

PEARSON IT
CERTIFICATION

vmware®

Official Cert Guide

Advance your IT career with hands-on learning

VCP-DCV for vSphere 7.x

Exam 2V0-21.20



Practice
Tests



Flash
Cards



Review
Exercises



Study
Planner



JOHN A. DAVIS
STEVE BACA
OWEN THOMAS

Table of Contents

Foreword xxv

Introduction xxvi

Part I: vSphere Architecture, Integration, and Requirements

Chapter 1 vSphere Overview, Components, and Requirements 3

“Do I Know This Already?” Quiz 3

Foundation Topics 6

vSphere Components and Editions 6

 vSphere Components 6

 Editions and Licenses 8

vCenter Server Topology 10

 Single Sign-On (SSO) Domain 12

 Enhanced Linked Mode 12

 vCenter HA 13

Infrastructure Requirements 14

 Compute and System Requirements 14

 Storage Requirements 16

 Network Requirements 17

 Infrastructure Services 21

Other Requirements 23

 Additional Requirements 23

 vSphere Replication Requirements 24

 vCenter High Availability Requirements 24

 SDDC Requirements 25

VMware Cloud vs. VMware Virtualization 27

 Server Virtualization 27

 VMware SDDC 27

 vCloud Suite and Private Clouds 28

 VCF and Hybrid Clouds 28

 VMC on AWS 28

 VMware vCloud Director 28

 Cloud Automation 28

Exam Preparation Tasks 29

Review All the Key Topics 29

Complete Tables and Lists from Memory 29

Define Key Terms 29

Answer Review Questions 30

Chapter 2 Storage Infrastructure 33

“Do I Know This Already?” Quiz 33

Foundation Topics 36

Storage Models and Datastore Types 36

How Virtual Machines Access Storage 36

Storage Virtualization: The Traditional Model 36

Software-Defined Storage Models 40

Datastore Types 41

Storage in vSphere with Kubernetes 45

VMware NVMe 46

vSAN Concepts 49

vSAN Characteristics 50

vSAN Terminology 51

What Is New in vSAN 7.0 53

vSAN Deployment Options 54

vSAN Limitations 59

vSAN Space Efficiency 59

vSAN Encryption 61

vSAN File Services 62

vSAN Requirements 63

Other vSAN Considerations 67

vSphere Storage Integration 68

VASA 69

VAAI 70

Virtual Volumes (VVols) 72

Storage Multipathing and Failover 73

Multipathing Overview 74

Pluggable Storage Architecture (PSA) 74

Storage Policies 78

Storage Policy Based Management (SPBM) 79

Virtual Disk Types 79

vSAN-Specific Storage Policies 79

Storage DRS (SDRS) 81

Initial Placement and Ongoing Balancing 81

Space Utilization Load Balancing 81

I/O Latency Load Balancing 81

SDRS Automation Level 82

SDRS Thresholds and Behavior 82

SDRS Recommendations 83

Anti-affinity Rules 83

Datastore Cluster Requirements 83

NIOC, SIOC, and SDRS 84

Exam Preparation Tasks 85

- Review All Key Topics 85
- Complete Tables and Lists from Memory 85
- Define Key Terms 85
- Review Questions 86

Chapter 3 Network Infrastructure 89

- “Do I Know This Already?” Quiz 89

Foundation Topics 92

- Networking Terms and Concepts 92
 - Traditional Networking Terminology 92
 - Virtual NICs 93
 - Virtual Switch Concepts 94
 - VLANs 94
- vSphere Standard Switch (vSS) 95
 - MTU 97
 - vSS Network Policies 98
 - NIC Teaming Policies 98
 - Network Security Policies 100
 - Traffic Shaping Policy 101
 - VLAN Policies 101
- Distributed Virtual Switch (vDS) 102
 - Distributed Port Groups 103
 - Uplink Port Groups 103
 - vSS and vDS Comparison 103
 - vDS Network Policies 104
 - Inbound Traffic Shaping 105
 - Port-Blocking Policies 105
 - Load-Based NIC Teaming 105
 - Resource Allocation Policy 105
 - NetFlow and Monitoring Policy 108
 - Traffic Filtering and Marking Policy 109
- vDS Settings and Features 110
 - Private VLANs 110
 - Data Center–Level Management 111
 - Port State Monitoring 111
 - Port State with vMotion 111
 - Port Mirroring 111
 - Port Binding and Allocation 112
 - LACP Support 113
 - vDS Health Check 115

Other vSphere Networking Features	116
Multicast Filtering Mode	116
Discovery Protocol	117
TCP Segmentation Offload	118
DirectPath I/O	118
Single Root I/O Virtualization (SR-IOV)	119
VMkernel Networking and TCP/IP Stacks	121

Exam Preparation Tasks 123

Review All Key Topics	123
Complete Tables and Lists from Memory	123
Define Key Terms	123
Review Questions	124

Chapter 4 Clusters and High Availability 127

“Do I Know This Already?” Quiz	127
--------------------------------	-----

Foundation Topics 130

Cluster Concepts and Overview	130
Enhanced vMotion Compatibility (EVC)	131
vSAN Services	134
Distributed Resource Scheduler (DRS)	134
Recent DRS Enhancements	134
DRS Rules	137
DRS Migration Sensitivity	138
Resource Pools	139
vSphere High Availability (HA)	143
vSphere HA Requirements	145
vSphere HA Response to Failures	145
Heartbeats	146
vSphere HA Admission Control	146
vSphere HA Advanced Options	148
Virtual Machine Settings	149
VM Component Protection (VMCP)	150
Virtual Machine and Application Monitoring	150
vSphere HA Best Practices	151
Proactive HA	151
Other Resource Management and Availability Features	151
Predictive DRS	152
Distributed Power Management (DPM)	152
Fault Tolerance (FT)	153
vCenter Server High Availability	157
VMware Service Lifecycle Manager	157

Exam Preparation Tasks 158

- Review All Key Topics 158
- Complete Tables and Lists from Memory 158
- Define Key Terms 158
- Review Questions 159

Chapter 5 vCenter Server Features and Virtual Machines 163

- “Do I Know This Already?” Quiz 163

Foundation Topics 166

- vCenter Server and vSphere 166
 - vSphere Managed Inventory Objects 166
 - Host Profiles 170
 - Content Libraries 171
 - vSphere with Tanzu 173
- Virtual Machine File Structure 173
 - Configuration File 174
 - Virtual Disk Files 175
 - Snapshot Files 175
- Virtual Machine Snapshots 175
 - Snapshot Use Cases 177
 - What a Snapshot Preserves 177
 - Parent Snapshots 178
 - Snapshot Behavior 178
 - Limitations 179
- Virtual Machine Settings 180
 - VM Hardware/Compatibility 180
 - Virtual Disk Provisioning 183
 - VMware Tools 183
 - Virtual Machine Options 183
 - Virtual Machine Advanced Settings 184
- Virtual Machine Migration 185
 - Virtual Machine Migration 185
 - vMotion Details 189
 - Storage vMotion Details 192
- Virtual Machine Cloning 194
 - Clones 194
 - Rapid Provisioning with Templates 195
 - Instant Clones 195
- Exam Preparation Tasks 197**
 - Review All Key Topics 197
 - Complete Tables and Lists from Memory 197
 - Define Key Terms 198
 - Review Questions 198

Chapter 6 VMware Product Integration 201

“Do I Know This Already?” Quiz 201

Foundation Topics 204

vSphere Add-ons 204

vSphere with Tanzu 204

vCenter Converter 205

VMware vSphere Replication 206

VMware SkyLine 206

vRealize Suite 207

vRealize Operations Manager (vROps) 207

vRealize Log Insight (vRLI) 208

vRealize Automation (vRA) 209

vRealize Orchestrator (vRO) 213

vRealize Network Insight (vRNI) 214

Desktop and Application Virtualization 215

VMware Horizon 215

App Volumes 217

Replication and Disaster Recovery 219

vSphere Replication 219

Site Recovery Manager (SRM) 221

Private, Public, and Hybrid Clouds 222

VMware Cloud Foundation (VCF) 223

VMware Hybrid Cloud Extension (HCX) 224

VMware on AWS 226

Azure VMware Solution 226

Networking and Security 227

AppDefense 227

NSX 228

Exam Preparation Tasks 230

Review All Key Topics 230

Complete Tables and Lists from Memory 230

Define Key Terms 230

Review Questions 231

Chapter 7 vSphere Security 233

“Do I Know This Already?” Quiz 233

Foundation Topics 236

vSphere Certificates 236

vSphere Certificates Overview 236

Certificate Requirements 238

ESXi Host Certificates 241

vSphere Permissions	242
Authentication and Authorization	242
Inventory Hierarchy and Objects	243
Privileges and Roles	244
Permissions	246
Global Permissions	247
Best Practices for Roles and Permissions	248
Required Privileges for Common Tasks	248
How Permissions Are Applied by vCenter Server	251
ESXi and vCenter Server Security	253
Built-in Security Features	254
Security Profiles	254
ESXi Password Hardening	256
Joining an ESXi Host to a Directory Service	257
vSphere Authentication Proxy	257
ESXi Host Access	257
Control MOB Access	257
ESXi Secure Boot and TPM	258
vSphere Trust Authority (vTA)	258
vCenter Server Security	259
vSphere Network Security	262
Virtual Machine Security	265
Virtual Machine Hardening Best Practices	265
Configuring UEFI Boot	266
Disabling Unexposed Features	266
Other Common Settings	267
Virtual Machine Risk Profiles	268
Protecting Virtual Machines Against Denial-of-Service Attacks	269
Controlling VM Device Connections	269
Virtual Machine Encryption	270
Encrypted vSphere vMotion	272
virtual Trusted Platform Module (vTPM)	273
virtual Intel Software Guard Extension (vSGX)	274
Available Add-on Security	275
Compliance Using vRealize Operations Manager	275
VMware NSX	276
AppDefense	277
Exam Preparation Tasks	279
Review All the Key Topics	279
Complete Tables and Lists from Memory	279
Define Key Terms	280
Review Questions	280

Part II: vSphere Installation/Configuration

Chapter 8 vSphere Installation 283

“Do I Know This Already?” Quiz 283

Foundation Topics 286

Installing ESXi Hosts 286

 Installing ESXi Interactively 286

 Scripted ESXi Installation 288

 Using Auto Deploy 292

Deploying vCenter Server Components 297

 vCenter Server Database 297

 Platform Services Controller (PSC) 297

 vCenter Server Appliance 298

 Configuring and Managing VMware Certificate Authority (VMCA) 303

Configuring Single Sign-On (SSO) 305

 SSO and Identity Sources Overview 305

 Adding, Editing, and Removing SSO Identity Sources 306

 Adding an Active Directory Identity Source 307

 Adding an LDAP Authentication Source 309

 Enabling and Disabling Single Sign-On (SSO) Users 310

 Configuring SSO Policies 311

 Configuring Identity Federation 313

Initial vSphere Configuration 315

 Implementing vSphere Client 315

 Implementing VMware vSphere Lifecycle Manager 315

 Configuring the vCenter Server Inventory 315

 Implementing vCenter HA 316

 Using Host Profiles 317

 VMware Tools 320

 Advanced ESXi Host Options 321

Exam Preparation Tasks 323

Review All the Key Topics 323

Complete Tables and Lists from Memory 323

Define Key Terms 323

Review Questions 324

Chapter 9 Configuring and Managing Virtual Networks 327

“Do I Know This Already?” Quiz 327

Foundation Topics 330

vSphere Standard Switches (vSS) 330

 Creating and Configuring vSphere Standard Switches 330

 Creating and Configuring Standard Port Groups 332

vSphere Distributed Switches (vDS)	334
Creating and Configuring vSphere Distributed Switches	334
Creating and Configuring Distributed Port Groups	337
VMkernel Networking	338
Configuring and Managing VMkernel Adapters	338
Configuring TCP/IP Stacks	339
Configuring and Managing Networking Features	340
Configuring Network I/O Control (NIOC)	340
Creating a Network Resource Pool	341
Using Private VLANs	342
Using DirectPath I/O	343
Single Root I/O Virtualization (SR-IOV)	343
Configuring and Managing Port Mirroring	345
Configuring and Managing Link Aggregation Groups (LAGs)	346
Managing Host Networking with vDS	350
Adding Hosts to a vDS	350
Managing Host Physical Network Adapters on a vDS	351
Migrating VMkernel Network Adapters to a vDS	352
Removing Hosts from a vDS	352
Migrating Virtual Machines to a vDS	353
Monitoring the State of Ports in a Distributed Port Group	353
Using the vDS Health Check	354
Networking Policies and Advanced Features	355
Exam Preparation Tasks	357
Review All the Key Topics	357
Complete Tables and Lists from Memory	357
Define Key Terms	357
Review Questions	358

Part III: vSphere Management and Optimization

Chapter 10 Managing and Monitoring Clusters and Resources 361

“Do I Know This Already?” Quiz	361
Foundation Topics	364
Creating and Configuring a vSphere Cluster	364
Creating a Cluster	364
Configuring a Cluster with Quickstart	365
EVC Mode	367
Creating and Configuring a vSphere DRS Cluster	368
Creating a vSphere DRS Cluster	368
Creating a Resource Pool	368
Configuring Advanced DRS Options	369

Creating and Configuring a vSphere HA Cluster	370
Creating a vSphere HA Cluster	370
Configuring Advanced vSphere HA Options	370
Configuring vSphere HA Admission Control	371
Configuring VMCP	371
Configuring Virtual Machine and Application Monitoring	372
Configuring Proactive HA	372
Configuring vSphere Fault Tolerance	373
Monitoring and Managing vSphere Resources	373
Metrics	374
vSphere Client Performance Charts	375
Troubleshooting and Optimizing Performance	379
Monitoring and Managing Cluster Resources	384
Monitoring and Managing Resource Pool Resources	385
Monitoring and Managing Host Resources and Health	386
Monitoring and Managing Virtual Machine Resources	388
ESXTOP	393
VIMTOP	396
vCenter Server Appliance Management Interface (VAMI)	396
Events, Alarms, and Automated Actions	396
Events	396
Viewing Events in the vSphere Client	397
Viewing the System Event Log	397
Streaming Events to a Remote Syslog Server	398
Alarms	399
Viewing and Acknowledging Triggered Alarms	399
Creating Alarm Definitions	400
Alarm Actions	401
Advanced Use Cases for Alarms	401
Logging in vSphere	401
ESXi Logs	402
vCenter Server Logs	404
Uploading System Logs to VMware	404
Log Levels	404
Configuring Syslog on ESXi Hosts	405
vRealize Log Insight (vRLI)	407
Exam Preparation Tasks	408
Review All the Key Topics	408
Complete Tables and Lists from Memory	408
Define Key Terms	408
Review Questions	409

Chapter 11 Managing Storage 411

“Do I Know This Already?” Quiz 411

Foundation Topics 414

Configuring and Managing vSAN 414

Preparing for vSAN 414

Creating a vSAN Cluster with Quickstart 415

Manually Enabling vSAN 416

Editing vSAN Settings 417

Licensing vSAN 418

Viewing a vSAN Datastore 418

Configuring vSAN and vSphere HA 419

Disabling vSAN 421

Shutting Down and Restarting vSAN 421

Deploying vSAN with vCenter Server 422

Expanding a vSAN Cluster 422

Working with Maintenance Mode 424

Managing vSAN Fault Domains 426

Extending a vSAN Datastore Across Two Sites 427

Managing Devices in a vSAN Cluster 429

Increasing Space Efficiency in a vSAN Cluster 430

Using Encryption in a vSAN Cluster 432

Using vSAN Policies 435

Viewing vSAN Storage Providers 436

Using vSAN File Service 436

Managing Datastores 438

Managing VMFS Datastores 438

Managing Raw Device Mappings (RDMs) 443

Managing NFS Datastores 444

Storage DRS and SIOC 447

Configuring and Managing Storage DRS 447

Configuring and Managing SIOC 449

NVMe and PMem 451

Managing VMware NVMe 451

Managing PMem 454

Multipathing, Storage Policies, and vVols 456

Managing Multipathing 456

Managing Storage Policies 459

Configuring and Managing vVols 463

Exam Preparation Tasks 465

Review All the Key Topics 465

Complete Tables and Lists from Memory 465

Define Key Terms	465
Review Questions	466

Chapter 12 Managing vSphere Security 469

“Do I Know This Already?” Quiz	469
Foundation Topics	472
Configuring and Managing Authentication and Authorization	472
Managing SSO	472
Users and Groups	474
Privileges and Roles	475
Permissions	475
Global Permissions	476
Editing Permissions	476
Configuring and Managing vSphere Certificates	477
Managing vSphere Client Certificates	477
Using Custom Certificates	478
Managing ESXi Certificates	479
General ESXi Security Recommendations	481
Configuring ESXi Using Host Profiles	482
Using Scripts to Manage Host Configuration Settings	483
ESXi Passwords and Account Lockout	485
SSH and ESXi Shell Security	487
PCI and PCIe Devices and ESXi	489
Disabling the Managed Object Browser	490
ESXi Networking Security Recommendations	490
ESXi Web Proxy Settings	490
vSphere Auto Deploy Security Considerations	491
Controlling CIM Access	491
Configuring and Managing ESXi Security	492
Configuring the ESXi Firewall	492
Customizing ESXi Services	493
Using Lockdown Mode	494
Managing the Acceptance Levels of Hosts and VIBs	496
Assigning Privileges for ESXi Hosts	496
Using Active Directory to Manage ESXi Users	497
Configuring vSphere Authentication Proxy	498
Configuring Smart Card Authentication for ESXi	499
Configuring UEFI Secure Boot for ESXi Hosts	499
Securing ESXi Hosts with Trusted Platform Module	500
Securing ESXi Log Files	501

Additional Security Management	501
Key Management Server	502
Changing Permission Validation Settings	502
Configuring and Managing vSphere Trust Authority (vTA)	502
Securing Virtual Machines with Intel Software Guard Extensions (SGX)	505
Encrypting a Virtual Machine	505

Exam Preparation Tasks 507

Review All the Key Topics	507
Complete Tables and Lists from Memory	507
Define Key Terms	507
Review Questions	508

Chapter 13 Managing vSphere and vCenter Server 511

“Do I Know This Already?” Quiz	511
--------------------------------	-----

Foundation Topics 514

vCenter Server Backup	514
Upgrading to vSphere 7.0	517
vCenter Server Data Transfer	519
Upgrading vCenter Server Appliance	519
Migrating vCenter Server for Windows to vCenter Server Appliance	522
Upgrading ESXi and Virtual Machines	524
Using Update Planner	524
Using vSphere Lifecycle Manager	526
About VMware Update Manager	529
VMware Update Manager Download Service (UMDS)	529
Baselines and Images	530
ESXi Quick Boot	535
ESXi Firmware Updates	536
Hardware Compatibility Checks	537
Exporting and Importing Cluster Images	538
Backup and Restore Scenarios	538
Upgrading Virtual Machines	539
Managing ESXi Hosts	540
Monitoring and Managing vCenter Server	542
Monitoring and Managing vCenter Server with the VAMI	543
Monitoring and Managing vCenter Server with the vSphere Client	547
Updating the vCenter Server	554
Managing a vCenter HA Cluster	557
Repointing a vCenter Server to Another Domain	558

Exam Preparation Tasks 561

Review All the Key Topics	561
---------------------------	-----

Complete Tables and Lists from Memory	561
Define Key Terms	562
Review Questions	562

Chapter 14 Managing Virtual Machines 565

“Do I Know This Already?” Quiz	565
Foundation Topics	568
Creating and Configuring Virtual Machines	568
Creating a New Virtual Machine	568
Powering on a VM	569
Opening a Console to a VM	569
Installing and Upgrading VMware Tools	570
Shutting Down a Guest	572
Cloning a Virtual Machine	572
Converting Between a VM and a Template	573
Deploying a Virtual Machine from a Template	574
Customizing the Guest OS	574
Deploying OVF/OVA Templates	577
Managing Virtual Machines	578
Configuring Virtual Machine Hardware	578
Editing Virtual Machine Options	583
Configuring Guest User Mappings	585
Editing OVF Details	585
Creating and Managing Virtual Machine Snapshots	586
Migrating Virtual Machines	587
Advanced Virtual Machine Management	589
Managing OVF Templates	589
Virtualization-Based Security	590
Managing VMs by Using PowerCLI	590
Configuring VMs to Support vGPUs	592
Content Libraries	594
Introduction to Content Libraries	594
Creating a Content Library	595
Publishing a Content Library	596
Subscribing to a Content Library	596
Content Library Permissions	597
Content Library Synchronization Options	598
Adding Items to a Content Library	598
Deploying VMs by Using a Content Library	599
Exam Preparation Tasks	600
Review All the Key Topics	600
Complete Tables and Lists from Memory	600

Define Key Terms 600

Review Questions 601

Chapter 15 Final Preparation 603

Getting Ready 603

Taking the Exam 604

Appendix A Answers to the “Do I Know This Already?” Quizzes and Review Questions 607

Index 627

Online Elements:

Appendix B Memory Tables

Appendix C Memory Table Answer Key

Appendix D Study Planner

Glossary

Sample pages

Clusters and High Availability

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
Other Resource Management and Availability Features	8–10

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.
L2	Intel Nehalem	Includes the Intel Penryn feature set and exposes additional CPU features, including SSE4.2 and POPCOUNT.
L3	Intel Westmere	Includes the Intel Nehalem feature set and exposes additional CPU features, including AES and PCLMULQDQ.
L4	Intel Sandy Bridge	Includes the Intel Westmere feature set and exposes additional CPU features, including AVX and XSAVE.
L5	Intel Ivy Bridge	Includes the Intel Sandy Bridge feature set and exposes additional CPU features, including RDRAND, ENFSTRG, FSGSBASE, SMEP, and F16C.
L6	Intel Haswell	Includes the Intel Ivy Bridge feature set and exposes additional CPU features, including ABMX2, AVX2, MOVBE, FMA, PERMD, RORX/MULX, INVPCID, and VMFUNC.
L7	Intel Broadwell	Includes the Intel Haswell feature set and exposes additional CPU features, including Transactional Synchronization Extensions, Supervisor Mode Access Prevention, Multi-Precision Add-Carry Instruction Extensions, PREFETCHW, and RDSEED.

Level	EVC Mode	Description
L8	Intel Skylake	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.
L9	Intel Cascade Lake	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

Level	EVC Mode	Description
A0	AMD Opteron Generation 1	Smallest AMD feature set for EVC mode.
A1	AMD Opteron Generation 2	Includes the AMD Generation 1 feature set and exposes additional CPU features, including CPMXCHG16B and RDTSCP.
A3	AMD Opteron Generation 3	Includes the AMD Generation 2 feature set and exposes additional CPU features, including SSE4A, MisAlignSSE, POPCOUNT, and ABM (LZCNT).
A2, B0	AMD Opteron Generation 3 (without 3DNow!)	Includes the AMD Generation 3 feature set without 3DNow support.
B1	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.
B2	AMD Opteron Piledriver	Includes the AMD Generation 4 feature set and exposes additional CPU features, including FMA, TBM, BMI1, and F16C.
B3	AMD Opteron Steamroller	Includes the AMD Piledriver feature set and exposes additional CPU features, including XSAVEOPT RDFSBASE, RDGSBASE, WRFSBASE, WRGSBAS, and FSGSBASE.
B4	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.
B5	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

Key Topic

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

Option	Description
Shares	<p>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. Shares can be thought of as priority under contention.</p> <p>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).</p> <p>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.</p> <p>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%.</p>
Reservations	<p>A reservation specifies the guaranteed minimum allocation for a virtual machine or a resource pool. A CPU reservation is expressed in megahertz, and a memory reservation is expressed in megabytes. You can power on a virtual machine only if there are enough unreserved resources to satisfy the reservation of the virtual machine. If the virtual machine starts, then it is guaranteed that amount, even when the physical server is heavily loaded.</p> <p>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.</p> <p>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.</p>

Option	Description
Expandable reservations	<p>You can enable expandable reservations to effectively allow a child resource pool to borrow from its parent. Expandable reservations, which are enabled by default, are considered during admission control. When powering on a virtual machine, if the resource pool does not have sufficient unreserved resources, the resource pool can use resources from its parent or ancestors.</p> <p>For example, say that in a resource pool where 8 GHz is reserved and expandable reservations is disabled, you try to start nine virtual machines each with 1 GHz, but the last virtual machine does not start. If you enable expandable reservation in the resource pool, and its parent pool (or cluster) has sufficient unreserved CPU resources, you can start the ninth virtual machine.</p>
Limits	<p>A limit specifies an upper bound for CPU or memory resources that can be allocated to a virtual machine or a resource pool.</p> <p>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.</p>

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

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

For example, the share values for a resource pool configured with normal CPU shares and high memory shares are 4000 (that is, 4×1000) 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

**Key
Topic**

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.
- 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.




FIGURE 4-2 vSphere HA Host Failover

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 > Default VM Restart Priority	You can indicate the order in which virtual machines are restarted when the host fails (higher priority machines first).

Option	Description
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.