowCoreDumpOnUSB=TRUE/FALSE (default)**  
**t**  
**FALSE**  
**E)**

If this option is set to TRUE, ESXi can write kernel crash core dumps to the VMFS-L Locker volume on a USB boot device.

---

**dumpSize (default:0 (automatically sized))**

This option sets the size of the core dump file (in megabytes) created on the system VMFS-L volume. This is limited to one-half of the space available on the VMFS-L volume.

---

**autoCreatedDumpFile=TRUE/FALSE (default)**  
**t**  
**TRUE**  
**)**

This option, when set to TRUE, automatically creates a core dump file. This is attempted in the following order:

- VMFS-L ESX-OSData
- USB VMFS-L

- Local VMFS

---

## Chapter 9

**Table 9-2** VLAN ID Details

VLAN ID	VLAN Tagging Mode	Description
0	External switch tagging (EST)	The virtual switch does not pass traffic associated with a VLAN.
1 to 4094	Virtual switch tagging (VST)	The virtual switch tags traffic with the entered tag.
4095	Virtual guest tagging (VGT)	Virtual machines handle VLANs. The virtual switch passes traffic from any VLAN.

Enhanced LACP support for vDS supports the following load-balancing modes (hashing algorithms):

- Destination IP address
- Destination IP address and TCP/UDP port
- Destination IP address and VLAN
- Destination IP address, TCP/UDP port, and VLAN
- Destination MAC address
- Destination TCP/UDP port
- Source IP address
- Source IP address and TCP/UDP port
- Source IP address and VLAN
- Source IP address, TCP/UDP port, and VLAN
- Source MAC address
- Source TCP/UDP port

- Source and destination IP address
- Source and destination IP address and TCP/UDP port
- Source and destination IP address and VLAN
- Source and destination IP address, TCP/UDP port, and VLAN
- Source and destination MAC address
- Source and destination TCP/UDP port
- Source port ID
- VLAN

## Chapter 10

**Table 10-4** Performance Chart Types

Chart Type	Description	Example
Line chart	Displays metrics for a single inventory object, where data for each metric is represented by a separate line.	For example, a network chart for a host can contain one line showing the number of packets received and another line showing the number of packets transmitted.
Bar chart	Displays metrics for objects, where each bar represents metrics for an object.	A bar chart can display metrics for datastores, where each datastore is represented as a bar. Each bar displays metrics based on the file type, such as virtual disk or snapshot.
Pie chart	Displays metrics for a single object, where each slice represents a	A datastore pie chart can display the amount of storage space occupied by each virtual machine or by each file type.

category or child object.

Displays metrics for child objects.

A host's stacked CPU usage chart displays metrics for the 10 virtual machines on the host that are consuming the most CPU. The Other amount displays the total CPU usage of the remaining virtual machines.

**Table 10-6** CPU Performance Analysis

Symptoms	Likely Causes	Potential Solutions
Host: CPU usage is consistently high.	The host has insufficient CPU resources to meet the demand.	Add the host to a DRS cluster.
Virtual machine: CPU usage is above 90%. CPU ready is above 20%. Application performance is poor.	Too many virtual CPUs are running on the host.  Storage or network operations are placing the CPU in a wait state.  The guest OS generates too much load for the CPU.	Increase the number of hosts in the DRS cluster.  Migrate one or more virtual machines to other hosts.  Upgrade the physical CPUs of the host.  Upgrade ESXi to



the latest version.

Enable CPU-saving features such as TCP segmentation offload, large memory pages, and jumbo frames.

Increase the amount of memory allocated to the virtual machines, which may improve cached I/O and reduce CPU utilization.

Reduce the number of virtual CPUs assigned to virtual machines.

Ensure that VMware Tools is installed.

Compare the CPU usage of troubled virtual machines with that of other virtual machines on the host or in the resource pool.

(Hint: Use a stacked graph.)

Increase the CPU limit, shares, or reservation on the troubled virtual machine.

---

Host: Memory usage is consistently 94% or higher. Free memory is 6% or less.

The host has insufficient memory resources to meet the demand.

Ensure that VMware Tools is installed and that the balloon driver is enabled for all virtual machines.

Virtual machine: Swapping is occurring. (Memory usage may be high or low.)

Reduce the memory size on oversized virtual machines.

Reduce the memory reservation of virtual machines where it is set higher than needed.

Add the host to a DRS cluster.

Increase the number of hosts in the DRS cluster.

Migrate one or more virtual machines to other hosts.

Add physical memory to the host.

---

Virtual machine: Memory usage is high.

The guest OS is not provided sufficient memory by the virtual machine.

Increase the memory size of the virtual machine.

Guest OS: Memory usage is high. Paging is occurring.

---

Virtual machine: CPU ready is low.

The guest OS is not provided sufficient CPU resources by the virtual machine.

Increase the number of CPUs for the virtual machine.

Guest OS: CPU utilization is high.

Migrate the virtual machine to a host with faster CPUs.

---

Datastore: Space utilization is high.

Snapshot files are consuming a lot of datastore space.

Delete or consolidate virtual machine snapshots.

Some virtual machines are provisioned with more storage space than required.

Convert some

The datastore has insufficient storage space to meet the demand.

virtual disks to thin provisioned.

Migrate one or more virtual machines (or virtual disks) to other datastores.

Add the datastore to a Storage DRS datastore cluster.

Add datastores with available space to the datastore cluster.

Add more storage space to the datastore.

---

Disk: Device latency is greater than 15 ms.

Problems are occurring with the storage array.

Migrate the virtual machines to datastores backed by other storage arrays.

---

Disk: VMkernel latency is greater than 4 ms. Queue latency is greater than zero.

The maximum throughput of a storage device is not sufficient to meet the demand of the current workload.

Migrate the virtual machines to datastores backed by storage devices (LUNs) with more spindles.

Balance virtual machines and their

disk I/O across the available physical resources. Use Storage DRS I/O balancing.

Add more disks (spindles) to the storage device backing the datastore.

Configure the queue depth and cache settings on the RAID controllers. Adjust the Disk.SchedNumReqOutstanding parameter.

Configure multipathing.

Increase the memory size of the virtual machine to eliminate any guest OS paging. Increase the guest OS caching of disk I/O.

Ensure that no

virtual machine swapping or ballooning is occurring.

Defragment guest file systems.

Use eager zeroed thick provisioned virtual disks.

---

<p>Network: The number of packets dropped is greater than zero. Latency is high. The transfer rate is low.</p>	<p>The maximum throughput of a physical network adapter is not sufficient to meet the demand of the current workload.</p>	<p>Install VMware Tools on each virtual machine and configure the guest OS to use the best-performing network adapter driver (such as vmxnet3).</p>
	<p>Virtual machine network resource shares are too few.</p>	
	<p>Network packet size is too large, which results in high network latency. Use the VMware AppSpeed performance monitoring application or a third-party application to check network latency.</p>	<p>Migrate virtual machines to other hosts or to other physical network adapters.</p>
	<p>Network packet size is too small, which increases the demand for the CPU resources needed for processing each packet. Host CPU, or possibly virtual machine CPU, resources are not enough to handle the load.</p>	<p>Verify that all NICs are running in full duplex mode.</p>
		<p>Implement TCP Segmentation</p>

Offload (TSO) and jumbo frames.

Assign additional physical adapters as uplinks for the associated port groups.

Replace physical network adapters with high-bandwidth adapters.

Place sets of virtual machines that communicate with each other regularly on the same ESXi host.

---

Performance charts are empty.	Some metrics are not available for pre-ESXi 5.0 hosts.	Upgrade hosts to a later version of ESXi.
	Data is deleted when you remove objects to vCenter Server or remove them.	Allow time for data collection on objects that were recently added, migrated, or recovered to the vCenter Server.
	Performance chart data for inventory objects that were moved to a new site by VMware vCenter Site Recovery Manager is deleted from the old site and not copied to the new site.	
	Performance chart data is deleted when	Power on all hosts and allow time for

you use VMware vMotion across vCenter Server instances.

real-time statistics to collect.

Real-time statistics are not available for disconnected hosts or powered-off virtual machines.

Allow time for the required roll-ups for non-real-time statistics.

Non-real-time statistics are rolled up at specific intervals. For example, 1-day statistics might not be available for 30 minutes after the current time, depending on when the sample period began.

The 1-day statistics are rolled up to create one data point every 30 minutes. If a delay occurs in the roll-up operation, the 1-week statistics might not be available for 1 hour after the current time. It takes 30 minutes for the 1-week collection interval, plus 30 minutes for the 1-day collection interval.

The 1-week statistics are rolled up to create one data point every two hours. If a delay occurs in the roll-up operations, the 1-month statistics might not be available for 3 hours. It takes 2 hours for the 1-month collection interval, plus 1 hour for the 1-week collection interval.

The 1-month statistics are rolled up to create one data point every day. If a delay occurs in the roll-up operations, the statistics might not be available for 1 day and 3 hours. It takes 1 day for the past year collection interval, plus 3 hours



for the past month collection interval.  
During this time, the charts are empty.

**Table 10-9** Key ESXTOP Panels and Metrics

Panel	Statistic	Description
CPU	%USED	Percentage of physical CPU core cycles used by the virtual machine.
CPU	%RUN	Percentage of total time scheduled for the virtual machine without accounting for hyperthreading, system time, co-stopping, and waiting:  $\%RUN = 100\% - \%RDY - \%CSTP - \%WAIT$
CPU	%RDY	Percentage of time the virtual machine was ready to run but was not provided CPU resources on which to execute. Indicator of CPU contention on the host.
CPU	%WAIT	Percentage of time the virtual machine spent in the blocked or busy wait state, including idle time. %WAIT includes %SWPWT.
CPU	%CSTP	Percentage of time a virtual machine spends in a ready, co-deschedule state. A high value indicates that the virtual machine's multiple CPUs are in contention.
CPU	%SWPWT	Percentage of time a virtual machine spends waiting for the host to swap memory.
Memory	MEM	Amount of physical memory allocated to a virtual machine:

	M SZ	MEMSZ = GRANT + MCTLSZ + SWCUR + “Never Touched”
Memory	GRANT	Amount of guest physical memory mapped to a virtual machine
Memory	CNSM	Amount of the memory consumed by the virtual machine: CNSM = GRANT – Shared Memory
Memory	SWCUR	Amount of memory swapped by the virtual machine.
Memory	SWR/s	Rate at which the host swaps in memory from disk for the virtual machine.
Memory	OVHD	Amount of memory used for virtual machine overhead, which is memory charged to the virtual machine that is not used by the guest OS.
Virtual Machine Storage	READS/s	Number of read commands issued per second.
Virtual Machine Storage	WRITES/s	Number of write commands issued per second.
Virtual Machine	MB/s	Megabytes read per second.

e Storage	R E A D/ s	
Virtual Machin e Storage	L AT /rd	Average latency (in milliseconds) per read.
Networ k	P K R R X/ s	Number of packets received per second.
Networ k	M bT X/ s	Megabits transmitted per second.
Networ k	% D RP T X	Percentage of transmit packets dropped. Indicates that the physical network adapter cannot meet the demand, perhaps due to load from other virtual machines.
Networ k	% D RP R X	Percentage of receive packets dropped. Indicates that insufficient CPU resources are available for network processing.

**Table 10-13** ESXi Log Files

<b>Component</b>	<b>Location</b>	<b>Description</b>
VMkernel	/var/log/vmkernel.log	Data related to virtual machines and ESXi

VMkernel warnings	/var/log/vmkwarning.log	Data related to virtual machines
VMkernel summary	/var/log/vmksummary.log	Data related to uptime and availability statistics for ESXi
ESXi host agent	/var/log/hostd.log	Data related to the agent that manages and configures the ESXi host and its virtual machines
vCenter agent	/var/log/vpxa.log	Data related to the agent that communicates with vCenter Server
ESXi Shell	/var/log/shell.log	Data related to each command typed into the ESXi Shell as well as shell events
Authentication	/var/log/auth.log	Data related to event authentication for the local system
System messages	/var/log/syslog.log	General log messages that can be used for troubleshooting
Virtual machines	vmware.log located in the same folder as the virtual machine configuration file.	Data related to virtual machine power events, system failure information, tool status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and more
Trusted infrastructure	/var/run/log/kmxa.log	Data related to the client service on the ESXi trusted host

ture  
agen  
t

Key provi der servi ce	/var/run/log/kmxd.l og	Data related to the vSphere Trust Authority key provider service
Attes tatio n servi ce	/var/run/log/attestd .log	Data related to the vSphere Trust Authority attestation service
ESX toke n servi ce	/var/run/log/esxtok end.log	Data related to the vSphere Trust Authority ESXi token service
ESX API forw arder	/var/run/log/esxapi adapter.log	Data related to the vSphere Trust Authority API forwarder
Quic k Boot	/var/log/loadESX.l og	Data related to restarting an ESXi host through Quick Boot

**Table 10-14** vCenter Server Logging Options

<b>Logging Option</b>	<b>Description</b>
None (Disable Logging)	No vCenter Server logging occurs.
Error (Errors Only)	The vCenter Server collects only error entries in its log files.
Warning (Warning and	The vCenter Server collects warning and error entries in its log files.

Errors)

Info (Normal Logging)	The vCenter Server collects information, warning, and error entries in its log files.
Verbose (Verbose)	The vCenter Server collects verbose, information, warning, and error entries in its log files.
Trivia (Extended Verbose)	The vCenter Server collects trivia, verbose, information, warning, and error entries in its log files.

## Chapter 11

**Table 11-2** Network Differences in vSAN and non-vSAN Clusters

<b>Factor</b>	<b>vSAN Is Enabled</b>	<b>vSAN Is Not Enabled</b>
Network used by vSphere HA	vSAN network	Management network
Heartbeat datastores	Any datastore, other than a vSAN datastore, that is mounted to multiple hosts in the cluster	Any datastore that is mounted to multiple hosts in the cluster
Host isolation criteria	Isolation addresses not pingable and vSAN storage network inaccessible	Isolation addresses not pingable and management network inaccessible

**Table 11-4** Datastore Browser Options

<b>Option</b>	<b>Description</b>
Upload Files	Upload a local file to the datastore.
Upload Folder	Upload a local folder to the datastore.
Download	Download a file from the datastore to the local machine.
New Folder	Create a folder on the datastore.

Copy to	Copy selected folders or files to a new location on the datastore or on another datastore.
Move to	Move selected folders or files to a new location on the datastore or on another datastore.
Rename to	Rename selected files.
Delete	Delete selected folders or files.
Inflate	Convert a selected thin virtual disk to thick.

**Table 11-5** Storage Filters

<b>Filter</b>	<b>Description</b>
config.vpxd.filter.vmf sFilter	Hides storage devices (LUNs) that are used by a VMFS datastore on any host managed by vCenter Server.
(VMFS filter)	
config.vpxd.filter.rdm Filter	Hides storage devices (LUNs) that are used by an RDM on any host managed by vCenter Server.
(RDM filter)	
config.vpxd.filter.sameHostsAndTransports Filter	Hides storage devices (LUNs) that are ineligible for use as VMFS datastore extents because of incompatibility with the selected datastore. Hides LUNs that are not exposed to all hosts that share the original datastore. Hides LUNs that use a storage type (such as Fibre Channel, iSCSI, or local) that is different from the original datastore.

(Same Hosts and Transports filter)

---

config.vpxd.filter.host Rescan Filter Automatically rescans and updates VMFS datastores following datastore management operations. If you present a new LUN to a host or a cluster, the hosts automatically perform a rescan, regardless of this setting.

(Host Rescan filter)

**Table 11-7** SCSI over Fabric and NVMe over Fabric Comparison

<b>Shared Storage Capability</b>	<b>SCSI over Fabric</b>	<b>NVMe over Fabric</b>
RDM	Supported	Not supported
Coredump	Supported	Not supported
SCSI-2 reservations	Supported	Not supported
Shared VMDK	Supported	Not supported
vVols	Supported	Not supported
Hardware acceleration with VAAI plug-ins	Supported	Not supported
Default MPP	NMP	HPP (NVMe-oF targets cannot be claimed by NMP.)
Limits	LUNs=1024, paths=4096	Namespaces=32, paths=128 (maximum 4 paths per namespace in a host)



## Chapter 12

**Table 12-2** Sample ESXCLI Commands

<b>Command</b>	<b>Description</b>
<b>esxcli system account add</b>	Creates an ESXi host local user account
<b>esxcli system account set</b>	Configures an ESXi host local user account
<b>esxcli system account list</b>	Lists ESXi host local user accounts
<b>esxcli system account remove</b>	Deletes an ESXi host local user accounts
<b>esxcli network ip dns server list</b>	Lists the host's DNS servers
<b>esxcli network nic list</b>	Lists the ESXi host's physical network adapters
<b>esxcli system settings advanced get /UserVars/ESXiShellTimeOut</b>	Displays the shell interactive timeout for the host

**Table 12-4** ESXi Lockdown Mode Behavior

<b>Service</b>	<b>Normal Mode</b>	<b>Normal Lockdown Mode</b>	<b>Strict Lockdown Mode</b>
vSphere	All users, based on permissions	vCenter (vpxuser)	vCenter (vpxuser)
Web Services API		Exception users, based on permissions	Exception users, based on permissions
		vCloud Director (vslauser, if available)	vCloud Director (vslauser, if available)

CIM providers	Users with administrator privileges on the host	vCenter (vpxuser)  Exception users, based on permissions  vCloud Director (vslauser, if available)	vCenter (vpxuser)  Exception users, based on permissions  vCloud Director (vslauser, if available)
DCUI	Users with administrator privileges on the host and users defined in the DCUI.Access advanced option	Users defined in the DCUI.Access advanced option  Exception users with administrator privileges on the host	DCUI service is stopped
ESXi Shell (if enabled)	Users with administrator privileges on the host	Users defined in the DCUI.Access advanced option  Exception users with administrator privileges on the host	Users defined in the DCUI.Access advanced option  Exception users with administrator privileges on the host
SSH (if enabled)	Users with administrator privileges on the host	Users defined in the	Users defined in the

enable d)	DCUI.Access advanced option	DCUI.Access advanced option
	Exception users with administrator privileges on the host	Exception users with administrator privileges on the host

## Chapter 13

**Table 13-4** Lifecycle Manager Definitions

<b>Term</b>	<b>Definition</b>
Update	A software release that makes small changes to the current version, such as vSphere 7.0 Update 1, 7.0 Update 2, and so on.
Upgrade	A software release that introduces major changes to the software. For example, you can upgrade from vSphere 6.5 to 6.7 and 7.0.
Patch	A small software update that provides bug fixes or enhancements to the current version of the software, such as 7.0a, 7.0 Update 1a, and so on.
VIB (vSphere Installation Bundle)	The smallest installable software package (metadata and binary payload) for ESXi.
VIB metadata	An XML file that describes the contents of the VIB, including dependency information, textual descriptions, system requirements, and information about bulletins.
Standalone VIB	A VIB that is not included in a component.
Depot	The hosted version of updates provided by VMware, OEMs, and third-party software vendors, containing the metadata and

the actual VIBs.

Offline bundle/offline depot	An archive (ZIP file) that contains VIBs and metadata that you use for offline patching and updates. A single offline bundle might contain multiple base images, vendor add-ons, or components.
OEM (original equipment manufacturer)	A VMware partner, such as Dell, HPE, or VMware Cloud on AWS.
Third-party software provider	A provider of I/O filters, device drivers, CIM modules, and so on.

**Table 13-8** Collection Intervals

<b>Collection Interval (Archive Length)</b>	<b>Collection Frequency</b>	<b>Default Behavior</b>
1 day	5 minutes	Real-time (20-second) statistics are rolled up to create one data point every 5 minutes. The result is 288 data points every day.  You can change the interval duration and archive length of the 1-day collection interval by configuring the statistics settings.
1 week	30 minutes	1-day statistics are rolled up to create one data point every 30 minutes. The result is 336 data points every week.  You cannot change the default settings of the 1-week collection interval.
1 month	2 hours	1-week statistics are rolled up to create one data point every 2 hours. The result is 360 data points

every month.

You cannot change the default settings of the 1-month collection interval.

---

1 year	1 day	1-month statistics are rolled up to create one data point every day. The result is 365 data points each year.
--------	-------	---

You can change the archive length of the 1-year collection interval by configuring the statistics settings.

---

## APPENDIX D

# Study Planner

Practice Test	Reading	Review
---------------	---------	--------

Element	Task	Goal Date	First Date Completed	Second Date Completed (Optional)	Notes
Introduction	Read Introduction				
Your Study Plan	Read Your Study Plan				
1. vSphere Overview, Components and Requirements	Read Foundation Topics				
1. vSphere Overview, Components and Requirements	Review Key Topics using the book or companion website				
1. vSphere Overview, Components and Requirements	Define Key Terms using the book or companion website				
1. vSphere Overview, Components and	Repeat DIKTA questions using the book or PTP exam engine				

Requirements					
1. vSphere Overview, Components and Requirements	Complete all memory tables in this chapter using the companion website				
1. vSphere Overview, Components and Requirements	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
2. Storage Infrastructure	Read Foundation Topics				
2. Storage Infrastructure	Review Key Topics using the book or companion website				
2. Storage Infrastructure	Define Key Terms using the book or companion website				
2. Storage Infrastructure	Repeat DIKTA questions using the book or PTP exam engine				
2. Storage Infrastructure	Complete all memory tables in this chapter using the companion website				
2. Storage Infrastructure	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				

3. Network Infrastructure	Read Foundation Topics				
3. Network Infrastructure	Review Key Topics using the book or companion website				
3. Network Infrastructure	Define Key Terms using the book or companion website				
3. Network Infrastructure	Repeat DIKTA questions using the book or PTP exam engine				
3. Network Infrastructure	Complete all memory tables in this chapter using the companion website				
3. Network Infrastructure	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
4. Clusters and High Availability	Read Foundation Topics				
4. Clusters and High Availability	Review Key Topics using the book or companion website				
4. Clusters and High Availability	Define Key Terms using the book or companion website				
4. Clusters and High Availability	Repeat DIKTA questions using the book or PTP exam engine				
4. Clusters and High Availability	Complete all memory tables in this chapter using				



	the companion website				
4. Clusters and High Availability	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
5. vCenter Server Features and Virtual Machines	Read Foundation Topics				
5. vCenter Server Features and Virtual Machines	Review Key Topics using the book or companion website				
5. vCenter Server Features and Virtual Machines	Define Key Terms using the book or companion website				
5. vCenter Server Features and Virtual Machines	Repeat DIKTA questions using the book or PTP exam engine				
5. vCenter Server Features and Virtual Machines	Complete all memory tables in this chapter using the companion website				
5. vCenter Server Features and Virtual Machines	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for				

	this chapter				
6. VMWare Product Integration	Read Foundation Topics				
6. VMWare Product Integration	Review Key Topics using the book or companion website				
6. VMWare Product Integration	Define Key Terms using the book or companion website				
6. VMWare Product Integration	Repeat DIKTA questions using the book or PTP exam engine				
6. VMWare Product Integration	Complete all memory tables in this chapter using the companion website				
6. VMWare Product Integration	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
7. vSphere Security	Read Foundation Topics				
7. vSphere Security	Review Key Topics using the book or companion website				
7. vSphere Security	Define Key Terms using the book or companion website				
7. vSphere Security	Repeat DIKTA questions using the book or PTP				

	exam engine				
7. vSphere Security	Complete all memory tables in this chapter using the companion website				
7. vSphere Security	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
Part I. vSphere Architecture, Integration, Requirements	Complete all exercises in Part I Review				
Practice Test	Take practice test in study mode using Part Review exam in practice test software for this part				
8. vSphere Installation	Read Foundation Topics				
8. vSphere Installation	Review Key Topics using the book or companion website				
8. vSphere Installation	Define Key Terms using the book or companion website				
8. vSphere Installation	Repeat DIKTA questions using the book or PTP exam engine				
8. vSphere Installation	Complete all memory tables in this chapter using the companion website				
8. vSphere	Complete Review				

Installation	Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
9. Configure and Manage Virtual Networks	Read Foundation Topics				
9. Configure and Manage Virtual Networks	Review Key Topics using the book or companion website				
9. Configure and Manage Virtual Networks	Repeat DIKTA questions using the book or PTP exam engine				
9. Configure and Manage Virtual Networks	Complete all memory tables in this chapter using the companion website				
9. Configure and Manage Virtual Networks	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
Part II. vSphere Installation/Configuration	Complete all exercises in Part II Review				
Practice Test	Take practice test in study mode using Part Review exam in practice test software for this part				
10. Monitoring and Managing	Read Foundation Topics				

Clusters and Resources					
10. Monitoring and Managing Clusters and Resources	Review Key Topics using the book or companion website				
10. Monitoring and Managing Clusters and Resources	Define Key Terms using the book or companion website				
10. Monitoring and Managing Clusters and Resources	Repeat DIKTA questions using the book or PTP exam engine				
10. Monitoring and Managing Clusters and Resources	Complete all memory tables in this chapter using the companion website				
10. Monitoring and Managing Clusters and Resources	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
11. Manage Storage	Read Foundation Topics				
11. Manage Storage	Review Key Topics using the book or companion website				
11. Manage Storage	Define Key Terms using the book or companion website				

11. Manage Storage	Repeat DIKTA questions using the book or PTP exam engine				
11. Manage Storage	Complete all memory tables in this chapter using the companion website				
11. Manage Storage	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
12. Manage vSphere Security	Read Foundation Topics				
12. Manage vSphere Security	Review Key Topics using the book or companion website				
12. Manage vSphere Security	Define Key Terms using the book or companion website				
12. Manage vSphere Security	Repeat DIKTA questions using the book or PTP exam engine				
12. Manage vSphere Security	Complete all memory tables and practice exercises in this chapter using the companion website				
12. Manage vSphere Security	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in				

	practice test software for this chapter				
13. Manage vSphere and vCenter Server	Read Foundation Topics				
13. Manage vSphere and vCenter Server	Review Key Topics using the book or companion website				
13. Manage vSphere and vCenter Server	Define Key Terms using the book or companion website				
13. Manage vSphere and vCenter Server	Repeat DIKTA questions using the book or PTP exam engine				
13. Manage vSphere and vCenter Server	Complete all memory tables and practice exercises in this chapter using the companion website				
13. Manage vSphere and vCenter Server	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
14. Virtual Machine Management/Provision, Migrate, Replication	Read Foundation Topics				
14. Virtual Machine Management/Pro	Review Key Topics using the book or companion website				

vision, Migrate, Replication					
14. Virtual Machine Management/Provision, Migrate, Replication	Define Key Terms using the book or companion website				
14. Virtual Machine Management/Provision, Migrate, Replication	Repeat DIKTA questions using the book or PTP exam engine				
14. Virtual Machine Management/Provision, Migrate, Replication	Complete Review Questions using the book or PTP exam engine				
Practice Test	Take practice test in study mode using Exam #1 in practice test software for this chapter				
Part III. vSphere Management and Optimization	Complete all exercises in Part II Review				
Practice Test	Take practice test in study mode using Part Review exam in practice test software for this part				
Final Review	Take practice test in study mode for all Book Questions in practice test software				
Final Review	Review all Key Topics in all chapters or in the Key				



	Topics App using the companion website				
Final Review	Review all Key Terms in all chapters or using the Key Terms Flashcards on the companion website				
Final Review	Complete all memory tables and practice exercises for all chapters using the companion website				
Final Review	Complete all config checklists for all chapters using the companion website				
Final Review	Take practice test in practice exam mode using Exam Bank #1 questions for all chapters				
Final Review	Take practice test in practice exam mode using Exam Bank #2 questions for all chapters				

# Glossary

## A

**add-on:** In vSphere Lifecycle Manager, a collection of components that you can use to customize an ESXi image with OEM content and drivers.

**App Volumes:** A set of application and user management solutions for VMware Horizon, Citrix Virtual Apps and Desktops, and Remote Desktop Services Host (RDSH) virtual environments.

**AppDefense:** A data center endpoint security product that protects applications running in vSphere.

## B

**baseline:** In vSphere Lifecycle Manager, a set of bulletins.

## C

**Certificate Manager:** A command-line utility that you can use to generate certificate signing requests (CSRs) and replace certificates for machine and solution users.

**client performance charts:** vSphere charts that enable you to view performance metrics in different ways, depending on the selected object and metric type.

**cluster:** A set of ESXi hosts that are intended to work together as a unit.

**Common Information Model (CIM):** An open standard that defines a framework for agentless, standards-based monitoring of ESXi host hardware resources. The framework consists of a CIM broker and a set of CIM providers.

**content library:** A repository (container) for objects such as virtual machine templates, vApp templates, ISO images, and files that you might want to share

among multiple vCenter Servers in a vSphere environment.

**CPU Ready Time:** A metric that indicates the amount of time a VCPU is ready to work (that is, has a workload and is ready to be scheduled) but is waiting to be scheduled on hardware. High CPU Ready Time is a sign of CPU contention.

## D

**data center:** A container object in the vSphere inventory that is an aggregation of all the different types of objects used to work in virtual infrastructure.

**disk group:** A group of local disks on an ESXi host that contributes to the vSAN datastore.

**Distributed Resource Scheduler (DRS):** A vSphere feature that balances VM workload in a cluster based on compute usage. It includes live (vMotion) migrations of VMs, when necessary.

## E-F

**ESXi base image:** The ESXi image that VMware provides with each release of ESXi, which is a complete set of components that can boot up a server.

**ESXTOP:** A utility that provides a detailed real-time look at resource usage from the ESXi Shell.

**EtherChannel:** A logical channel formed by bundling together two or more links to aggregate bandwidth and provide redundancy. Other acceptable names for EtherChannel (an IOS term) are port channel (an NX-OS term) and link aggregation group (LAG).

## G

**graphics processing unit (GPU):** A specialized processor developed for parallel processing, primarily for rendering graphical images.

## H

**High-Performance Plug-in (HPP):** The default plug-in that claims NVMe-oF targets.

**host profile:** A feature that enables you to encapsulate the configuration of one

host and apply it to other hosts.

**hybrid cloud:** A cloud that is a combination of a private cloud, a public cloud, and on-premises infrastructure.

## I-L

**I/O filter:** A software component that can be installed on ESXi hosts and can offer additional data services to virtual machines.

**image:** In vSphere Lifecycle Manager, a description of which software, drivers, and firmware to run on a host.

**Intel Software Guard Extension (SGX):** A processor-specific technology that enables application developers to protect code and data from disclosure or modification.

## M

**managed object browser (MOB):** A web-based interface that provides you with a means to explore the VMkernel object model.

**microsegmentation:** A type of network segmentation that decreases the level of risk and increases the security posture of a data center by providing granular control and distributed stateful firewalling. It effectively allows you to place a firewall on each VM network connection.

**Microsoft virtualization-based security (VBS):** A Microsoft feature for Windows 10 and Windows Server 2016 operating systems that uses hardware and software virtualization to enhance system security by creating an isolated, hypervisor-restricted, specialized subsystem.

## N

**NetFlow:** A switch feature that collects IP network traffic as it enters or exits an interface. NetFlow data provides an overview of traffic flows, based on the network source and destination.

**network resource pool:** A mechanism that enables you to apply a part of the bandwidth that is reserved for virtual machine system traffic to a set of distributed port groups.

**Non-Volatile Memory Express (NVMe) device:** A high-performance alternative to SCSI storage.

## O

**Open Virtual Appliance (OVA) template:** A single-file distribution of an OVF package.

**Open Virtual Format (OVF) template:** A set of files with the OVF, VMDK, and MF file extensions

## P-Q

**PMem device:** A non-volatile dual in-line memory module (NVDIMM) on the ESXi host that resides in a normal memory slot.

**port mirroring:** A process that allows administrators to duplicate everything that is happening on one distributed port to then be visible on another distributed port.

**Predictive DRS:** A feature that leverages the predictive analytics of vRealize Operations (vROps) Manager and vSphere DRS to provide workload balancing prior to the occurrence of resource utilization spikes and resource contention.

**private VLAN (PVLAN):** An extension of the VLAN standard that is not double encapsulated but that allows a VLAN to effectively be subdivided into other VLANs

**Proactive HA:** A vSphere feature that minimizes VM downtime by proactively detecting hardware failures and placing the host in Quarantine Mode or Maintenance Mode.

**Proactive High Availability (Proactive HA):** A feature that integrates with select hardware partners to detect degraded components and evacuate VMs from affected vSphere hosts before an incident causes a service interruption.

## R

**raw device mapping (RDM):** A mapping file that contains metadata that resides in a VMFS datastore and acts as a proxy for a physical storage device (LUN), allowing a virtual machine to access the storage device directly.

**resource pool:** A container object in the vSphere inventory that is used to compartmentalize the CPU and memory resources of a host or cluster.

## S

**Single Root I/O Virtualization (SR-IOV):** A feature that allows a single Peripheral Component Interconnect Express (PCIe) device to appear as multiple devices to the hypervisor (ESXi) or to a virtual machine's guest operating system.

**Site Recovery Manager (SRM):** A VMware business continuity solution that you can use to orchestrate planned migrations, test recoveries, and disaster recoveries.

**Skyline:** A proactive support technology developed by VMware Global Services that is available to customers with an active Production Support or Premier Services agreement.

**stateless caching:** A type of caching in which Auto Deploy does not store ESXi configuration or state data within the host. Instead, during subsequent boots, the host must connect to the Auto Deploy server to retrieve its configuration.

**Storage I/O Control (SIOC):** A vSphere feature that allows you to prioritize storage access during periods of contention, ensuring that the more critical virtual machines obtain more I/O than less critical VMs.

**Storage vMotion:** The hot cross-datastore migration of a virtual machine.

## T-U

**template:** An object in the vSphere inventory that is effectively a non-executable virtual machine.

**Trusted Platform Module (TPM):** A chip that is a secure cryptoprocessor that enhances host security by providing a trust assurance rooted in hardware as opposed to software.

## V

**vApp:** A container object in vSphere that provides a format for packaging and managing applications.

**vCenter Converter:** A free solution that automates the process of converting existing Windows and Linux machines into virtual machines running in a vSphere environment.

**vCenter HA:** A native high availability solution for vCenter Server Appliance.

**vCenter Single Sign-On (SSO):** An authentication broker and security token exchange infrastructure.

**vCenter Single Sign-On Security Token Service (STS):** A web service that issues, validates, and renews security tokens.

**VIMTOP:** A tool you can run in vCenter Server Appliance to see resource usage for services that are running.

**virtual LAN (VLAN):** A logical partition of a physical network at the data link layer (Layer 2).

**Virtual Machine Component Protection (VMCP):** A vSphere HA feature that can detect datastore accessibility issues and provide remediation for impacted virtual machines.

**virtual machine snapshot:** A copy that captures the state of a virtual machine and the data in the virtual machine at a specific point in time.

**virtual PMem disk (vPMemDisk):** A regular virtual disk that is assigned a PMem storage policy, which forces it to be placed on a host-local PMem datastore.

**Virtual Trusted Platform Module (vTPM):** A software-based representation of a physical TPM 2.0 chip.

**virtual volume:** An encapsulation of virtual machine files, virtual disks, and their derivatives that are stored natively inside a storage system.

**Virtual Volumes (vVOLs):** An integration and management framework that virtualizes SAN/NAS arrays, enabling a more efficient operational model.

**vMotion:** The hot cross-host migration of a virtual machine.

**VMware Certificate Authority (VMCA):** A certificate authority (CA) that is responsible for issuing certificates for VMware solution users, certificates for machines running required services, and certificates for ESXi hosts. It provisions each ESXi host, each machine in the environment, and each solution user with a certificate signed by VMCA.

**VMware Cloud (VMC):** An integrated cloud offering jointly developed by

AWS and VMware that provides a highly scalable, secure service that allows organizations to seamlessly migrate and extend their on-premises vSphere-based environments to the AWS cloud.

**VMware Directory Service (vmdir):** An identity source that handles SAML certificate management for authentication with vCenter Single Sign-On.

**VMware Endpoint Certificate Service (VECS):** A local (client-side) repository for certificates, private keys, and other certificate information that can be stored in a keystore. VECS is a mandatory component that is used when VMCA is not signing certificates.

**VMware Enhanced Authentication Plug-in:** A plug-in that provides integrated Windows authentication and Windows-based smart card functionality. In the vSphere 6.5 release, the VMware Enhanced Authentication Plug-in replaced the Client Integration Plug-in.

**VMware Horizon:** A platform for securely delivering virtual desktops and applications in private and hybrid clouds.

**VMware Lifecycle Manager:** A service that replaces VMware Update Manager and adds features and capabilities for ESXi lifecycle management at the cluster level.

**VMware PowerCLI:** A command-line and scripting tool built on Windows PowerShell that provides cmdlets for managing and automating VMware products, including vSphere.

**VMware Service Lifecycle Manager:** A vCenter Server service that monitors the health of services and takes preconfigured remediation action when it detects a failure.

**VMware Tools:** A set of software modules and services, including services that can communicate with the VMkernel.

**VMware Update Manager Download Service (UMDS):** An optional module of vSphere Lifecycle Manager, whose primary function is to download data when Lifecycle Manager does not have Internet connectivity.

**VMware vSphere Lifecycle Manager:** A service that runs in vCenter Server and provides simple, centralized lifecycle management for ESXi hosts and clusters by using images and baselines.

**vRealize Automation (vRA):** An automation platform for private and multi-cloud environments.



**vRealize Log Insight (vRLI):** A software product that provides intelligent log management for infrastructure and applications for any environment.

**vSAN fault domain:** A set of hosts that are members of a vSAN cluster. Fault domains provide additional protection against outage in the event of a rack or blade chassis failure.

**vSAN file service:** A file service on a vSAN cluster, which enables you to create file shares. It provides vSAN-backed file shares that virtual machines can access as NFS Version 3 and NFS Version 4.1 file shares.

**vSphere Client:** An HTML5-based GUI used for administration in vSphere.

**vSphere Distributed Switch (vDS):** A single virtual switch for all associated hosts in a data center. It provides centralized provisioning, monitoring, and management of virtual networks for associated hosts and virtual machines

**vSphere Fault Tolerance (FT):** A feature that provides continuous availability for a virtual machine (the primary VM) by ensuring that the state of a secondary VM is identical at any point in the instruction execution of the virtual machine.

**vSphere HA:** A vSphere feature that provides automated failover protection for VMs against host, hardware, network, and guest OS issues. In the event of host system failure, it performs cold migrations and restarts failed VMs on surviving hosts.

**vSphere Installation Bundle (VIB):** An ESXi software package, created and signed by VMware and its partners, that contains solutions, drivers, CIM providers, and applications.

**vSphere inventory:** A collection of managed virtual and physical objects.

**vSphere Replication:** An extension to VMware vCenter Server that provides hypervisor-based virtual machine replication and recovery.

## W-Z

**witness host:** A stretched vSAN component that consists only of metadata and acts as a tiebreaker.



## Where are the companion content files?

Register this digital version of [VCP-DCV for vSphere 7.x \(Exam 2V0-21.20\) Official Cert Guide](#) to access important downloads.

Register this eBook to unlock the companion files. Follow these steps:

1. Go to [pearsonITcertification.com/account](https://pearsonITcertification.com/account) and log in or create a new account.
2. Enter the ISBN: **9780135898192** (NOTE: Please enter the print book ISBN provided to register the eBook you purchased.)
3. Answer the challenge question as proof of purchase.
4. Click on the “Access Bonus Content” link in the Registered Products section of your account page, to be taken to the page where your downloadable content is available.

This eBook version of the print title does not contain the practice test software that accompanies the print book.

You May Also Like—Premium Edition eBook and Practice Test. To learn about the Premium Edition eBook and Practice Test series, visit [pearsonITcertification.com/practicetest](https://pearsonITcertification.com/practicetest)

## The Professional and Personal Technology Brands of Pearson



**Cisco Press**

**informIT**

PEARSON IT Certification

**QUE**

**SAMS**

# **Official Cert Guide VCP-DCV for vSphere 7.X Exam 2V0-21.20**

ISBN: 978-0-13-589819-2

**See inside ►►►**

for your Pearson Test Prep activation code and special offers

# **Official Cert Guide VCP-DCV for vSphere 7.x Exam 2V0-21.20**

## **Premium Edition eBook and Practice Test**

To enhance your preparation, Pearson IT Certification also sells a digital Premium Edition of this book. The Premium Edition provides you with three eBook files (PDF, EPUB, and MOBI/Kindle) as well as an enhanced edition of the Pearson Test Prep practice test software. The Premium Edition includes two additional practice exams with links for every question mapped to the PDF eBook.

## **Special Offer—Save 80%**

This single-use coupon code will allow you to purchase a copy of the Premium Edition at an 80% discount. Simply go to the URL below, add the Premium Edition to your cart, and apply the coupon code at checkout.

[www.pearsonITcertification.com/title/9780136484264](http://www.pearsonITcertification.com/title/9780136484264)

Coupon Code:

**DO NOT DISCARD THIS NUMBER**

You will need this activation code to activate your practice test in the Pearson Test Prep practice test software.

To access the online version, go to [www.PearsonTestPrep.com](http://www.PearsonTestPrep.com). Select **Pearson IT Certification** as your product group. Enter your email/password for your account. If you don't have an account on [PearsonITCertification.com](http://PearsonITCertification.com) or [CiscoPress.com](http://CiscoPress.com), you will need to establish one by going to [PearsonITCertification.com/join](http://PearsonITCertification.com/join). In the My Products tab, click the **Activate New Product button**. Enter the access code printed on this insert card to activate your product. The product will now be listed in your My Products page.

If you wish to use the Windows desktop offline version of the application, simply register your book at [www.pearsonITcertification.com/register](http://www.pearsonITcertification.com/register), select the **Registered Products** tab on your account page, click the **Access Bonus Content** link, and download and install the software from the companion website.

This activation code can be used to register your exam in both the online and the offline versions.

Activation Code:

# Special Offer

## **Save 80% on Premium Edition eBook and Practice Test**

The *VCP-DCV for vSphere 7.x (Exam 2V0-21.20) Official Cert Guide* provides three eBook files (PDF, EPUB, and MOBI/Kindle) to read on your preferred device and an enhanced edition of the Pearson Test Prep practice test software. You also receive two additional practice exams with links for every question mapped to the PDF eBook.

**See the card insert in the back of the book for your Pearson Test Prep activation code and special offers.**

# Official Cert Guide VCP-DCV for vSphere 7.x

Exam 2V0-21.20

Official Cert Guide **Companion Website**

---

Access interactive study tools on this book's companion website, including practice test software, memory tables, review exercises, Key Term flash card application, study planner, and more!

To access the companion website, simply follow these steps:

1. Go to [www.pearsonitcertification.com/register](http://www.pearsonitcertification.com/register).
2. Enter the print book ISBN: **9780135898192**..
3. Answer the security question to validate your purchase.
4. Go to your account page.
5. Click on the **Registered Products** tab.
6. Under the book listing, click on the **Access Bonus Content** link.

If you have any issues accessing the companion website, you can contact our support team by going to <http://pearsonitp.ehelp.org>.

## Code Snippets

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a “Click here to view code image” link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.



```
vmkfstools -configATSOOnly 1 [storage path]
```

```
displayName = "server1"  
guestOS = "centos-64"  
nvram = "server1.nvram"  
scsi0:0.fileName = "server1.vmdk"
```

```
ddb.geometry.cylinders = "1305"  
ddb.geometry.heads = "255"  
ddb.geometry.sectors = "63"
```

```
isolation.tools.diskWiper.disable = "TRUE"  
isolation.tools.diskShrink.disable = "TRUE"
```

```
isolation.tools.copy.disable = "TRUE"  
isolation.tools.paste.disable = "TRUE"
```

```
isolation.device.connectable.disable = "TRUE"
```

```
isolation.device.edit.disable = "TRUE"
```

```
tools.setInfo.sizeLimit = "1048576"
```

```
tools.guestlib.enableHostInfo = "FALSE"
```



```
isolation.device.connectable.disable = "TRUE"  
isolation.device.edit.disable = "TRUE"
```

```
ks=http://192.168.100.10/kickstart/ks-script-01.cfg  
nameserver=192.168.1.100 ip=192.168.100.101  
netmask=255.255.255.0 gateway=192.168.100.101
```

```
#
# Sample scripted installation file
#
# Accept the VMware End User License Agreement
vmaccepteula
# Set the root password for the DCUI and Tech Support Mode
rootpw myp@ssw0rd
# Install on the first local disk available on machine
install --firstdisk --overwritevmfs
# Set the network to DHCP on the first network adapter
network --bootproto=dhcp --device=vmnic0
# A sample post-install script
%post --interpreter=python --ignorefailure=true
import time
stampFile = open('/finished.stamp',
```

```
New-DeployRule -Name Rule-01 -Item "Auto-deployed Hosts" -allhosts
```

```
Set-DeployRule -DeployRule Rule-01 -Pattern  
"ipv4=192.168.100.101-192.168.100.112"
```

```
vcsa-deploy install --verify-only path-to-JSON-file
```

```
vcsa-deploy install --accept-eula --acknowledge-ceip  
optional_arguments path-to-JSON-file
```

```
esxcli network ip netstack add -N="stack_name"
```



```
<syslog-prefix> : Event [eventId] [partInfo] [createdTime]  
[eventType] [severity] [user] [target] [chainId] [desc]
```

```
esxcfg-advcfg -s 1 /VSAN/IgnoreClusterMemberListUpdates
```

```
esxcli system maintenanceMode set -e true -m noAction
```

```
esxcli system maintenanceMode set -e false
esxcli vsan cluster get
```

```
esxcfg-advcfg -s 0 /VSAN/IgnoreClusterMemberListUpdates
```

```
scsiX:Y.ignoreDeviceInquiryCache = "true"
```

```
esxcli storage core device inquirycache set --device deviceID  
--ignore true
```

```
esxcli storage core plugin list -plugin-class=MP
```



```
esxcli storage nmp device list --device=mpx.vmbha0:C0:T0:L0
```

```
/opt/vmware/bin/sso-config.sh -set_signing_cert -t vsphere.local  
~/newsts/newsts.pem
```

`/usr/lib/vmware-vmca/bin/certificate-manager`

retry=3 min=disabled,disabled,disabled,7,7

retry=3 min=disabled,disabled,16,7,7,passphrase=3

```
vifs --server hostname --username username --put filename /host/  
ssh_host_dsa_key_pub
```

```
esxcli --server=<server_name> software acceptance set
```

```
/usr/lib/vmware/secureboot/bin/secureBoot.py -c
```



```
vmware-umds -S --add-url https://web1.vendor1.com/index.html  
/index.xml --url-type HOST
```

```
vmware-umds -E -export-store repositoryPath
```

```
/usr/lib/vmware/loadesx/bin/loadESXCheckCompat.py
```

com.vmware.appliance.system.storage.resize

com.vmware.appliance.version1.resources.storage.stats.list

```
cmsso-util unregister --node-pnid Target_vCenter_FQDN  
--username username --passwd password
```

```
cmsso-util domain-repoint -m pre-check --src-emb-admin  
Administrator --replication-partner-fqdn Replication_  
Partner_FQDN --replication-partner-admin Existing_Domain_  
Admin --dest-domain-name Existing_Domain
```

```
cmsso-util domain-repoint -m execute --src-emb-admin  
Administrator --replication-partner-fqdn Replication_  
Partner_FQDN --replication-partner-admin Existing_Domain_  
Admin --dest-domain-name Existing_Domain
```



```
Install-Module VMware.PowerCLI -Scope CurrentUser
```

Set-ExecutionPolicy RemoteSigned

```
Connect-VIServer -Server server1.vsphere.local -Protocol http  
-User 'administrator@vsphere.local' -Password 'VMware1!'
```

```
Get-VM
```

```
<CreateVM>
  <VM>
    <Name>MyVM1</Name>
    <HDDCapacity>100</HDDCapacity>
  </VM>
  <VM>
    <Name>MyVM2</Name>
    <HDDCapacity>100</HDDCapacity>
  </VM>
</CreateVM>
```

```
[xml]$s = Get-Content myVM.xml
$s.CreateVM.VM | foreach {New-VM -VMHost $vmHost1 -Name $_.Name
-DiskGB $_.HDDCapacity}
```

```
Get-VM MyVM1 | Move-VM -Datastore DS2
```

```
esxcli software vib install -v /vmfs/volumes/ARL-ESX14-DS1/  
NVIDIA/NVIDIA-VMware_ESXi_6.7_Host_Driver_390.42-1OEM.  
670.0.0.7535516.vib
```

# Contents

[Cover Page](#)

[About This eBook](#)

[Title Page](#)

[Copyright Page](#)

[Credits](#)

[We Want to Hear from You!](#)

[Reader Services](#)

[Dedications](#)

[About the Authors](#)

[About the Reviewer](#)

[Acknowledgments](#)

[Contents at a Glance](#)

[Contents](#)

[Foreword](#)

[Introduction](#)

[Part I: vSphere Architecture, Integration, and Requirements](#)

[Chapter 1. vSphere Overview, Components, and Requirements](#)

[“Do I Know This Already?” Quiz](#)

[Foundation Topics](#)

[vSphere Components and Editions](#)

[vCenter Server Topology](#)

[Infrastructure Requirements](#)

[Other Requirements](#)

[VMware Cloud vs. VMware Virtualization](#)



Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Answer Review Questions

## Chapter 2. Storage Infrastructure

“Do I Know This Already?” Quiz

Foundation Topics

Storage Models and Datastore Types

vSAN Concepts

vSphere Storage Integration

Storage Multipathing and Failover

Storage Policies

Storage DRS (SDRS)

Exam Preparation Tasks

Review All Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Chapter 3. Network Infrastructure

“Do I Know This Already?” Quiz

Foundation Topics

Networking Terms and Concepts

vSphere Standard Switch (vSS)

Distributed Virtual Switch (vDS)

vDS Settings and Features

Other vSphere Networking Features

Exam Preparation Tasks

Review All Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

#### Chapter 4. Clusters and High Availability

“Do I Know This Already?” Quiz

Foundation Topics

Cluster Concepts and Overview

Distributed Resource Scheduler (DRS)

vSphere High Availability (HA)

Other Resource Management and Availability Features

Exam Preparation Tasks

Review All Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

#### Chapter 5. vCenter Server Features and Virtual Machines

“Do I Know This Already?” Quiz

Foundation Topics

vCenter Server and vSphere

Virtual Machine File Structure

Virtual Machine Snapshots

Virtual Machine Settings

Virtual Machine Migration

Virtual Machine Cloning

Exam Preparation Tasks

Review All Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Chapter 6. VMware Product Integration

“Do I Know This Already?” Quiz

Foundation Topics

vSphere Add-ons

vRealize Suite

Desktop and Application Virtualization

Replication and Disaster Recovery

Private, Public, and Hybrid Clouds

Networking and Security

Exam Preparation Tasks

Review All Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Chapter 7. vSphere Security

“Do I Know This Already?” Quiz

Foundation Topics

vSphere Certificates

vSphere Permissions

ESXi and vCenter Server Security

vSphere Network Security

Virtual Machine Security

Available Add-on Security

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Part II: vSphere Installation/Configuration

### Chapter 8. vSphere Installation

“Do I Know This Already?” Quiz

Foundation Topics

Installing ESXi Hosts

Deploying vCenter Server Components

Configuring Single Sign-On (SSO)

Initial vSphere Configuration

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

### Chapter 9. Configuring and Managing Virtual Networks

“Do I Know This Already?” Quiz

Foundation Topics

vSphere Standard Switches (vSS)

vSphere Distributed Switches (vDS)

VMkernel Networking

Configuring and Managing Networking Features

Managing Host Networking with vDS

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

### Part III: vSphere Management and Optimization

#### Chapter 10. Managing and Monitoring Clusters and Resources

“Do I Know This Already?” Quiz

Foundation Topics

Creating and Configuring a vSphere Cluster

Creating and Configuring a vSphere DRS Cluster

Creating and Configuring a vSphere HA Cluster

Monitoring and Managing vSphere Resources

Events, Alarms, and Automated Actions

Logging in vSphere

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

#### Chapter 11. Managing Storage

“Do I Know This Already?” Quiz

Foundation Topics

Configuring and Managing vSAN

Managing Datastores

Storage DRS and SIOC

NVMe and PMem

Multipathing, Storage Policies, and vVols

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Chapter 12. Managing vSphere Security

“Do I Know This Already?” Quiz

Foundation Topics

Configuring and Managing Authentication and Authorization

Configuring and Managing vSphere Certificates

General ESXi Security Recommendations

Configuring and Managing ESXi Security

Additional Security Management

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Chapter 13. Managing vSphere and vCenter Server

“Do I Know This Already?” Quiz

Foundation Topics

vCenter Server Backup

Upgrading to vSphere 7.0

Using vSphere Lifecycle Manager

Managing ESXi Hosts

Monitoring and Managing vCenter Server

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Chapter 14. Managing Virtual Machines

“Do I Know This Already?” Quiz

Foundation Topics

Creating and Configuring Virtual Machines

Managing Virtual Machines

Advanced Virtual Machine Management

Content Libraries

Exam Preparation Tasks

Review All the Key Topics

Complete Tables and Lists from Memory

Define Key Terms

Review Questions

## Chapter 15. Final Preparation

Getting Ready

Taking the Exam

Appendix A: Answers to the “Do I Know This Already?”  
Quizzes and Review Questions

Index

Appendix B: Memory Tables

Appendix C: Memory Table Answer Key

Appendix D: Study Planner

Glossary

Where are the companion content files? - Register

Inside Front Cover

Inside Back Cover

Code Snippets

1. [i](#)
2. [ii](#)
3. [iii](#)
4. [iv](#)
5. [v](#)
6. [vi](#)
7. [vii](#)
8. [viii](#)
9. [ix](#)
10. [x](#)
11. [xi](#)
12. [xii](#)
13. [xiii](#)
14. [xiv](#)
15. [xv](#)
16. [xvi](#)
17. [xvii](#)
18. [xviii](#)
19. [xix](#)
20. [xx](#)
21. [xxi](#)
22. [xxii](#)
23. [xxiii](#)



24. [xxiv](#)
25. [xxv](#)
26. [xxvi](#)
27. [xxvii](#)
28. [xxviii](#)
29. [xxix](#)
30. [xxx](#)
31. [xxxii](#)
32. [xxxiii](#)
33. [xxxiiii](#)
34. [xxxv](#)
35. [xxxvi](#)
36. [xxxvii](#)
37. [xxxviii](#)
38. [xxxix](#)
39. [xl](#)
40. [xli](#)
41. [xlii](#)
42. [xliii](#)
43. [xliv](#)
44. [xlv](#)
45. [xlvi](#)
46. [xlvii](#)
47. [xlviii](#)
48. [xlix](#)
49. [l](#)

- 51. [li](#)
- 52. [2](#)
- 53. [3](#)
- 54. [4](#)
- 55. [5](#)
- 56. [6](#)
- 57. [7](#)
- 58. [8](#)
- 59. [9](#)
- 60. [10](#)
- 61. [11](#)
- 62. [12](#)
- 63. [13](#)
- 64. [14](#)
- 65. [15](#)
- 66. [16](#)
- 67. [17](#)
- 68. [18](#)
- 69. [19](#)
- 70. [20](#)
- 71. [21](#)
- 72. [22](#)
- 73. [23](#)
- 74. [24](#)
- 75. [25](#)
- 76. [26](#)
- 77. [27](#)

- 78. [28](#)
- 79. [29](#)
- 80. [30](#)
- 81. [31](#)
- 82. [32](#)
- 83. [33](#)
- 84. [34](#)
- 85. [35](#)
- 86. [36](#)
- 87. [37](#)
- 88. [38](#)
- 89. [39](#)
- 90. [40](#)
- 91. [41](#)
- 92. [42](#)
- 93. [43](#)
- 94. [44](#)
- 95. [45](#)
- 96. [46](#)
- 97. [47](#)
- 98. [48](#)
- 99. [49](#)
- 100. [50](#)
- 101. [51](#)
- 102. [52](#)
- 103. [53](#)
- 104. [54](#)

- 105. 55
- 106. 56
- 107. 57
- 108. 58
- 109. 59
- 110. 60
- 111. 61
- 112. 62
- 113. 63
- 114. 64
- 115. 65
- 116. 66
- 117. 67
- 118. 68
- 119. 69
- 120. 70
- 121. 71
- 122. 72
- 123. 73
- 124. 74
- 125. 75
- 126. 76
- 127. 77
- 128. 78
- 129. 79
- 130. 80
- 131. 81

- 132. 82
- 133. 83
- 134. 84
- 135. 85
- 136. 86
- 137. 87
- 138. 88
- 139. 89
- 140. 90
- 141. 91
- 142. 92
- 143. 93
- 144. 94
- 145. 95
- 146. 96
- 147. 97
- 148. 98
- 149. 99
- 150. 100
- 151. 101
- 152. 102
- 153. 103
- 154. 104
- 155. 105
- 156. 106
- 157. 107
- 158. 108

- 159. [109](#)
- 160. [110](#)
- 161. [111](#)
- 162. [112](#)
- 163. [113](#)
- 164. [114](#)
- 165. [115](#)
- 166. [116](#)
- 167. [117](#)
- 168. [118](#)
- 169. [119](#)
- 170. [120](#)
- 171. [121](#)
- 172. [122](#)
- 173. [123](#)
- 174. [124](#)
- 175. [125](#)
- 176. [126](#)
- 177. [127](#)
- 178. [128](#)
- 179. [129](#)
- 180. [130](#)
- 181. [131](#)
- 182. [132](#)
- 183. [133](#)
- 184. [134](#)
- 185. [135](#)

- 186. [136](#)
- 187. [137](#)
- 188. [138](#)
- 189. [139](#)
- 190. [140](#)
- 191. [141](#)
- 192. [142](#)
- 193. [143](#)
- 194. [144](#)
- 195. [145](#)
- 196. [146](#)
- 197. [147](#)
- 198. [148](#)
- 199. [149](#)
- 200. [150](#)
- 201. [151](#)
- 202. [152](#)
- 203. [153](#)
- 204. [154](#)
- 205. [155](#)
- 206. [156](#)
- 207. [157](#)
- 208. [158](#)
- 209. [159](#)
- 210. [160](#)
- 211. [161](#)
- 212. [162](#)

- 213. 163
- 214. 164
- 215. 165
- 216. 166
- 217. 167
- 218. 168
- 219. 169
- 220. 170
- 221. 171
- 222. 172
- 223. 173
- 224. 174
- 225. 175
- 226. 176
- 227. 177
- 228. 178
- 229. 179
- 230. 180
- 231. 181
- 232. 182
- 233. 183
- 234. 184
- 235. 185
- 236. 186
- 237. 187
- 238. 188
- 239. 189



240. [190](#)  
241. [191](#)  
242. [192](#)  
243. [193](#)  
244. [194](#)  
245. [195](#)  
246. [196](#)  
247. [197](#)  
248. [198](#)  
249. [199](#)  
250. [200](#)  
251. [201](#)  
252. [202](#)  
253. [203](#)  
254. [204](#)  
255. [205](#)  
256. [206](#)  
257. [207](#)  
258. [208](#)  
259. [209](#)  
260. [210](#)  
261. [211](#)  
262. [212](#)  
263. [213](#)  
264. [214](#)  
265. [215](#)  
266. [216](#)

267. [217](#)  
268. [218](#)  
269. [219](#)  
270. [220](#)  
271. [221](#)  
272. [222](#)  
273. [223](#)  
274. [224](#)  
275. [225](#)  
276. [226](#)  
277. [227](#)  
278. [228](#)  
279. [229](#)  
280. [230](#)  
281. [231](#)  
282. [232](#)  
283. [233](#)  
284. [234](#)  
285. [235](#)  
286. [236](#)  
287. [237](#)  
288. [238](#)  
289. [239](#)  
290. [240](#)  
291. [241](#)  
292. [242](#)  
293. [243](#)

- 294. [244](#)
- 295. [245](#)
- 296. [246](#)
- 297. [247](#)
- 298. [248](#)
- 299. [249](#)
- 300. [250](#)
- 301. [251](#)
- 302. [252](#)
- 303. [253](#)
- 304. [254](#)
- 305. [255](#)
- 306. [256](#)
- 307. [257](#)
- 308. [258](#)
- 309. [259](#)
- 310. [260](#)
- 311. [261](#)
- 312. [262](#)
- 313. [263](#)
- 314. [264](#)
- 315. [265](#)
- 316. [266](#)
- 317. [267](#)
- 318. [268](#)
- 319. [269](#)
- 320. [270](#)

- 321. [271](#)
- 322. [272](#)
- 323. [273](#)
- 324. [274](#)
- 325. [275](#)
- 326. [276](#)
- 327. [277](#)
- 328. [278](#)
- 329. [279](#)
- 330. [280](#)
- 331. [281](#)
- 332. [282](#)
- 333. [283](#)
- 334. [284](#)
- 335. [285](#)
- 336. [286](#)
- 337. [287](#)
- 338. [288](#)
- 339. [289](#)
- 340. [290](#)
- 341. [291](#)
- 342. [292](#)
- 343. [293](#)
- 344. [294](#)
- 345. [295](#)
- 346. [296](#)
- 347. [297](#)

- 348. [298](#)
- 349. [299](#)
- 350. [300](#)
- 351. [301](#)
- 352. [302](#)
- 353. [303](#)
- 354. [304](#)
- 355. [305](#)
- 356. [306](#)
- 357. [307](#)
- 358. [308](#)
- 359. [309](#)
- 360. [310](#)
- 361. [311](#)
- 362. [312](#)
- 363. [313](#)
- 364. [314](#)
- 365. [315](#)
- 366. [316](#)
- 367. [317](#)
- 368. [318](#)
- 369. [319](#)
- 370. [320](#)
- 371. [321](#)
- 372. [322](#)
- 373. [323](#)
- 374. [324](#)

375. [325](#)  
376. [326](#)  
377. [327](#)  
378. [328](#)  
379. [329](#)  
380. [330](#)  
381. [331](#)  
382. [332](#)  
383. [333](#)  
384. [334](#)  
385. [335](#)  
386. [336](#)  
387. [337](#)  
388. [338](#)  
389. [339](#)  
390. [340](#)  
391. [341](#)  
392. [342](#)  
393. [343](#)  
394. [344](#)  
395. [345](#)  
396. [346](#)  
397. [347](#)  
398. [348](#)  
399. [349](#)  
400. [350](#)  
401. [351](#)

- 402. [352](#)
- 403. [353](#)
- 404. [354](#)
- 405. [355](#)
- 406. [356](#)
- 407. [357](#)
- 408. [358](#)
- 409. [359](#)
- 410. [360](#)
- 411. [361](#)
- 412. [362](#)
- 413. [363](#)
- 414. [364](#)
- 415. [365](#)
- 416. [366](#)
- 417. [367](#)
- 418. [368](#)
- 419. [369](#)
- 420. [370](#)
- 421. [371](#)
- 422. [372](#)
- 423. [373](#)
- 424. [374](#)
- 425. [375](#)
- 426. [376](#)
- 427. [377](#)
- 428. [378](#)

429. [379](#)  
430. [380](#)  
431. [381](#)  
432. [382](#)  
433. [383](#)  
434. [384](#)  
435. [385](#)  
436. [386](#)  
437. [387](#)  
438. [388](#)  
439. [389](#)  
440. [390](#)  
441. [391](#)  
442. [392](#)  
443. [393](#)  
444. [394](#)  
445. [395](#)  
446. [396](#)  
447. [397](#)  
448. [398](#)  
449. [399](#)  
450. [400](#)  
451. [401](#)  
452. [402](#)  
453. [403](#)  
454. [404](#)  
455. [405](#)



- 456. [406](#)
- 457. [407](#)
- 458. [408](#)
- 459. [409](#)
- 460. [410](#)
- 461. [411](#)
- 462. [412](#)
- 463. [413](#)
- 464. [414](#)
- 465. [415](#)
- 466. [416](#)
- 467. [417](#)
- 468. [418](#)
- 469. [419](#)
- 470. [420](#)
- 471. [421](#)
- 472. [422](#)
- 473. [423](#)
- 474. [424](#)
- 475. [425](#)
- 476. [426](#)
- 477. [427](#)
- 478. [428](#)
- 479. [429](#)
- 480. [430](#)
- 481. [431](#)
- 482. [432](#)

483. [433](#)  
484. [434](#)  
485. [435](#)  
486. [436](#)  
487. [437](#)  
488. [438](#)  
489. [439](#)  
490. [440](#)  
491. [441](#)  
492. [442](#)  
493. [443](#)  
494. [444](#)  
495. [445](#)  
496. [446](#)  
497. [447](#)  
498. [448](#)  
499. [449](#)  
500. [450](#)  
501. [451](#)  
502. [452](#)  
503. [453](#)  
504. [454](#)  
505. [455](#)  
506. [456](#)  
507. [457](#)  
508. [458](#)  
509. [459](#)

- 510. [460](#)
- 511. [461](#)
- 512. [462](#)
- 513. [463](#)
- 514. [464](#)
- 515. [465](#)
- 516. [466](#)
- 517. [467](#)
- 518. [468](#)
- 519. [469](#)
- 520. [470](#)
- 521. [471](#)
- 522. [472](#)
- 523. [473](#)
- 524. [474](#)
- 525. [475](#)
- 526. [476](#)
- 527. [477](#)
- 528. [478](#)
- 529. [479](#)
- 530. [480](#)
- 531. [481](#)
- 532. [482](#)
- 533. [483](#)
- 534. [484](#)
- 535. [485](#)
- 536. [486](#)

- 537. [487](#)
- 538. [488](#)
- 539. [489](#)
- 540. [490](#)
- 541. [491](#)
- 542. [492](#)
- 543. [493](#)
- 544. [494](#)
- 545. [495](#)
- 546. [496](#)
- 547. [497](#)
- 548. [498](#)
- 549. [499](#)
- 550. [500](#)
- 551. [501](#)
- 552. [502](#)
- 553. [503](#)
- 554. [504](#)
- 555. [505](#)
- 556. [506](#)
- 557. [507](#)
- 558. [508](#)
- 559. [509](#)
- 560. [510](#)
- 561. [511](#)
- 562. [512](#)
- 563. [513](#)

- 564. [514](#)
- 565. [515](#)
- 566. [516](#)
- 567. [517](#)
- 568. [518](#)
- 569. [519](#)
- 570. [520](#)
- 571. [521](#)
- 572. [522](#)
- 573. [523](#)
- 574. [524](#)
- 575. [525](#)
- 576. [526](#)
- 577. [527](#)
- 578. [528](#)
- 579. [529](#)
- 580. [530](#)
- 581. [531](#)
- 582. [532](#)
- 583. [533](#)
- 584. [534](#)
- 585. [535](#)
- 586. [536](#)
- 587. [537](#)
- 588. [538](#)
- 589. [539](#)
- 590. [540](#)

- 591. [541](#)
- 592. [542](#)
- 593. [543](#)
- 594. [544](#)
- 595. [545](#)
- 596. [546](#)
- 597. [547](#)
- 598. [548](#)
- 599. [549](#)
- 600. [550](#)
- 601. [551](#)
- 602. [552](#)
- 603. [553](#)
- 604. [554](#)
- 605. [555](#)
- 606. [556](#)
- 607. [557](#)
- 608. [558](#)
- 609. [559](#)
- 610. [560](#)
- 611. [561](#)
- 612. [562](#)
- 613. [563](#)
- 614. [564](#)
- 615. [565](#)
- 616. [566](#)
- 617. [567](#)

- 618. [568](#)
- 619. [569](#)
- 620. [570](#)
- 621. [571](#)
- 622. [572](#)
- 623. [573](#)
- 624. [574](#)
- 625. [575](#)
- 626. [576](#)
- 627. [577](#)
- 628. [578](#)
- 629. [579](#)
- 630. [580](#)
- 631. [581](#)
- 632. [582](#)
- 633. [583](#)
- 634. [584](#)
- 635. [585](#)
- 636. [586](#)
- 637. [587](#)
- 638. [588](#)
- 639. [589](#)
- 640. [590](#)
- 641. [591](#)
- 642. [592](#)
- 643. [593](#)
- 644. [594](#)

- 645. [595](#)
- 646. [596](#)
- 647. [597](#)
- 648. [598](#)
- 649. [599](#)
- 650. [600](#)
- 651. [601](#)
- 652. [602](#)
- 653. [603](#)
- 654. [604](#)
- 655. [605](#)
- 656. [606](#)
- 657. [607](#)
- 658. [608](#)
- 659. [609](#)
- 660. [610](#)
- 661. [611](#)
- 662. [612](#)
- 663. [613](#)
- 664. [614](#)
- 665. [615](#)
- 666. [616](#)
- 667. [617](#)
- 668. [618](#)
- 669. [619](#)
- 670. [620](#)
- 671. [621](#)



672. [622](#)  
673. [623](#)  
674. [624](#)  
675. [625](#)  
676. [626](#)  
677. [627](#)  
678. [628](#)  
679. [629](#)  
680. [630](#)  
681. [631](#)  
682. [632](#)  
683. [633](#)  
684. [634](#)  
685. [635](#)  
686. [636](#)  
687. [637](#)  
688. [638](#)  
689. [639](#)  
690. [640](#)  
691. [641](#)  
692. [642](#)  
693. [643](#)  
694. [644](#)  
695. [645](#)  
696. [646](#)  
697. [647](#)  
698. [648](#)

- 699. 649
- 700. 650
- 701. 651
- 702. 652
- 703. 653
- 704. 654
- 705. 655
- 706. 656
- 707. 657
- 708. 658
- 709. 659
- 710. 660
- 711. 661
- 712. 662
- 713. 663
- 714. 664
- 715. 665
- 716. 666
- 717. 667
- 718. 668
- 719. 669
- 720. 670
- 721. B-1
- 722. B-2
- 723. B-3
- 724. B-4
- 725. B-5

- 726. [B-6](#)
- 727. [B-7](#)
- 728. [B-8](#)
- 729. [B-9](#)
- 730. [B-10](#)
- 731. [B-11](#)
- 732. [B-12](#)
- 733. [B-13](#)
- 734. [B-14](#)
- 735. [B-15](#)
- 736. [B-16](#)
- 737. [B-17](#)
- 738. [B-18](#)
- 739. [B-19](#)
- 740. [B-20](#)
- 741. [B-21](#)
- 742. [B-22](#)
- 743. [B-23](#)
- 744. [B-24](#)
- 745. [B-25](#)
- 746. [B-26](#)
- 747. [B-27](#)
- 748. [B-28](#)
- 749. [B-29](#)
- 750. [B-30](#)
- 751. [B-31](#)
- 752. [B-32](#)

- 753. B-33
- 754. B-34
- 755. B-35
- 756. B-36
- 757. B-37
- 758. B-38
- 759. B-39
- 760. B-40
- 761. B-41
- 762. B-42
- 763. B-43
- 764. B-44
- 765. C-1
- 766. C-2
- 767. C-3
- 768. C-4
- 769. C-5
- 770. C-6
- 771. C-7
- 772. C-8
- 773. C-9
- 774. C-10
- 775. C-11
- 776. C-12
- 777. C-13
- 778. C-14
- 779. C-15

- 780. C-16
- 781. C-17
- 782. C-18
- 783. C-19
- 784. C-20
- 785. C-21
- 786. C-22
- 787. C-23
- 788. C-24
- 789. C-25
- 790. C-26
- 791. C-27
- 792. C-28
- 793. C-29
- 794. C-30
- 795. C-31
- 796. C-32
- 797. C-33
- 798. C-34
- 799. C-35
- 800. C-36
- 801. C-37
- 802. C-38
- 803. C-39
- 804. C-40
- 805. C-41
- 806. C-42

807. C-43

808. C-44

809. C-44

810. G-1

811. G-2

812. G-3

813. G-4

814. G-5

815. G-6

816. G-7

817. G-8