

**VMware vSphere:
Install, Configure, Manage**
Lecture Manual
ESXi 6.5 and vCenter Server 6.5



VMware® Education Services
VMware, Inc.
www.vmware.com/education

VMware vSphere:

Install, Configure, Manage

ESXi 6.5 and vCenter Server 6.5

Part Number EDU-EN-VSICM65-LECT

Lecture Manual

Copyright © 2017 VMware, Inc. All rights reserved. This manual and its accompanying materials are protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

The training material is provided “as is,” and all express or implied conditions, representations, and warranties, including any implied warranty of merchantability, fitness for a particular purpose or noninfringement, are disclaimed, even if VMware, Inc., has been advised of the possibility of such claims. This training material is designed to support an instructor-led training course and is intended to be used for reference purposes in conjunction with the instructor-led training course. The training material is not a standalone training tool. Use of the training material for self-study without class attendance is not recommended.

These materials and the computer programs to which it relates are the property of, and embody trade secrets and confidential information proprietary to, VMware, Inc., and may not be reproduced, copied, disclosed, transferred, adapted or modified without the express written approval of VMware, Inc.

CONTENTS

MODULE 1	Course Introduction	1
	Importance	2
	Learner Objectives (1)	3
	Learner Objectives (2)	4
	You Are Here	5
	Typographical Conventions	6
	References	7
	VMware Online Resources	8
	VCP-Foundations Certification Alignment	9
	VMware Education Overview	10
	VMware Certification Overview	11
MODULE 2	Introduction to vSphere and the Software-Defined Data Center	13
	You Are Here	14
	Importance	15
	Module Lessons	16
	Lesson 1: Overview of vSphere and the Software-Defined Data Center	17
	Learner Objectives	18
	Traditional Architecture	19
	Virtual Architecture	20
	About Virtual Machines	21
	Benefits of Using Virtual Machines	22
	Types of Virtualization	24
	About the Software-Defined Data Center	25
	vSphere and Cloud Computing	26
	Review of Learner Objectives	28
	Lesson 2: Abstraction: Overview of a Virtual Machine	29
	Learner Objectives	30
	Virtual Machine: Guest and Consumer of ESXi Host	31
	Physical and Virtual Architecture	32
	Physical Resource Sharing	33
	CPU Virtualization	34
	Physical and Virtualized Host Memory Usage	35
	Physical and Virtual Networking	36
	Physical File Systems and VMFS	38
	Encapsulation	39
	vSphere Clients	40
	vSphere Web Client	41
	vSphere Client	42
	VMware Host Client	43
	Review of Learner Objectives	44
	Lesson 3: Shared Resources: Overview of ESXi	45

Learner Objectives	46
About ESXi Hosts	47
Physical and Virtual Architecture	49
Configuring an ESXi Host.	50
Configuring an ESXi Host: Root Access	51
Configuring an ESXi Host: Management Network.	52
Configuring an ESXi Host: Other Settings	53
Remote Access Settings: Security Profile	54
Configuring Lockdown Mode	55
Strict Lockdown Mode	56
Managing User Accounts: Best Practices	57
ESXi Host as an NTP Client	58
Labs	59
Lab 1: Installing ESXi	60
Lab 2: Configuring ESXi Hosts	61
Review of Learner Objectives.	62
Lesson 4: Centralized Management: Overview of vCenter Server.	63
Learner Objectives.	64
About the vCenter Server Management Platform	65
vCenter Server Architecture	66
Additional vCenter Server Services and Interfaces.	67
Platform Services Controller.	68
vCenter Server Services	69
Monitoring the Health and Status of Services and Nodes Across vCenter Server Systems	70
ESXi and vCenter Server Communication	71
Review of Learner Objectives.	72
Key Points	73
MODULE 3	
Creating Virtual Machines	75
You Are Here	76
Importance.	77
Module Lessons.	78
Lesson 1: Virtual Machine Concepts	79
Learner Objectives.	80
About Virtual Machine Files.	81
About Virtual Machine Virtual Hardware.	83
Virtual Hardware Versions	85
About Virtual Hardware Version 13	86
About CPU and Memory.	87
About Virtual Storage	89
About Virtual Disks.	90
About Thick-Provisioned Virtual Disks	91
About Thin-Provisioned Virtual Disks	92

About Virtual Networks	94
About Virtual Network Adapters (1)	95
About Virtual Network Adapters (2)	97
About Miscellaneous Devices	99
About the Virtual Machine Console	100
Review of Learner Objectives	101
Lesson 2: Creating a Virtual Machine	102
Learner Objectives	103
About Provisioning Virtual Machines	104
Creating Virtual Machines with the New Virtual Machine Wizard (1)	106
Creating Virtual Machines with the New Virtual Machine Wizard (2)	107
New Virtual Machine Wizard Settings	108
Installing the Guest Operating System	109
Deploying OVF Templates	110
Removing a Virtual Machine	111
About VMware Tools	112
Managing VMware Tools	114
VMware Tools: Supported ISO Images	115
Troubleshooting OS Installation Failures in Virtual Machines	116
Troubleshooting a Failed VMware Tools Installation on a Guest Operating System	117
Lab 3: Deploying and Configuring a Virtual Machine	118
Review of Learner Objectives	119
Key Points	120
MODULE 4	
vCenter Server	121
You Are Here	122
Importance	123
Module Lessons	124
Lesson I: vCenter Server Architecture	125
Learner Objectives	126
Overview of vCenter Server Appliance (1)	127
Overview of vCenter Server Appliance (2)	128
About the vCenter Server Management Platform	129
vCenter Server Services	130
vCenter Server Appliance Architecture	131
vCenter Server Appliance Scalability	132
vCenter Server Deployment Options	133
Platform Services Controller Deployment Recommendations (1)	134
Platform Services Controller Deployment Recommendations (2)	135
Platform Services Controller Deployment Recommendations (3)	136

vCenter Server APIs	137
High Availability for vCenter Server Appliance.....	138
Review of Learner Objectives.....	139
Lesson 2: Deploying, Backing Up, and Restoring vCenter Server Appliance	140
Learner Objectives.....	141
Preparing for vCenter Server Appliance Deployment (1).....	142
Preparing for vCenter Server Appliance Deployment (2).....	143
vCenter Server Appliance Native UI Installer.....	144
vCenter Server Install, Upgrade, Migrate, and Restore.....	145
vCenter Server Appliance Deployment.....	146
vCenter Server Appliance Two-Stage Deployment	147
vCenter Server Deployment Wizard	148
Getting Started with vCenter Server	149
Adding License Keys to vCenter Server.....	150
Configuring vCenter Server Settings.....	151
Logging In to the vCenter Server Appliance Management UI	152
vCenter Server Appliance Management Home.....	153
Native vCenter Server Backup and Restore	154
Lab 4: Working with vCenter Server	155
Review of Learner Objectives.....	156
Lesson 3: vSphere Clients.....	157
Learner Objectives.....	158
Accessing vSphere Clients	159
vSphere Web Client Home Page.....	160
Using the vSphere Web Client Navigator	161
vCenter Server Views: Hosts and Clusters, VMs, and Templates	162
vCenter Server Views: Storage and Networks	163
Viewing Object Information	164
Using Quick Filters in vSphere Web Client	165
Using Drag-and-Drop Functionality in vSphere Web Client	166
Using Pin and Unpin Functionality.....	167
Lab 5: Navigating the vSphere Clients	168
Review of Learner Objectives.....	169
Lesson 4: Managing the vCenter Server Inventory.....	170
Learner Objectives.....	171
About Data Center Objects	172
Organizing Inventory Objects into Folders.....	173
Adding a Data Center and Organizational Objects to vCenter Server	175
Adding ESXi Hosts to vCenter Server	176
Creating Custom Tags for Inventory Objects	177
vCenter Server Events.....	178

vCenter Server System Logs	179
Output vCenter Server Events and Logs to Syslog Collector	180
vCenter Server Database Health	181
Managing the vCenter Server Services	182
Lab 6: Creating Folders in vCenter Server Appliance	183
Review of Learner Objectives	184
Lesson 5: vCenter Server Roles and Permissions	185
Learner Objectives	186
Access Control Overview	187
vCenter Server Permissions	188
Adding Permissions to the vCenter Server Inventory	189
About Roles	191
About Objects	193
About Global Permissions	194
Assigning Permissions	195
Viewing Roles and Assignments	196
Applying Permissions: Scenario 1	197
Applying Permissions: Scenario 2	198
Applying Permissions: Scenario 3	199
Applying Permissions: Scenario 4	200
Creating a Role	201
Review of Learner Objectives	202
Key Points	203

MODULE 5

Configuring and Managing Virtual Networks	205
You Are Here	206
Importance	207
Module Lessons	208
Lesson 1: Introduction to vSphere Standard Switches	209
Learner Objectives	210
Types of Virtual Switch Connections	211
Adding ESXi Networking	212
Virtual Switch Connection Examples	213
Standard Switch Components	214
Viewing the Standard Switch Configuration	215
About VLANs	216
Network Adapter Properties	217
Types of Virtual Switches	218
Distributed Switch Architecture	219
Standard Switch and Distributed Switch Feature Comparison	220
Review of Learner Objectives	221
Lesson 2: Configuring Standard Switch Policies	222
Learner Objectives	223
Network Switch and Port Policies	224

Configuring Security Policies	225
Traffic-Shaping Policy	226
Configuring Traffic Shaping	227
NIC Teaming and Failover Policies	228
Load-Balancing Method: Originating Virtual Port ID	230
Load-Balancing Method: Source MAC Hash	231
Load-Balancing Method: Source and Destination IP Hash	232
Detecting and Handling Network Failure	233
Physical Network Considerations	235
Lab 7: Using Standard Switches	236
Review of Learner Objectives	237
Key Points	238

MODULE 6

Configuring and Managing Virtual Storage	239
You Are Here	240
Importance	241
Module Lessons	242
Lesson 1: Storage Concepts	243
Learner Objectives	244
Basic Storage Overview	245
Storage Protocol Overview	247
About Datastores	249
About VMFS6	250
About NFS	252
vSAN Overview	253
About vSphere Virtual Volumes	254
About Raw Device Mapping	255
Physical Storage Considerations	256
Review of Learner Objectives	257
Lesson 2: Fibre Channel Storage	258
Learner Objectives	259
About Fibre Channel	260
Fibre Channel SAN Components	261
Fibre Channel Addressing and Access Control	262
Multipathing with Fibre Channel	264
FCoE Adapters	266
Configuring Software FCoE: Creating a VMkernel Port	267
Configuring Software FCoE: Activating the Software FCoE Adapter	268
Multipathing with Software FCoE	269
Reviewing Learner Objectives	270
Lesson 3: iSCSI Storage	271
Learner Objectives	272
iSCSI Components	273

iSCSI Addressing	274
Storage Device Naming Conventions	275
iSCSI Adapters	276
Setting Up iSCSI Adapters	277
ESXi Network Configuration for IP Storage.....	278
Creating Datastores and Discovering iSCSI Targets.....	279
iSCSI Security: CHAP	280
Multipathing with iSCSI Storage	282
Lab 8: Accessing iSCSI Storage	284
Review of Learner Objectives.....	285
Lesson 4: VMFS Datastores	286
Learner Objectives.....	287
Using VMFS Datastores with ESXi Hosts	288
Creating and Viewing VMFS Datastores	289
Browsing Datastore Contents	290
Managing Overcommitted Datastores.....	291
Increasing the Size of a VMFS Datastore	292
Deleting or Unmounting a VMFS Datastore.....	293
Multipathing Algorithms.....	295
Configuring Storage Load Balancing	296
Lab 9: Managing VMFS Datastores	297
Review of Learner Objectives.....	298
Lesson 5: NFS Datastores	299
Learner Objectives.....	300
NFS Components.....	301
Configuring an NFS Datastore	302
NFS v3 and NFS v4.1	303
NFS Version Compatibility with Other vSphere Technologies	304
NFS Datastore Best Practices	306
NFS Datastore Name and Configuration.....	307
Configuring AD and NFS Servers to Use Kerberos	308
Configuring ESXi Host Authentication and NFS Kerberos Credentials	309
Considerations When Using Kerberos for NFS	310
Configuring a Datastore to Use Kerberos	311
Viewing IP Storage Information	312
Unmounting an NFS Datastore	313
Multipathing and NFS v4.1 Storage	314
Enabling Session Trunking and Multipathing.....	316
Lab 10: Accessing NFS Storage	317
Review of Learner Objectives.....	318
Lesson 6: vSAN Datastores.....	319
Learner Objectives.....	320

About vSAN	321
vSAN Requirements	323
Configuring a vSAN Datastore	325
Disk Groups	326
Viewing vSAN Cluster Summary	327
Using vSAN	328
Objects in vSAN Datastores	329
Virtual Machine Storage Policies	330
Configuring Virtual Machine Storage Policies	331
Viewing a Virtual Machine's vSAN Datastore	332
vSAN Cluster Member Maintenance Mode Options	333
Removing a Host from a vSAN Cluster	334
Review of Learner Objectives	335
Key Points	336

MODULE 7

Virtual Machine Management	337
You Are Here	338
Importance	339
Module Lessons	340
Lesson 1: Creating Templates and Clones	341
Learner Objectives	342
Using a Template	343
Creating a Template	344
Deploying a Virtual Machine from a Template	345
Updating a Template	346
Cloning a Virtual Machine	347
Customizing the Guest Operating System	348
Review of Learner Objectives	350
Lesson 2: Working with Content Libraries	351
Learner Objectives	352
About the Content Library	353
Adding Templates to a Content Library	354
Deploying VMs from Templates in a Content Library	355
Benefits of Content Libraries	356
Types of Content Library	358
Lab 11: Using Templates and Clones	359
Review of Learner Objectives	360
Lesson 3: Modifying Virtual Machines	361
Learner Objectives	362
Modifying Virtual Machine Settings	363
Hot-Pluggable Devices	364
Creating an RDM	365
Dynamically Increasing a Virtual Disk's Size	366

Inflating a Thin-Provisioned Disk	367
Virtual Machine Options	368
VMware Tools Options	369
Boot Options	370
Lab 12: Modifying Virtual Machines	371
Review of Learner Objectives	372
Lesson 4: Migrating Virtual Machines	373
Learner Objectives	374
Migrating Virtual Machines	375
Comparison of Migration Types	376
vSphere vMotion Migration	377
vSphere vMotion Migration Workflow	378
vSphere vMotion Migration Requirements	380
Host Requirements for vSphere vMotion Migration	381
CPU Constraints on vSphere vMotion Migration	382
Other Cluster Settings: EVC for vSphere vMotion	383
CPU Baselines for an EVC Cluster	384
EVC Cluster Requirements	385
Hiding or Exposing NX/XD	386
Identifying CPU Characteristics	387
Checking vSphere vMotion Errors	388
vSphere Storage vMotion in Action	389
vSphere Storage vMotion Guidelines and Limitations	390
Shared-Nothing vSphere vMotion Migration	391
Shared-Nothing vSphere vMotion Migration Considerations	392
Migration Between vCenter Server Instances	393
VMkernel Networking Layer and TCP/IP Stacks	394
vSphere vMotion TCP/IP Stacks	396
Long-Distance vSphere vMotion Migration	397
Networking Requirements for Long-Distance vSphere vMotion Migration	398
Network Checks for Migrations Between vCenter Server Instances	399
Encrypted vSphere vMotion	400
Lab 13: Migrating Virtual Machines	401
Review of Learner Objectives	402
Lesson 5: Creating Virtual Machine Snapshots	403
Learner Objectives	404
Virtual Machine Snapshots	405
Virtual Machine Snapshot Files	406
Taking a Snapshot	408
Managing Snapshots	409
Deleting a Virtual Machine Snapshot (1)	410
Deleting a Virtual Machine Snapshot (2)	411

Deleting a Virtual Machine Snapshot (3)	412
Deleting All Virtual Machine Snapshots.	413
About Snapshot Consolidation	414
Discovering When to Consolidate.	415
Performing Snapshot Consolidation	416
Lab 14: Managing Virtual Machines.	417
Review of Learner Objectives.	418
Key Points	419

MODULE 8

Resource Management and Monitoring	421
You Are Here	422
Importance.	423
Module Lessons.	424
Lesson 1: Virtual CPU and Memory Concepts.	425
Learner Objectives.	426
Memory Virtualization Basics.	427
Virtual Machine Memory Overcommitment.	428
Memory Reclamation Techniques.	429
Virtual SMP.	431
Hyperthreading	432
CPU Load Balancing.	433
Review of Learner Objectives.	434
Lesson 2: Resource Controls and Resource Pools	435
Learner Objectives.	436
Shares, Limits, and Reservations.	437
Defining Resource Settings for Individual VMs.	438
Resource Sharing by Virtual Machines.	439
About Resource Pools	440
Resource Pool Attributes.	441
Reasons to Use Resource Pools.	442
Resource Pool Example.	443
Resource Pools Example: CPU Shares	444
Resource Pools Example: CPU Contention	445
Expandable Reservation	446
Admission Control for CPU and Memory Reservations.	447
Resource Pool Summary Tab	448
Resource Reservation Tab.	449
Scheduling Changes to Resource Settings	450
Lab 15: Managing Resource Pools	451
Review of Learner Objectives.	452
Lesson 3: Creating vApps	453
Learner Objectives.	454
Managing Virtual Machines with a vApp	455
vApp Characteristics	456

Exporting and Deploying vApps	457
Lab 16: Managing vApps	458
Review of Learner Objectives	459
Lesson 4: Monitoring Resource Use	460
Learner Objectives	461
Performance-Tuning Methodology	462
Resource-Monitoring Tools	463
Guest Operating System Monitoring Tools	464
Using Perfmon to Monitor Virtual Machine Resources	465
Using esxtop to Monitor Virtual Machine Resources	466
About Monitoring Inventory Objects with Performance Charts	467
Working with Overview Performance Charts	468
Working with Advanced Performance Charts	469
Chart Options: Real-Time and Historical	470
Chart Types	471
Saving Charts	472
Objects and Counters	473
Statistics Type	474
Rollup	475
Setting Log Levels	477
Interpreting Data from the Tools	478
CPU-Constrained Virtual Machine	479
Memory-Constrained Virtual Machine	481
Memory-Constrained Host	482
Monitoring Active Memory of a Virtual Machine	483
Disk-Constrained Virtual Machines	484
Monitoring Disk Latency	485
Network-Constrained Virtual Machines	486
Lab 17: Monitoring Virtual Machine Performance	487
Review of Learner Objectives	488
Lesson 5: Using Alarms	489
Learner Objectives	490
About Alarms	491
Alarm Settings	492
Alarm Triggers	493
Configuring Condition Triggers	494
Configuring Event Triggers	495
Configuring Actions	496
Configuring vCenter Server Notifications	498
Viewing and Acknowledging Triggered Alarms	499
Lab 18: Using Alarms	500
Review of Learner Objectives	501
Key Points	502

MODULE 9

vSphere HA, vSphere Fault Tolerance, and Protecting Data	503
You Are Here	504
Importance.	505
Module Lessons.	506
Lesson 1: Introduction to vSphere HA	507
Learner Objectives.	508
Protection at Every Level	509
vCenter Server Availability: Recommendations.	511
About vSphere HA	512
vSphere HA Scenario: ESXi Host Failure.	513
vSphere HA Scenario: Guest Operating System Failure.	514
vSphere HA Scenario: Application Failure.	515
Importance of Redundant Heartbeat Networks.	516
Redundancy Using NIC Teaming	517
Redundancy Using Additional Networks	518
Review of Learner Objectives.	519
Lesson 2: vSphere HA Architecture	520
Learner Objectives.	521
vSphere HA Architecture: Agent Communication	522
vSphere HA Architecture: Network Heartbeats	524
vSphere HA Architecture: Datastore Heartbeats.	525
vSphere HA Failure Scenarios	526
Failed Slave Host.	527
Failed Master Host	528
Isolated Host	529
Design Considerations.	530
Virtual Machine Storage Failures	531
Virtual Machine Component Protection	532
Review of Learner Objectives.	533
Lesson 3: Configuring vSphere HA	534
Learner Objectives.	535
About Clusters.	536
vSphere HA Prerequisites	537
Configuring vSphere HA Settings.	538
vSphere HA Settings: Failure and Responses	539
vSphere HA Settings: Virtual Machine Monitoring	540
vSphere HA Settings: Heartbeat Datastores	541
vSphere HA Settings: Admission Control.	542
Example: Admission Control Using Cluster Resources Percentage.	543
Example: Admission Control Using Slots.	544
vSphere HA Settings: Performance Degradation VMs Tolerate.	545
vSphere HA Settings: Advanced Options	546


vSphere HA Orchestrated Restart	547
Configuring vSphere HA Orchestrated Restart	548
VM Dependencies in Orchestrated Restart	549
Network Configuration and Maintenance	550
Cluster Resource Reservation	551
Monitoring Cluster Status	552
Lab 19: Using vSphere HA	553
Review of Learner Objectives	554
Lesson 4: Introduction to vSphere Fault Tolerance	555
Learner Objectives	556
vSphere Fault Tolerance	557
vSphere Fault Tolerance Features	558
vSphere Fault Tolerance with vSphere HA and vSphere DRS	559
Redundant VMDKs	560
vSphere Fault Tolerance Checkpoint	561
vSphere Fault Tolerance: Precopy	562
vSphere Fault Tolerance Fast Checkpointing	563
Shared Files	564
shared.vmf File	565
Enabling vSphere Fault Tolerance on a Virtual Machine	566
Review of Learner Objectives	567
Lesson 5: vSphere Replication and vSphere Data Protection	568
Learner Objectives	569
About vSphere Replication	570
vSphere Replication Appliance	571
Replication Functions	572
Configuring vSphere Replication for a Single Virtual Machine	573
Configuring MPIT	574
Recovering Virtual Machines	576
About vSphere Data Protection	577
vSphere Data Protection Requirements and Architecture	578
Offloaded Backup Processing	580
Changed Block Tracking	581
Data Deduplication	582
vSphere Data Protection Deployment and Configuration	583
Virtual Machine Backup	584
Performing Restores with vSphere Data Protection	585
Review of Learner Objectives	586
Key Points	587
MODULE 10	
vSphere DRS	589
You Are Here	590
Importance	591
Learner Objectives	592

vSphere DRS Cluster Prerequisites	593
vSphere DRS Cluster Settings: Automation Level	594
vSphere DRS Cluster Settings: Forecasted Metrics	596
Other Cluster Settings: Swap File Location	597
vSphere DRS Cluster Settings: Virtual Machine Affinity	598
vSphere DRS Cluster Settings: DRS Groups	599
vSphere DRS Cluster Settings: VM-Host Affinity Rules	600
VM-Host Affinity Rule: Preferential	601
VM-Host Affinity Rule: Required	602
vSphere DRS Cluster Settings: Automation at the Virtual Machine Level	603
Adding a Host to a Cluster	604
Viewing vSphere DRS Cluster Information	605
Viewing vSphere DRS Recommendations	607
Monitoring Cluster Status	608
Maintenance Mode and Standby Mode	609
Removing a Host from the vSphere DRS Cluster	610
Disabling vSphere DRS and Restoring a Resource Pool Tree	611
Improving Virtual Machine Performance Methods	612
Using vSphere HA with vSphere DRS	613
Lab 20: Implementing a vSphere DRS Cluster	614
Review of Learner Objectives	615
Key Points	616
MODULE 11	
vSphere Update Manager	617
You Are Here	618
Importance	619
Learner Objectives	620
About vSphere Update Manager	621
VMware Tools Integration with vSphere Update Manager	622
VMware Tools Product Locker	623
vSphere Update Manager Capabilities	624
vSphere Update Manager Components	625
Configuring vSphere Update Manager Settings	627
Baseline and Baseline Groups	629
Creating and Editing Patch or Extension Baselines	631
Attaching a Baseline	632
Scanning for Updates	633
Viewing Compliance for vSphere Objects	634
Remediating Objects	635
Patch Recall Notification	637
Lab 21: Using vSphere Update Manager	638
Review of Learner Objectives	639
Key Points	640

MODULE 1

Course Introduction

Slide 1-1



Importance

Slide 1-2

Administrators must have the knowledge, skills, and abilities to build and run a VMware vSphere® environment.

You must know how to install and configure VMware ESXi™ hosts and VMware vCenter Server®. You must also know how to manage ESXi hosts and virtual machines with vCenter Server.

Learner Objectives (1)

Slide 1-3

By the end of this course, you should be able to meet the following objectives:

- Describe the software-defined data center
- Deploy an ESXi host and create virtual machines
- Describe the vCenter Server architecture
- Deploy VMware vCenter® Server Appliance™
- Back up and restore vCenter Server
- Configure the VMware vCenter Server® High Availability Basic mode
- Use vCenter Server to manage an ESXi host
- Configure and manage the vSphere infrastructure with VMware vSphere® Client™ and VMware vSphere® Web Client
- Configure virtual networks with vSphere standard switches
- Use vSphere distributed switches to improve network scalability
- Use vCenter Server to manage various types of storage
- Manage virtual machines, templates, clones, and snapshots

Learner Objectives (2)

Slide 1-4

By the end of this course, you should be able to meet the following objectives:

- Create a vApp
- Describe and use the content library
- Migrate virtual machines with VMware vSphere® vMotion®
- Use VMware vSphere® Storage vMotion® to migrate virtual machine storage
- Monitor resource usage and manage resource pools
- Manage VMware vSphere® High Availability and VMware vSphere® Fault Tolerance
- Use VMware vSphere® Replication™ and VMware vSphere® Data Protection™ to replicate virtual machines and perform data recovery
- Use VMware vSphere® Distributed Resource Scheduler™ clusters to improve host scalability
- Use VMware vSphere® Update Manager™ to apply patches and perform upgrades
- Perform basic troubleshooting of ESXi hosts, virtual machines, and vCenter Server

You Are Here

Slide 1-5

- 1. Course Introduction**
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Typographical Conventions

Slide 1-6

The following typographical conventions are used in this course.

Monospace

Filenames, folder names, path names, and command names:
Navigate to the `VMS` folder.

Monospace bold

What the user types:
Enter `ipconfig /release`.

Boldface

User interface controls:
Click the **Configuration** tab.

Italic

Book titles and placeholder variables:

- *vSphere Virtual Machine Administration*
- *ESXi_host_name*

References

Slide 1-7

Title	Location
<i>vSphere Installation and Setup</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
<i>vCenter Server and Host Management</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
<i>vSphere Virtual Machine Administration</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
<i>vSphere Networking</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
<i>vSphere Security</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
<i>vSphere Resource Management</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
<i>vSphere Availability</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
<i>vSphere Monitoring and Performance</i>	https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html

VMware Online Resources

Slide 1-8

Documentation for vSphere: <http://www.vmware.com/support/pubs>

VMware vSphere Blog: <http://blogs.vmware.com/vsphere/>

VMware Communities: <http://communities.vmware.com>

VMware Support: <http://www.vmware.com/support>

VMware Education: <http://www.vmware.com/education>

VMware Certification: <http://mylearn.vmware.com/portals/certification>

VMware Education and Certification Blog:
<http://blogs.vmware.com/education/>

VCP-Foundations Certification Alignment

Slide 1-9

VMware vSphere: Install, Configure, Manage aligns with the VCP-Foundations certification:

- The VCP-Foundations exam blueprint served as the basis for the design of this course.
- You should use the VCP-Foundations exam blueprint as a reference when preparing for the test.
- This course should not be used as the only resource for exam preparation.
- VMware certification details can be found at

https://mylearn.vmware.com/mgrReg/plan.cfm?plan=64178&ui=www_cert

VMware Education Overview


Slide 1-10

VMware Education provides training and certification programs to grow your skills with VMware technology.



Learning paths help you find the course you need based on the product, your role, and your level of experience.

For example, this course is part of the Data Center Virtualization Infrastructure learning path.



Examples of delivery methods:

- On-Demand:
 - Self-paced learning solution that combines modular training with hands-on practice labs, and is an alternative to traditional classroom training.
- Lab Connect:
 - Self-paced, technical lab environment designed to enhance your learning experience.
 - These cloud-based, on-demand labs let you practice skills learned during instructor-led training and get extra hands-on practice before applying your new skills in an operational environment.


VMware Certification Overview

Slide 1-11

VMware Certification sets the standard for IT professionals and validates critical skills with VMware technology.



This course can be used to fulfill the training requirement for the VMware Certified Professional 6 – Data Center Virtualization (VCP6-DCV) certification.



VMware Learning Zone provides a cloud-based video training library that can help you prepare efficiently and give you confidence to pass the certification exam.

MODULE 2

Introduction to vSphere and the Software-Defined Data Center

Slide 2-1



You Are Here

Slide 2-2

1. Course Introduction
2. **Introduction to vSphere and the Software-Defined Data Center**
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Importance

Slide 2-3

A vSphere administrator should be familiar with the many components on which vSphere is based. You should also understand the following concepts and best practices:

- Virtualization, ESXi, and the virtual machine
- Fundamental vSphere components and use of vSphere in the software-defined data center
- Use of vSphere clients to administer and manage vSphere environments

Module Lessons

Slide 2-4

- | | |
|-----------|--|
| Lesson 1: | Overview of vSphere and the Software-Defined Data Center |
| Lesson 2: | Abstraction: Overview of a Virtual Machine |
| Lesson 3: | Shared Resources: Overview of ESXi |
| Lesson 4: | Centralized Management: Overview of vCenter Server |

Lesson 1: Overview of vSphere and the Software-Defined Data Center

Slide 2-5



Lesson 1: Overview of vSphere and the Software-Defined Data Center

Learner Objectives

Slide 2-6

By the end of this lesson, you should be able to meet the following objectives:


- Understand similarities and differences between physical and virtual machines
- Describe benefits of using virtual machines
- Identify virtual machine files and file extensions
- Describe how a virtual machine is a guest and consumer of host resources
- Explain how vSphere interacts with CPUs, memory, networks, and storage
- Use vSphere clients
- Describe how vSphere fits into the software-defined data center and the cloud infrastructure

Traditional Architecture

Slide 2-7

Traditional architecture has inherent challenges:

- Poor use of physical resources
- High management and maintenance costs
- High physical infrastructure costs
- Provisioning challenges
- Insufficient failover and poor disaster protection



Traditionally, operating systems and software were on a physical computer. Large physical infrastructures pose several challenges in a data center. The model depicted in the diagram is not flexible and can be inefficient. The planning and costs of proper infrastructure (square footage, rack space, power, cooling, cabling, and server provisioning) are but a few of the challenges that IT staff must address.

In this physical model, a one-to-one relationship exists between a physical computer and the software running on it. This relationship leaves most computers vastly underutilized. Often, between only 5 and 10 percent of physical server capacity is in use. The cost of space, power, and cooling can be expensive.


Further, provisioning physical servers is a time-consuming process. In nonvirtualized environments, time must be allotted to procure new hardware, place it in the data center, install an operating system, patch the operating system, and so forth. Installing and configuring the required applications can take weeks. This process also includes a myriad of other tasks to integrate the system into the infrastructure, for example, configuring firewall rules, enabling switch ports, and provisioning storage.

Virtual Architecture

Slide 2-8

Virtual architecture has inherent benefits:

- Expanded use of physical resources
- Reduced management and maintenance costs
- Improved desktop manageability and security
- Increased availability of applications
- Increased operational flexibility



Using virtualization technology changes the way servers are utilized. Instead of the one-to-one relationship in a physical infrastructure you can have a many-to-one relationship with virtualization. This relationship saves significantly on both capital and operational expenditures. This relationship is also called consolidation.

The benefits do not stop at consolidation. You can easily back up and restore VMs, migrate VMs from one host to another, rapidly restart VMs if a hardware failure occurs, as well as clone and deploy VMs to or from templates to name a few additional benefits.

About Virtual Machines

Slide 2-9

A virtual machine is a software representation of a physical computer and its components.

The virtualization software converts the physical machine and its components into files.

Virtual Machine



Virtual Machine Components

- Operating system
- VMware Tools™
- Virtual resources, such as:
 - CPU and memory
 - Network adapters
 - Disks and controllers
 - Parallel and serial ports



The virtual machine includes a set of specification and configuration files, and is backed by the physical resources of a host. Every virtual machine has virtual devices that provide the same functionality of physical hardware but are more portable, more secure, and easier to manage.

Virtual machines typically have an operating system, applications, VMware Tools™, and virtual resources and hardware that you manage in much the same way as you manage a physical computer.

VMware Tools is a suite of utilities that you install in the operating system of a virtual machine. VMware Tools improves the performance and management of a virtual machine.

Benefits of Using Virtual Machines

Slide 2-10

Physical machines	Virtual machines
<p>Difficult to move or copy</p> <p>Bound to a specific set of hardware components</p> <p>Often has a short lifecycle</p> <p>Requires personal contact to upgrade hardware</p>	<p>Easy to move or copy</p> <ul style="list-style-type: none">• Encapsulated into files• Independent of physical hardware <p>Easy to manage</p> <ul style="list-style-type: none">• Isolated from other virtual machines running on same physical hardware• Insulated from physical hardware changes
	

In a physical machine, the operating system (for example, Windows, UNIX, or Linux) is installed directly on the hardware. The operating system requires specific device drivers to support specific hardware. If the computer is upgraded with new hardware, new device drivers are required.

If applications interface directly with hardware drivers, an upgrade to the hardware, drivers, or both can have significant repercussions if incompatibilities exist. These potential repercussions put the burden of testing hardware upgrades against a wide variety of application suites and operating systems on the hands-on technical support personnel.

Virtualizing these systems saves on this cost because virtual machines are 100 percent software.

Multiple virtual machines are isolated from one another. You can have a database server and an email server running on the same physical computer. The isolation between the virtual machines means that software-dependency conflicts are not a problem. Even users with system administrator privileges on a virtual machine's guest operating system cannot breach this layer of isolation to access another virtual machine, unless they have been explicitly granted access by the VMware ESXi™ system administrator. As a result of virtual machine isolation, if a guest operating system running in a virtual machine fails, other virtual machines on the same host are unaffected and continue to run.

A guest operating system failure has no effect on the following items:

- The ability of users to access the other virtual machines
- The ability of the operational virtual machines to access the resources that they must have
- The performance of the other virtual machines

Virtual machines enable you to consolidate your physical servers and make more efficient use of your hardware. Because a virtual machine is a set of files, features not available or not as efficient on physical architectures are available to you, for example:

- You can rapidly and consistently provision virtual machines.
- With virtual machines, you can use live migration, fault tolerance, high availability, and improved disaster recovery scenarios, for example, that increase uptime and reduce recovery time when failures happen.
- Multitenancy enables the ability to mix virtual machines into specialized configurations, such as a DMZ.

With virtual machines, you can support legacy applications and operating systems on newer hardware when maintenance contracts on the existing hardware expire.

Types of Virtualization

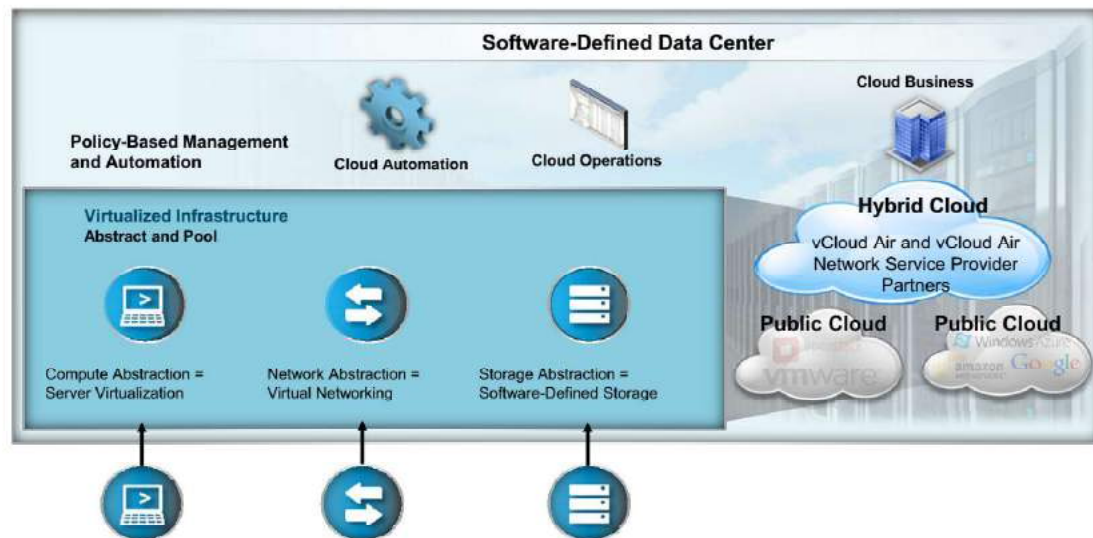
Slide 2-11



About the Software-Defined Data Center

Slide 2-12

In a software-defined data center, all infrastructure is virtualized and the control of the data center is entirely automated by software. vSphere is the foundation of the software-defined data center.



A software-defined virtual data center is deployed with isolated computing, storage, networking, and security resources faster than the traditional, hardware-based data center. All of the resources, CPU, memory, disk and network, of a software-defined virtual data center are abstracted into files thereby enabling all the benefits of virtualization at all levels of the infrastructure independent of the physical infrastructure below.

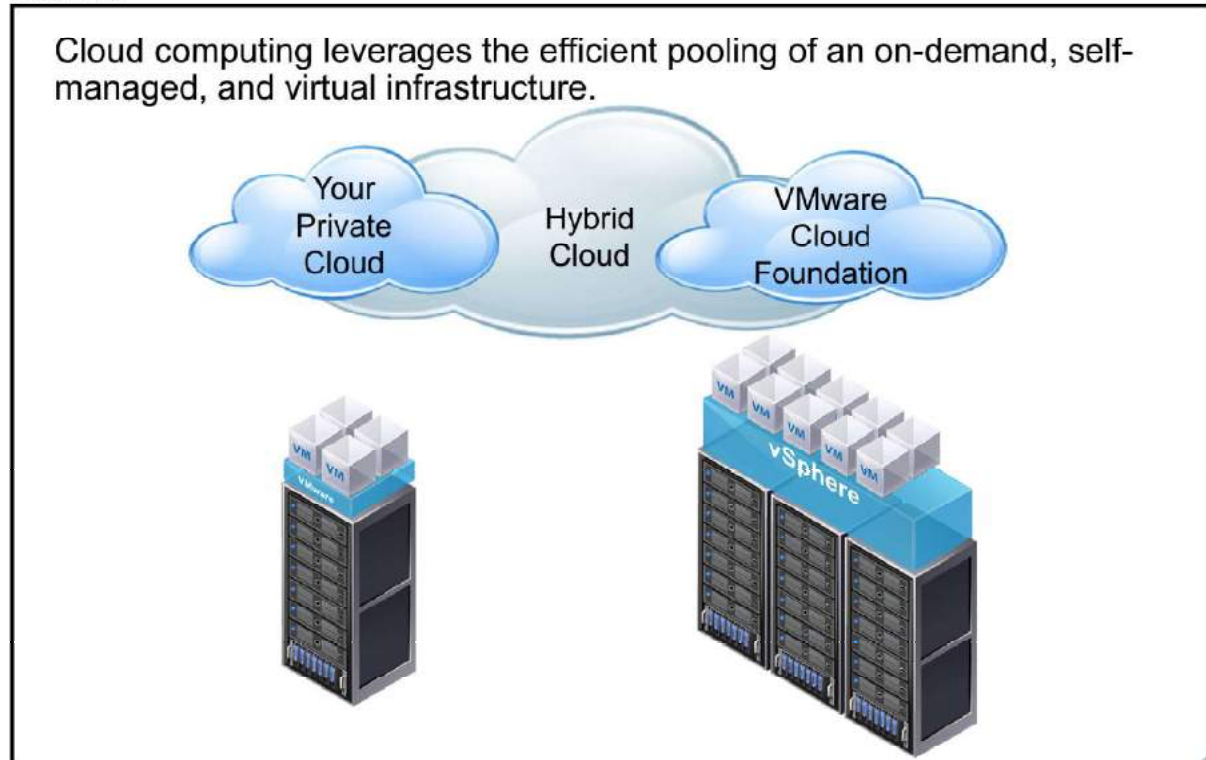
The software-defined data center can include the following components:

- Virtualized infrastructure policy-based management
- Automation hybrid cloud services, such as VMware vCloud® Air™

VMware vSphere® provides features such as hardware abstraction, network virtualization, and resource pooling. These features are critical for the success of a software-defined data center deployment.

vSphere and Cloud Computing

Slide 2-13



As defined by the National Institute of Standards and Technology (NIST), cloud computing is a model for enabling ubiquitous, convenient, and on-demand network access to a shared pool of configurable computing resources. For example, networks, servers, storage, applications, and services can be rapidly provisioned and released with minimal management effort or little service provider interaction. vSphere is the foundation of the technology that enables shared and configurable resource pools. vSphere abstracts the physical resources of the data center to separate the workload from the physical hardware. VMware vCloud® Air™ enables you to abstract the virtual resources managed by vSphere to easily allocate them to the resource consumers in the form of secure, high-performance virtual data centers, which can offer multiple tiers of service and performance. A software user interface can provide the framework for managing and maintaining this abstraction and allocation.

The purpose of vSphere in cloud computing is to hide the complexities of the physical resources from the consumer of those resources, and to provide managed access to those resources on which virtual machines can run. Meanwhile, the providers of the cloud resource still benefit from all the operational and maintenance advantages of virtualization. As such, the cloud can be deployed as a private (or community) cloud, public cloud, or hybrid cloud to serve the needs of one or more resource consumers in one or more legal business entities.

The software-defined data center gives you the basis for building a private, public, or hybrid cloud, enabling you to deliver IT as a service. The software-designed data center architecture provides a

common management, orchestration, networking, and security model across on-premise and off-premise environments.

Individual departments or internal corporate organizations (divisions) are able to deploy and manage IT infrastructure through virtual systems as needed. The private cloud has many advantages. Groups in a corporation can manage their own IT services. Policy-based self-service and automation enable IT resources to be deployed quickly when they are needed. Because customers are in control, their infrastructure matches their needs. Using VMware vCloud Suite®, you can build an on-premise private cloud that is based on vSphere. vCloud Suite is an integrated offering for building and managing a vSphere private cloud that is based on the software-defined data center architecture.

In the same way that Internet service providers can host Web sites for businesses, cloud service providers host IT operations for multiple businesses. Typically, a public cloud is owned and operated by a third party. A VPN connection is also typically available, so using a public cloud can be perceived as an off-premise extension of a private enterprise. Companies using a public cloud receive all of the advantages that a private cloud offers. A small company might be able to entirely outsource its IT.

A private cloud enables companies to move applications and data to virtualized platforms in their environment. After the application is virtualized, the company can reap additional cost savings by moving the application to an externally available public cloud. Applications can be shifted between private clouds and public clouds as desired. vCloud Air aids in constructing and managing hybrid clouds. You can build a data center that is based on vSphere with a hybrid cloud built on VMware technologies and operated by VMware with vCloud Air. You can also use the extensive VMware ecosystem or certified vCloud partners worldwide. vCloud Air supports existing workloads, third-party applications, and new application development. This solution allows virtual machines to be migrated between private and public clouds without conversion.

To learn more about VMware cloud computing, go to <http://www.vmware.com/cloud-computing/overview.html>.

Review of Learner Objectives


Slide 2-14

You should be able to meet the following objectives:

- Understand similarities and differences between physical and virtual machines
- Describe benefits of using virtual machines
- Identify virtual machine files and file extensions
- Describe how a virtual machine is a guest and consumer of host resources
- Explain how vSphere interacts with CPUs, memory, networks, and storage
- Use vSphere clients
- Describe how vSphere fits into the software-defined data center and the cloud infrastructure

Lesson 2: Abstraction: Overview of a Virtual Machine

Slide 2-15



Learner Objectives

Slide 2-16


By the end of this lesson, you should be able to meet the following objectives:

- Describe similarities and differences between a physical machine and a virtual machine
- Identify benefits of using virtual machines
- Highlight that a virtual machine is a set of specification and configuration files
- Recognize that a virtual machine is a guest and consumer of a host and its resources
- Explain how vSphere interacts with CPUs, memory, networks, and storage
- Navigate vSphere clients and examine VM settings
- Use vSphere Web Client to access and manage your vCenter Server system and ESXi host

Virtual Machine: Guest and Consumer of ESXi Host

Slide 2-17

Any application in any OS can run in a virtual machine (guest) and consume CPU, memory, disk, and network from host-based resources.

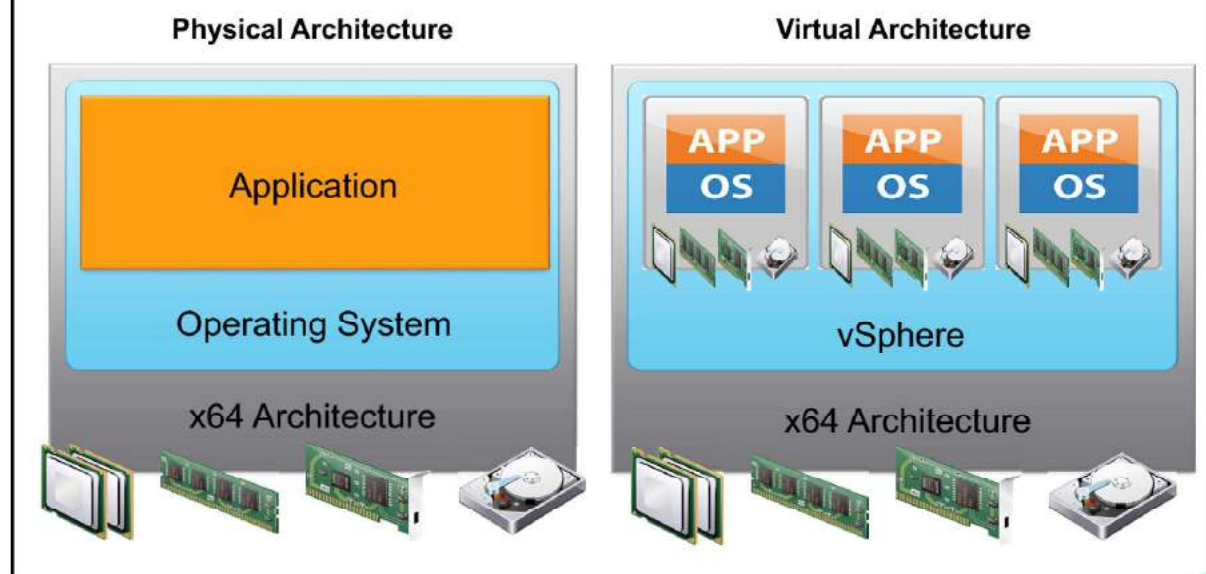


A virtual machine is an abstraction in software of a physical machine turning components into files that act like physical components.

Physical and Virtual Architecture

Slide 2-18

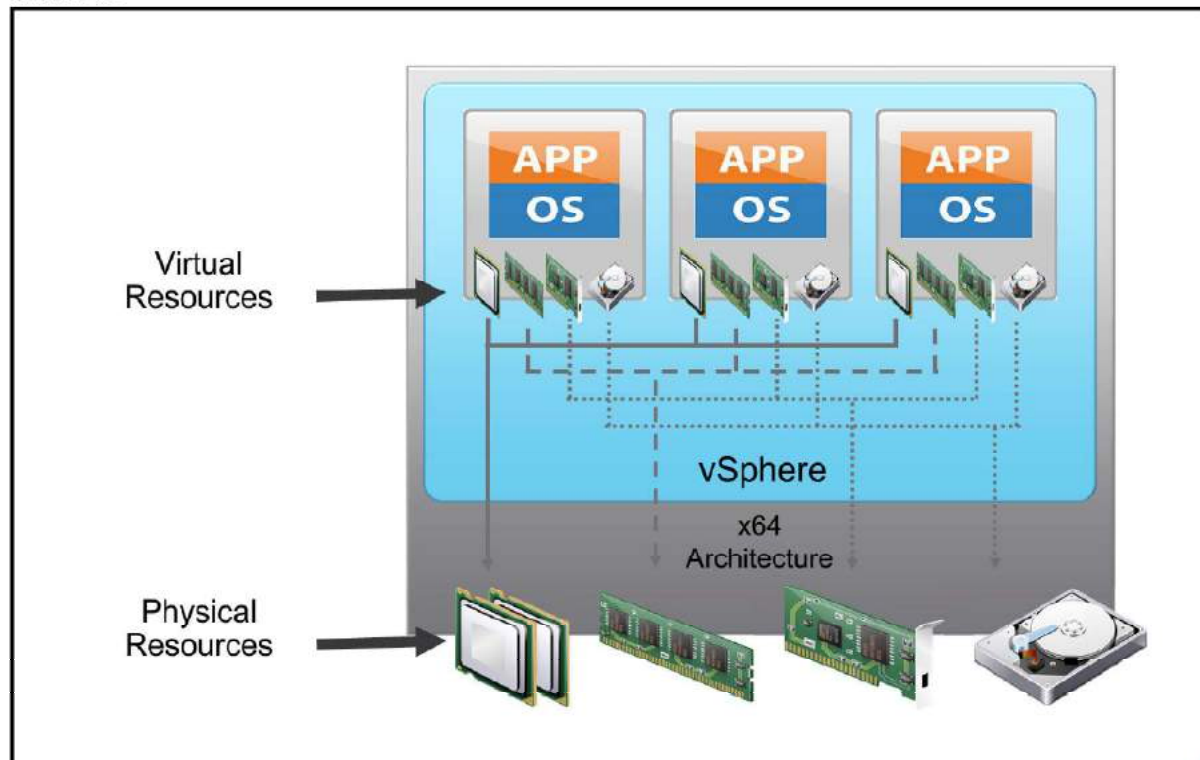
Virtualization is a technology that abstracts physical components into software components and provides solutions to many problems that are faced by IT staff.



Virtualization enables you to consolidate and run multiple workloads as virtual machines on a single computer. A virtual machine is a computer that is created by software that, like a physical computer, runs an operating system and applications. The slide illustrates the differences between a virtualized and a nonvirtualized host. In traditional architectures, the operating system interacts directly with the installed hardware. The operating system schedules processes to run, allocates memory to applications, sends and receives data on network interfaces, and reads from and writes to attached storage devices. In comparison, a virtualized host interacts with installed hardware through a thin layer of software called the virtualization layer or hypervisor. The hypervisor provides physical hardware resources dynamically to virtual machines as needed to support the operation of the virtual machines. The hypervisor enables virtual machines to operate with a degree of independence from the underlying physical hardware. For example, a virtual machine can be moved from one physical host to another. In addition, its virtual disks can be moved from one type of storage to another without affecting the functioning of the virtual machine.

Physical Resource Sharing

Slide 2-19



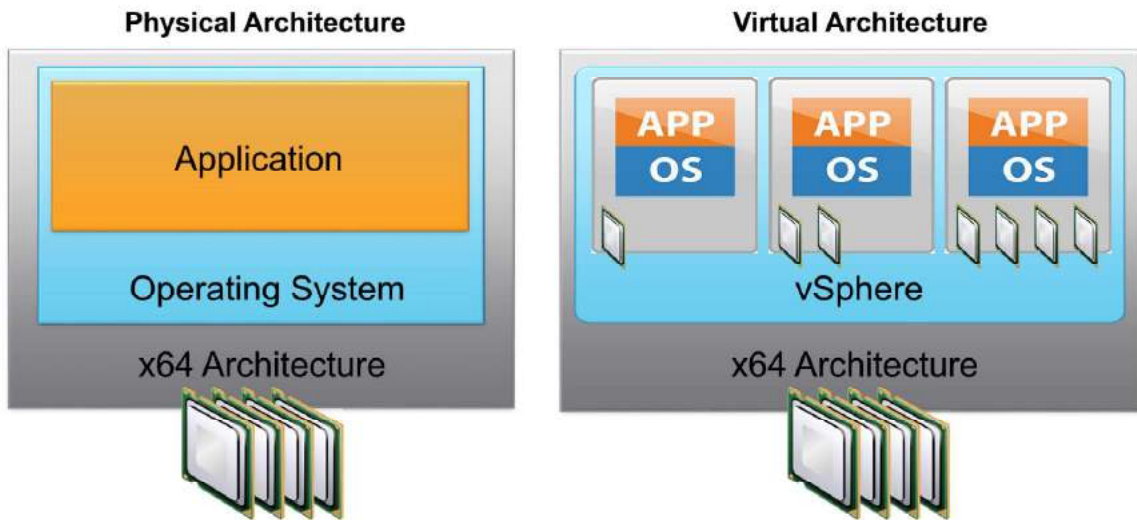
With virtualization, you can run multiple virtual machines on a single physical host, with each virtual machine sharing the resources of one physical computer across multiple environments. Virtual machines share access to CPUs and are scheduled to run by the hypervisor. In addition, virtual machines are assigned their own region of memory to use and share access to the physical network cards and disk controllers. Different virtual machines can run different operating systems and applications on the same physical computer. When multiple virtual machines run on an ESXi host, each virtual machine is allocated a portion of the physical resources. The hypervisor schedules virtual machines like a traditional operating system allocates memory for and schedules applications. These virtual machines run on various CPUs. Virtual machines, like applications, use network and disk bandwidth. However, virtual machines are managed with elaborate control mechanisms to manage how much access is available for each virtual machine. With the default resource allocation settings, all virtual machines associated with the same ESXi host receive an equal share of available resources.

CPU Virtualization

Slide 2-20

In a physical environment, the operating system assumes the ownership of all the physical CPUs in the system.

CPU virtualization emphasizes performance and runs directly on the available CPUs.



The virtualization layer runs instructions only when needed to make virtual machines operate as if they were running directly on a physical machine. CPU virtualization is not emulation. A software emulator enables programs to run on a computer system other than the one for which they were originally written. Emulation provides portability, but might negatively affect performance. CPU virtualization is not emulation because the supported guest operating systems are those designed for x64 processors. The hypervisor allows the operating systems to run natively on the hosts' physical x64 processors.

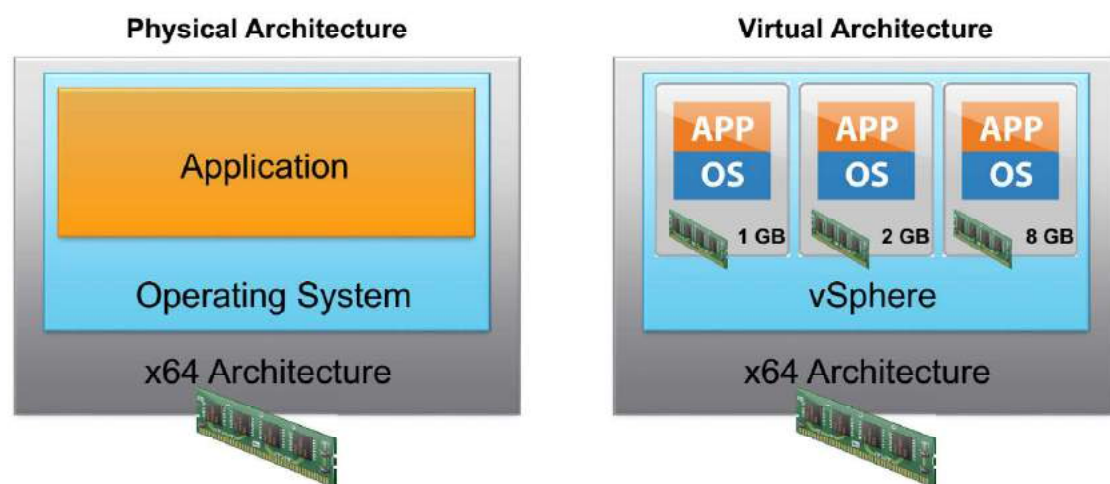
When many virtual machines are running on an ESXi host, those virtual machines might compete for CPU resources. When CPU contention occurs, the ESXi host time-slices the physical processors across all virtual machines so that each virtual machine runs as if it has a specified number of virtual processors.

Physical and Virtualized Host Memory Usage

Slide 2-21

In a physical environment, the operating system assumes the ownership of all physical memory in the system.

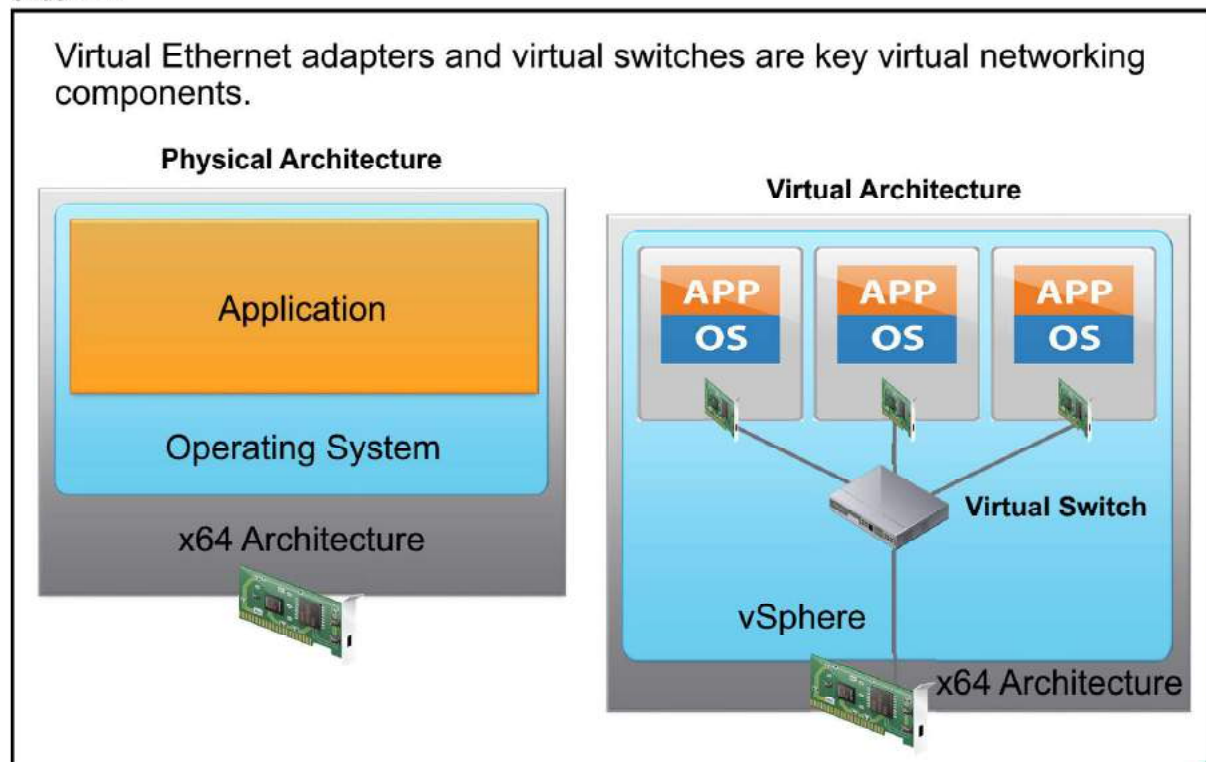
Memory virtualization emphasizes performance and runs directly on the available RAM.



When an application starts, it uses the interfaces provided by the operating system to allocate or release virtual memory pages during the execution. Virtual memory is a decades-old technique used in most general-purpose operating systems. Virtual memory enables operating systems to present more memory to applications than they physically have access to. Almost all modern processors have hardware to support it. Virtual memory creates a uniform virtual address space for applications and enables the operating system and hardware to handle the address translation between the virtual address space and the physical address space. This technique adapts the execution environment to support large address spaces, process protection, file mapping, and swapping in modern computer systems. In a virtualized environment, the VMware virtualization layer creates a contiguous addressable memory space for the virtual machine when it is started. The memory space allocated is configured when the virtual machine is created and has the same properties as the virtual address space. This configuration enables the hypervisor to run multiple virtual machines simultaneously while protecting the memory of each virtual machine from being accessed by others.

Physical and Virtual Networking

Slide 2-22

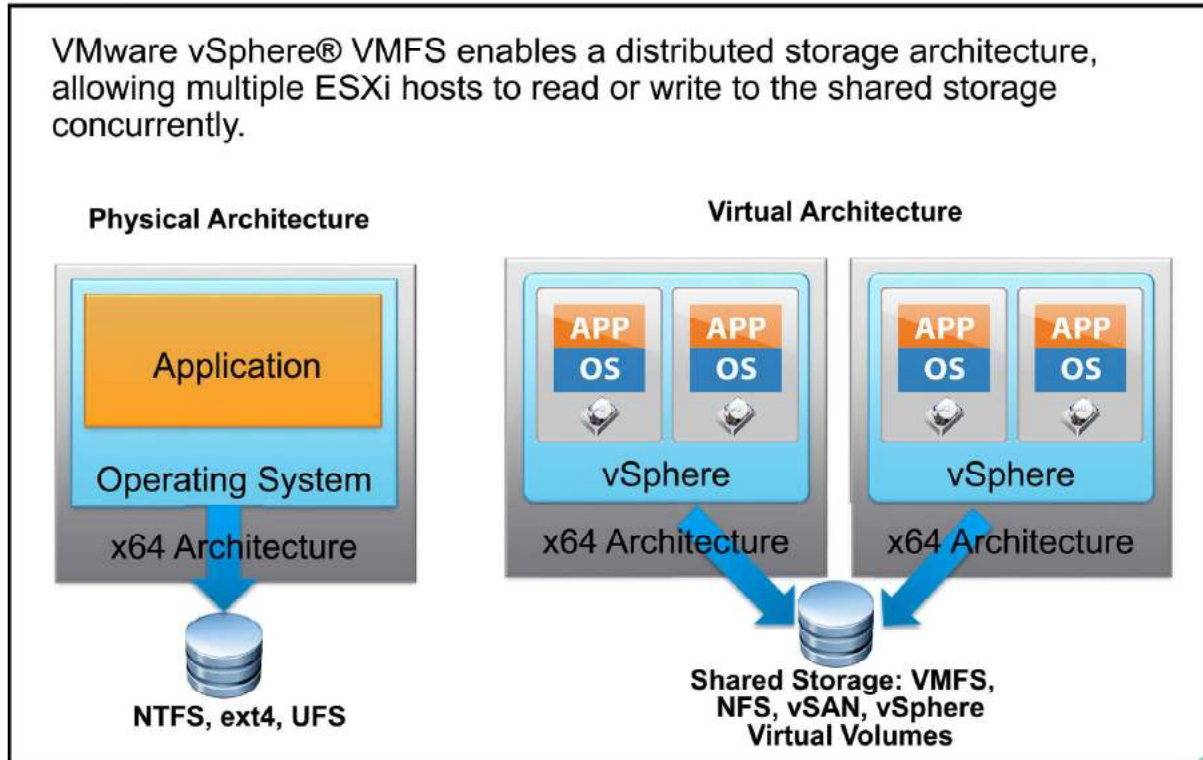


A virtual machine can be configured with one or more virtual Ethernet adapters. Virtual switches enable virtual machines on the same ESXi host to communicate with one another by using the same protocols that are used over physical switches, without the need for additional hardware. Virtual switches also support VLANs that are compatible with standard VLAN implementations from other networking equipment vendors. With the VMware virtual networking, you can link local virtual machines together and link local virtual machines to the external network through a virtual switch. A virtual switch, like a physical Ethernet switch, forwards frames at the data link layer. An ESXi host might contain multiple virtual switches. The virtual switch connects to the external network through outbound Ethernet adapters, called vmnics. The virtual switch is capable of binding multiple vmnics together, like NIC teaming on a traditional server, offering greater availability and bandwidth to the virtual machines using the virtual switch. Virtual switches are similar to modern physical Ethernet switches in many ways. Like a physical switch, each virtual switch is isolated and has its own forwarding table. So every destination that the switch looks up can match only ports on the same virtual switch where the frame originated. This feature improves security, making it difficult for hackers to break virtual switch isolation. Virtual switches also support VLAN segmentation at the port level, so that each port can be configured as an access or trunk port, providing access to either single or multiple VLANs.

However, unlike physical switches, virtual switches do not require the Spanning Tree Protocol, because a single-tier networking topology is enforced. Multiple virtual switches cannot be interconnected and network traffic cannot flow directly from one virtual switch to another virtual switch on the same host. Virtual switches provide all the ports that you need in one switch. Virtual switches do not need to be cascaded because virtual switches do not share physical Ethernet adapters and leaks do not occur between virtual switches.

Physical File Systems and VMFS

Slide 2-23



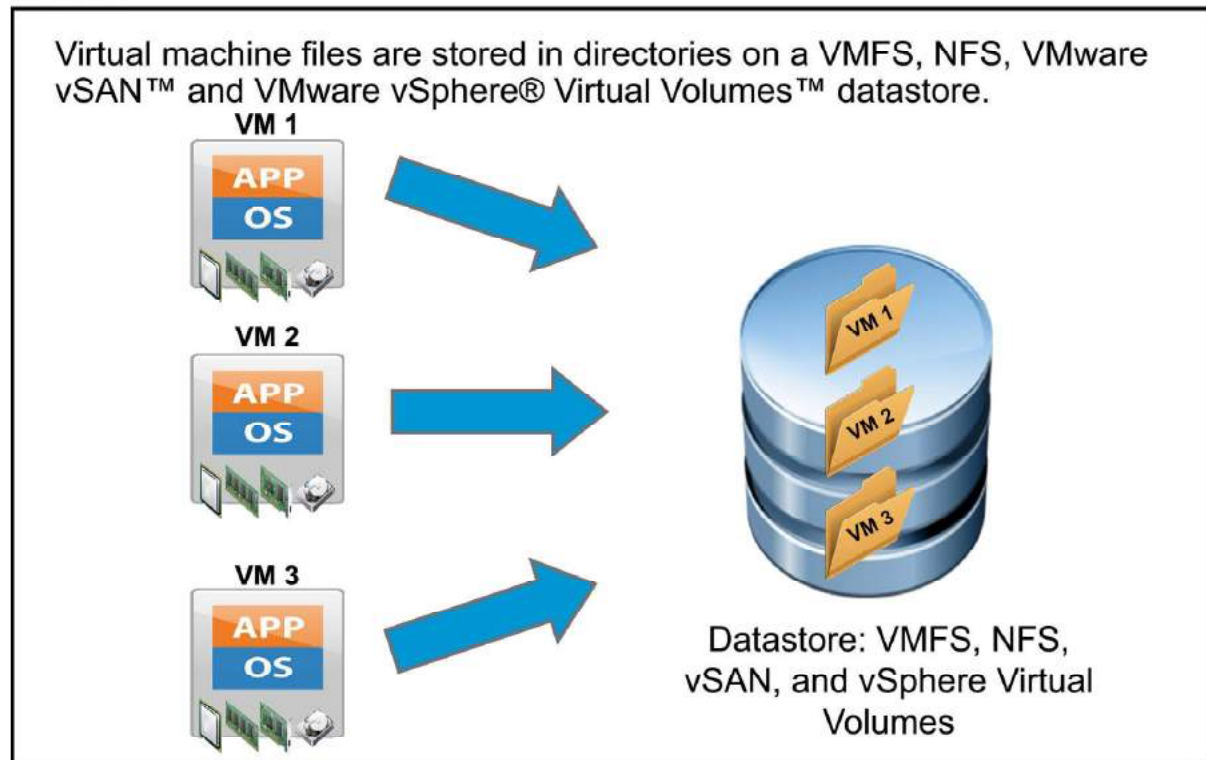
VMware vSphere® VMFS is designed, constructed, and optimized for a virtualized environment. VMFS is a high-performance cluster file system designed for virtual machines. VMFS uses distributed journaling of its file system metadata changes to enable fast and resilient recovery if a hardware failure occurs. VMFS increases resource utilization by providing multiple virtual machines with shared access to a consolidated pool of clustered storage. VMFS is also the foundation of distributed infrastructure services, such as live migration of virtual machines and virtual machine files, dynamically balanced workloads across available compute resources, automated restart of virtual machines, and fault tolerance.

VMFS provides an interface to storage resources so that several storage protocols (Fibre Channel, Fibre Channel over Ethernet, and iSCSI) can be used to access datastores on which virtual machines can reside. Dynamic growth of VMFS data stores through aggregation of storage resources and dynamic expansion of a VMFS datastore enables you to increase a shared storage resource pool with no downtime.

With the distributed locking methods, VMFS forges the link between the virtual machine and the underlying storage resources. The unique capabilities of VMFS enable virtual machines to join a cluster seamlessly, with no management overhead.

Encapsulation

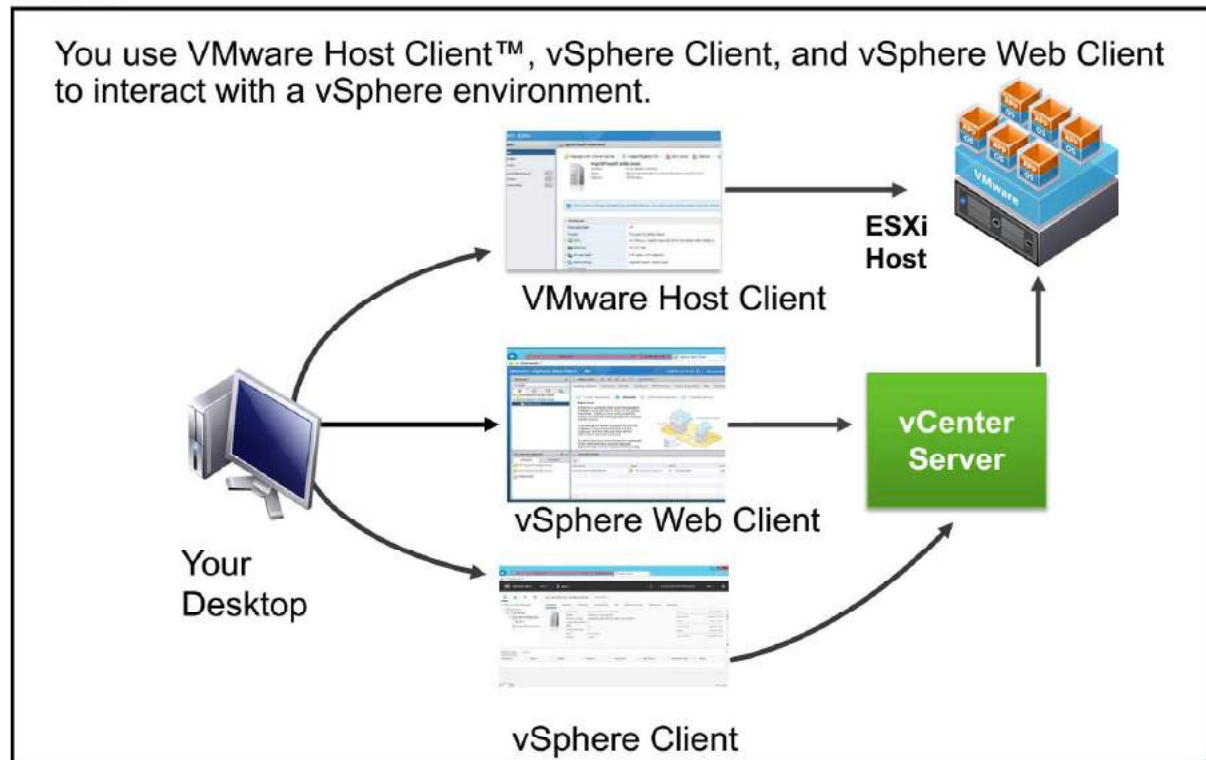
Slide 2-24



VMFS provides encapsulation of the entire virtual machine so that VMFS can easily become part of a business continuity or disaster recovery solution.

vSphere Clients

Slide 2-25



VMware vSphere® Web Client is a browser-based, fully extensible, platform-independent user interface. vSphere Web Client is based on Adobe Flex. All operations necessary for working with vSphere, ESXi, and VMware vCenter Server® are possible with vSphere Web Client.

VMware vSphere® Client™ is a Windows-installed application. vSphere Client is the legacy user interface for vSphere and has been deprecated. vSphere 6.5 introduces two new HTML5-based clients for the ESXi host and the vCenter Server system, while still retaining vSphere Web Client. The two new HTML5-based clients are designed with the following benefits:

- Clean, modern UI built on the new VMware Clarity UI standards
- No browser plug-ins to install or manage
- Integrated into vCenter Server and ESXi

vSphere Web Client

Slide 2-26


vSphere Web Client has the following components:

- Adobe Flex client application running in a browser
- Java server embedded in vCenter Server Appliance 6.5

No dedicated installation-time configuration is needed.

You access vSphere Web Client from vCenter Server Appliance at https://your_vCenter_Server_Appliance/vsphere-client.

Client Integration Plug-in is not required.



The blue image on the slide is the unique label that you will see on the upper-left corner of the screen when viewing vSphere Web Client. You can use this image to quickly differentiate vSphere Web Client from the other clients.

vSphere Web Client in vSphere 6.5 provides a complete set of functionality to enable you to manage VMware vCenter® Server Appliance™ through a Web browser. vSphere Web Client was designed to replace the client functionality of the original vSphere Client for Windows.

When you use https://your_vCSA/vsphere-client to access vSphere Web Client, it will internally redirect to port 9443 on your vCenter Server system.

The vSphere Web Client UI includes new features such as custom attributes, object tabs, and live refresh, presented with other performance and usability improvements.

Most vSphere objects tabbing structure have been reorganized to be more familiar and easier to use.

The Client Integration Plug-in is no longer required. The Client Integration Plug-in was previously necessary for a certain set of functions in vSphere Web Client. vSphere Web Client was redesigned to remove any dependency on the following functions:

- Datastore File Upload/Download
- OVF Export, Deploy
- Content Library Import/Export

The only remaining function that has dependencies is Windows Session Authentication. For this exception you install the Enhanced Authentication Plug-in.

vSphere Client


Slide 2-27

HTML-based vSphere Client has no dependence on installing Adobe Flex.

vSphere Client uses the same Java server as vSphere Web Client.

No dedicated installation-time configuration is needed.

You access vSphere Client from vCenter Server Appliance at https://your_vCenter_Server_Appliance/ui.



The dark image in the slide is the unique label that you will see on the upper-left corner of the screen when viewing vSphere Client. You can use this client to quickly differentiate vSphere Client from the other clients.

When you use https://your_vCSA/ui to access vSphere Client, it will internally redirect to port 9443 on your vCenter Server.


vSphere Client in vSphere 6.5 enables you to manage vCenter Server Appliance through a Web browser but has no requirement for Adobe Flex to be enabled in the browser.

VMware Host Client

Slide 2-28

With VMware Host Client, no dedicated installation-time configuration is needed.

VMware Host Client is served from ESXi 6.5: https://your_ESXi_host/ui.



The blue image in the slide is the unique label that you will see on the upper-left corner of the screen when viewing VMware Host Client. You can use this image to quickly differentiate VMware Host Client from the other clients.

The HTML5-embedded VMware Host Client is a new product designed to replace the host client functionality of the original vSphere Client for Windows. The layout of the HTML5 UI is similar to vSphere Web Client to remain consistent in workflow navigation while simplified in areas where vCenter Server functionality is not required.

Review of Learner Objectives


Slide 2-29

You should be able to meet the following objectives:

- Describe similarities and differences between a physical machine and a virtual machine
- Identify benefits of using virtual machines
- Highlight that a virtual machine is a set of specification and configuration files
- Recognize that a virtual machine is a guest and consumer of a host and its resources
- Explain how vSphere interacts with CPUs, memory, networks, and storage
- Navigate vSphere clients and examine VM settings
- Use vSphere Web Client to access and manage your vCenter Server system and ESXi host

Lesson 3: Shared Resources: Overview of ESXi

Slide 2-30



Learner Objectives

Slide 2-31

By the end of this lesson, you should be able to meet the following objectives:

- Describe the ESXi host architecture
- Navigate the Direct Console User Interface (DCUI) to configure an ESXi host
- Use the new VMware Host Client to administer an ESXi host
- Configure ESXi host settings
- Discuss user account best practices

About ESXi Hosts

Slide 2-32

An ESXi host has the following availability and features:

- Available for purchase with vSphere or as a free, downloadable version
- High security:
 - Host-based firewall
 - Memory hardening
 - Kernel module integrity
 - Trusted Platform Module
 - UEFI secure boot
 - Lockdown modes
- Small disk footprint
- Installable on hard disks, SAN LUNs, USB devices, SD cards, and diskless hosts

You can get a free version of ESXi, called VMware vSphere® Hypervisor, or you can purchase a licensed version with vSphere. ESXi can be installed on a hard disk, USB device, or SD card. ESXi can also be installed on diskless hosts (directly into memory) with VMware vSphere® Auto Deploy™.

ESXi has a small disk footprint for added security and reliability. ESXi provides additional protection with the following features:

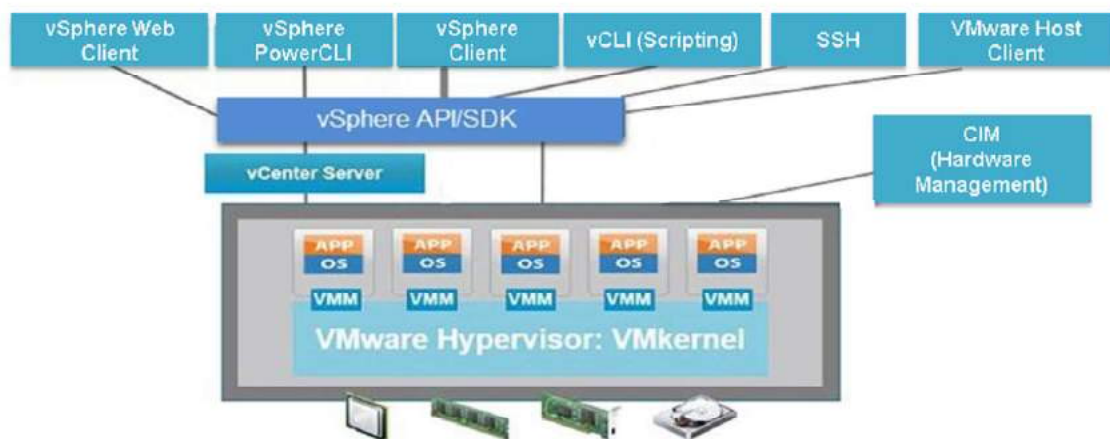
- Host-based firewall: To minimize the risk of an attack through the management interface, ESXi includes a firewall between the management interface and the network.
- Memory hardening: The ESXi kernel, user-mode applications, and executable components such as drivers and libraries are located at random, nonpredictable memory addresses. Combined with the non-executable memory protections made available by microprocessors, this provides protection that makes it difficult for malicious code to use memory exploits to take advantage of vulnerabilities.
- Kernel module integrity: Digital signing ensures the integrity and authenticity of modules, drivers, and applications as they are loaded by the VMkernel.
- Trusted Platform Module: A hardware element that creates a trusted platform. This element affirms that the boot process and all drivers loaded are genuine.

- **UEFI secure boot:** This feature is for systems that support UEFI secure boot firmware which contains a digital certificate that the VMware Infrastructure Bundles (VIBs) chain to. At boot time, a verifier is started before other processes and it will check the VIB's chain to the certificate in the firmware.
- **Lockdown modes:** Lockdown mode is a vSphere feature that disables login and API functions from being executed directly on an ESXi host. Before vSphere 6.0, only one lockdown mode existed. Feedback from customers indicated that this lockdown mode was inflexible in some use cases. With vSphere 6.0, the introduction of normal and strict lockdown modes aims to improve this feature.

Physical and Virtual Architecture

Slide 2-33

The ESXi hypervisor provides a virtualization layer that abstracts the processor, memory, storage, and networking resources of the physical host and allocates them to multiple virtual machines.



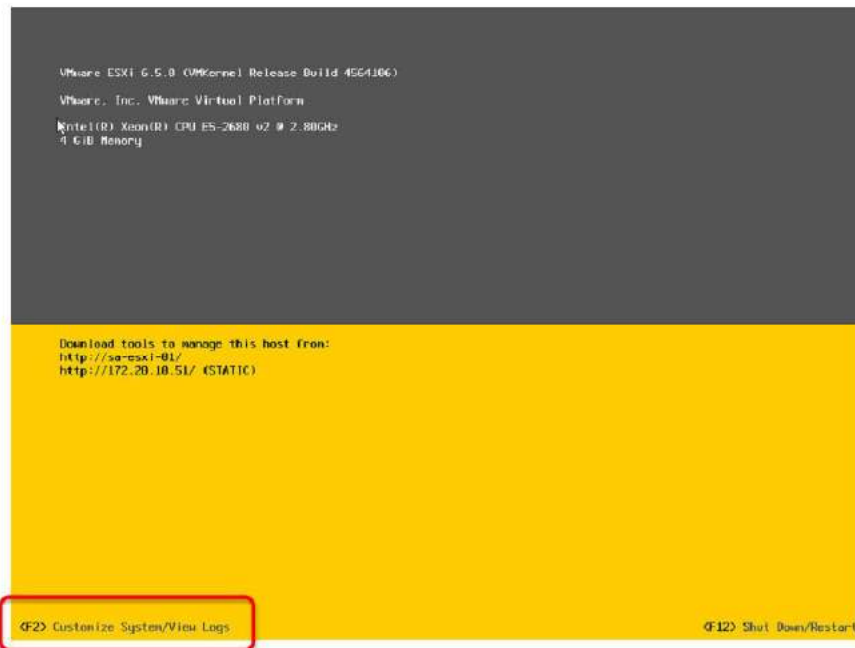
ESXi is a bare-metal hypervisor that creates the foundation for a dynamic and automated data center. In the ESXi architecture, applications running in virtual machines access CPU, memory, disk, and network interfaces without direct access to the underlying hardware. The ESXi hypervisor is called the VMkernel. The VMkernel receives requests from virtual machines for resources from the virtual machine monitor (VMM) and presents the requests to the physical hardware. Each powered-on virtual machine has its own dedicated VMM that is responsible for presenting virtual hardware to the virtual machine and receiving requests. An ESXi host can be accessed through several interfaces:

- vSphere Web Client (connected to vCenter Server)
- vSphere HTML5-based clients (connected directly to the host or to vCenter Server)
- VMware vSphere® Command-Line Interface
- VMware vSphere® API and VMware vSphere® Management SDK
- Common Information Model

Configuring an ESXi Host

Slide 2-34

The DCUI is a text-based user interface with keyboard-only interaction.



You use the Direct Console User Interface (DCUI) to configure certain settings for ESXi hosts. The DCUI is a low-level configuration and management interface, accessible through the console of the server, used primarily for initial basic configuration. You press F2 to start customizing system settings.

Configuring an ESXi Host: Root Access

Slide 2-35

DCUI enables an administrator to configure root access settings:

- Set a root password (complex passwords only).
- Enable or disable lockdown mode:
 - Limits management of the host to vCenter Server.
 - Enabled only for hosts managed by vCenter Server.



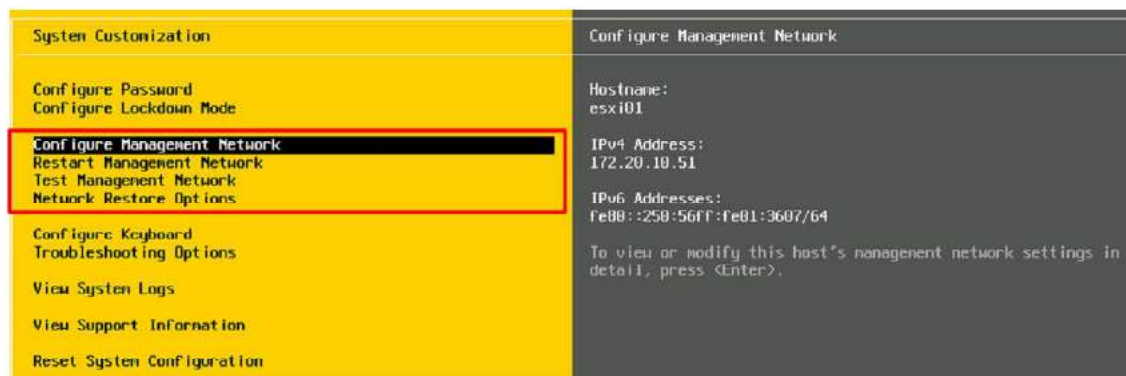
The administrative user name for the ESXi host is root. By default, the administrative password is not set. If you do not set a root password, the root user is allowed to log in to the ESXi host without providing a password. A configuration issues message appears on the ESXi host Summary page, with a reminder that the default password for the root user was not changed.

Configuring an ESXi Host: Management Network

Slide 2-36

The DCUI enables you to modify network settings:

- Host name
- IP configuration (IP address, subnet mask, default gateway)
- DNS servers

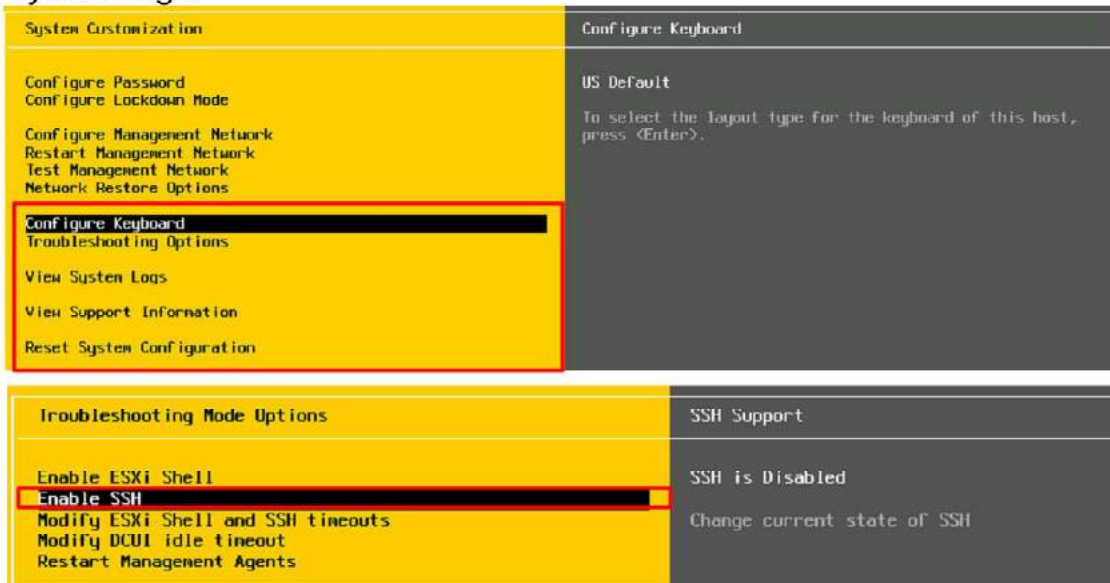


You must set up your IP address before your ESXi host is operational. By default, a DHCP-assigned address is configured for the ESXi host. To change or configure basic network settings, use DCUI. From DCUI, you can change the host name, the IP settings (such as IP address, subnet mask, default gateway), and the DNS servers. You can also modify the network adapter used for the management network, configure VLAN settings, configure IPv6 addressing, and set custom DNS suffixes. You can restart the management network (without rebooting the system), test the management network (using ping and DNS requests), and disable a management network.

Configuring an ESXi Host: Other Settings

Slide 2-37

The DCUI enables an administrator to configure the keyboard layout, enable troubleshooting services, view support information, and view system logs.



The DCUI enables you to change the keyboard layout, view support information, such as the host's license serial number, and view system logs. The default keyboard layout is U.S. English.

The troubleshooting options allow you to enable or disable troubleshooting services. By default, they are disabled:

- VMware vSphere® ESXi™ Shell: For troubleshooting issues locally
- SSH: For troubleshooting issues remotely by using an SSH client, for example, PuTTY

The best practice is to keep troubleshooting services disabled until they are necessary, for example, when you are working with VMware technical support to resolve a problem.

The **Reset System Configuration** option on this screen enables you to reset the system configuration to its software defaults and remove custom extensions or packages that you added to the host.

Remote Access Settings: Security Profile

Slide 2-38

The security profile controls remote access to an ESXi host:

- ESXi includes a firewall that is enabled by default.
- The ESXi firewall blocks incoming and outgoing traffic, except for the traffic that is enabled in the host's security profile.
- You can customize many essential security settings for an ESXi host through the Security Profile panel in vSphere Web Client.
- Some services can be managed by the administrator. Some daemons, such as the DCUI and NTP client processes, can start and stop automatically with the ESXi host.

An ESXi host includes a firewall as part of the default installation. On ESXi hosts, remote clients are typically prevented from accessing services on the host. Similarly, local clients are typically prevented from accessing services on remote hosts. To ensure the integrity of the host, few ports are open by default. To provide or prevent access to certain services or clients, you must modify the properties of the security profile.

You can configure firewall settings for incoming and outgoing connections for a service or a management agent. For some services, you can manage service details. For example, you can use the **Start**, **Stop**, or **Restart** buttons to change the status of a service temporarily. Alternatively, you can change the startup policy so that the service starts with the host or with port usage. For some services, you can explicitly specify IP addresses from which connections are allowed.

Configuring Lockdown Mode

Slide 2-39

To increase the security of your ESXi hosts, you can put your hosts in lockdown mode:

- Two lockdown modes are available: Normal and strict.

When you enable normal lockdown mode, no users but vpxuser have authentication permissions. Also, users cannot perform operations against the host directly.



Lockdown mode forces all operations to be performed through vCenter Server. The host can only be accessed using vSphere Client or vSphere Web Client.

When a host is in lockdown mode, you cannot run commands from VMware vSphere® Command-Line Interface, from an administration server, or from a script. External software or management tools might not be able to retrieve or modify information from the ESXi host.

The root user is still authorized to log in to the DCUI when lockdown mode is enabled.

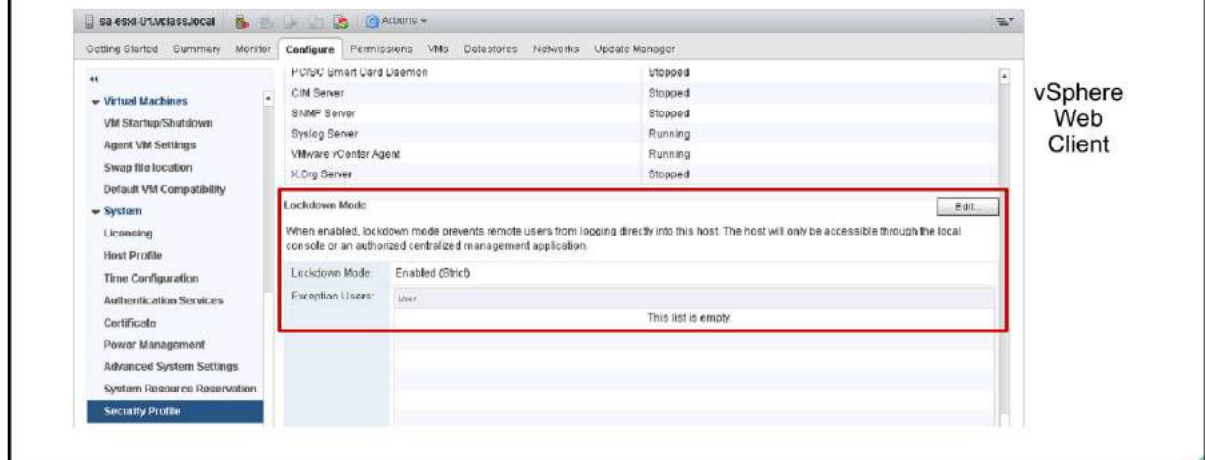
Strict Lockdown Mode

Slide 2-40

In strict lockdown mode, the DCUI service is also stopped.

If the connection to the vCenter Server system is lost and vSphere Web Client is no longer available, the ESXi host becomes unavailable.

The host can be accessed in this situation only if the VMware vSphere® ESXi™ Shell and SSH services are enabled and authorized users are added to the Exception Users list.



In normal lockdown mode, the ESXi host can be accessed through vCenter Server. Only users who are on the exception users list and have administrator privileges can log in to the DCUI. If SSH or vSphere ESXi Shell are enabled, access might be possible through these services.

In strict lockdown mode, the ESXi host can be accessed only through vCenter Server. If SSH or vSphere ESXi Shell is enabled, running sessions for accounts in the DCUI . Access advanced option that have administrator privileges remain enabled. Sessions for users in the Exception Users list that have administrator privileges also remain enabled. All other DCUI sessions are terminated.

Users defined in the Exception Users list keep their permissions when the ESXi host enters lockdown mode. The accounts are meant to be used by third-party solutions and external applications that must continue functioning when an ESXi host is in lockdown mode. Only users associated with these applications should be added to the Exception Users list.

Managing User Accounts: Best Practices

Slide 2-41

Exercise care when assigning user accounts to access ESXi hosts or vCenter Server systems:

- Strictly control root privileges to ESXi hosts.
- Create strong root account passwords that have at least eight characters. Use special characters, case changes, and numbers. Change passwords periodically.
- Manage ESXi hosts centrally through the vCenter Server system by using the appropriate vSphere client.

On an ESXi host, the root user account is the most powerful user account on the system. The user root has access to all files and all commands. This user has almost unlimited capabilities. Securing this account is the most important step that you can take to secure an ESXi host.

Whenever possible, use vSphere Web Client or an HTML-5-based client to log in to the vCenter Server system and manage your ESXi hosts. In some unusual circumstances, for example, when the vCenter Server system is down, you use VMware Host Client™ to connect directly to the ESXi host. Although you can log in to your ESXi host through vSphere CLI or vSphere ESXi Shell, these access methods should be reserved for troubleshooting or configuration that cannot be accomplished by using VMware Host Client. If a host needs to be managed directly, avoid creating local users on the host. If possible, join the host to a Windows domain and log in with domain credentials instead.

ESXi Host as an NTP Client

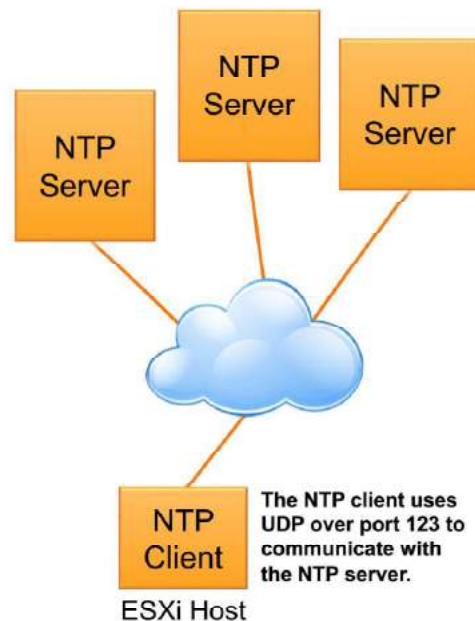
Slide 2-42

Network Time Protocol (NTP) is a client-server protocol used to synchronize a computer's clock to a time reference.

NTP is important:

- For accurate performance graphs
- For accurate time stamps in log messages
- So that virtual machines have a source to synchronize with

An ESXi host can be configured as an NTP client. It can synchronize time with an NTP server on the Internet or your corporate NTP server.



NTP is an Internet standard protocol that is used to synchronize computer clock times in a network. The benefits to synchronizing an ESXi host's time include:

- Performance data can be displayed and interpreted properly.
- Accurate time stamps appear in log messages, which make audit logs meaningful.
- Virtual machines can synchronize their time with the ESXi host. Time synchronization is beneficial to applications, such as database applications, running on the virtual machines.

NTP is a client-server protocol. When you configure the ESXi host to be an NTP client, the host synchronizes its time with an NTP server, which can be a server on the Internet or your corporate NTP server.

For information about NTP, see <http://www.ntp.org>.

For more about timekeeping, see VMware knowledge base articles 1318 at <http://kb.vmware.com/kb/1318> and 1006427 at <http://kb.vmware.com/kb/1006427>.

Labs

Slide 2-43

Lab 1: Installing ESXi

Lab 2: Configuring ESXi Hosts

Lab 1: Installing ESXi

Slide 2-44

Install ESXi on a VM using your student desktop

1. Access Your Student Desktop
2. Install ESXi

Lab 2: Configuring ESXi Hosts

Slide 2-45

Configure an ESXi host

1. Examine the Options in the DCUI
2. Configure the Management Network
3. Enable SSH
4. View System Logs
5. Clean Up for the Next Lab

Review of Learner Objectives


Slide 2-46

You should be able to meet the following objectives:

- Describe the ESXi host architecture
- Navigate the DCUI to configure an ESXi host
- Use the new VMware Host Client to administer an ESXi host
- Configure ESXi host settings
- Discuss user account best practices

Lesson 4: Centralized Management: Overview of vCenter Server

Slide 2-47



Learner Objectives

Slide 2-48

By the end of this lesson, you should be able to meet the following objectives:


- Describe the vCenter Server architecture
- Discuss how ESXi hosts communicate with vCenter Server
- Identify the vCenter Server services, components, and modules
- Explain VMware Platform Services Controller™
- Describe the vCenter Server security

About the vCenter Server Management Platform

Slide 2-49

vCenter Server is a service that acts as a central administration point for ESXi hosts and their virtual machines connected on a network:

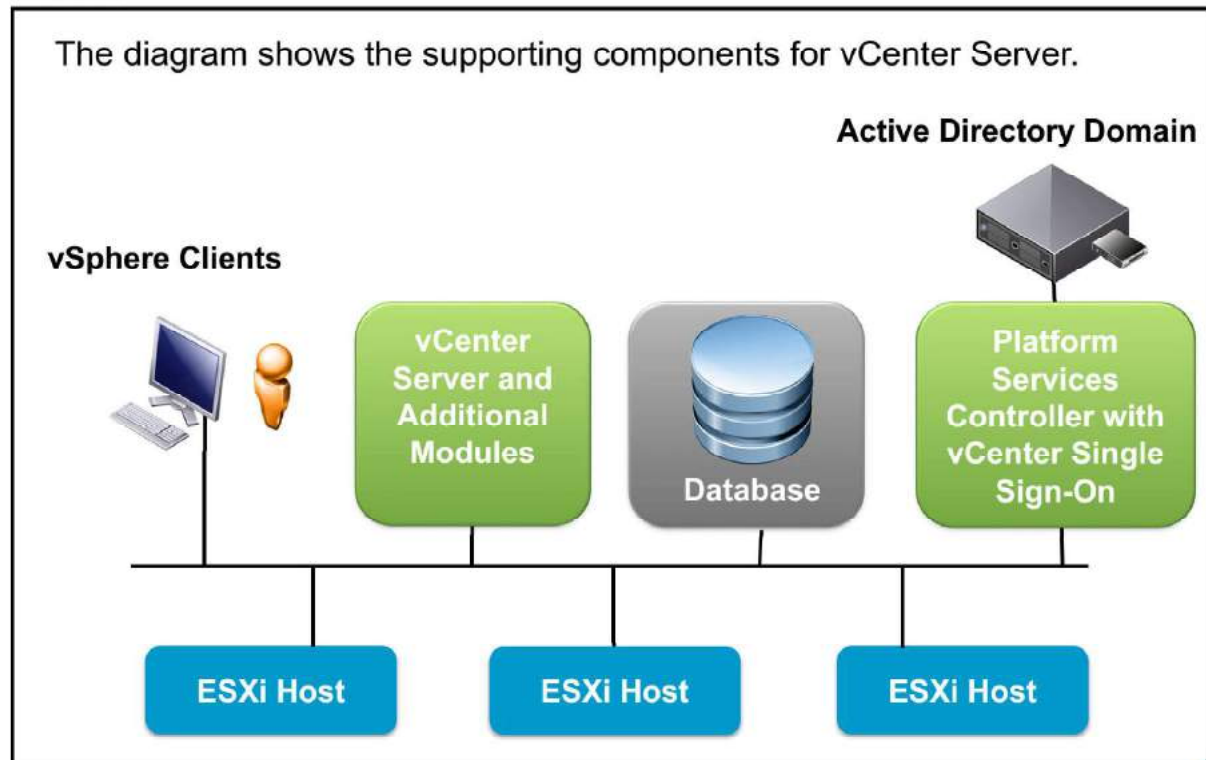
- Runs on Windows or on a Linux-based appliance
- Directs the actions of virtual machines and hosts



vCenter Server is a service that acts as a central administrator for ESXi hosts connected in a network. vCenter Server enables you to pool and manage the resources of multiple hosts. You can deploy vCenter Server Appliance on an ESXi host in your infrastructure. vCenter Server Appliance is a preconfigured Linux-based virtual machine that is optimized for running vCenter Server and the vCenter Server components. vCenter Server Appliance provides advanced features, such as VMware vSphere® Distributed Resource Scheduler™, VMware vSphere® High Availability, VMware vSphere® Fault Tolerance, VMware vSphere® vMotion®, and VMware vSphere® Storage vMotion®.

vCenter Server Architecture

Slide 2-50

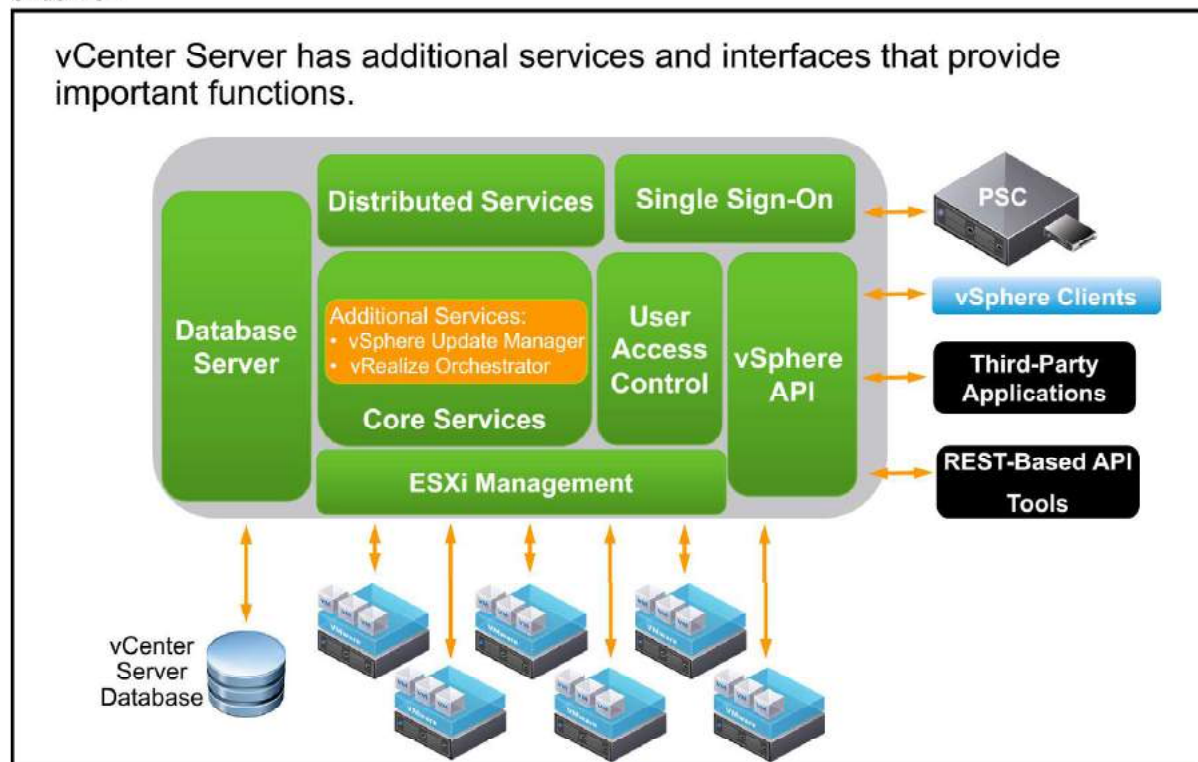


The vCenter Server architecture relies on the following components:

- vSphere Web Client and VMware Host Client: VMware Host Client is used to connect directly to ESXi hosts. vSphere Web Client connects directly to vCenter Server. You can also use the new vSphere Client to connect to the vCenter Server system. When an ESXi host is managed by vCenter Server, administrators should always use vCenter Server and the vSphere clients to manage that host.
- vCenter Server database: The vCenter Server database is the most critical component. The database stores the inventory items, security roles, resource pools, performance data, and other critical information for vCenter Server.
- VMware vCenter® Single Sign-On™: Provides a security domain defined in your vSphere environment. Authentication is performed by the vCenter Single Sign-On server. The vCenter Single Sign-On server can be configured to authenticate against multiple user repositories, also called identity sources, such as an Active Directory domain.
- Managed hosts: vCenter Server enables you to manage ESXi hosts and the virtual machines that run on them.

Additional vCenter Server Services and Interfaces

Slide 2-51



vCenter Server includes these services and interfaces:

- Core services include management of resources and virtual machines by the Inventory service, task scheduling, statistics logging, management of alarms and events, virtual machine provisioning, and host and virtual machine configuration.
- The vCenter Lookup Service contains topology information about the vSphere infrastructure, enabling vSphere components to connect to each other securely. Services, such as the Inventory Service and vCenter Server, register with the vCenter Lookup Service so that other vSphere components, like vSphere Web Client, can find them.
- Distributed services include vSphere vMotion, vSphere DRS, and vSphere HA, which are installed with vCenter Server.
- Additional services are packaged as part of the base product, for example, VMware vSphere® Update Manager™ and VMware vRealize® Orchestrator™. No additional license is necessary.
- A database interface provides access to the vCenter Server database.
- vCenter Server provides access to the ESXi host through a vCenter Server agent, which is started on the host when it is added to the vCenter Server inventory.

Platform Services Controller

Slide 2-52

vCenter Server includes Platform Services Controller:

- Platform Services Controller includes a set of common infrastructure services:
 - VMware vCenter® Single Sign-On
 - VMware License Server
 - Lookup Service
 - VMware Certificate Authority
 - Certificate Store
 - VMware Directory Services
- Other features are installed under the vCenter Server component.
- You can install vCenter Server and Platform Services Controller on the same or different machines.



vCenter Server has VMware Platform Services Controller™. This system includes a set of common infrastructure services. Some of these, such as vCenter Single Sign-on, were present under vCenter Server in vSphere 5.x. Other functions like an internal certificate authority and certificate store are present in vSphere 6.0 and are subsystems designed to make vCenter Server more robust.

vCenter Server Services

Slide 2-53

The vCenter Server group of services contains:

- vCenter Server
- vSphere Web Client (server)
- VMware Inventory Service
- vSphere Update Manager
- VMware vSphere® Auto Deploy™
- VMware vSphere® ESXi™ Dump Collector
- VMware vSphere® Syslog Collector

Platform Services
Controller

vCenter Server

You cannot distribute these vCenter Server functions across multiple servers. When you deploy vCenter Server Appliance, all of these features are included.

The vCenter Server component includes the following systems that were present in vSphere 5:

- VMware Inventory Service
- vSphere Web Client
- vCenter Server

vSphere Web Client is not a client, but is a Web server that is based on the Apache Tomcat Web server.

In addition, the vCenter Server component also includes systems that were optional before:

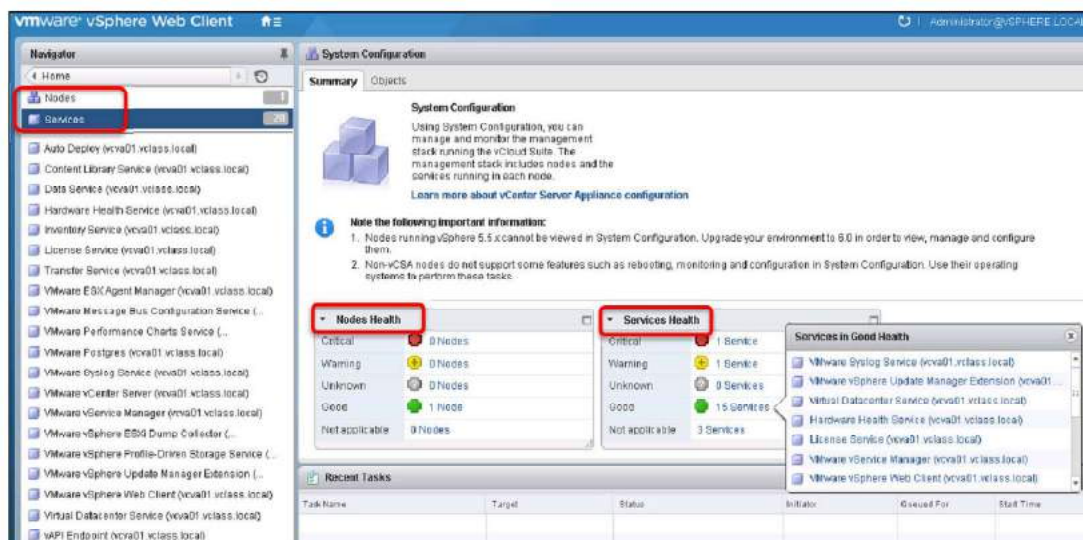
- VMware vSphere® Syslog Collector
- VMware vSphere® ESXi™ Dump Collector
- VMware vSphere® Auto Deploy™

Although installation of these services is not optional, the utilization of this functionality is up to the administrator. Some services such as vSphere ESXi Dump Collector are installed in a disabled state when vCenter Server is installed on a Windows server.

Monitoring the Health and Status of Services and Nodes Across vCenter Server Systems

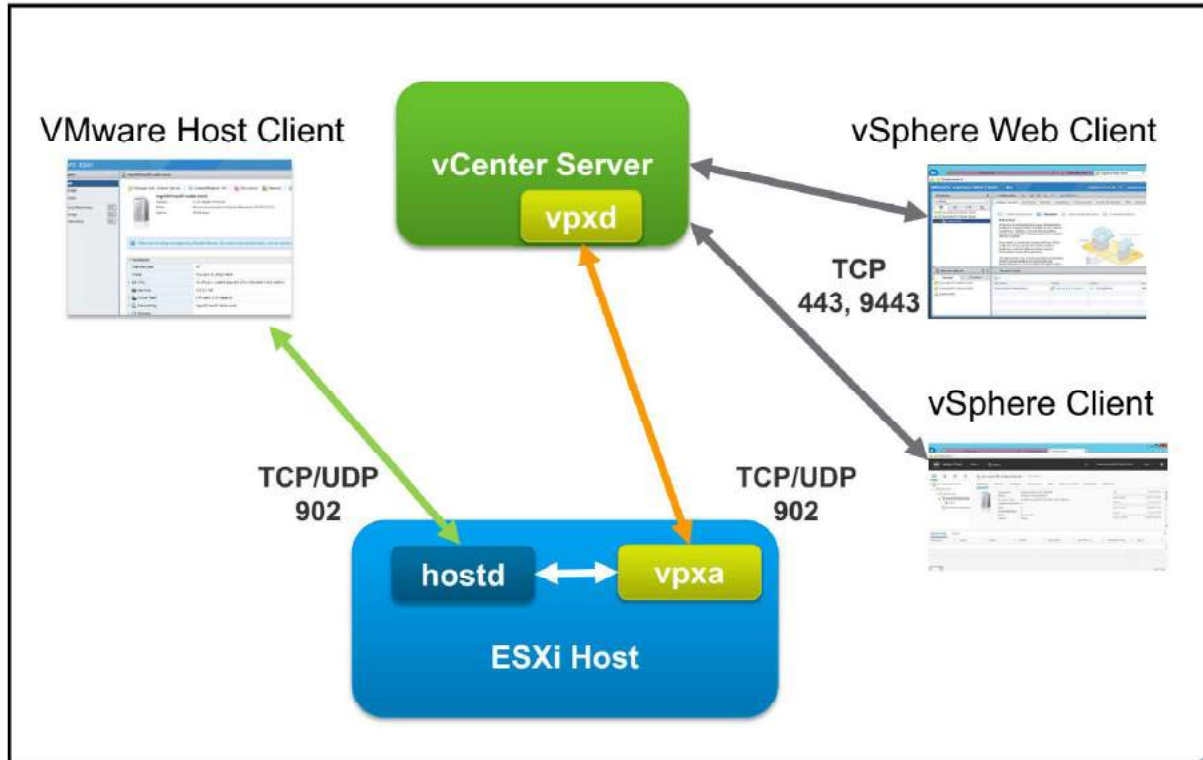
Slide 2-54

vSphere Web Client enables you to monitor the status of all manageable services and nodes across vCenter Server systems. A list of default services is available in each vCenter Server instance.



ESXi and vCenter Server Communication

Slide 2-55



vSphere Web Client and the new vSphere Client communicate directly with vCenter Server. If you must communicate directly with an ESXi host, then you must use VMware Host Client.

vCenter Server provides direct access to the ESXi host through a vCenter Server agent named virtual provisioning X agent (vpxa). The vpxa process is started on the host when it is added to the vCenter Server inventory. The vCenter Server service (vpxd) communicates with the ESXi host daemon (hostd) through the vCenter Agent (vpxa).

For clients communicating directly with the host, bypassing vCenter, they converse with hostd. The hostd process runs directly on the ESXi host and is responsible for managing most of the operations on the ESXi host. It is aware of all virtual machines that are registered on the ESXi host, the storage volumes visible to the ESXi host, and the status of all virtual machines. Most commands or operations come from vCenter Server through vpxa. Examples include creating, migrating, and powering on virtual machines. vpxa acts as an intermediary between the vpxd process, which runs on vCenter Server, and the hostd process to relay the tasks to perform on the host. When you are logged in to the vCenter Server system through vSphere Web Client, vCenter Server passes commands to the ESXi host through the vpxa process. The vCenter Server database is also updated. If you are using VMware Host Client to communicate directly with an ESXi host, communications go directly to the hostd process and the vCenter Server database is not updated.

Review of Learner Objectives

Slide 2-56

You should be able to meet the following objectives:

- Describe the vCenter Server architecture
- Discuss how ESXi hosts communicate with vCenter Server
- Identify the vCenter Server services, components, and modules
- Explain Platform Services Controller
- Describe the vCenter Server security

Key Points

Slide 2-57


- Using virtual machines solves many data center problems.
- Virtual machines are hardware independent.
- Virtual machines share the physical resources of the ESXi host on which they reside.
- A virtual machine is a set of files that is easy to transfer and back up.
- Virtual machine files are encapsulated into a folder and placed on a datastore.
- The ESXi hypervisor runs directly on the host.
- vSphere abstracts CPU, memory, storage, and networking for virtual machine use.

Questions?

MODULE 3

Creating Virtual Machines

Slide 3-1



You Are Here

Slide 3-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. **Creating Virtual Machines**
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Importance

Slide 3-3

You can create a virtual machine in several ways. Choosing the correct method for your task can help you save time and make the deployment process manageable and scalable.

Module Lessons

Slide 3-4

- | | |
|-----------|----------------------------|
| Lesson 1: | Virtual Machine Concepts |
| Lesson 2: | Creating a Virtual Machine |

Lesson 1: Virtual Machine Concepts

Slide 3-5



Lesson 1: Virtual Machine Concepts

Learner Objectives

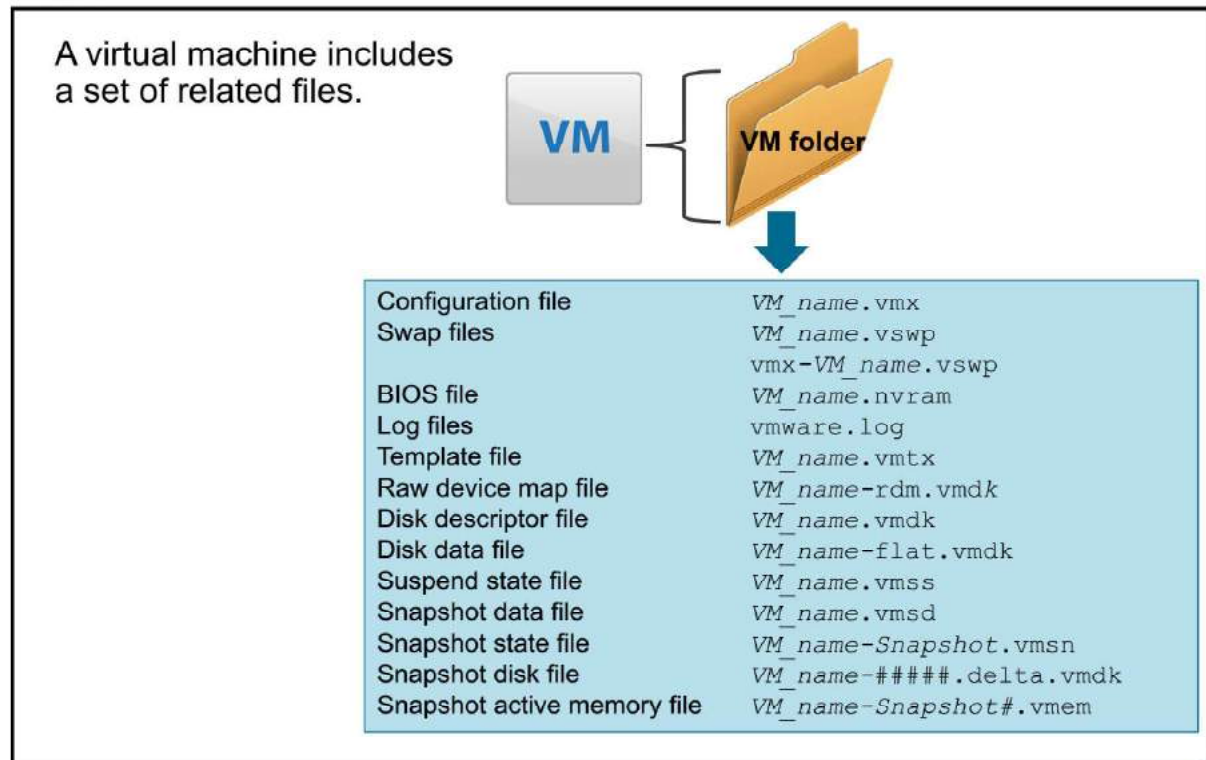
Slide 3-6

By the end of this lesson, you should be able to meet the following objectives:

- Identify files that make up a virtual machine
- Compare virtual machine hardware version 13 to other versions
- Describe components of a virtual machine
- Identify the various methods to access a virtual machine console
- Identify the virtual network adapters and highlight the enhanced VMXNET3
- Discuss the vRDMA and vNVMe features
- Compare and contrast the types of virtual disk provisioning

About Virtual Machine Files

Slide 3-7



The slide lists some of the files that make up a virtual machine. Except for the log files, the name of each file starts with the virtual machine's name *VM_name*. A virtual machine consists of the following files:

- A configuration file (*.vmx*).
- One or more virtual disk files. The first virtual disk has files *VM_name.vmdk* and *VM_name-flat.vmdk*.
- A file containing the virtual machine's BIOS settings (*.nvram*).
- A virtual machine's current log file (*.log*) and a set of files used to archive old log entries (*-#.log*).
- Swap files (*.vswp*) used to reclaim memory during periods of contention.
- A snapshot description file (*.vmsd*). This file is empty if the virtual machine has no snapshots.

If the virtual machine is converted to a template, a virtual machine template configuration file (*.vmtx*) replaces the virtual machine configuration file (*.vmx*). A virtual machine template is a master copy of the virtual machine.


If the virtual machine has more than one disk file, the file pair for the second disk file and later is named `VM_name_#.vmdk` and `VM_name_#-flat.vmdk`, where # is the next number in the sequence, starting with 1. For example, if the virtual machine named Test01 has two virtual disks, this virtual machine has the `Test01.vmdk`, `Test01-flat.vmdk`, `Test01_1.vmdk`, and `Test01_1-flat.vmdk` files.

In addition to the current log file, `vmware.log`, up to six archive log files are maintained at one time. For example, `-1.log` to `-6.log` might exist at first. The next time an archive log file is created, for example, when the virtual machine is powered off and powered back on, the following action occurs: `-2.log` to `-7.log` are maintained, `-1.log` is deleted, `-3.log` to `-8.log`, and so on.

A virtual machine can have other files, for example, if one or more snapshots were taken or if raw device mappings (RDMs) were added. A virtual machine has an additional lock file if it resides on an NFS datastore. A virtual machine has a change block tracking file (`.ctk`) if it is backed up with the VMware vSphere® Data Protection™ appliance or other backup software that has enabled the CDP feature.

About Virtual Machine Virtual Hardware

Slide 3-8



A virtual machine uses virtual hardware. Each guest operating system sees ordinary hardware devices. The guest operating system does not know that these devices are virtual. All virtual machines have uniform hardware, except for a few variations that the system administrator can apply. Uniform hardware makes virtual machines portable across VMware virtualization platforms.

You can configure virtual machine memory and CPU settings. vSphere supports many of the latest CPU features, including virtual CPU performance counters. You can add virtual hard disks and NICs. You can also add and configure virtual hardware, such as CD/DVD drives, floppy drives, and SCSI devices. All devices are not available to add and configure. For example, you cannot add video devices, but you can configure available video devices and video cards.

You can add multiple USB devices, such as security dongles and mass storage devices, to a virtual machine that resides on an ESXi host to which the devices are physically attached. When you attach a USB device to a physical host, the device is available only to virtual machines that reside on that host. Those virtual machines cannot connect to a device on another host in the data center. A USB device is available to only one virtual machine at a time. When you remove a device from a virtual machine, it becomes available to other virtual machines that reside on the host.

You can add up to 16 PCI VMware vSphere® DirectPath I/O™ devices to a virtual machine. The devices must be reserved for PCI passthrough on the host on which the virtual machine runs. Snapshots are not supported with vSphere DirectPath I/O passthrough devices.

Virtual Machine Communication Interface (VMCI) provides a high-speed communication channel between a virtual machine and the hypervisor. You cannot add or remove VMCI devices.

The SATA controller provides access to virtual disks and DVD/CD-ROM devices. The SATA virtual controller appears to a virtual machine as an AHCI SATA Controller.

The VMCI is an infrastructure that provides fast and efficient communication between a virtual machine and the host operating system. The VMCI SDK facilitates the development of applications that use the VMCI infrastructure. Without VMCI, virtual machines communicate with the host using the network layer. Using the network layer adds overhead to the communication. With VMCI, communication overhead is minimal and tasks that require that communication can be optimized. An internal network can transmit an average of slightly over 2 Gbit/s using VMXNET3. VMCI can go up to nearly 10 Gbit/s with twelve 8k sized queue pairs.

The following types of communication exist:

- Datagrams: Connectionless and similar to UDP queue pairs
- Connection oriented: Similar to TCP

VMCI provides socket APIs that are very similar to APIs that are already used for TCP/UDP applications. IP addresses are replaced with VMCI ID numbers. For example, you can port netperf to use VMCI sockets instead of TCP/UDP. VMCI is disabled by default.

For more information about the virtual hardware, see *vSphere Virtual Machine Administration Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Virtual Hardware Versions

Slide 3-9

The virtual hardware version determines the operating system functions that a virtual machine supports. Do not use a later version that is not supported by the VMware product.

Compatibility	Hardware Version
ESXi 6.5 and later	13
ESXi 6 and later	11
ESXi 5.5 and later	10
ESXi 5.1 and later	9
ESXi/ESX 5.0 and later	8

Each release of a VMware product has a corresponding virtual machine hardware version included. The table shows the highest hardware version level that each ESX/ESXi version supports. Each virtual machine compatibility level supports at least five major or minor vSphere releases. For example, a virtual machine with ESX/ESXi 4.0 and later compatibility can run on ESX/ESXi 4.0, ESX/ESXi 4.1, ESXi 5.0, ESXi 5.1, ESXi 5.5, and ESXi 6.

For more information about configuring virtual machine hardware, see *vSphere Virtual Machine Administration* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

For a complete list of virtual machine configuration maximums, see *Configuration Maximums* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

About Virtual Hardware Version 13

Slide 3-10

Virtual hardware version 13 provides several features and benefits.

Features	Benefits
Increased RAM capacity	Hardware version 13 virtual machines support up to 6 TB of RAM.
NVMe in Guest (vNVMe)	Provides high performance guest block I/O and up to a 50% reduction in software overhead. vNVMe also supports back-end vSAN and vSphere Virtual Volumes storage object-backed disk support.
RDMA in Guest (vRDMA)	vRDMA supports Remote Direct Memory Access providing OS bypass, zero-copy, low latency, and high bandwidth with less power usage and faster data access.

The increased RAM capacity supports the VMware vision of any application in a virtual machine. Improving the infrastructure performance also fits that vision. vNVME minimizes or eliminates VMkernel overhead in various cases, such as maintaining drivers for all operating systems, lower cost per I/O in guests, and support for storage based on VMware vSAN™ and VMware vSphere® Virtual Volumes™. Virtual RDMA (vRDMA) also helps accelerate the VMkernel network performance by allowing multiple guests simultaneous access to virtual devices providing low latency and high bandwidth.

For further information, see VMware knowledge base article 2051652 at <http://kb.vmware.com/kb/2051652>.

About CPU and Memory

Slide 3-11

You can add, change, or configure CPU and memory resources to improve virtual machine performance.

The maximum number of vCPUs that you can assign to a virtual machine depends on:


- The number of logical CPUs on the host
- The host license
- The type of installed guest operating system

A virtual machine running on an ESXi 6 host can have up to 128 vCPUs.

Maximum memory size for a virtual machine depends on:

- The host's physical memory
- The virtual machine's compatibility setting

The maximum memory size of a virtual machine with ESXi 6.5 compatibility running on ESXi 6.5 is 6 TB.



You size the virtual machine's CPU and memory according to the applications and the guest operating system.

The maximum number of virtual CPUs that you can assign to a virtual machine depends on the number of logical CPUs on the host, the host license, and the type of guest operating system that is installed on the virtual machine. The multicore vCPU feature enables you to control the number of cores per virtual socket in a virtual machine. This capability enables operating systems with socket restrictions to use more of the host CPU's cores, which increases overall performance.

A virtual machine cannot have more virtual CPUs than the number of logical CPUs on the host. The number of logical CPUs is the number of physical processor cores, or two times that number if hyperthreading is enabled. For example, if a host has 128 logical CPUs, you can configure the virtual machine for 128 vCPUs.

You can set most of the memory parameters during virtual machine creation or after the guest operating system is installed. Some actions require that you power off the virtual machine before changing the settings. The memory resource settings for a virtual machine determine how much of the host's memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size.

ESXi hosts limit the memory resource use to the maximum amount useful for the virtual machine, so that you can accept the default of Unlimited memory resources. You can reconfigure the amount of memory allocated to a virtual machine to enhance performance. Maximum memory size for a virtual machine depends on the host's physical memory and the virtual machine's compatibility setting.

About Virtual Storage

Slide 3-12

Virtual disks are connected to virtual SCSI adapters. The ESXi host offers several choices in storage adapters to a virtual machine:

- BusLogic Parallel: The latest Mylex (BusLogic) BT/KT-958 compatible host bus adapter.
- LSI Logic Parallel: LSI Logic LSI53C10xx Ultra320 SCSI I/O controller is supported.
- LSI Logic SAS: LSI Logic SAS adapter has a serial interface.
- VMware Paravirtual SCSI: PVSCSI adapters are high-performance storage adapters that can provide greater throughput and lower CPU utilization.
- AHCI SATA controllers: A standard defined by Intel that specifies the operation of Serial ATA (SATA) host bus adapters.
- vNVME: Virtualized Non-volatile Memory host controller interface specification (NVMHCI) is a logical device interface for accessing nonvolatile storage media through a PCI Express bus.

The Typical Configuration option of the Virtual Machine Creation wizard in vSphere Client selects the type of virtual SCSI adapter, based on the choice of guest operating system. The virtual disk is stored in the same folder as the virtual machine configuration file. However, you can select to place a virtual disk in an alternate location, for example, when separating boot and data disks. You select a VMFS datastore to hold the new, blank virtual disk, and specify the disk's size. The names of the virtual disk files contain the name of the virtual machine. On the slide, the virtual machine name is Server1. You can also site the disk at a specific virtual SCSI target ID and logical unit number (LUN).


Remote Direct Memory Access (RDMA) enables the access between host memory directly without involving the CPU on the systems. RDMA enables high-throughput and low-latency networking between two hosts and is used in parallel computing clusters. Thus it allows the transfer of data to or from application memory and requires no work to be done by the CPUs. In essence, the application data is written directly to the network bypassing the host.

vRDMA is the VMware implementation of RDMA inside a guest virtual machine, thereby improving networking performance.

About Virtual Disks

Slide 3-13

A virtual machine usually has at least one virtual disk.



Sample virtual disk definition:

Virtual disk size:	8 GB
Datastore:	MyVMFS
Virtual disk node:	0:0
Virtual storage adapter:	LSI Logic SAS
Virtual disk files:	Server1.vmdk and Server1-flat.vmdk
Default disk mode:	Snapshots allowed
Optional disk mode:	Independent: Persistent or nonpersistent
Disk provisioning policy:	Thick provision lazy zeroed, thick provision eager zeroed, or thin provision

About Thick-Provisioned Virtual Disks


Slide 3-14

Thick provisioning uses all the defined disk space at the creation of the virtual disk:

- Virtual machine disks consume all the capacity, as defined at creation, regardless of the amount of data in the guest operating system file system.

Eager-zeroed or lazy-zeroed:

- Every block in an eager-zeroed thick-provisioned disk is prefilled with a zero.
- Every block in a lazy-zeroed thick-provisioned disk is filled with a zero when data is written to the block.



When you create a virtual disk, these virtual disk types are available:

- **Thick Provision Lazy Zeroed:** Space required for the virtual disk is allocated during creation. Data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine. This type is the default disk type.
- **Thick Provision Eager Zeroed:** Space required for the virtual disk is allocated during creation. Data remaining on the physical device is zeroed out when the disk is created.
- **Thin Provision:** A thin-provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can expand to the maximum capacity allocated to it.

About Thin-Provisioned Virtual Disks

Slide 3-15

Thin provisioning enables virtual machines to use storage space as needed:


- Thin-provisioned virtual machine disks consume only the capacity needed to hold the current files.
- A virtual machine sees the full allocated disk size at all times.

You can mix thick and thin formats.

Full reporting and alerts help manage allocations and capacity.

More efficient use of storage:

- Virtual disk allocation: 140 GB
- Available datastore capacity: 100 GB
- Used storage capacity: 80 GB



Thin provisioning provides alarms and reports that track allocation versus current usage of storage capacity. Thin provisioning enables storage administrators to optimize the allocation of storage for virtual environments. Thin provisioning enables users to optimally but safely use available storage space through overallocation.

Thin provisioning is often used with storage array deduplication to improve storage use and to back up virtual machines.

The following table identifies the differences between the virtual disk options. The differences between the time it takes to create the virtual disk type, how block allocation and zeroing are performed, and how the virtual disk is to be laid out on disk are compared.

	Thick Provisioned Lazy Zeroed	Thick Provisioned Eager Zeroed	Thin Provision
Creation time	Fast.	Slow and proportional to disk size.	Fastest.
Block allocation	Fully preallocated.	Fully preallocated.	Allocated and zeroed out on demand upon first write to block.
Virtual disk layout	Higher chance of contiguous file blocks.	Highest chance of contiguous file blocks.	Layout varies according to dynamic state of the volume at time of block allocation.
Zeroing of allocated file blocks	File blocks are zeroed out when each block is first written to.	File blocks are allocated and zeroed out when disk is created.	File blocks are zeroed out when blocks are allocated.


About Virtual Networks

Slide 3-16

A virtual network enables communication between virtual machines and physical machines.

When you configure networking for a virtual machine, you select or change the following items:

- The network adapter type
- The network connection
- Whether to connect to the network when the virtual machine powers on



For information about vSphere networking, see *vSphere Networking* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.


About Virtual Network Adapters (1)

Slide 3-17

When you configure a virtual machine, you can add network adapters (NICs) and specify the adapter type. Whenever possible, select VMXNET3.

Supported network adapter types:

- Flexible: Can function as either a Vlance or VMXNET adapter.
- E1000-E1000E: High-performance adapter available for only some guest operating systems.
- VMXNET, VMXNET2, and VMXNET3 are VMware drivers that are available only with VMware Tools.



virtual machine

The types of network adapters that are available depend on the following factors:

- The virtual machine compatibility level (or hardware version), which depends on the host that created or most recently updated it. For example, the VMXNET3 virtual NIC requires hardware version 7 (ESX/ESXi 4.0 or later).
- Whether the virtual machine compatibility is updated to the latest version for the current host.
- The guest operating system

The following NIC types are supported:

- E1000E: Emulated version of the Intel 82574 Gigabit Ethernet NIC. E1000E is the default adapter for Windows 8 and Windows Server 2012.
- E1000: Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in most newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later
- Flexible: Identifies itself as a Vlance adapter when a virtual machine boots, but initializes itself and functions as either a Vlance or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlance adapter to the higher performance VMXNET adapter.


- **Vlance:** Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in 32-bit legacy guest operating systems. A virtual machine configured with this network adapter can use its network immediately.
- **VMXNET2 (Enhanced):** Based on the VMXNET adapter but provides high-performance features commonly used on modern networks, such as jumbo frames and hardware offloads. VMXNET2 (Enhanced) is available only for some guest operating systems on ESX/ESXi 3.5 and later.
- **VMXNET3:** A paravirtualized NIC designed for performance. VMXNET3 offers all the features available in VMXNET2 and adds several new features, such as multiqueue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery.

About Virtual Network Adapters (2)

Slide 3-18

Supported network adapter types:

- SR-IOV passthrough: The virtual machine and the physical adapter exchange data without using the VMkernel as an intermediary:
 - Limited guest operating system support
- VMware vSphere® DirectPath I/O™: vSphere DirectPath I/O allows virtual machine access to physical PCI network functions on platforms with an I/O memory management unit.
- vRDMA: vRDMA is a paravirtualized device that provides improved virtual device performance. It provides an RDMA-like interface for vSphere guests.



virtual machine

- SR-IOV passthrough: Representation of a virtual function on a physical NIC with SR-IOV support. This adapter type is suitable for virtual machines that require more CPU resources or where latency might cause failure. The virtual machine and the physical adapter exchange data without using the VMkernel as an intermediary. If certain virtual machines are sensitive to network delay, SR-IOV can provide direct access to the virtual functions of supported physical NICs, bypassing the virtual switches and hence reducing overhead.
- SR-IOV passthrough is available in ESXi 5.5 and later for Red Hat Enterprise Linux 6 and later, and Windows Server 2008 R2 with SP2. An operating system release might contain a default virtual function driver for certain NICs, while for others you must download and install it from a location provided by the vendor of the NIC or of the host.
- vSphere DirectPath I/O allows virtual machine access to physical PCI network functions on platforms with an I/O Memory Management Unit. Passthrough devices provide the means to more efficiently use resources and improve performance in your environment. You can configure a passthrough PCI device on a virtual machine in vSphere Web Client. The following features are unavailable for virtual machines configured with vSphere DirectPath I/O:
 - Hot adding and removing of virtual devices
 - Suspend and resume

- Record and replay
 - Fault tolerance
 - High availability
 - vSphere DRS: Limited availability. The virtual machine can be part of a cluster, but cannot migrate across hosts.
 - Snapshots
- RDMA has been exploited for accelerating ESXi hypervisor services like vSphere vMotion. VMware developed a paravirtual device driver called Virtual RDMA for RDMA-capable fabrics. vRDMA is a paravirtualized device that provides improved virtual device performance. vRDMA provides an RDMA-like interface for vSphere guests. vRDMA allows multiple guests to access the RDMA device by using Verbs API, an industry-standard interface. A set of these Verbs was implemented to expose an RDMA-capable guest device (vRDMA) to applications. The applications can use the vRDMA guest driver to communicate with the underlying physical device. vRDMA supports RDMA, providing the following functions:
 - OS bypass
 - Zero-copy
 - Low latency and high bandwidth
 - Less power usage and faster data access

About Miscellaneous Devices

Slide 3-19

A virtual machine must have a vCPU and virtual memory. The addition of other virtual devices makes the virtual machine more useful.

CD/DVD drive:

- Connect to CD, DVD, or ISO image.

USB 3.0:

- Smart-card readers

Floppy drive:


- Connect a virtual machine to a floppy drive or a floppy image.

Generic SCSI devices:

- A virtual machine can be connected to additional SCSI adapters.

vGPUs:

- Enable a virtual machine to use GPUs on the physical host for high-computation activities.



Virtual CPU and virtual memory are the minimum required virtual hardware. Having a virtual hard disk and virtual NICs makes the virtual machine more useful. Other virtual hardware for a virtual machine includes a virtual CD/DVD drive, a virtual floppy drive, and generic virtual SCSI devices. The virtual CD/DVD drive or floppy drive can point to the following devices:

- The CD/DVD drive on the ESXi host
- CD/DVD ISO image (.iso) or floppy (.flp) images
- The CD/DVD or floppy drive on your local system

You can map the virtual machine's CD/DVD drive either to a physical drive or to an ISO file for your CD/DVD drive. An ISO file is a byte-for-byte copy of a CD or a DVD that has been ripped. These virtual CDs or DVDs can be accessed remotely and are usually faster than physical CDs or DVDs.

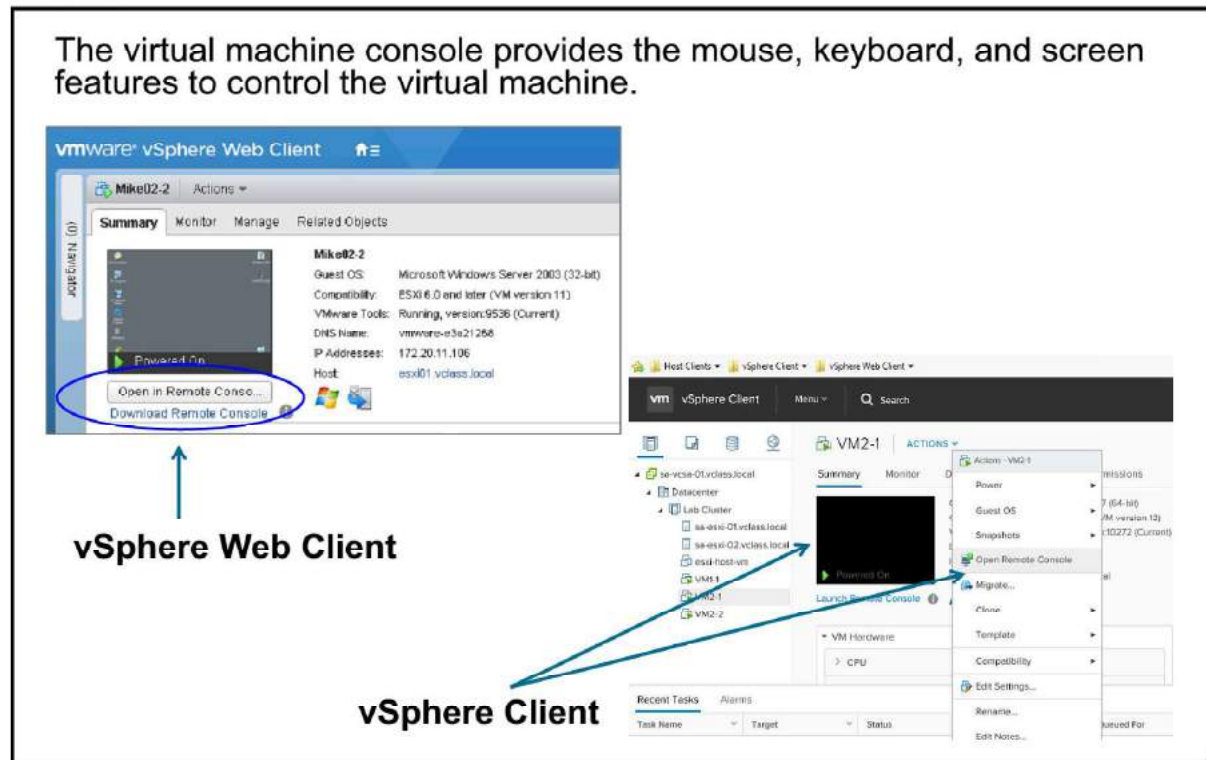
You can also add generic SCSI devices to your virtual machine. You can connect these devices to the virtual SCSI adapters on your virtual machine.

For information about creating a virtual machine, see *vSphere Virtual Machine Administration* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

About the Virtual Machine Console

Slide 3-20

The virtual machine console provides the mouse, keyboard, and screen features to control the virtual machine.



You can open the virtual machine console from vSphere Client and vSphere Web Client.

The virtual machine console, available in both UIs, provides the mouse, keyboard, and screen features to control a virtual machine. The console is launched through the methods shown in the slide.

vSphere Web Client requires the installation of a plug-in to enable the Windows Authentication option.

You use the virtual machine console to access the BIOS of the virtual machine, install an operating system on a virtual machine, power the virtual machine on and off, and reset the virtual machine.

The virtual machine console is normally not used to connect to the virtual machine for daily tasks. Remote Desktop Connection, Virtual Network Connection, or other options are normally used to connect to the virtual desktop. The virtual machine console is used for tasks such as power cycling, configuring hardware, and troubleshooting network issues.

Review of Learner Objectives

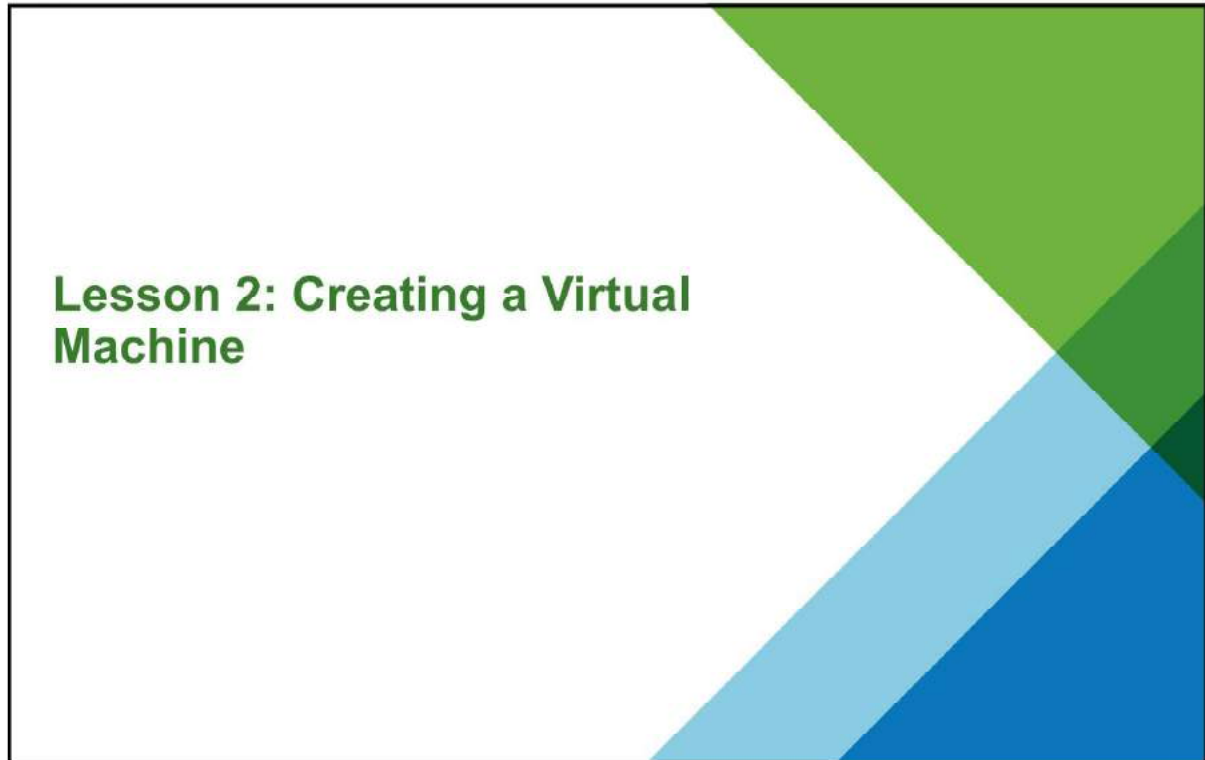
Slide 3-21

You should be able to meet the following objectives:

- Identify files that make up a virtual machine
- Compare virtual machine hardware version 13 to other versions
- Describe components of a virtual machine
- Identify the various methods to access a virtual machine console
- Identify the virtual network adapters and highlight the enhanced VMXNET3
- Discuss the vRDMA and vNVMe features
- Compare and contrast the types of virtual disk provisioning

Lesson 2: Creating a Virtual Machine

Slide 3-22



Learner Objectives

Slide 3-23

By the end of this lesson, you should be able to meet the following objectives:

- Create, provision, and remove a virtual machine
- Explain the importance of VMware Tools
- Describe how to import a virtual appliance Open Virtual Machine Format (OVF) template
- Manage VMware Tools
- Explain troubleshooting OS installation and VMware Tools problems

About Provisioning Virtual Machines

Slide 3-24

You can create virtual machines in several ways:

- Use the New Virtual Machine wizard to create virtual machines.
- Deploy virtual machines, virtual appliances, and vApps stored in OVF.
- Use a CentOS, Linux, or Windows template in a VMware vCloud® Air™ catalog to create virtual machines.

VMware provides several methods to provision vSphere virtual machines. The optimal method for your environment depends on factors such as the size and type of your infrastructure and the goals that you want to achieve.

You can use the New Virtual Machine wizard to create a single virtual machine if no other virtual machines in your environment have the requirements you are looking for, such as a particular operating system or hardware configuration. For example, you might need a virtual machine that is configured only for testing purposes. You can also create a single virtual machine, install an operating system on it, and use that virtual machine as a template from which to clone other virtual machines.

Deploy virtual machines, virtual appliances, and vApps stored in Open Virtual Machine Format (OVF) to use a preconfigured virtual machine. A virtual appliance is a virtual machine that typically has an operating system and other software preinstalled. You can deploy virtual machines from OVF templates that are located on local file systems, (for example, local disks, such as C:), removable media (for example, CDs or USB keychain drives), shared network drives, or from URLs.

The New Virtual Machine wizard that is run from vSphere Client enables you to select between the **Typical** and **Custom** configuration types. Running the wizard from vSphere Web Client does not offer a choice of configurations.

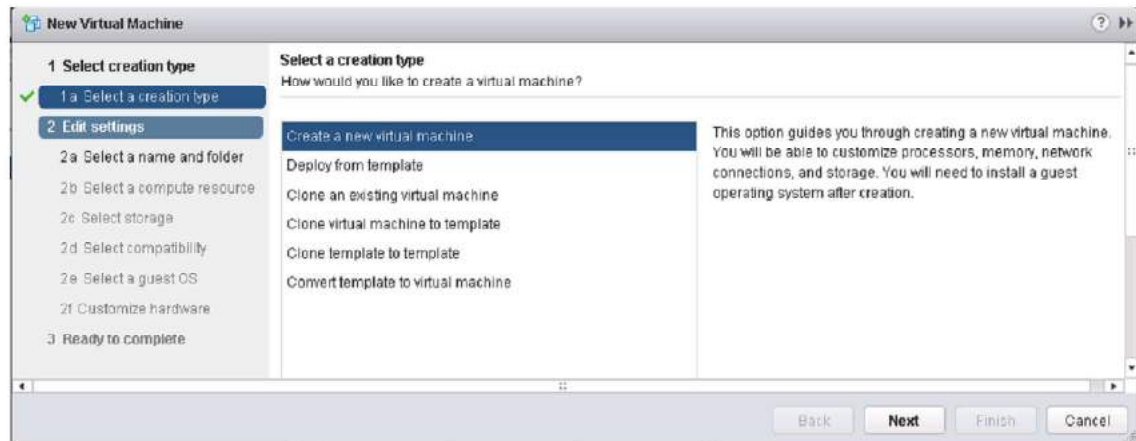
Regardless of where it is run, the New Virtual Machine wizard prompts you for standard information:

- The virtual machine name and folder. The resource on which the virtual machine will run: a host, a cluster, a vApp, or a resource pool. The virtual machine will have access to the resources of the selected object.
- The datastore on which to store the virtual machine's files. Each datastore might have a different size, speed, availability, and other properties. The available datastores are accessible from the destination resource that you selected.
- The guest operating system to be installed into the virtual machine.
- The number of NICs, the network to connect to, and the network adapter type.
- Virtual disk provisioning choice.

Creating Virtual Machines with the New Virtual Machine Wizard (1)

Slide 3-25

You can use the New Virtual Machine wizard in vSphere Web Client to create a virtual machine.

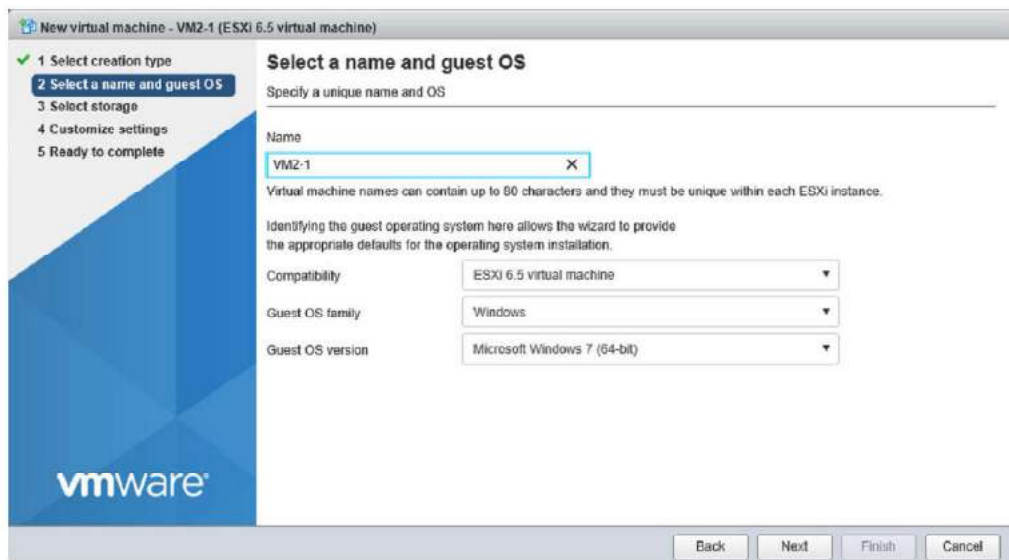


Regardless of which client UI you are using, a wizard is available to walk you through the deployment of a new virtual machine. The slide includes a screen capture of the wizard in vSphere Web Client.

Creating Virtual Machines with the New Virtual Machine Wizard (2)

Slide 3-26

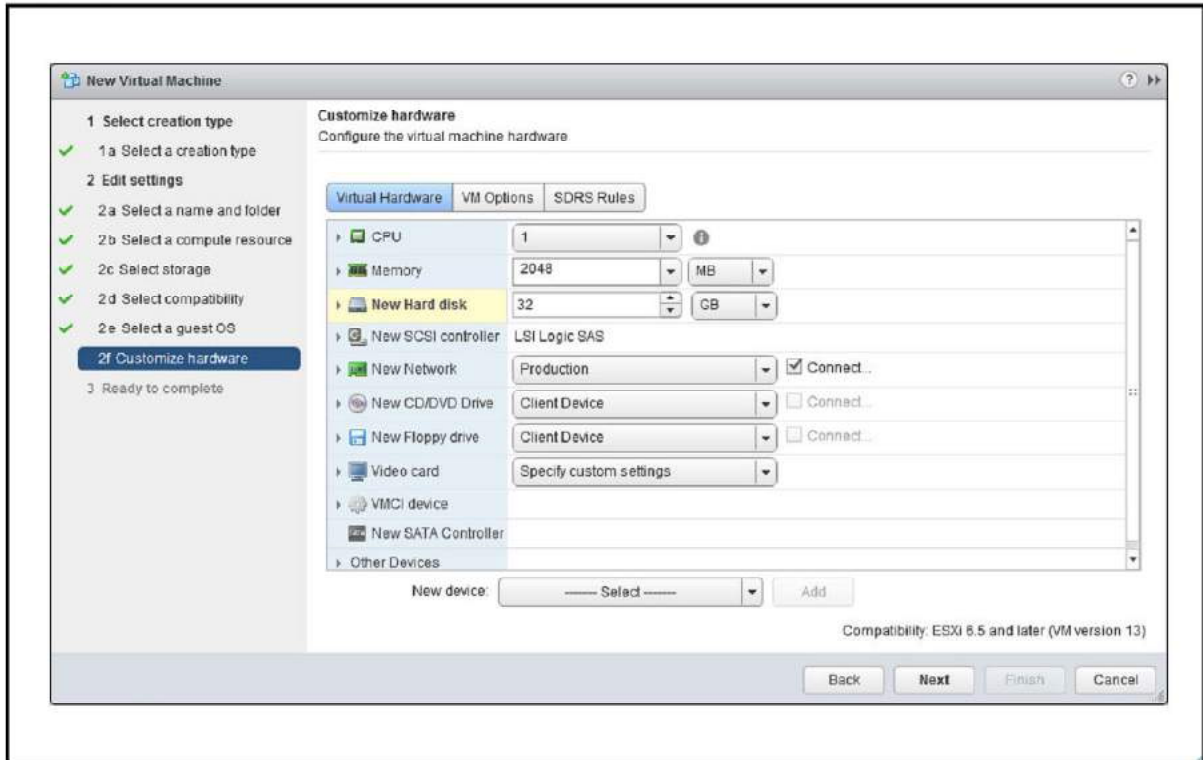
You can use the New Virtual Machine wizard in VMware Host Client to create a virtual machine.



The slide shows the New Virtual Machine wizard in VMware Host Client. This wizard is also available in vSphere Client.

New Virtual Machine Wizard Settings

Slide 3-27




The slide depicts the configuration settings for the virtual machine that are deployed. These values are populated by the previous choices made for the operating system.

Installing the Guest Operating System

Slide 3-28

Installing a guest operating system in your virtual machine is like installing it on a physical computer.



Installing a guest operating system is the same in a virtual machine as it is in a physical computer. To install the guest operating system, you interact with the virtual machine through the virtual machine console. Using vSphere Web Client, you can attach a CD, DVD, or ISO image containing the installation image to the virtual CD/DVD drive. On the slide, the Windows Server 2008 guest operating system is being installed. You can use vSphere Web Client to install a guest operating system. You can install a guest operating system from a CD or from an ISO image. Installing from an ISO image is typically faster and more convenient than a CD installation. For more information about installing guest operating systems, see *vSphere Virtual Machine Administration* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>. For more about the supported guest operating systems, see *Hardware and Guest Operating System Compatibility Guides* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

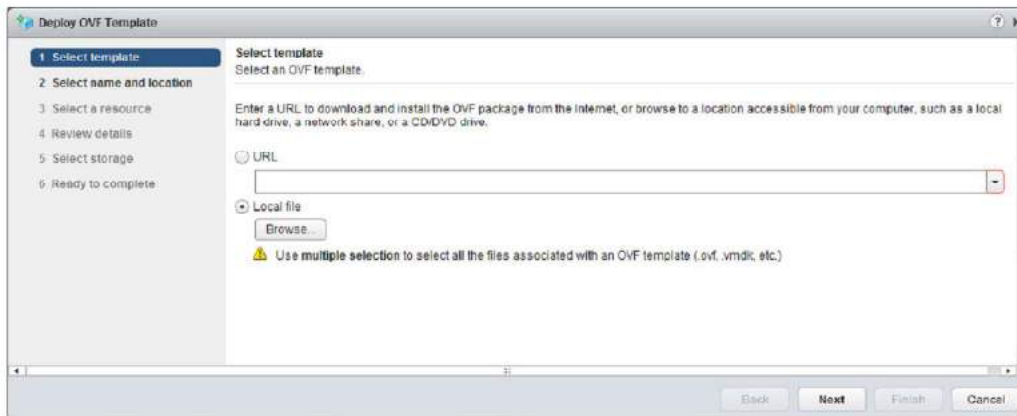
Deploying OVF Templates

Slide 3-29

You can deploy any virtual machine or a virtual appliance stored in OVF.

Virtual appliances are preconfigured virtual machines:

- They are usually designed for a single purpose, for example, a safe browser or firewall.
- They are available from VMware Solution Exchange.



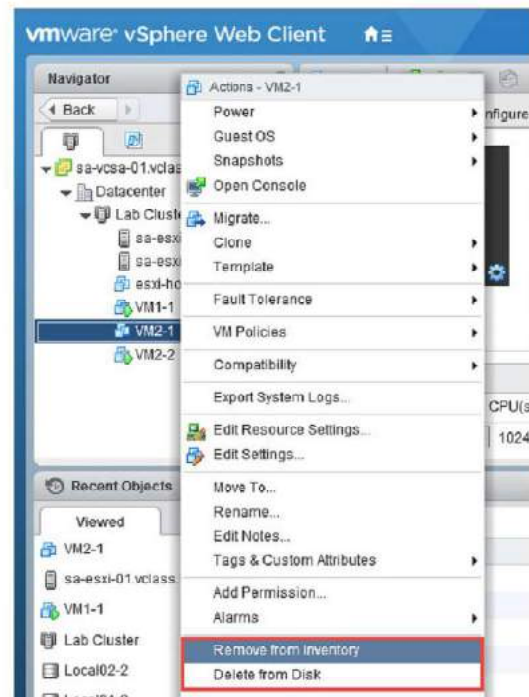
A virtual appliance is a preconfigured virtual machine that typically includes a preinstalled guest operating system and other software. A virtual appliance is usually designed for a specific purpose, for example, to provide a secure Web browser, a firewall, or a backup-and-recovery utility. A virtual appliance can be added, or imported, to your vCenter Server system inventory or ESXi inventory. Virtual appliances can be imported from Web sites such as the VMware Virtual Appliance Marketplace at https://solutionexchange.vmware.com/store/category_groups/19. Virtual appliances are deployed as OVF templates. OVF is a platform-independent, efficient, extensible, and open packaging and distribution format for virtual machines. OVF files are compressed, enabling faster downloads. vSphere Web Client validates an OVF file before importing it and ensures that it is compatible with the intended destination server. If the appliance is incompatible with the selected host, you cannot import it.

Removing a Virtual Machine

Slide 3-30

You can remove a virtual machine in the following ways:

- Remove from the inventory:
 - This type of removal unregisters the virtual machine.
 - The virtual machine's files remain on the disk.
 - The virtual machine can later be registered (added) to the inventory.
- Delete from disk:
 - All virtual machine files are permanently deleted from the virtual machine datastore.



Removing a virtual machine from the vCenter Server inventory unregisters the virtual machine from the host and vCenter Server. This process does not delete the virtual machine files from the datastore. Virtual machine files remain at the same storage location, and the virtual machine can be re-registered in the Datastore Browser.

However, the delete function permanently removes the virtual machine files from the data store as well as unregistering the virtual machine from host and vCenter Server.

About VMware Tools

Slide 3-31

VMware Tools is a suite of utilities that enhance the performance of the virtual machine's guest operating system.

VMware Tools benefits:

- Device drivers:
 - SVGA display
 - VMXNET/VMXNET3
 - Balloon driver for memory management
 - Sync driver for quiescing I/O
- Increased graphics performance
- Improved mouse performance

VMware Tools features:

- Copying and pasting text, graphics, and files between the virtual machine and the client desktop
- Time synchronization
- Ability to shut down the virtual machine
- Guest authentication (vCenter Single Sign-On)

VMware Tools improves management of the virtual machine by replacing generic operating system drivers with VMware drivers tuned for virtual hardware. You install VMware Tools into the guest operating system. Although the guest operating system can run without VMware Tools, you lose important features and convenience. VMware recommends installing VMware Tools. When you install VMware Tools, you install these items:

- The VMware Tools service: This service synchronizes the time in the guest operating system with the time in the host operating system.
- A set of VMware device drivers, with additional Perfmon monitoring options.
- A set of scripts that helps you automate guest operating system operations. You can configure the scripts to run when the virtual machine's power state changes.

VMware Tools enhances the performance of a virtual machine and makes possible many of the ease-of-use features in VMware products:

- Significantly faster graphics performance and Windows Aero on operating systems that support Aero.
- The Unity feature, which enables an application in a virtual machine to appear on the host desktop like any other application window.

- Shared folders between host and guest file systems.
- Copying and pasting text, graphics, and files between the virtual machine and the host or client desktop.
- Improved mouse performance. •
- Synchronization of the clock in the virtual machine with the clock on the host or client desktop.
- Scripting that helps automate guest operating system operations.
- The ability to shut down the virtual machine.

Although the guest operating system can run without VMware Tools, many VMware features are not available until you install VMware Tools. For example, if you do not have VMware Tools installed in your virtual machine, you cannot use the shutdown or restart options from the toolbar. You can use only the power options.

Managing VMware Tools

Slide 3-32

The version of VMware Tools distributed with vSphere 6.5 is 10.1.

VMware Tools 10.1 provides the following features:

- Digital signature verification
- Three supported guest operating system ISO images
- Product locker for storing ISOs

Additional ISO images for other operating systems can be downloaded from VMware.

vSphere 6.5 provides some enhancements to VMware Tools by implementing security, reducing distribution media volume requiring less space, offloading ISOs to a on-disk product locker and providing access for other ISO images needed through download from the VMware Web site.

VMware Tools: Supported ISO Images

Slide 3-33

The following ISOs are included with vSphere 6.5:

- `windows.iso`: For Vista and later guests
- `winPreVista.iso`: For Windows 2000, XP, and Server 2003 guests
- `linux.iso`: For Linux 3.x and later guests, CentOS 6 and later, Debian 6 and later, Ubuntu 10 and later

VMware Tools for other guest operating systems, such as FreeBSD, Solaris, and Mac OS X, can be downloaded from My VMware at <https://download.vmware.com>.

VMware has reduced the number of VMware Tools ISOs that are distributed to only three.

These enhancements greatly reduce the number of ISO images required to be retained and managed, thereby making it easier to manage and reduce your host deployment size by both the number provided and through the availability of a disk-based product locker.

If you require tools for other operating systems, they are available at the VMware Web site for download and deployment to your product locker.

VMware Tools can contain a hash digital signature. The VMware Tools hash and signature are used to verify the ISO file authenticity, integrity, and origin. When installing VMware Tools, the Secure Hash Algorithm is used to validate the integrity of the ISO file and the associated digital signature is used to validate the author:

- `iso.sha`: Used to enforce and verify ISO integrity
- `iso.sig`: Used to enforce and verify ISO authentication

You must restart the host for the changes to take effect.

Troubleshooting OS Installation Failures in Virtual Machines

Slide 3-34

Problems:

- The installation of a 64-bit operating system cannot start.
- The installation of a 64-bit guest operating system stops responding at the Setup is starting the Windows screen.
- The installation of a 64-bit operating system cannot complete.

Resolutions:

1. Verify that the guest operating system that you are attempting to install is fully certified by VMware.
2. Verify that your ESX/ESXi host meets the hardware and firmware requirements for running 64-bit virtual machines.
3. If your ESX/ESXi host uses Intel processors, verify that virtualization technology is enabled in the BIOS.
4. Verify that the correct guest operating system is selected.

Troubleshooting a Failed VMware Tools Installation on a Guest Operating System

Slide 3-35

Problems:

- VMware Tools installation errors before completion.
- VMware Tools installation fails to complete.
- Unable to complete VMware Tools for Windows or Linux installation.
- VMware Tools hangs when installing or reinstalling.

Solutions:

1. Verify that the guest operating system that you are trying to install is fully certified by VMware.
2. Verify that the correct operating system is selected.
3. Verify that the ISO image is not corrupted.
4. If installing on a Windows operating system, ensure that you are not experiencing problems with your Windows registry.
5. If installing on a 64-bit Linux guest operating system, verify that no dependencies are missing.

Validate that each resolution shown on the slide is true for your environment. The steps are ordered in the most appropriate sequence to isolate the issue and identify the proper resolution. For more detailed information, see VMware knowledge base article 1003908 at <http://kb.vmware.com/kb/1003908>.

Lab 3: Deploying and Configuring a Virtual Machine

Slide 3-36

Create and prepare a virtual machine for use

1. Create a Virtual Machine
2. Install a Guest Operating System
3. Install VMware Tools
4. Install Files

Review of Learner Objectives

Slide 3-37

You should be able to meet the following objectives:

- Create, provision, and remove a virtual machine
- Explain the importance of VMware Tools
- Describe how to import a virtual appliance OVF template
- Manage VMware Tools
- Explain troubleshooting OS installation and VMware Tools problems

Key Points

Slide 3-38


- Virtual machines can be provisioned by using various methods:
 - You can use the New Virtual Machine wizard in vSphere Client, vSphere Web Client, and VMware Host Client to create and clone virtual machines.
 - You can create a virtual machine by deploying an OVF template.
- VMware Tools increases the performance of the virtual machine's guest operating system.

Questions?

MODULE 4

vCenter Server

Slide 4-1



You Are Here

Slide 4-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. **vCenter Server**
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Importance

Slide 4-3

vCenter Server enables you to centrally manage multiple ESXi hosts and their virtual machines. Failure to properly install, configure, and manage vCenter Server might result in reduced administrative efficiency or possible ESXi host and virtual machine downtime.

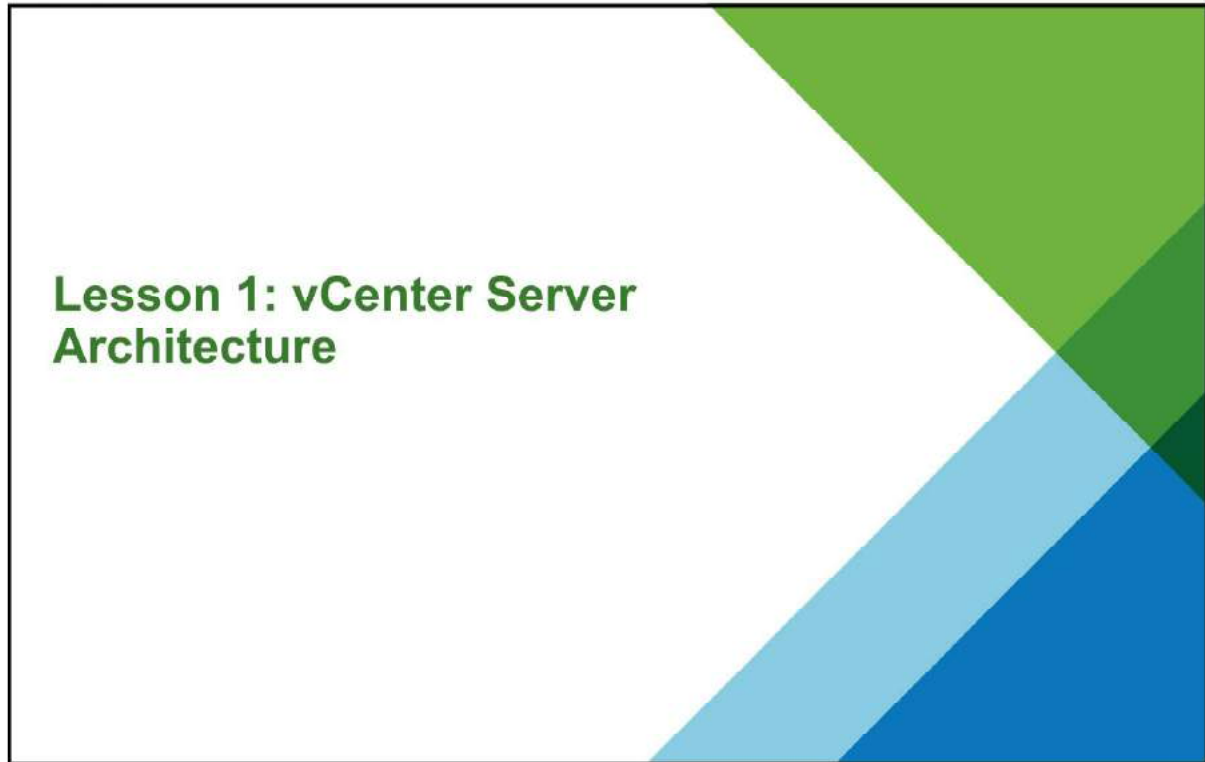
Module Lessons

Slide 4-4

- | | |
|-----------|---|
| Lesson 1: | vCenter Server Architecture |
| Lesson 2: | Deploying, Backing Up, and Restoring vCenter Server Appliance |
| Lesson 3: | vSphere Clients |
| Lesson 4: | Managing the vCenter Server Inventory |
| Lesson 5: | vCenter Server Roles and Permissions |

Lesson 1: vCenter Server Architecture

Slide 4-5



Lesson 1: vCenter Server Architecture

Learner Objectives

Slide 4-6

By the end of this lesson, you should be able to meet the following objectives:

- Describe the vCenter Server architecture
- Discuss how ESXi hosts communicate with vCenter Server
- Identify the vCenter Server services, components, and modules
- Explain Platform Services Controller
- Discuss the REST-based API
- Describe vCenter Server High Availability

Overview of vCenter Server Appliance (1)

Slide 4-7

vCenter Server Appliance is a preconfigured Linux-based virtual machine that is optimized for running vCenter Server and the associated services.

vCenter Server Appliance reduces the deployment time of vCenter Server and the associated services, and provides a low-cost alternative to the Windows-based vCenter Server installation.

The vCenter Server Appliance package contains the following software:

- VMware Photon™ OS 1.0
- The Platform Services Controller group of infrastructure services
- The vCenter Server group of services
- PostgreSQL

vCenter Server is a service that acts as a central administrator for ESXi hosts connected in a network. vCenter Server enables you to pool and manage the resources of multiple hosts. You can deploy the vCenter Server system as a vCenter Server Appliance system. vCenter Server Appliance is a preconfigured Linux-based virtual machine optimized for running vCenter Server and the vCenter Server components. You can deploy vCenter Server Appliance on hosts running ESXi 5.5 or later.

Overview of vCenter Server Appliance (2)

Slide 4-8

vCenter Server Appliance uses the embedded PostgreSQL database that has the scalability of up to 2,000 hosts and 35,000 virtual machines. During the deployment, you can choose the vCenter Server Appliance size for your vSphere environment size and the storage size for your database requirements.

Starting with vSphere 6.5, the vCenter Server services in vCenter Server Appliance include:


- vSphere Update Manager extension.
- vCenter Server Appliance supports high availability.
- vCenter Server Appliance and Platform Services Controller support file-based backup and restore.

vCenter Server provides advanced features, such as vSphere DRS, vSphere HA, vSphere Fault Tolerance, vSphere vMotion, and vSphere Storage vMotion.

About the vCenter Server Management Platform

Slide 4-9

vCenter Server is an application service that acts as a central administration point for ESXi hosts and their virtual machines connected on a network. This service directs the actions of virtual machines and hosts.



vCenter Server Services

Slide 4-10

The vCenter Server group of services contains the following functions:

- vCenter Server
- PostgreSQL
- vSphere Web Client (server)
- vSphere Auto Deploy
- vSphere ESXi Dump Collector
- vSphere Syslog Collector
- vSphere Update Manager



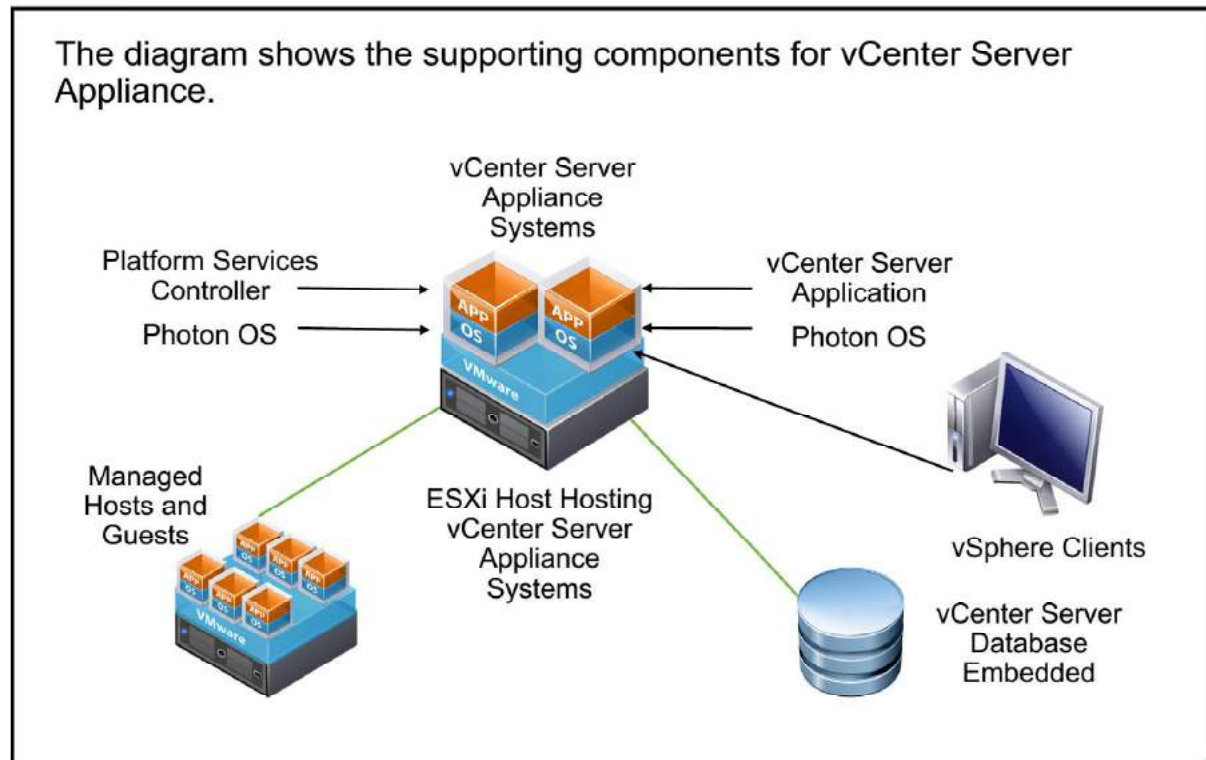
You cannot distribute these vCenter Server functions across multiple servers. When you deploy vCenter Server Appliance, all of these features are included.

vCenter Server includes these services and interfaces:

- Core services include management of resources and virtual machines by the Inventory service, task scheduling, statistics logging, management of alarms and events, virtual machine provisioning, and host and virtual machine configuration.
- The vCenter Lookup Service contains topology information about the vSphere infrastructure, enabling vSphere components to connect to each other securely. Services, such as the Inventory Service and vCenter Server, register with the vCenter Lookup Service so that other vSphere components, like vSphere Web Client, can find them.
- Distributed services include vSphere vMotion, vSphere DRS, and vSphere HA, which are installed with vCenter Server.
- Additional services are packaged separately from the base product and require separate installation, for example, vSphere Update Manager and vRealize Orchestrator. No additional license is necessary.
- A database interface provides access to the vCenter Server database.
- vCenter Server provides access to the ESXi host through a vCenter Server agent, which is started on the host when it is added to the vCenter Server inventory.
- vCenter Single Sign-On provides access to domain user accounts.
- VMware vSphere® API with VMware vSphere® Management SDK provides an interface for writing custom applications that access vCenter Server functionality.

vCenter Server Appliance Architecture

Slide 4-11



The vCenter Server architecture relies on the following components:

- vSphere Web Client and vSphere Client: vSphere Client is used to connect directly to ESXi hosts. vSphere Web Client connects directly to vCenter Server. When an ESXi host is managed by vCenter Server, administrators should always use vCenter Server and vSphere Web Client to manage that host.
- vCenter Server database: The most critical component is the vCenter Server database. The database stores the inventory items, security roles, resource pools, performance data, and other critical information for vCenter Server.
- vCenter Single Sign-On provides a security domain defined in your vSphere environment. Authentication is performed by the vCenter Single Sign-On server. The vCenter Single Sign-On server can be configured to authenticate against multiple user repositories, also called identity sources, such as an Active Directory domain.
- Managed hosts: vCenter Server enables you to manage ESXi hosts and the virtual machines that run on them.

vCenter Server Appliance Scalability

Slide 4-12

vCenter Server Appliance scales beyond the capacity of vCenter Server installed on a Windows machine.

Metric	Windows vCenter Server	vCenter Server Appliance
Hosts per vCenter Server system	1,000	2,000
Virtual machines per vCenter Server system	10,000	35,000
Powered-on virtual machines per vCenter Server	25,000	25,000
Hosts per cluster	64	64
Virtual machines per cluster	8,000	8,000
Database	Must be Oracle database or SQL for full scalability	Embedded Postgres
Linked Mode	Yes	Yes


vCenter Server Deployment Options

Slide 4-13

vCenter Server Appliance is functionally equivalent to vCenter Server installed on a Windows server:

- vCenter Server Appliance can be configured in the following ways:
 - As an embedded system with an internal Platform Services Controller instance
 - As a distributed system with an external Platform Services Controller instance

vCenter Server Appliance supports Enhanced Linked Mode.



You can deploy vCenter Server Appliance on an ESXi host. vCenter Server Appliance is a preconfigured Linux-based virtual machine that is optimized for running vCenter Server.

You can download the vCenter Server Appliance installer, install the Client Integration Plug-in, and deploy vCenter Server Appliance. During the deployment of the appliance, you select whether you want to deploy vCenter Server Appliance with an external Platform Services Controller or with an embedded Platform Services Controller:

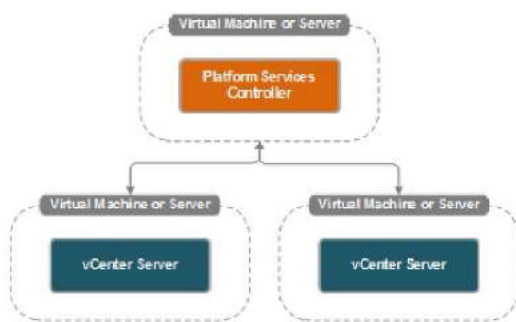
- vCenter Server with an embedded Platform Services Controller: All services bundled with the Platform Services Controller appliance are deployed on the same host machine as vCenter Server. vCenter Server with an embedded Platform Services Controller is suitable for smaller environments with eight or less product instances.
- vCenter Server with an external Platform Services Controller: You must first deploy the Platform Services Controller appliance on one virtual machine or host and then deploy vCenter Server on another virtual machine or host. Platform Services Controller can be shared across many products. This configuration is suitable for larger environments with nine or more product instances.

Platform Services Controller Deployment Recommendations (1)

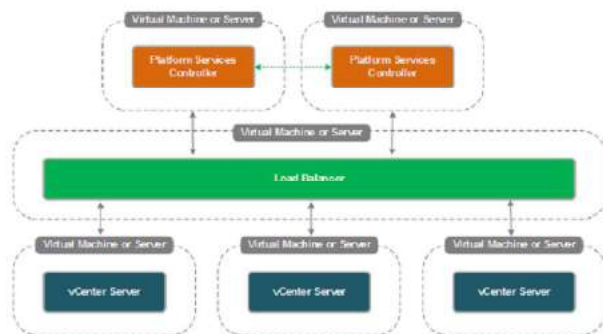
Slide 4-14

These deployment models are recommended for Platform Services Controller in Enhanced Linked Mode.

Platform Services Controller Without Load Balancer



Platform Services Controller with Load Balancer



For Enhanced Linked Mode with an external Platform Services Controller without vSphere HA, the Platform Services Controller is configured on a separate virtual machine and the vCenter Server systems are joined to that domain, providing the Enhanced Linked mode functionality.

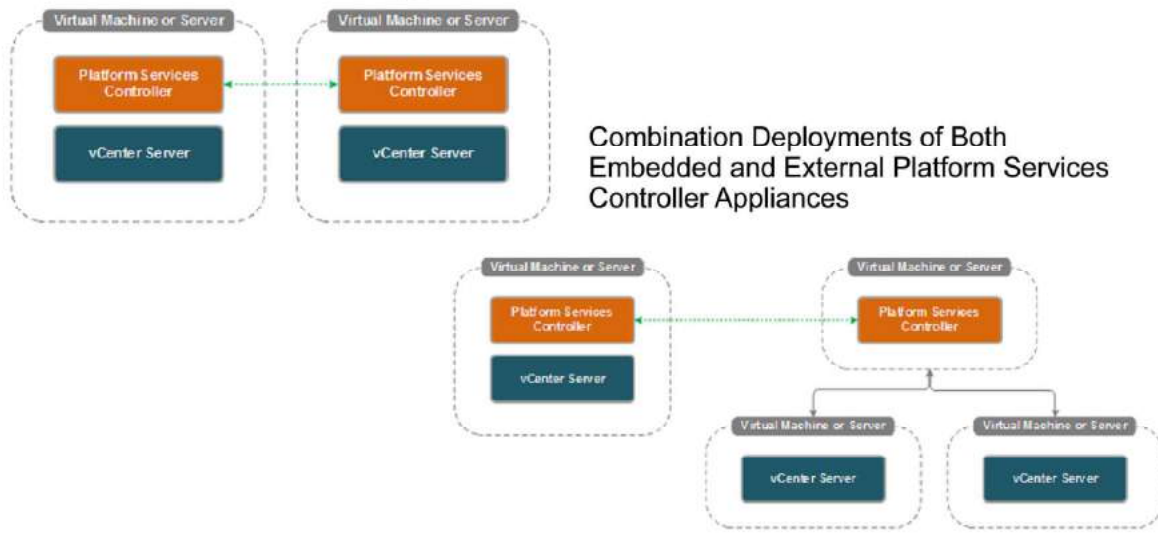
For Enhanced Linked Mode with an external Platform Services Controllers with vSphere HA, the Platform Services Controller appliances are configured on separate virtual machines and configured behind a load balancer to provide high availability to the configuration. The vCenter Server systems are joined to that domain by using the shared load balancer IP address, which provides the Enhanced Linked Mode functionality, but which is resilient to failures.

Platform Services Controller Deployment Recommendations (2)

Slide 4-15

These deployment models are not recommended for Platform Services Controller in Enhanced Linked Mode.

Enhanced Linked Mode with Embedded Platform Services Controller Appliances



For Enhanced Linked Mode with embedded Platform Services Controllers, the vCenter Server system is installed in an embedded configuration on the first server. Subsequent installations are configured in embedded mode but joined to an existing vCenter Single Sign-On domain. Linking embedded Platform Services Controller appliances is possible, but the configuration is not recommended. An external configuration for the Platform Services Controller appliance is preferred.


In combination deployments, linking an embedded Platform Services Controller and an external Platform Services Controller is possible, but the configuration is not recommended. An external configuration for the Platform Services Controller appliance is preferred.

Platform Services Controller Deployment Recommendations (3)

Slide 4-16

This deployment model is not recommended for Platform Services Controller in Enhanced Linked Mode.

External vCenter Server System Linked to an Embedded Platform Services Controller



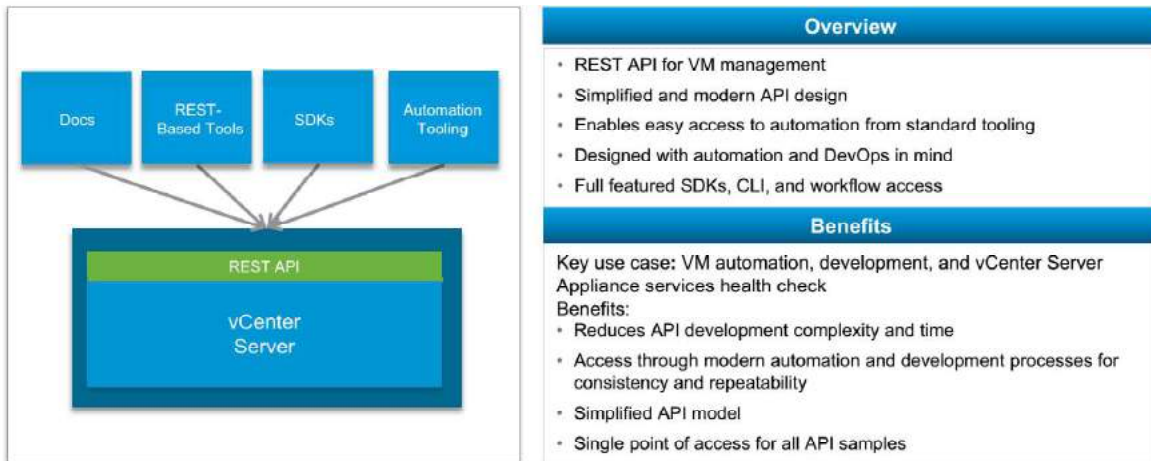
For an embedded Platform Services Controller and vCenter Server system linked with an external, standalone vCenter Server system, linking a second vCenter Server system to an existing embedded vCenter Server system and Platform Services Controller is possible. VMware does not recommend this configuration. An external configuration for the Platform Services Controller appliance is preferred.

For more information about Platform Services Controller recommendations, see VMware knowledge base article 2108548 at <http://kb.vmware.com/kb/2108548>. For more information about backing up and restoring vCenter Server 6 external deployment models, see VMware knowledge base article 2110294 at <http://kb.vmware.com/kb/2110294>.

vCenter Server APIs

Slide 4-17

vSphere 6.5 includes a developer and automation-friendly REST-based API and interfaces that simplify automation and development.



VMware provides SDKs in several programming languages that allow you to build service capability into your vCenter Server API client.

The following SDKs are included in this release:

- VMware vSphere Automation SDK for Java
- VMware vSphere Automation SDK for .NET
- VMware vSphere Automation SDK for REST
- VMware vSphere Automation SDK for Python
- VMware vSphere Automation SDK for Perl
- VMware vSphere Automation SDK for Ruby

Each of these SDKs contains:

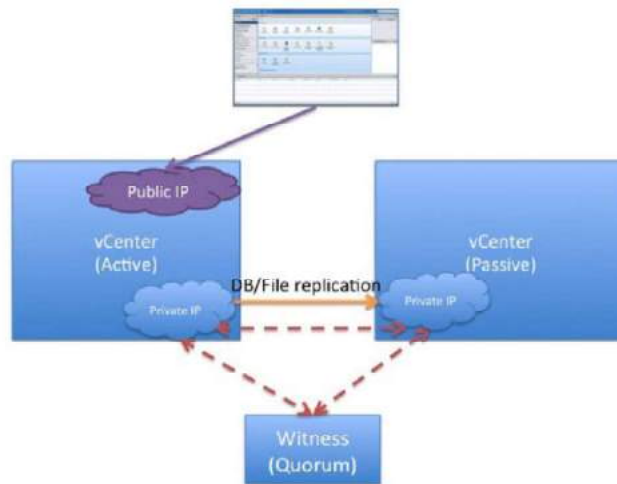
- Bindings to the vSphere Automation API
- vSphere Web Services API
- API Reference
- Programming guide
- Samples

High Availability for vCenter Server Appliance

Slide 4-18

High availability for vCenter Server Appliance protects against both hardware and software failures and ensures that your implementation can recover quickly.

- The protected node is called the active node.
- Two other appliance nodes are created: a passive node and a witness node.
- If the active node fails, the passive node takes over the role of the active node.
- The state of the active node is replicated to the passive node and captured in a PostgreSQL database and in the configuration files.



Downtime, whether planned or unplanned, brings considerable costs with it. However, solutions to ensure higher levels of availability have traditionally been costly, hard to implement, and difficult to manage.

VMware software makes it simpler and less expensive to provide higher levels of availability for important applications. With vSphere, organizations can easily increase the baseline level of availability provided for all applications and provide higher levels of availability more easily and cost effectively.

VMware vCenter Server® High Availability protects vCenter Server Appliance against host and hardware failures. The active-passive architecture of the solution can also help you reduce downtime significantly when you patch vCenter Server Appliance.

After some network configuration, you create a three-node cluster that contains Active, Passive, and Witness nodes. Different configuration paths are available. Your selection depends on your existing configuration.

Before you can configure vCenter Server High Availability, you must consider several factors. A vCenter Server Appliance deployment can use an internal or external Platform Services Controller. A brownfield deployment with components that use different versions of vSphere requires different considerations than a greenfield deployment that includes only vSphere 6.5 components. Resource and software requirements and the networking setup must also be considered carefully.

Review of Learner Objectives

Slide 4-19

You should be able to meet the following objectives:

- Describe the vCenter Server architecture
- Discuss how ESXi hosts communicate with vCenter Server
- Identify the vCenter Server services, components, and modules
- Explain Platform Services Controller
- Discuss the REST-based API
- Describe vCenter Server High Availability

Lesson 2: Deploying, Backing Up, and Restoring vCenter Server Appliance

Slide 4-20



Lesson 2: Deploying, Backing Up, and Restoring vCenter Server Appliance

Learner Objectives

Slide 4-21

By the end of this lesson, you should be able to meet the following objectives:

- Discuss the vCenter Server deployment models
- Deploy vCenter Server Appliance into an infrastructure
- Add license keys to vCenter Server
- Configure vCenter Server settings
- Create a vCenter Server backup
- Restore vCenter Server Appliance from backup

Preparing for vCenter Server Appliance Deployment (1)

Slide 4-22

Before deploying vCenter Server Appliance, you must complete several tasks:

- Verify that all vCenter Server Appliance system requirements are met.
- Prepare a vCenter Server database:
 - Use the included PostgreSQL database.
- For the first installation of vCenter Server Appliance, Platform Services Controller must be deployed before vCenter Server:
 - If you deploy vCenter Server Appliance with an embedded Platform Services Controller, this operation occurs automatically.
 - If you install vCenter Server Appliance with an external Platform Services Controller instance, you must first install Platform Services Controller and then install vCenter Server.

To ensure that your browser supports the Client Integration Plug-in, verify that you have one of the supported Web browsers. For more information, see *vSphere Installation and Setup* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Preparing for vCenter Server Appliance Deployment (2)

Slide 4-23

Before deploying vCenter Server Appliance, you must complete several tasks:


- You must provide the fully qualified domain name (FQDN) or the static IP of the host machine on which you are performing the install or upgrade. VMware recommends using the FQDN.
- You must verify that clocks on all machines on the vSphere network are synchronized.

vCenter Server Appliance Native UI Installer

Slide 4-24

With vSphere 6.5, a native application has been developed to facilitate the deployment of vCenter Server Appliance 6.5:

- A native application has been written for Windows, Linux, and Mac OS X and has no dependency on browsers or a plug-in.
- This GUI application performs validations and prechecks during the deployment to ensure that no mistakes are made and that a compatible environment is created.

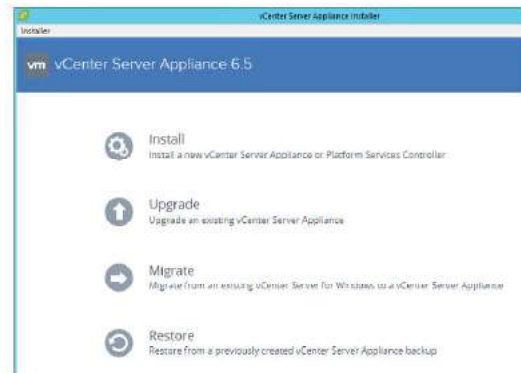


vCenter Server Install, Upgrade, Migrate, and Restore

Slide 4-25

The new UI deployment tool has the following featured options:

- **Install:** Installs a new vCenter Server Appliance or Platform Services Controller
- **Upgrade:** Upgrades an existing vCenter Appliance
- **Migrate:** Migrates an existing vCenter Server for Windows to a vCenter Server Appliance
- **Restore:** Restores from a previously created vCenter Server Appliance backup



vCenter Server Appliance Deployment

Slide 4-26

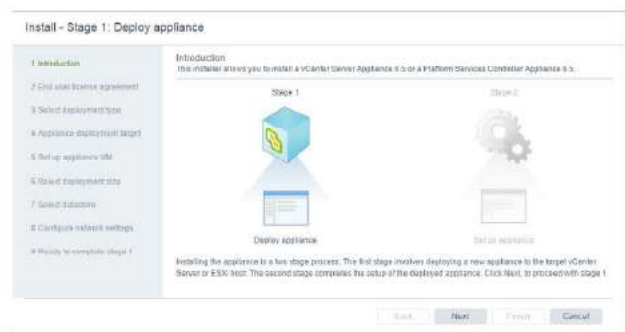
Installer support for Windows, Mac,
and Linux

vSphere Update Manager is included

vCenter Server Appliance and
Platform Services Controller
install is a two-stage process:

- Stage 1: Deploy OVF
- Stage 2: Configuration

Fully automatable by using
JSON templates, with Windows,
Linux, and Mac support



When you choose to deploy vCenter Server Appliance with an embedded Platform Services Controller, you deploy Platform Services Controller and vCenter Server as one appliance. A single Platform Services Controller appliance is suitable for deployments with eight or fewer vCenter Server systems.

When you choose to deploy vCenter Server Appliance with an external Platform Services Controller, the vCenter Server instance is deployed as two appliances:

1. You deploy Platform Services Controller.
2. You deploy vCenter Server and the vCenter Server components as another virtual appliance, and connect vCenter Server Appliance to the external Platform Services Controller.

vCenter Server Appliance Two-Stage Deployment

Slide 4-27

Stage 1: UI

- Accept the EULA.
- Select the deployment type.
- Connect to the target ESXi host or vCenter Server system to deploy vCenter Server Appliance.
- Define vCenter Server Appliance name and root password.
- Select deployment size (Mem/CPU) and storage size.
- Select datastore location (thin disk).
- Configure networking.

Stage 2: Deployment

- OVF is deployed to the ESXi host.
- Disks are configured.
- RPMs are installed (depending on Embedded, Platform Service Controller, vCenter Server deployment choice).
- Networking is configured.

vCenter Server Deployment Wizard

Slide 4-28

The vCenter Server Deployment wizard prompts for information depending on your choice of deployment methods.

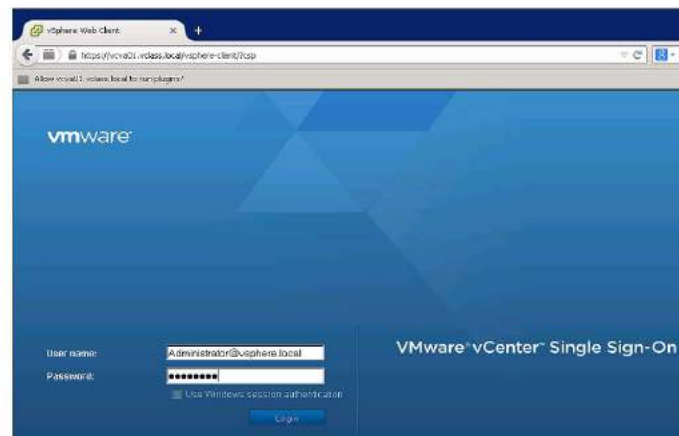
	Embedded	PSC	vCenter Server
Stage 1			
Deployment target	✓	✓	✓
Deployment type	✓	✓	✓
Deployment size	✓	✗	✓
Define VM name	✓	✓	✓
Define root password	✓	✓	✓
Define datastore location	✓	✓	✓
Define networking	✓	✓	✓
Stage 2			
Create SSO domain	✓	✓	✗
Join SSO domain	✗	✓	✓
Configure CEIP	✓	✓	✓

Getting Started with vCenter Server

Slide 4-29

After you deploy vCenter Server Appliance, log in to it by using one of the vSphere clients to manage your vSphere inventory:

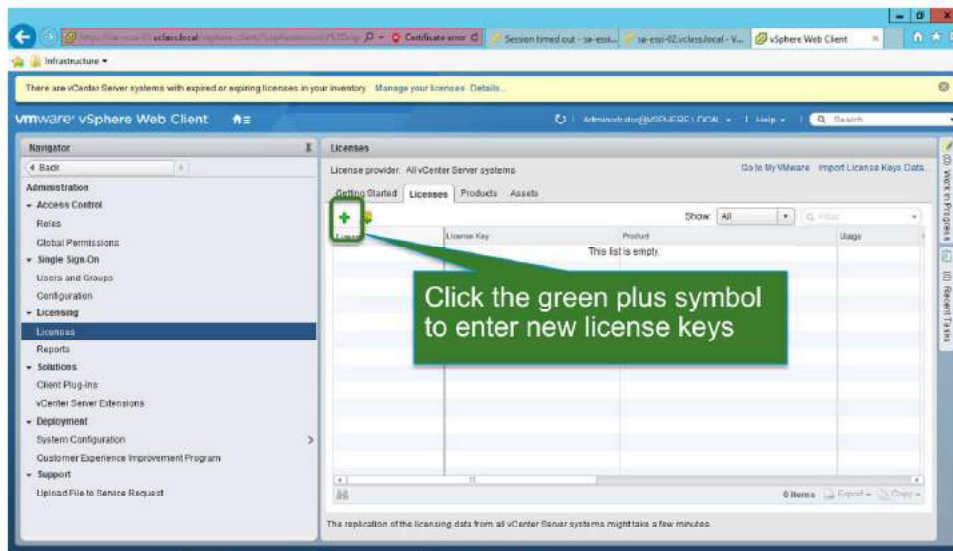
- vSphere Web Client: https://FQDN_for_vCenter_Server/vsphere-client
- vSphere Client: https://FQDN_for_vCenter_Server/ui



Adding License Keys to vCenter Server

Slide 4-30

Assign a license to vCenter Server before its 60-day evaluation period expires.

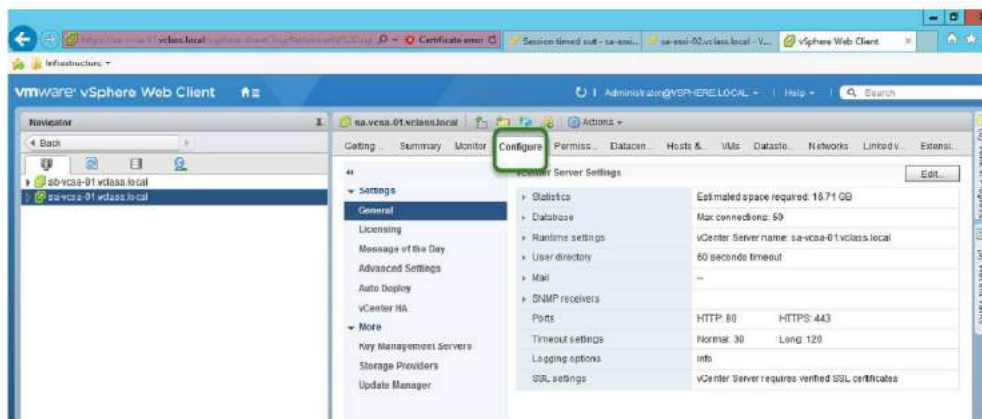


Configuring vCenter Server Settings

Slide 4-31

You can configure your vCenter Server system from vSphere Web Client, including settings such as licensing, statistics collection, logging, and other settings:

- To access the vCenter Server system settings, navigate to the vCenter Server system in vSphere Web Client and click the **Configure** > **Settings** tabs.



Logging In to the vCenter Server Appliance Management UI

Slide 4-32

To back up or restore vCenter Server Appliance, you must connect to the Appliance Management interface at `https://FQDN_or_IP_address:5480`.



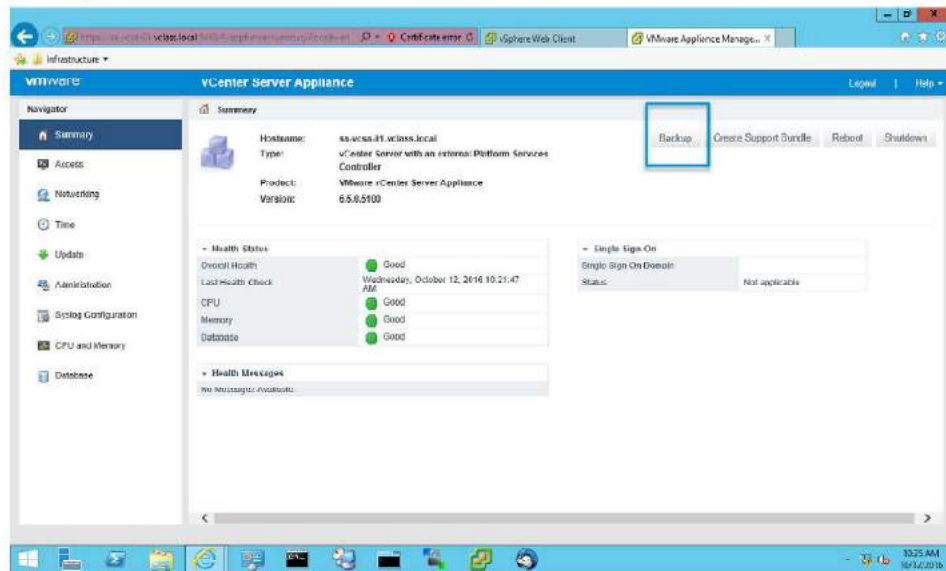
The Appliance Management UI is an HTML client specifically designed to configure vCenter Server Appliance.

The Appliance Management UI connects directly to port 5480 of the appliance and is used during the initial configuration of the appliance. The Appliance Management UI provides access to all the service APIs of the appliance.

vCenter Server Appliance Management Home

Slide 4-33

After logging in to the vCenter Server Appliance Management UI, you see the functions that you can perform. To back up the appliance, select the **Backup** tab.



Native vCenter Server Backup and Restore

Slide 4-34

Removes dependency on third-party backup solutions

Restores a vCenter Server instance to a brand new appliance

Supports backup or restore of vCenter Server Appliance and Platform Services Controller

Includes embedded and external deployments

Supports protocols, including:

- HTTP/S
- SCP
- FTP/S

Includes option for encryption

Restores directly from the vCenter Server Appliance ISO

vCenter Server Appliance 6.5 Restore - Stage 1: Deploy OVF

Enter backup details
Enter the backup location type, location, and credentials.

Backup location type:

Backup location:

Port:

User name:

Password:

Encryption password: optional

Install
Install a new vCenter Server Appliance or Platform Services Controller Appliance

Upgrade
Upgrade an existing vCenter Server Appliance

Migrate
Migrate from an existing vCenter Server for Windows to a vCenter Server Appliance

Restore
Restore from a previously created vCenter Server Appliance backup

The vCenter Server Appliance Management API supports backing up key parts of the appliance. You can protect vCenter Server data and minimize the time required to restore data center operations.

The backup process collects key files into a tar bundle and compresses the bundle to reduce network load. To minimize storage impact, the transmission is streamed without caching in the appliance. To reduce the total time required to complete the backup operation, the backup process handles the different components in parallel.

You can encrypt the compressed file before transmission to the backup storage location. When you choose encryption, you must supply a password that can be used to decrypt the file during restoration.

The backup operation always includes the vCenter Server database and system configuration files, so that a restore operation has all the data needed to recreate an operational appliance. Optionally, you can specify that a backup operation should include Statistics, Events, and Tasks from the current state of the data center. Current alarms are always included in a backup.

Lab 4: Working with vCenter Server

Slide 4-35

Install and use vCenter Server Appliance

1. Deploy vCenter Server Appliance
2. Access and Configure vCenter Server Appliance
3. Add Your ESXi Hosts to the vCenter Server Inventory
4. Configure the ESXi Hosts as NTP Clients
5. Back Up vCenter Server Appliance
6. Complete the vCenter Server Appliance Deployment

Review of Learner Objectives

Slide 4-36

You should be able to meet the following objectives:

- Discuss the vCenter Server deployment models
- Deploy vCenter Server Appliance into an infrastructure
- Add license keys to vCenter Server
- Configure vCenter Server settings
- Create a vCenter Server backup
- Restore vCenter Server Appliance from backup

Lesson 3: vSphere Clients

Slide 4-37



Learner Objectives

Slide 4-38

By the end of this lesson, you should be able to meet the following objectives:

- Access the vSphere clients
- Install the Enhanced Authentication Plug-In for Windows
- Navigate the vSphere clients

Accessing vSphere Clients

Slide 4-39

To access a vSphere client, you open a Web browser and enter the URL for the desired vSphere client.

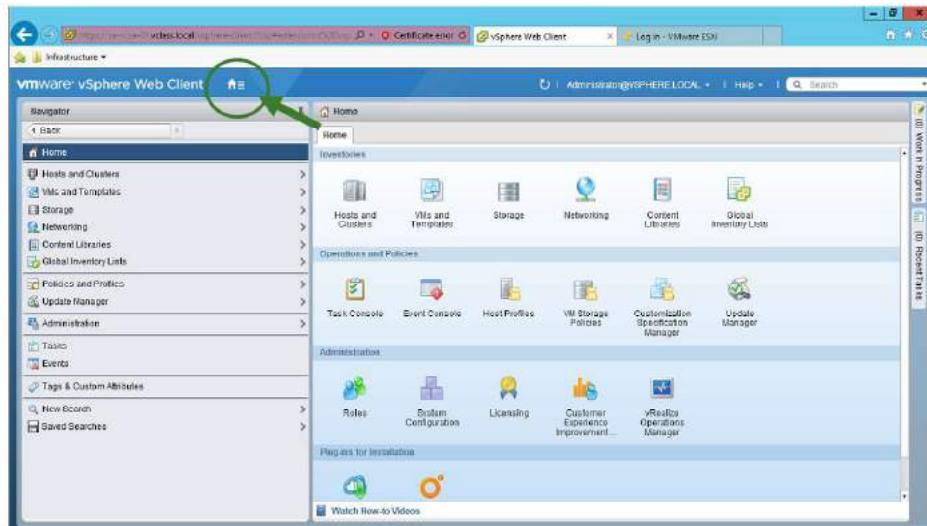
The screenshot shows a web browser window displaying the VMware vCenter Single Sign-On login page. The browser's address bar shows a URL with a certificate error. The page features the VMware logo at the top left. A green callout box points to the address bar area, containing the text: `https://FQDN_of_vCenter_Server/vsphere-client` and `https://FQDN_of_vCenter_Server/ui`. Below the logo, there is a login form with fields for 'User name:' (containing 'example@domain.local') and 'Password:', a checkbox for 'Use Windows session authentication', and a 'Login' button. At the bottom left, there is a link 'Download Enhanced Authentication Plugin'. A second green callout box points to this link, containing the text: 'Click here to download and install the Enhanced Authentication Plug-in.'

vSphere Web Client Home Page

Slide 4-40

Click the **Home** icon to reach the vCenter Server Home page.

The Home page has a Navigator pane on the left and Inventories, Monitoring, and Administration panes on the right.



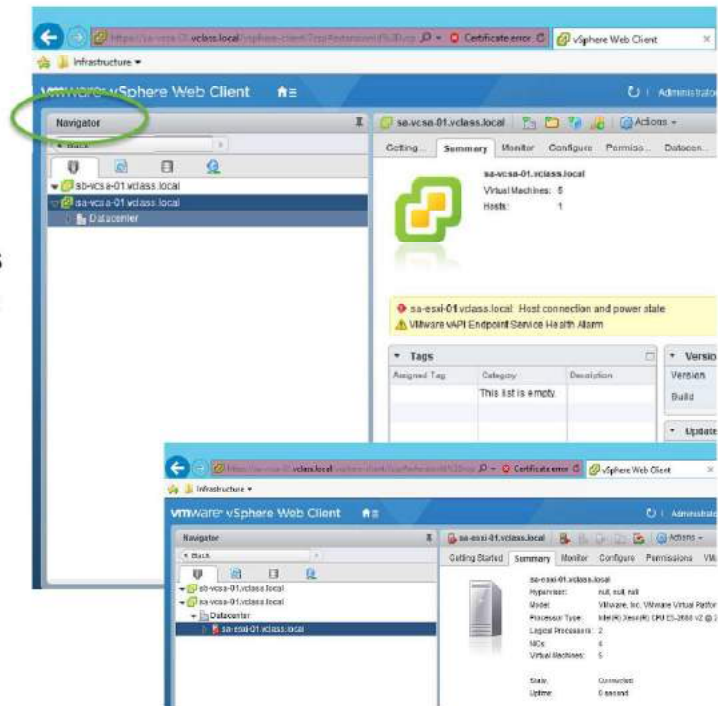
From the vSphere Web Client Home page, you can manage your vCenter Server system inventory, monitor your infrastructure environment, and complete system administration tasks.

Using the vSphere Web Client Navigator

Slide 4-41

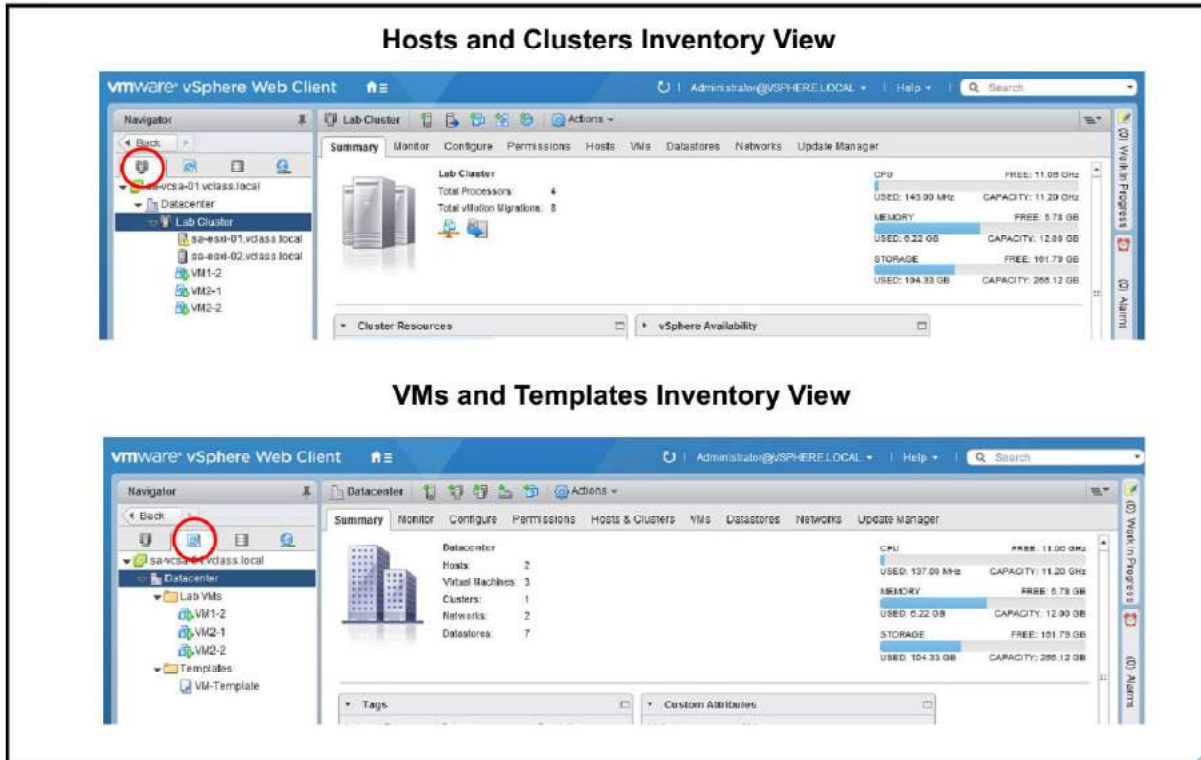
You can use the Navigator pane to browse and select objects in the vSphere Web Client inventory.

The navigator presents a list-based view of the inventory, which enables you to navigate inventory objects.



vCenter Server Views: Hosts and Clusters, VMs, and Templates

Slide 4-42

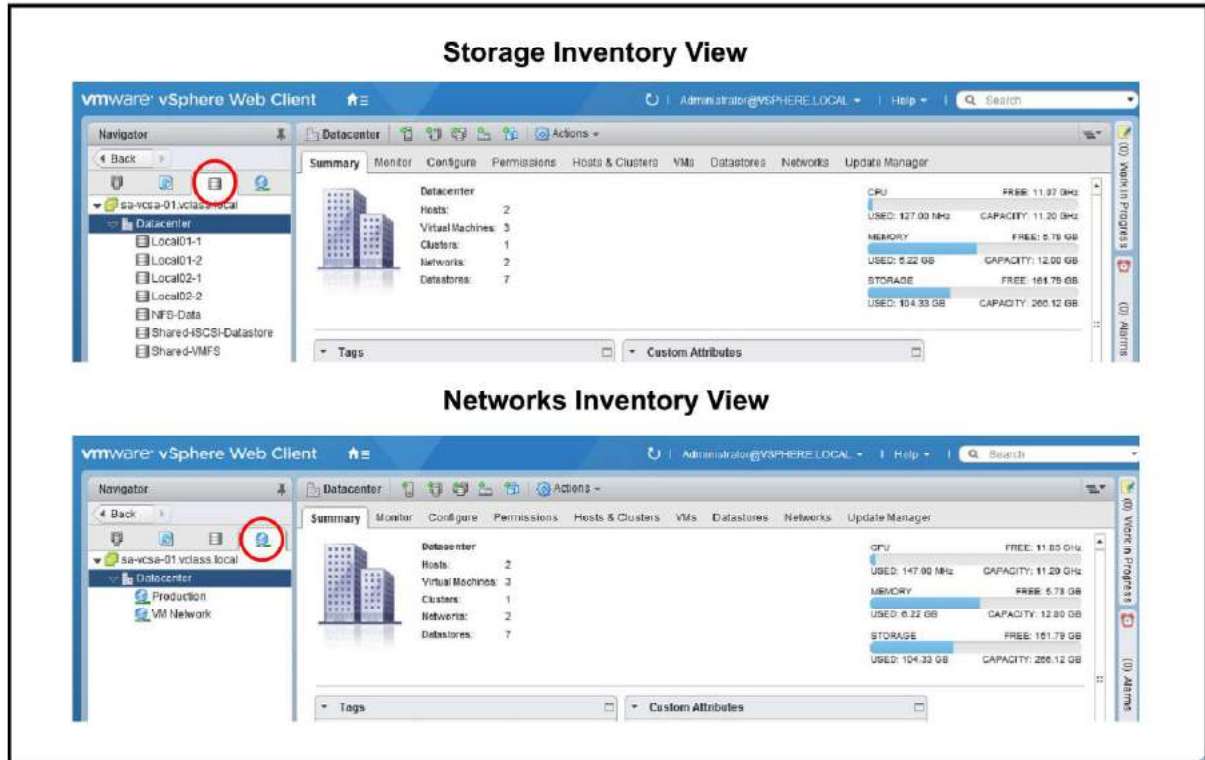


The Hosts and Clusters inventory view displays all host and cluster objects in a data center.

The VMs and Templates inventory view displays all virtual machine and template objects in a data center.

vCenter Server Views: Storage and Networks

Slide 4-43



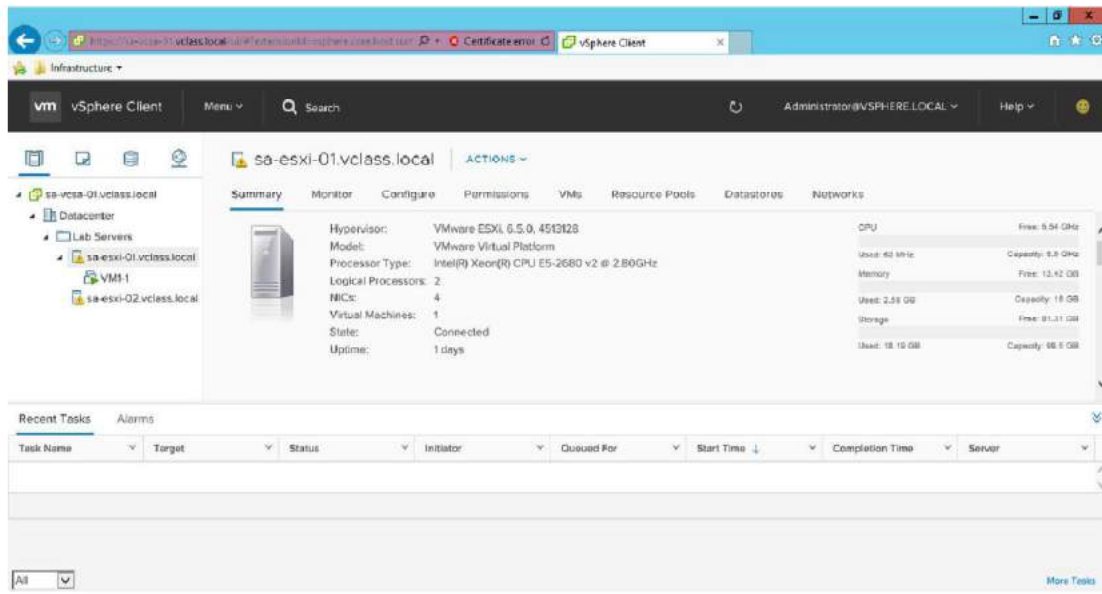
The Storage inventory view displays all the details for datastores in the data center.

The Networking inventory view displays all virtual machine port groups and distributed switches. Similar to the other inventory views, you can organize your datastore and network objects into folders.

Viewing Object Information

Slide 4-44

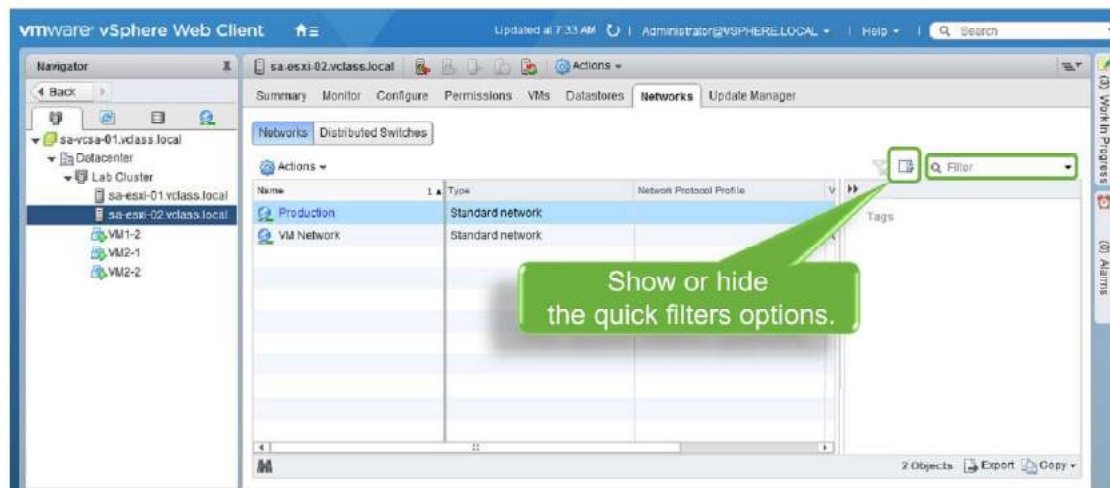
Because you can navigate to view object information and access related objects, monitoring and managing object properties is easy.



Using Quick Filters in vSphere Web Client

Slide 4-45

You can use quick filters to find an object or a set of objects in the vSphere inventory by using certain display criteria.



You can use quick filters to find an object or a set of objects in your vSphere Web Client inventory that fit certain criteria.

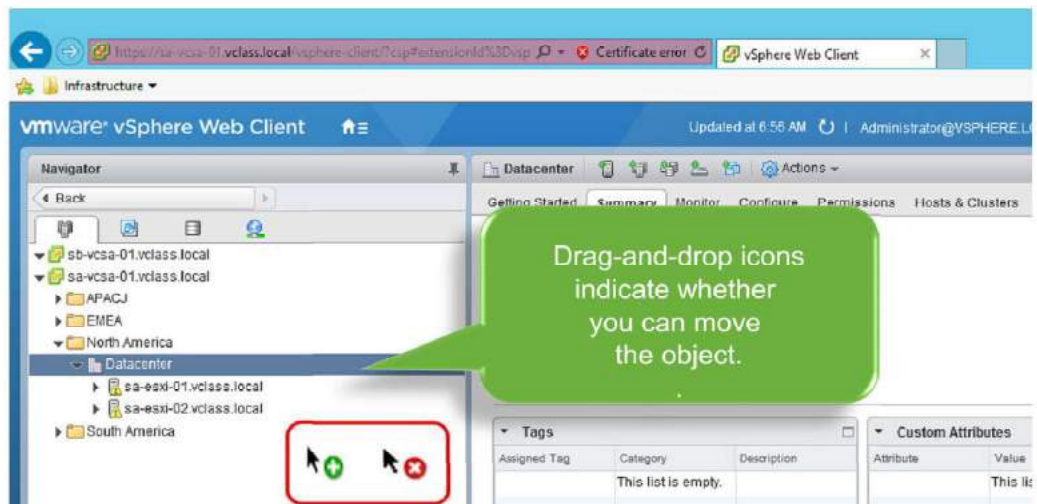
Quick filters are available in the list views, which appear in the **Objects** tab of an inventory list and in search results.

For example, you can use the quick filter options for virtual machines to find all virtual machines in your vSphere inventory that are powered on but do not have VMware Tools running.

Using Drag-and-Drop Functionality in vSphere Web Client

Slide 4-46

You can drag an inventory object to another location.

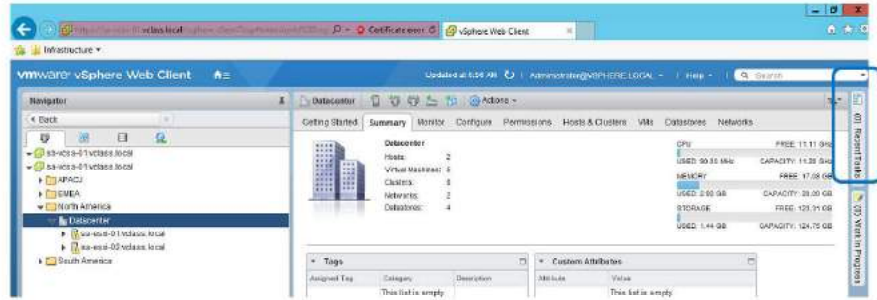


Using Pin and Unpin Functionality

Slide 4-47

You can pin and unpin display panes within the user interface.

Unpinned
Recent Tasks



Pinned
Tasks



Lab 5: Navigating the vSphere Clients

Slide 4-48

Become familiar with vSphere Client and vSphere Web Client

1. Navigate vSphere Client
2. Navigate vSphere Web Client

Review of Learner Objectives

Slide 4-49

You should be able to meet the following objectives:

- Access the vSphere clients
- Install the Enhanced Authentication Plug-In for Windows
- Navigate the vSphere clients

Lesson 4: Managing the vCenter Server Inventory

Slide 4-50



Lesson 4: Managing the vCenter Server Inventory

Learner Objectives

Slide 4-51

By the end of this lesson, you should be able to meet the following objectives:


- Create and organize vCenter Server inventory objects
- Add data center and organizational objects to vCenter Server
- Add hosts to vCenter Server
- Discuss how to create custom inventory tags for inventory objects
- Recognize how to view vCenter Server logs and events
- Manage the vCenter Server services
- Monitor vCenter Server Appliance

About Data Center Objects

Slide 4-52

A virtual data center is a container for all the inventory objects required to complete a fully functional environment for operating virtual machines:

- You can create multiple data centers to organize sets of environments.
- Each data center has its own hosts, virtual machines, templates, datastores, and networks.

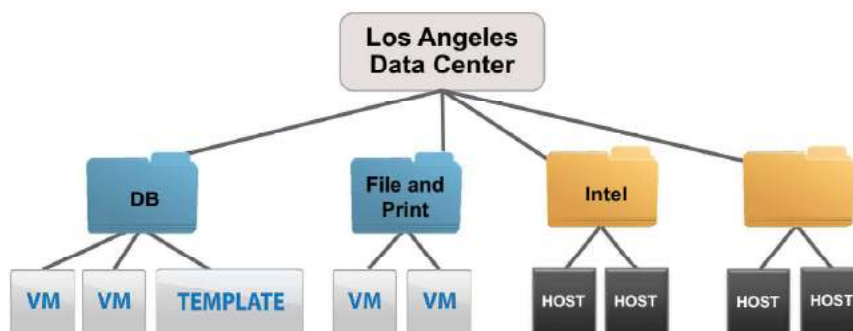


A virtual data center is a container for all the inventory objects required to complete a fully functional environment for operating virtual machines. You can create multiple data centers to organize sets of environments. For example, you might create a data center for each organizational unit in your enterprise or create some data centers for high-performance environments and other data centers for less demanding virtual machines.

Organizing Inventory Objects into Folders

Slide 4-53

Items in the data center can be placed into folders. Folders and subfolders can be created to better organize systems.



Plan the setup of your virtual environment. A large vSphere implementation might contain several virtual data centers with a complex arrangement of hosts, clusters, resource pools, and networks. It might involve multiple vCenter Server systems connected using Enhanced Linked Mode. Smaller implementations might require a single virtual data center with a less complex topology. Regardless of the scale of your virtual environment, consider how the virtual machines it will support are going to be used and administered.

Populating and organizing your inventory involves the following activities:

- Create data centers.
- Add hosts to the data centers.
- Organize inventory objects in folders.
- Set up networking by using vSphere standard switches or vSphere distributed switches. To use services such as vSphere vMotion, TCP/IP storage, virtual SAN, and fault tolerance, set up VMkernel networking for these services. For more information, see *vSphere Networking* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.
- Configure storage systems and create datastore inventory objects to provide logical containers for storage devices in your inventory. See *vSphere Storage* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

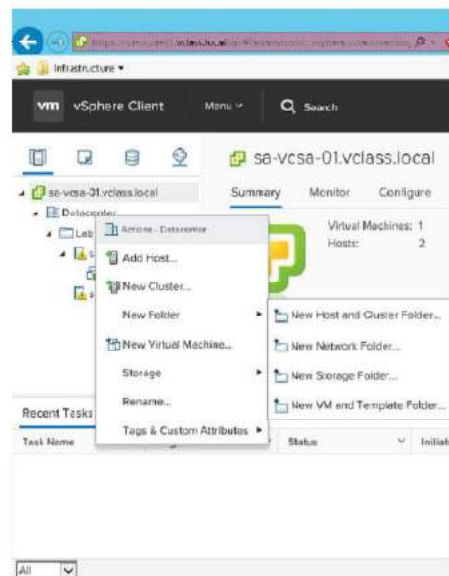
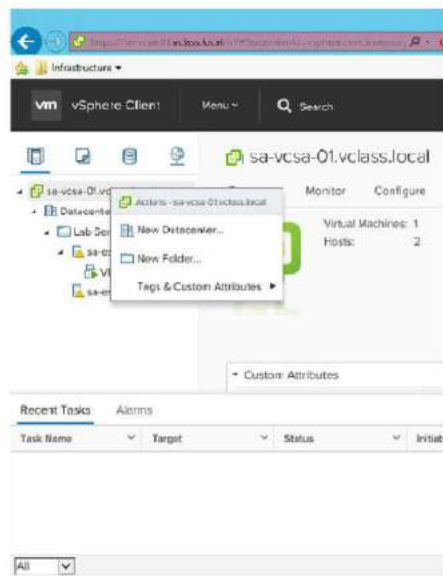
- Create clusters to consolidate the resources of multiple hosts and virtual machines. You can enable vSphere HA and vSphere DRS for increased availability and more flexible resource management. For information about configuring vSphere HA, see *vSphere Availability* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>. For information about configuring vSphere DRS, see *vSphere Resource Management* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.
- Create resource pools to provide logical abstraction and flexible management of the resources in vSphere. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources. For details, see *vSphere Resource Management* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Adding a Data Center and Organizational Objects to vCenter Server

Slide 4-54

You can add a data center and host and cluster folders:


- You can use folders to group objects of the same type for easier management.



Adding ESXi Hosts to vCenter Server

Slide 4-55

You can add an ESXi host.



Add Host

- 1 Name and location
- 2 Connection settings
- 3 Host summary
- 4 Assign license
- 5 Lockdown mode
- 6 VM location
- 7 Ready to complete

Name and location
Enter the name or IP address of the host to add to vCenter Server.

Host name or IP address:	<input type="text"/>
Location:	<input type="button" value="Datacenter"/>

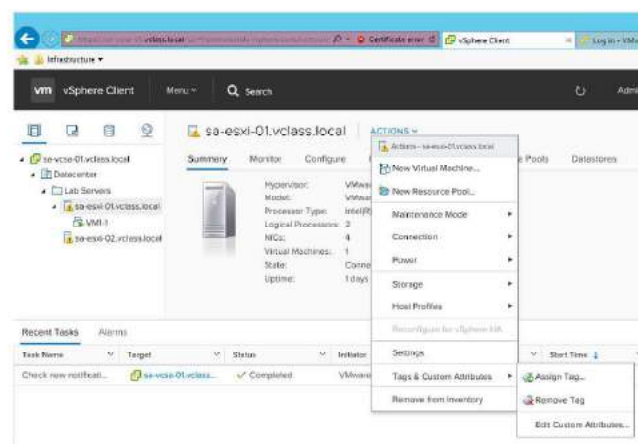
Creating Custom Tags for Inventory Objects

Slide 4-56

Tags enable you to attach metadata to objects in the vSphere inventory to make these objects more sortable.

You can associate a set of objects of the same type:

- Search for objects by that tag.
- Enable a business case where customers want to create groups of virtual machines, clusters, and datastores for ease of management.



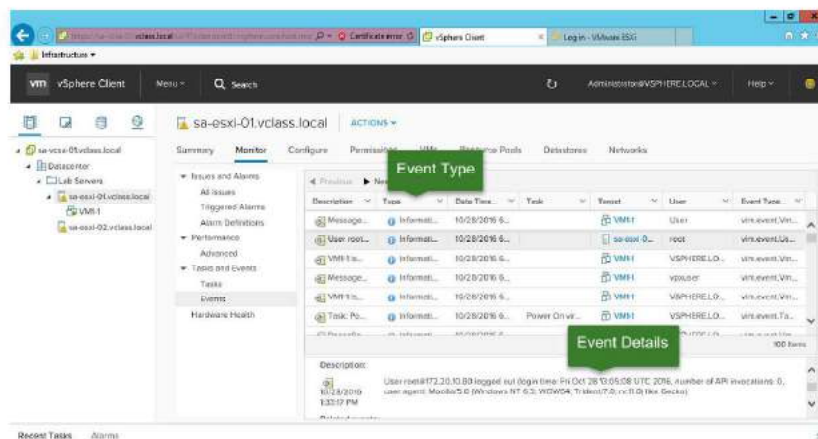
You use tags to add metadata to inventory objects. You can record information about your inventory objects in tags and use the tags in searches.

vCenter Server Events

Slide 4-57

The vCenter Server events and audit trail allows selectable retention periods in increments of 30 days:

- User-action information includes the user's account and specific event details.
- All actions are reported, including file ID, file path, source of operation, operation name, and date and time of operation.



An Event is a data object type that contains information about state changes of managed entities and other objects on the server. Events include user actions and system actions that occur on data centers, datastores, clusters, hosts, resource pools, virtual machines, networks, and distributed virtual switches. For example, these common system activities generate one or more Event data objects:

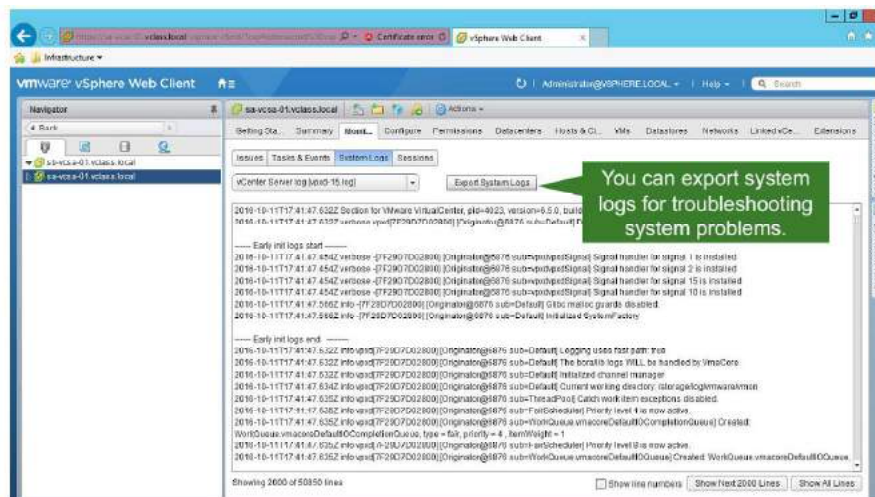
- Powering a virtual machine on or off
- Creating a virtual machine
- Installing VMware Tools on the guest OS of a virtual machine
- Reconfiguring a compute resource
- Adding a newly configured ESX/ESXi system to a vCenter Server system

vCenter Server System Logs

Slide 4-58

vSphere records events in the vCenter Server database:

- System log entries include information, such as who generated the event, when the event was created, and the type of event.



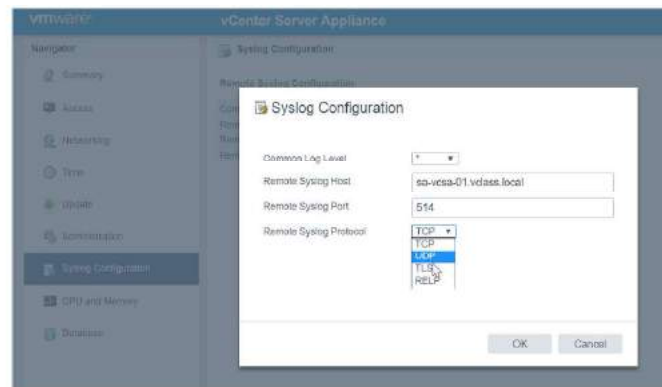
vSphere records events in the vCenter Server database. System log entries include information such as who generated the event, when the event was created, and the type of event.

Output vCenter Server Events and Logs to Syslog Collector

Slide 4-59

vCenter Server is capable of streaming its event and log information to a remote Syslog server:

- Enable in the vCenter Appliance VAMI (VMware Appliance Management Interface) – https://FQDN_of_vCenter:5480
- This feature can help prevent the filling up of the vCenter Server database by retaining events locally for 30 days and streaming older event data to a Syslog server.



You can stream the information about the events that your vSphere system generates to a remote Syslog server. Streaming events is supported only for vCenter Server Appliance.

The streaming of events to a remote Syslog server is disabled by default. You can enable and configure the streaming of vCenter Server events to a remote Syslog server from the vCenter Server Appliance Management interface. After you enable remote streaming, vCenter Server Appliance starts streaming and only the newly generated events are streamed to the remote Syslog server.

All Syslog messages begin with a specific prefix. You can distinguish the vCenter Server Appliance events from other Syslog messages by their Event prefix.

The Syslog protocol limits the length of Syslog messages to 1,024 characters. Messages that are longer than 1,024 characters split into multiple Syslog messages.

In the Syslog server, events have the following format:

```
<syslog-prefix> : Event [eventId] [partInfo] [createdTime] [eventType]
[severity] [user] [target] [chainId] [desc]
```

vCenter Server Database Health

Slide 4-60

vCenter Server checks the status of the database every 15 minutes:

- Database health warnings trigger an alarm when the volume (by default) reaches 80 percent.
- The alarm changes from warning to error when the free space reaches 95 percent and vCenter Server services shut down to allow the user to configure more disk space or remove unwanted content.

These features are available for the following databases:

- PostgreSQL and MSSQL
- Not available for Oracle

Managing the vCenter Server Services

Slide 4-61

You can manage vCenter Server services by selecting **Administration > System Configuration** from the Home page and selecting **Services**.



Lab 6: Creating Folders in vCenter Server Appliance

Slide 4-62

Create folders in vCenter Server Appliance

1. Create a Host and Cluster Folder
2. Create Virtual Machine and Template Folders

Review of Learner Objectives

Slide 4-63

You should be able to meet the following objectives:

- Create and organize vCenter Server inventory objects
- Add data center and organizational objects to vCenter Server
- Add hosts to vCenter Server
- Discuss how to create custom inventory tags for inventory objects
- Recognize how to view vCenter Server logs and events
- Manage the vCenter Server services
- Monitor vCenter Server Appliance

Lesson 5: vCenter Server Roles and Permissions

Slide 4-64



Lesson 5: vCenter Server Roles and Permissions

Learner Objectives

Slide 4-65

By the end of this lesson, you should be able to meet the following objectives:

- Define a permission
- Describe the rules for applying permissions
- Create a custom role
- Create a permission

Access Control Overview

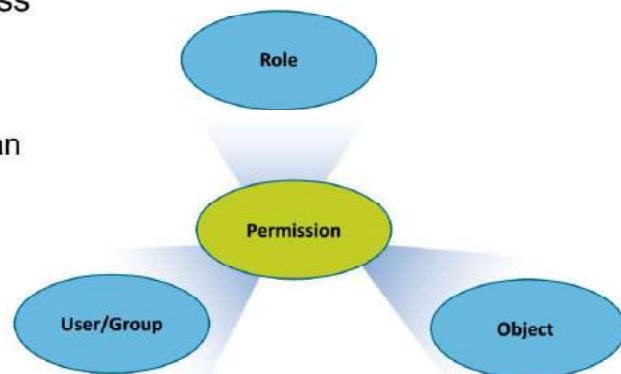
Slide 4-66

The access control system enables the vCenter Server administrator to define a user's privileges to access objects in the inventory.

Key concepts:

- Privilege: Defines an action that can be performed
- Role: A set of privileges
- Object: The target of the action
- User/group: Indicates who can perform the action

Together, a role, a user or group, and an object define a permission.



The authorization to perform tasks in vCenter Server is governed by an access control system. This system enables the vCenter Server administrator to specify in great detail which users or groups can perform which tasks on which objects. The access control system is defined with the following concepts:

- Privilege: The ability to perform a specific action or read a specific property. Examples include powering on a virtual machine and creating an alarm.
- Role: A collection of privileges. Roles provide a way to aggregate all the individual privileges that are required to perform a higher-level task, such as administering a virtual machine.
- Object: An entity upon which actions are performed.
- User or group: A user or group who can perform the action.

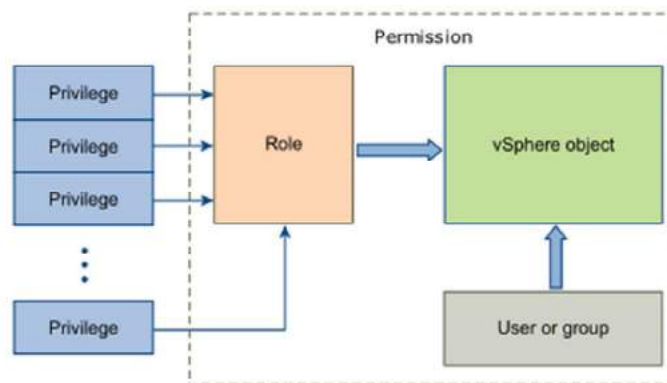
vCenter Server Permissions

Slide 4-67

The permission model for vCenter Server systems relies on assigning permissions to objects in the vSphere object hierarchy. Each permission gives one user or group a set of privileges, that is, a role for the selected object.

The following concepts are important:

- Permissions
- Users and groups
- Privileges
- Roles



A permission is set on an object in the vCenter Server object hierarchy. Each permission associates the object with a group or user and the group or user access roles. For example, you can select a virtual machine object, add one permission that gives the ReadOnly role to Group 1, and add a second permission that gives the Administrator role to User 2.

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.

Adding Permissions to the vCenter Server Inventory

Slide 4-68

Right-click the inventory object, select **Add Permission**, select the user and group, and click **OK**.

The image contains three numbered screenshots illustrating the process of adding permissions in vSphere Web Client:

- One:** A screenshot of the vSphere Web Client interface showing a right-click context menu over a folder in the inventory. The 'Add Permission' option is highlighted.
- Two:** A screenshot of the 'Lab Server Folder - Add Permission' dialog box. The 'Users and Groups' list is empty, and the 'Assigned Role' list has 'All Privileges' selected.
- Three:** A screenshot of the 'Select Users/Groups' dialog box. The 'Domain' is set to 'vsphere.local'. The 'Users and Groups' list shows 'vsphere.local/DCAdmins' selected in the 'Groups' field.

To manage permissions from vSphere Web Client, you must understand the following concepts:

- **Permission:** Each object in the vCenter Server object hierarchy has associated permissions. Each permission specifies for one group or user which privileges that group or user has on the object.
- **Users and Groups:** On vCenter Server systems, you can assign privileges only to authenticated users or groups of authenticated users. Users are authenticated through vCenter Single Sign-On. The users and groups must be defined in the identity source that vCenter Single Sign-On is using to authenticate. Define users and groups using the tools in your identity source, for example, Active Directory.
- **Privileges:** Privileges are fine-grained access controls. You can group those privileges into roles that you can map to users or groups.
- **Roles:** Roles are sets of privileges. Roles allow you to assign permissions on an object based on a typical set of tasks that users perform. Default roles, such as Administrator, are predefined on vCenter Server and cannot be changed. Other roles, such as Resource Pool Administrator, are predefined sample roles. You can create custom roles either from scratch or by cloning and modifying sample roles. For information about creating a custom role and cloning a role, see *vSphere Security* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

You can assign permissions to objects at different levels of the hierarchy, for example, you can assign permissions to a host object or to a folder object that includes all host objects. For information about hierarchical inheritance of permissions, see *vSphere Security* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>. You can also assign permissions to a global root object to apply the permissions to all objects in all solutions. For information about global permissions, see *vSphere Security* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

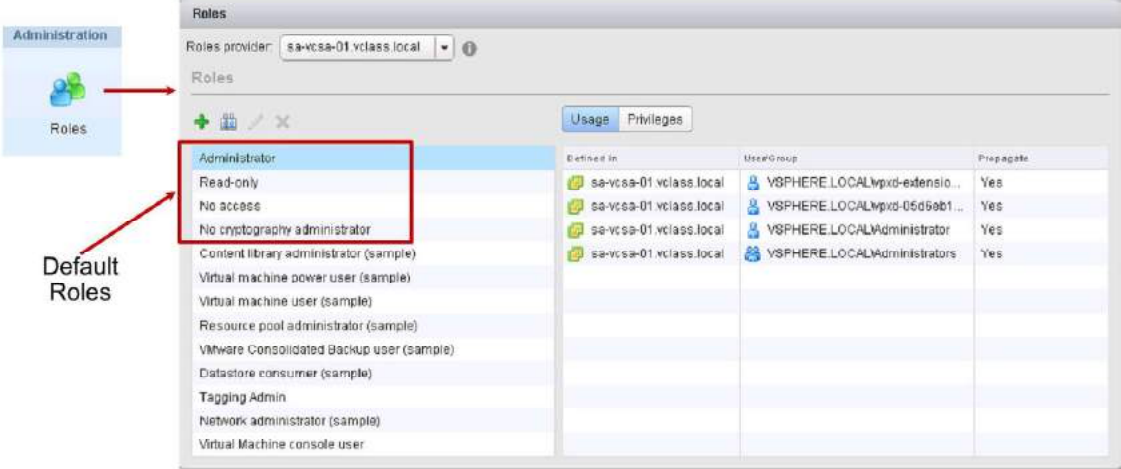
About Roles

Slide 4-69

Privileges are grouped into roles:

- They allow users to perform tasks.
- They are grouped into categories to make them simpler to configure.

vCenter Server provides a few default roles, which you cannot modify.



The screenshot shows the vCenter Server Roles configuration page. The 'Roles' tab is selected, and the 'Roles' list is displayed. A red box highlights the 'Default Roles' section, which includes the following roles:

- Administrator
- Read-only
- No access
- No cryptography administrator

The 'Usage' and 'Privileges' tabs are also visible. The 'Usage' tab shows a table of roles assigned to users or groups.

Defined in	User/Group	Propagate
sa-vc:sa-01 vclass.local	VSPHERE.LOCAL\wpzd-extensio...	Yes
sa-vc:sa-01 vclass.local	VSPHERE.LOCAL\wpzd-05d5eb1...	Yes
sa-vc:sa-01 vclass.local	VSPHERE.LOCAL\Administrator	Yes
sa-vc:sa-01 vclass.local	VSPHERE.LOCAL\Administrators	Yes

A role is a set of one or more privileges. A privilege allows access to a specific task and is grouped with other privileges related to it. For example, the Virtual Machine Power User sample role consists of several privileges in categories like Datastore and Global. A role is assigned to a user or group and determines the level of access of that user or group.

You cannot change the privileges associated with the default roles.

- Administrator role: Users with the Administrator role for an object are allowed to view and perform all actions on the object.
- No cryptography administrator role: Users with the No cryptography administrator role for an object have the same privileges as users with the Administrator role, except for cryptographic operations privileges.
- No access role: Users with the No access role for an object cannot view or change the object in any way.
- Read-only role: Users with the Read-only role for an object are allowed to view the state of the object and details about the object.

The default roles are organized as a hierarchy. Each role inherits the privileges of the previous role. For example, the Administrator role inherits the privileges of the Read-only role. Roles that you create do not inherit privileges from any of the system roles.

Except for the default roles, roles are not hierarchically organized. No role is superior or subordinate to another role. All roles are independent of one another.

About Objects

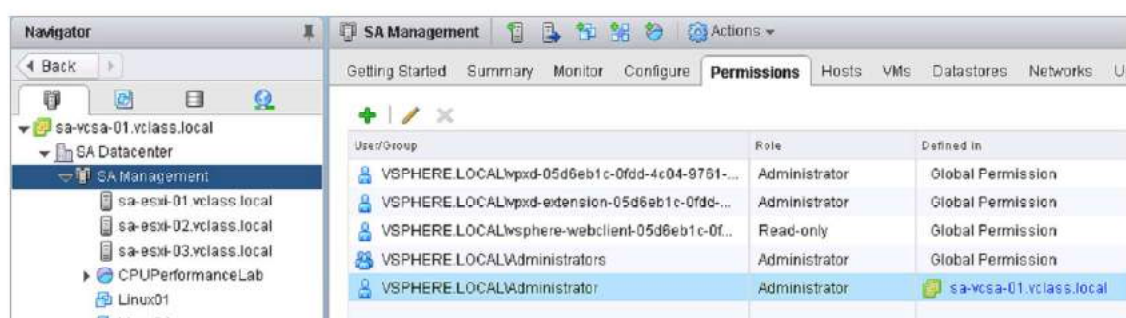
Slide 4-70

Objects are entities on which actions are performed:

- Objects include data centers, folders, resource pools, clusters, hosts, datastores, networks, and virtual machines.

All objects have a **Permissions** tab:

- The **Permissions** tab shows which user or group and role are associated with the selected object.



A user or group indicates who can perform the action. The object is the target of the action. Each combination of user or group, role, and object must be specified.

By using a role, one or more permissions can be assigned to any object in the vCenter Server inventory.

On the slide, the Administrator role was granted to the VSPHERE.LOCAL\Administrator user. The combination of user or group, plus role have been applied at the vCenter Server level. So this permission is allowed on all data centers in the vCenter Server inventory.

On the slide, the first three entries are for solution users. A solution user represents one or more vCenter Server services:

- vpzd: Solution user for vCenter Server
- vpzd-extension: Solution user for services such as Inventory Service and vSphere Auto Deploy
- vsphere-webclient: Solution user for services such as vSphere Web Client

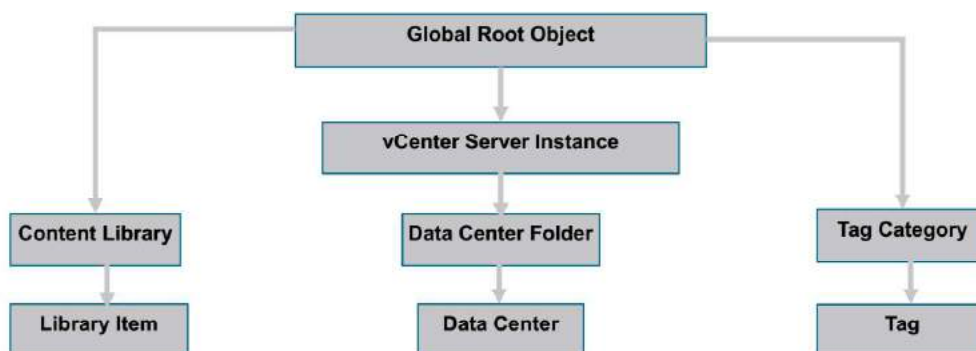
These solution users have the Administrator role, and have global permissions.

About Global Permissions

Slide 4-71

Global permissions support assigning privileges across solutions from a global root object:

- Global permissions span solutions such as vCenter Server and VMware vRealize® Orchestrator™.
- Global permissions give a user or group privileges for all objects in all object hierarchies.

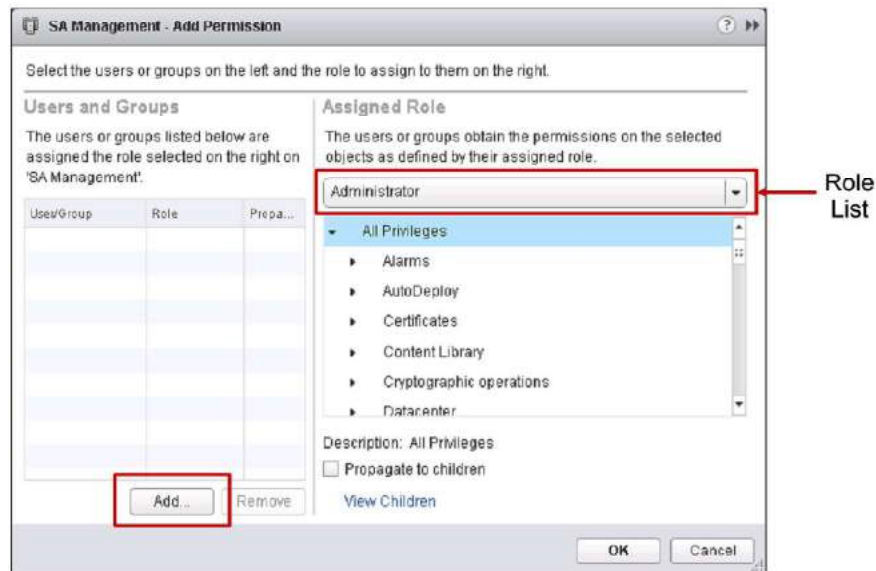


Often, you apply a permission to a vCenter Server inventory object such as an ESXi host or a virtual machine. When you apply a permission, you specify that a user or group has a set of privileges, called a role, on the object. Global permissions give a user or group privileges to view or manage all objects in each of the inventory hierarchies in your deployment. The example on the slide shows that the global root object has permissions over all vCenter Server objects, including content libraries, vCenter Server instances, and tags. vCenter Server permissions, on the other hand, are effective only on objects in the vCenter Server instance.

Assigning Permissions

Slide 4-72

Add permissions by adding a user or group, and selecting a role from the role list. You can also choose to propagate the permission to all child objects.



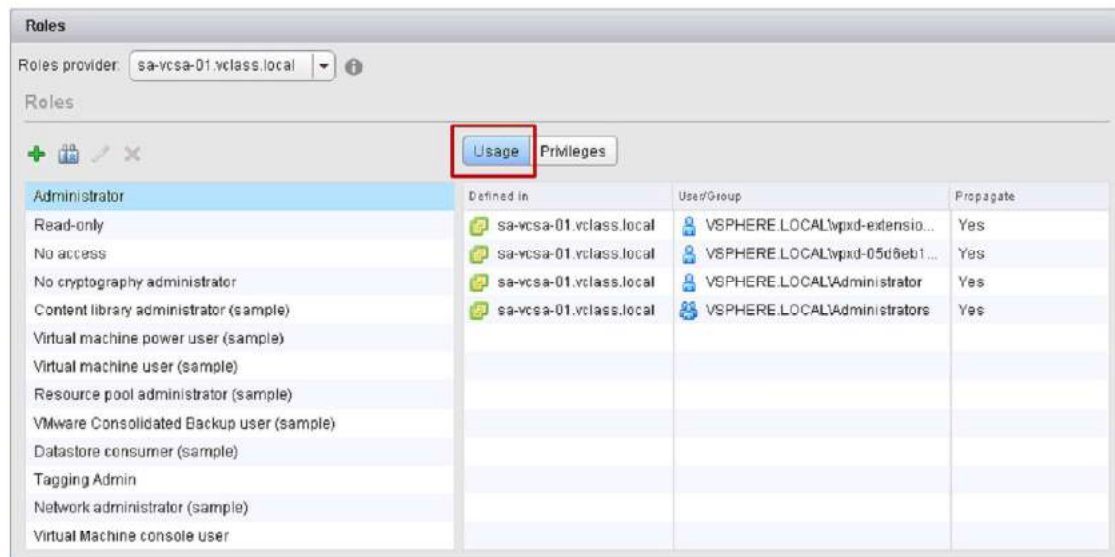
Role propagation is the act of passing along permissions. A role can be propagated to its child objects in the inventory.

For each permission, you can decide whether the permission propagates down the object hierarchy to all subobjects or if it applies only to the original object. For example, you can grant a user the Read-only role at the data center level and have the role propagate to the data center's child objects. Then you can grant another role, such as Virtual Machine Power User, to a virtual machine in the data center.

Viewing Roles and Assignments

Slide 4-73

The Roles pane shows which users are assigned the selected role on a particular object.



The screenshot shows the Roles pane in vSphere. The Roles provider is set to sa-vcsa-01.vclass.local. The Administrator role is selected. The Usage tab is active, showing a table of role assignments.

Role	Defined in	User/Group	Propagate
Administrator	sa-vcsa-01.vclass.local	VSPHERE.LOCAL\wpad-extensio...	Yes
Read-only	sa-vcsa-01.vclass.local	VSPHERE.LOCAL\wpad-05d0eb1...	Yes
No access	sa-vcsa-01.vclass.local	VSPHERE.LOCAL\Administrator	Yes
No cryptography administrator	sa-vcsa-01.vclass.local	VSPHERE.LOCAL\Administrators	Yes
Content library administrator (sample)			
Virtual machine power user (sample)			
Virtual machine user (sample)			
Resource pool administrator (sample)			
VMware Consolidated Backup user (sample)			
Datastore consumer (sample)			
Tagging Admin			
Network administrator (sample)			
Virtual Machine console user			

You can view all of the objects to which a role was assigned and all of the users or groups who were granted the role.

To view this information

1. Click **Usage** in the Roles pane.
2. Select a role in the roles list.

The information panel shows each object to which the role is assigned and the users and groups who were granted the role.

Applying Permissions: Scenario 1

Slide 4-74

A permission can propagate down the object hierarchy to all subobjects, or it can apply only to an immediate object.



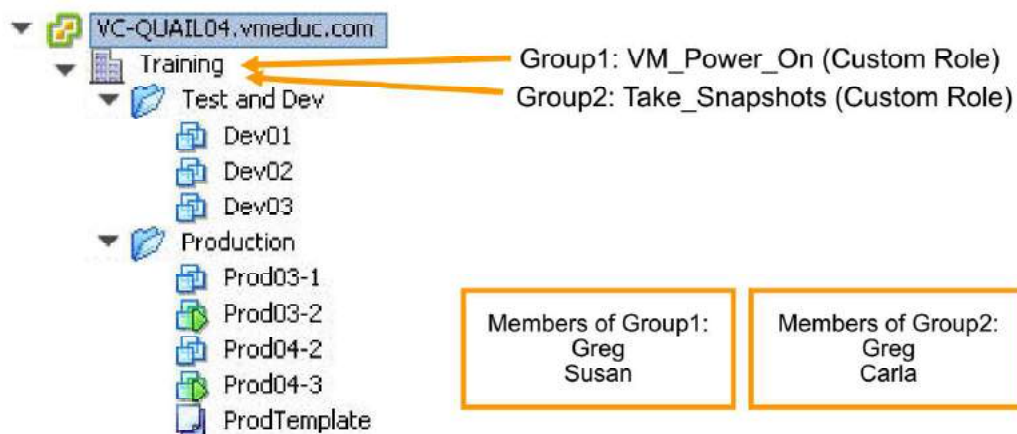
In addition to specifying whether permissions propagate downward, you can override permissions set at a higher level by explicitly setting different permissions for a lower-level object.

On the slide, user Greg is given Read-Only access in the Training data center. This role is propagated to all child objects except one, the Prod03-2 virtual machine. For this virtual machine, Greg is an administrator.

Applying Permissions: Scenario 2

Slide 4-75

When a user is a member of multiple groups with permissions on the same object, the user is assigned the union of privileges assigned to the groups for that object.



When a user is a member of multiple groups, and these groups have permissions on the same object in the inventory, the user is assigned the union of privileges assigned to the groups for that object.

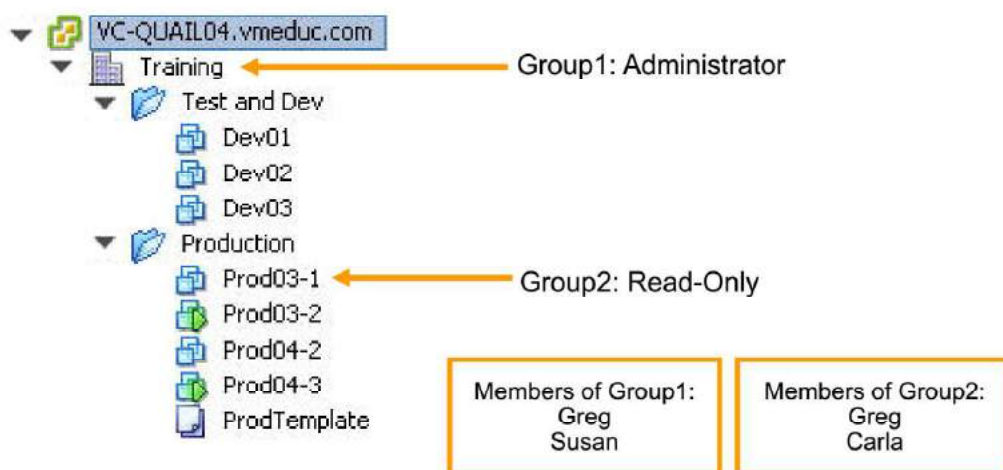
On the slide, Group1 is assigned the VM_Power_On role, a custom role that contains only one privilege: the ability to power on a virtual machine. Group2 is assigned the Take_Snapshots role, another custom role that contains the privileges to create and remove snapshots. Both roles propagate to the child objects.

Assume that Greg belongs to both Group1 and Group2. Greg gets both VM_Power_On and Take_Snapshots privileges for all objects in the Training data center.

Applying Permissions: Scenario 3

Slide 4-76

When a user is a member of multiple groups with permissions on different objects, for each object on which the group has permissions, the same permissions apply as if they were granted directly to the user.



A user can be a member of multiple groups and can have permissions on different objects in the inventory. For each object on which the group has permissions, the same permissions apply as if they were granted to the user directly. You can override permissions set for a higher-level object by explicitly setting different permissions for a lower-level object.

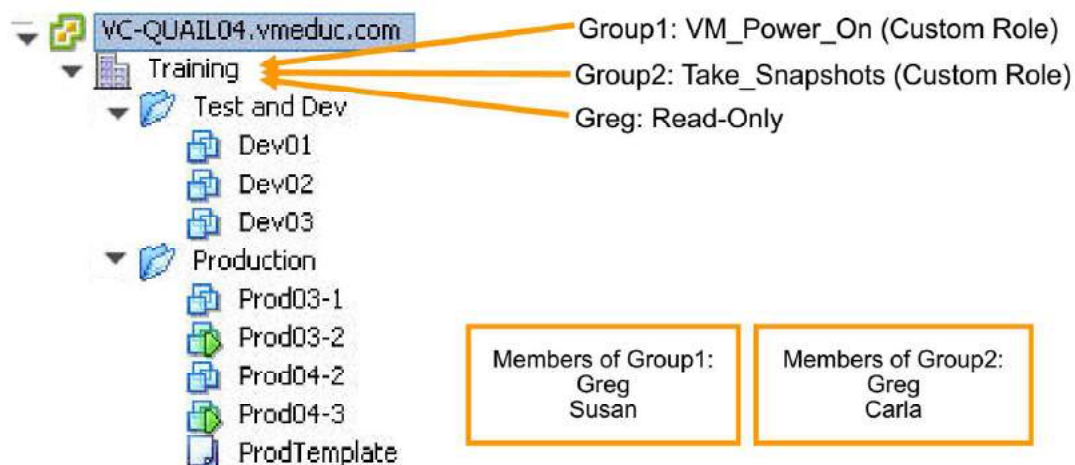
On the slide, Group1 is assigned the administrator role at the Training data center and Group2 is assigned the Read-only role on the virtual machine object, Prod03-1. The permission granted to Group1 is propagated to child objects.

Assume that user Greg is a member of both Group1 and Group2. Greg gets administrator privileges on the entire Training data center (the higher-level object), except for the virtual machine named Prod03-1 (the lower-level object). For this object, Greg gets read-only access.

Applying Permissions: Scenario 4

Slide 4-77

Permissions defined explicitly for the user on an object take precedence over all group permissions on that same object.



On the slide, three permissions are assigned to the Training data center:

- Group1 is assigned the VM_Power_On role.
- Group2 is assigned the Take_Snapshots role.
- User Greg is assigned the Read-only role.

Assume that Greg is a member of both Group1 and Group2. Assume also that propagation to child objects is enabled on all roles. Although Greg is a member of both Group1 and Group2, he gets the Read-only privilege to the Training data center and all objects under it. Greg gets the Read-only privilege because explicit user permissions on an object take precedence over all group permissions on that same object.

Creating a Role

Slide 4-78

Create roles that enable only the necessary tasks:

- Example: Virtual Machine Creator

Use folders to contain the scope of permissions:

- For example, assign the Virtual Machine Creator role to user Nancy and apply it to the Finance folder.

Virtual Machine Creator role:

Datastore > Allocate space

Network > Assign network

Resource > Assign virtual machine to resource pool

Virtual machine > Inventory > Create new

Virtual machine > Configuration > Add new disk

Virtual machine > Configuration > Add or remove device

The Virtual Machine Creator role is one of many examples of roles that can be created. As a best practice, define a role using the smallest number of privileges possible so that security and control over your environment can be maximized. Also, give the roles names that explicitly indicate what each role allows, to make its purpose clear.

Use folders to contain the scope of permissions. For example, to limit the creation of virtual machines, create a folder in the VMs and Templates inventory view. Apply the Virtual Machine Creator role on this folder for the users.

Review of Learner Objectives

Slide 4-79

You should be able to meet the following objectives:

- Define a permission
- Describe the rules for applying permissions
- Create a custom role
- Create a permission

Key Points

Slide 4-80


- The vCenter Server architecture includes the following components:
 - vCenter Server
 - vCenter Server database
 - Active Directory
 - Managed ESXi hosts
- vCenter Server has the following types of deployment models:
 - Embedded Platform Services Controller
 - External Platform Services Controller
- You use the vSphere clients to connect to vCenter Server systems and manage vSphere inventory objects.
- You can manage the vCenter Server inventory.
- You can back up and restore the vCenter Server system natively.

Questions?

MODULE 5

Configuring and Managing Virtual Networks

Slide 5-1



You Are Here

Slide 5-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. **Configuring and Managing Virtual Networks**
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Importance

Slide 5-3

Virtual machines must be able to communicate with other virtual machines and with physical machines. Remote host management and IP-based storage must be able to operate effectively.

Failure to properly configure ESXi networking can negatively affect the operations of your virtual infrastructure.

Module Lessons

Slide 5-4

- | | |
|-----------|---|
| Lesson 1: | Introduction to vSphere Standard Switches |
| Lesson 2: | Configuring Standard Switch Policies |

Lesson 1: Introduction to vSphere Standard Switches

Slide 5-5



Lesson 1: Introduction to vSphere Standard Switches

Learner Objectives

Slide 5-6

By the end of this lesson, you should be able to meet the following objectives:

- Describe the virtual switch connection types
- Configure and view standard switch configurations, such as virtual machine port group, VMkernel port, VLAN, and security features
- List the features comparison of standard and distributed switches


Types of Virtual Switch Connections

Slide 5-7

Virtual switches provide the connectivity between VMs on the same host or on different hosts. Virtual switches also support VMkernel network access for remote host management, vSphere vMotion, iSCSI, and NFS.

A virtual switch has specific connection types:

- Virtual machine port groups
- VMkernel port:
 - For IP storage, vSphere vMotion migration, vSphere Fault Tolerance, vSAN, and vSphere Replication
 - For the ESXi management network



A virtual switch provides the following connection types to hosts and virtual machines:

- Connecting virtual machines to the physical network.
- Connecting VMkernel services to the physical network. VMkernel services include access to IP storage, such as NFS or iSCSI, vSphere vMotion migrations, and access to the management network.

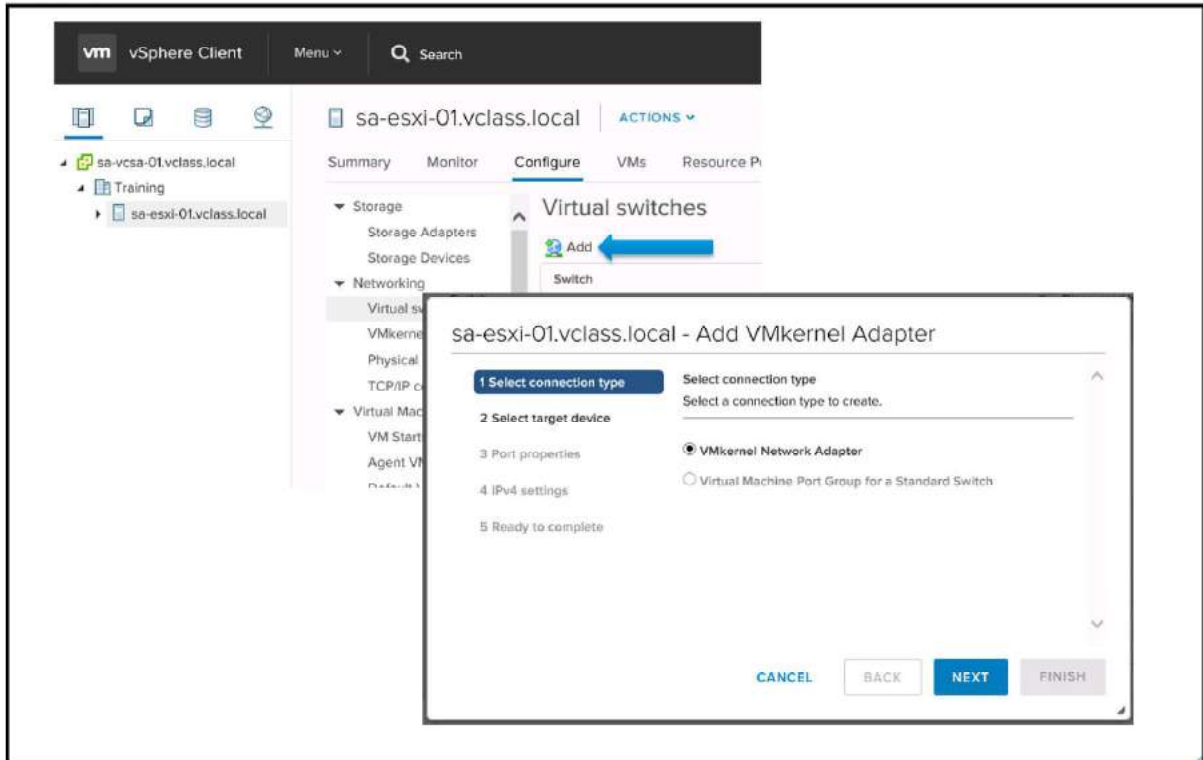
The ESXi management network port is used to connect to network or remote services, including vSphere Web Client. Each ESXi management network port and each VMkernel port must be configured with its own IP address, netmask, and gateway.

To help configure virtual switches, you can create port groups. A port group is a template that stores configuration information to create virtual switch ports on a virtual switch. Virtual machine port groups are used to connect virtual machines to one another with common networking properties.

Virtual machine port groups and VMkernel ports connect to the outside world through the physical Ethernet adapters that are connected to the virtual switch uplink ports.

Adding ESXi Networking

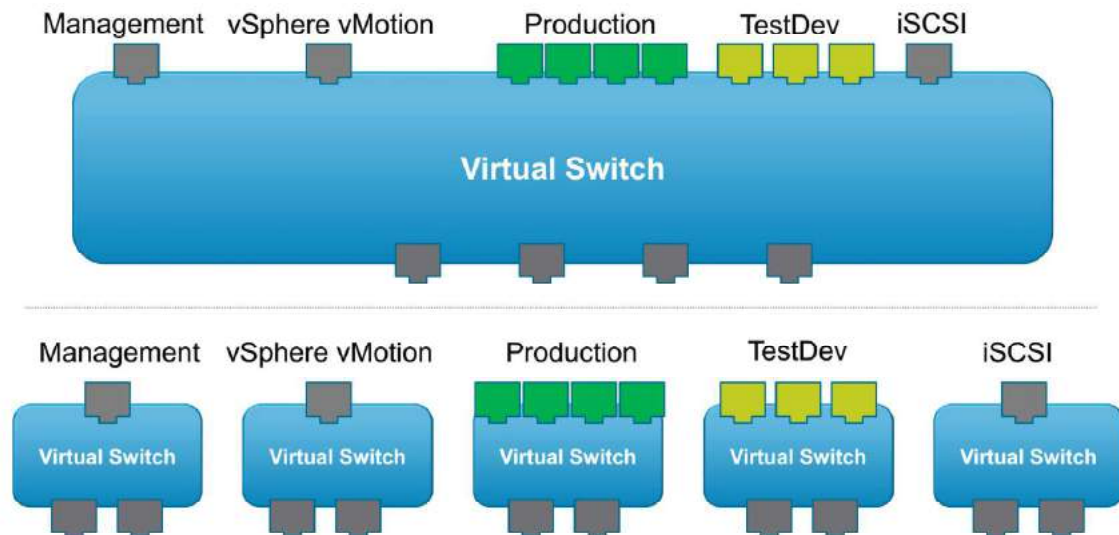
Slide 5-8



Virtual Switch Connection Examples

Slide 5-9

More than one network can coexist on the same virtual switch. Networks can also exist on separate virtual switches.



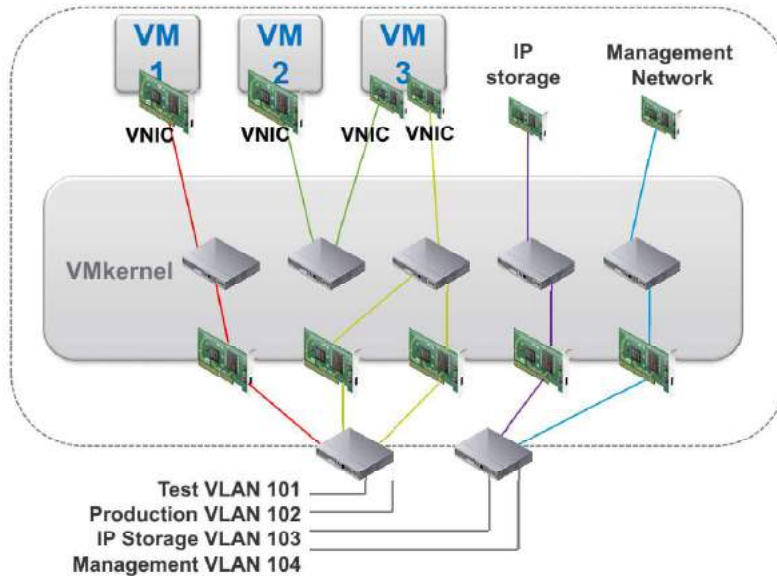
When you are designing your networking environment, vSphere enables you to place all your networks on a single virtual switch. Or you can opt for multiple virtual switches, each with a separate network. The decision partly depends on the layout of your physical networks. For example, you might not have enough network adapters to create a separate virtual switch for each network. Instead, you might team your network adapters in a single virtual switch and isolate the networks by using VLANs.

Because physical NICs are assigned at the virtual switch level, all ports and port groups that are defined for a particular switch share the same hardware.

Standard Switch Components

Slide 5-10

A standard switch provides connections for virtual machines to communicate with one another, whether they are on the same host or on different hosts.



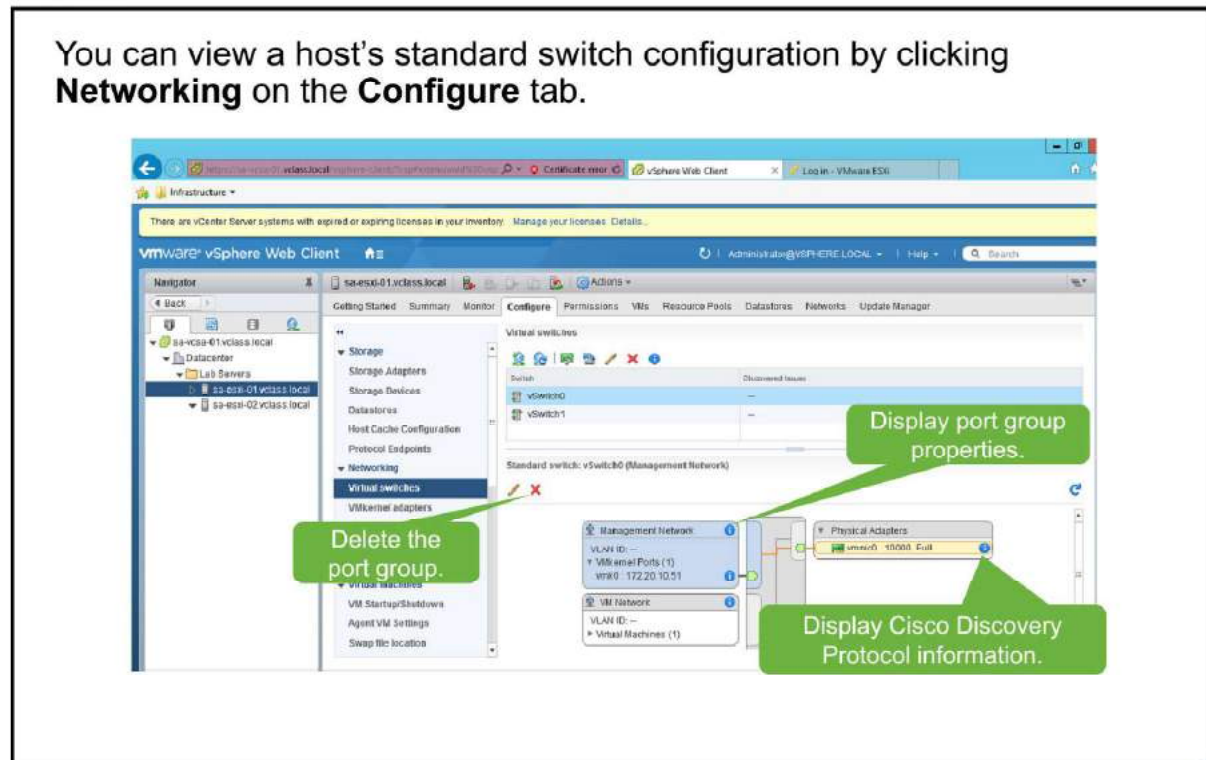
The slide shows five standard switches, each devoted to a different purpose. From left to right, the switches are in numerical order:

1. A standard switch with a single outbound adapter. This switch is used only by VM1.
2. An internal-only standard switch, which enables virtual machines in a single ESXi host to communicate directly with other virtual machines connected to the same standard switch. VM2 and VM3 can use this switch to communicate with each other.
3. A standard switch with teamed NICs. A NIC team provides automatic distribution of packets and failover.
4. A standard switch that is used by the VMkernel for accessing iSCSI- or NAS-based storage.
5. A standard switch that is used by the VMkernel to enable remote management capabilities.

Viewing the Standard Switch Configuration

Slide 5-11

You can view a host's standard switch configuration by clicking **Networking** on the **Configure** tab.



The slide shows the standard switch vSwitch0 on an ESXi host. By default, the ESXi installation creates a virtual machine port group named VM Network and a VMkernel port named Management Network. A good practice is to remove the VM Network virtual machine port group and keep virtual machine networks and management networks separated for performance and security reasons.

To remove a standard switch, click the red X next to the switch to be deleted. To display virtual switch properties, click the pencil icon above the virtual switch.

Port group properties for a port or port group can be displayed. If applicable, Cisco Discovery Protocol (CDP) information can be shown for a physical adapter.

CDP enables ESXi administrators to determine which Cisco switch port is connected to a given virtual switch. When CDP is enabled for a particular virtual switch, you can view properties of the Cisco switch from vSphere Web Client. Properties include device type, port ID, hardware capabilities, and so on.

About VLANs

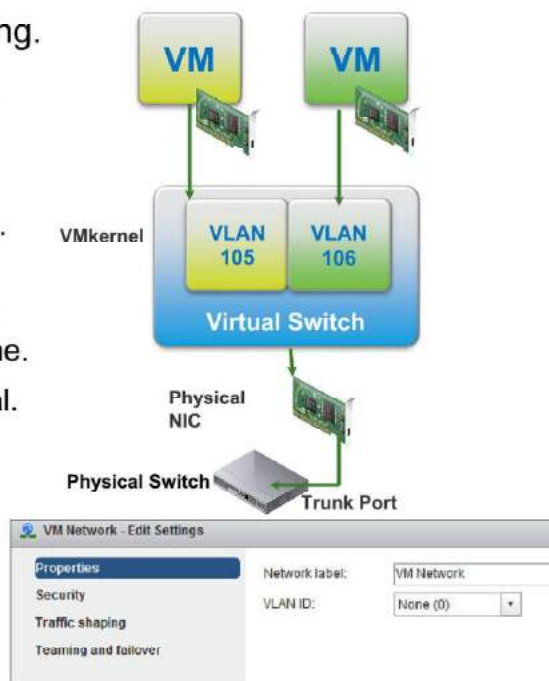
Slide 5-12

ESXi supports 802.1Q VLAN tagging.

Virtual switch tagging is one of the tagging policies supported:

- Frames from a virtual machine are tagged as they exit the virtual switch.
- Tagged frames arriving at a virtual switch are untagged before they are sent to the destination virtual machine.
- The effect on performance is minimal.

ESXi provides VLAN support by giving a VLAN ID to a port group.



VLANs provide for logical groupings of switch ports, enabling communications as if all virtual machines or ports in a VLAN were on the same physical LAN segment. A VLAN is a software-configured broadcast domain. Using a VLAN has the following benefits:

- Creation of logical networks that are not based on the physical topology
- Improved performance by confining broadcast traffic to a subset of ports on a switch
- Cost savings by partitioning the network without overhead of deploying new routers

VLANs can be configured at the port group level. The ESXi host provides VLAN support through virtual switch tagging, which is provided by giving a port group a VLAN ID. By default, a VLAN ID is optional. The VMkernel takes care of all tagging and untagging as the packets pass through the virtual switch.

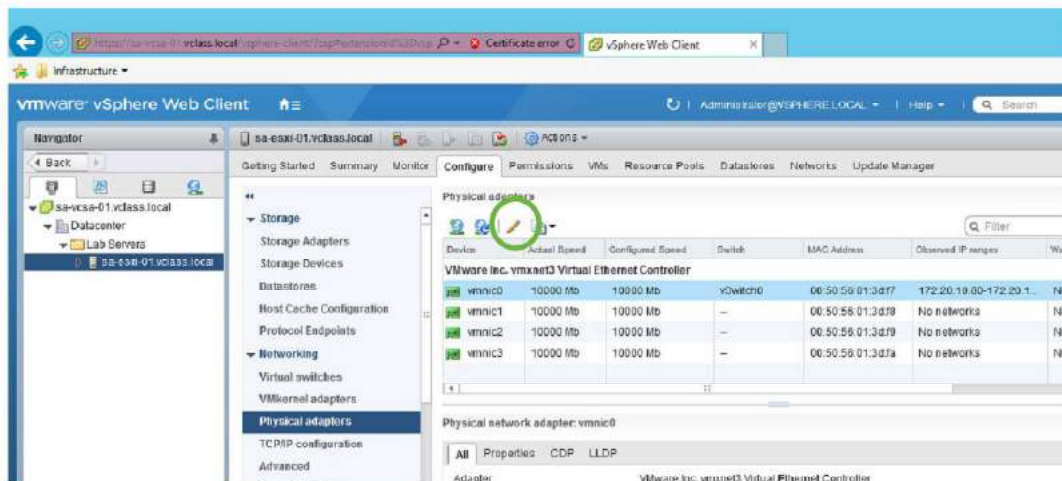
The port on a physical switch to which an ESXi host is connected must be defined as a static trunk port. A trunk port is a port on a physical Ethernet switch that is configured to send and receive packets tagged with a VLAN ID. No VLAN configuration is required in the virtual machine. In fact, the virtual machine does not know that it is connected to a VLAN.

For more information about how VLANs are implemented, see VMware knowledge base article 1003806 at <http://kb.vmware.com/kb/1003806>.

Network Adapter Properties

Slide 5-13

A physical adapter can become a bottleneck for network traffic if the adapter speed does not match application requirements.



You can change the connection speed and duplex of a physical adapter to transfer data in compliance with the traffic rate.

If the physical adapter supports SR-IOV, you can enable it and configure the number of virtual functions to use for virtual machine networking.

Types of Virtual Switches

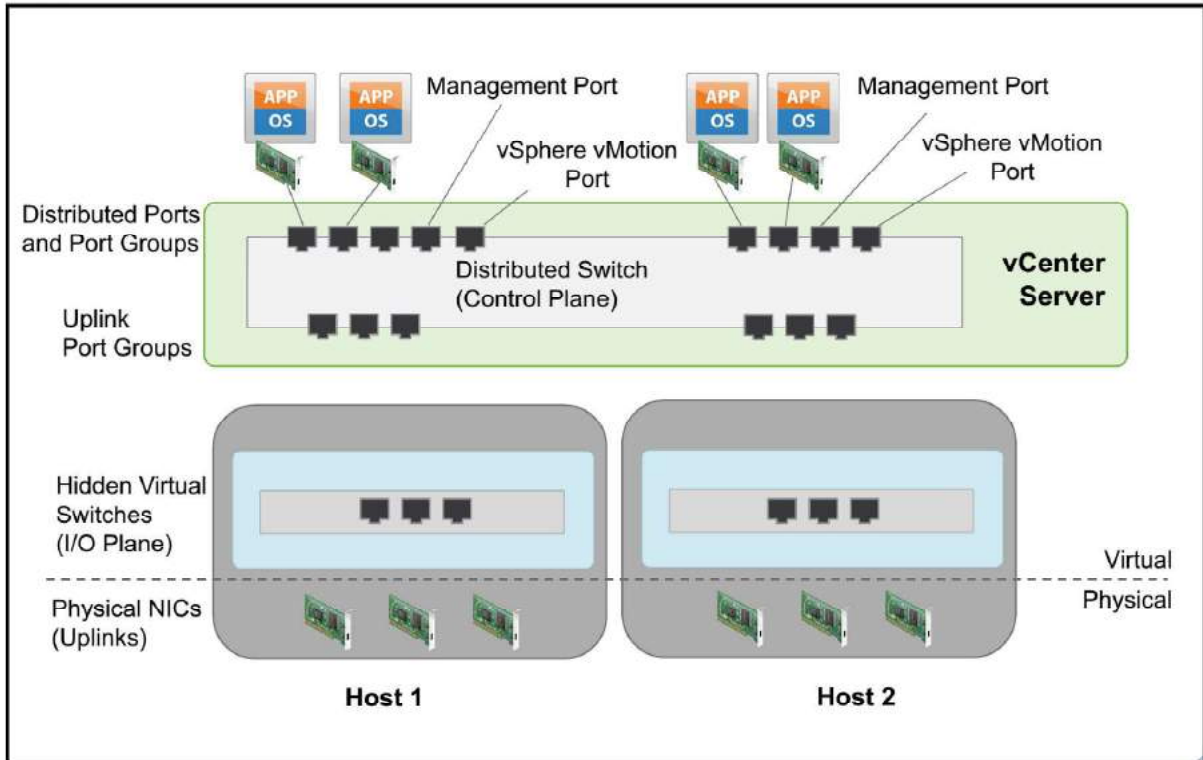
Slide 5-14

A virtual network supports these types of virtual switches:

- Standard switches:
 - Virtual switch configuration for a single host
- Distributed switches:
 - Virtual switches that provide a consistent network configuration for virtual machines as they migrate across multiple hosts.

Distributed Switch Architecture

Slide 5-15



vCenter Server owns the configuration of the distributed switch. The configuration is consistent across all hosts that use the distributed switch.

Standard Switch and Distributed Switch Feature Comparison

Slide 5-16

Feature	Standard Switch	Distributed Switch
Layer 2 switch	✓	✓
VLAN segmentation	✓	✓
IPv6 support	✓	✓
802.1Q tagging	✓	✓
NIC teaming	✓	✓
Outbound traffic shaping	✓	✓
Inbound traffic shaping		✓
VM network port block		✓
Private VLANs		✓
Load-based teaming		✓
Data center-level management		✓
vSphere vMotion migration over a network		✓
Per-port policy settings		✓
Port state monitoring		✓
NetFlow		✓
Port mirroring		✓

Review of Learner Objectives

Slide 5-17

You should be able to meet the following objectives:

- Describe the virtual switch connection types
- Configure and view standard switch configurations, such as virtual machine port group, VMkernel port, VLAN, and security features
- List the features comparison of standard and distributed switches

Lesson 2: Configuring Standard Switch Policies

Slide 5-18



Lesson 2: Configuring Standard Switch Policies

Learner Objectives

Slide 5-19

By the end of this lesson, you should be able to meet the following objectives:

- Explain how to set the security policies for a standard switch port group
- Explain how to set the traffic-shaping policies for a standard switch port group
- Explain how to set the NIC teaming and failover policies for a standard switch port group

Network Switch and Port Policies

Slide 5-20

Policies that are set at the standard switch level apply to all port groups on the standard switch by default.

Available network policies:

- Security
- Traffic shaping
- NIC teaming and failover

Policies are defined at the following levels:

- Standard switch level:
 - Default policies for all the ports on the standard switch.
- Port group level:
 - Effective policies: Policies defined at this level override the default policies that are set at the standard switch level.

Networking security policy provides protection against MAC address impersonation and unwanted port scanning.

Traffic shaping is useful when you want to limit the amount of traffic to a virtual machine or a group of virtual machines. You do traffic shaping to either protect a virtual machine or traffic in an oversubscribed network.

Use the teaming and failover policy to determine how the network traffic of virtual machines and VMkernel adapters that are connected to the switch is distributed between physical adapters, and how the traffic should be rerouted if an adapter fails.


These policies are defined for the entire standard switch and can also be defined for a VMkernel port or a virtual machine port group. When a policy is defined for an individual port or port group, the policy at this level overrides the default policies defined for the standard switch.

Configuring Security Policies

Slide 5-21

Administrators can define security policies at both the standard switch level and the port group level:

- **Promiscuous mode:** Allows a virtual switch or port group to forward all traffic regardless of the destination.
- **MAC address changes:** Accept or reject inbound traffic when the MAC address is altered by the guest.
- **Forged transmits:** Accept or reject outbound traffic when the MAC address is altered by the guest.



For a vSphere standard switch, you can configure security policy to reject the MAC address and promiscuous mode changes in the guest operating system of a virtual machine.

The network security policy contains the following exceptions:

- **Promiscuous mode:** Promiscuous mode allows a virtual switch or port group to forward all traffic regardless of their destinations. The default is **Reject**.
- **MAC address changes:** When set to **Reject**, if the guest attempts to change the MAC address assigned to the virtual NIC, it stops receiving frames. The default is **Accept**.
- **Forged transmits:** A frame's source address field may be altered by the guest, and contain a MAC address other than the assigned virtual NIC MAC address. You can set the Forged transmits parameter to accept or reject such frames. The default is **Accept**.

In general, these policies give you the option of disallowing certain behaviors that might compromise security. For example, a hacker might use a promiscuous mode device to capture network traffic for unscrupulous activities. Or, someone might impersonate a node and gain unauthorized access by spoofing its MAC address.

Set **Promiscuous mode** to **Accept** to use an application in a virtual machine that analyzes or sniffs packets, such as a network-based intrusion detection system.

Set **MAC Address Changes** and **Forged Transmits** to **Reject** to help protect against certain attacks launched by a rogue guest operating system.

Leave **MAC Address Changes** and **Forged Transmits** at their default values (**Accept**) if your applications change the mapped MAC address, as do some guest operating system-based firewalls.

Traffic-Shaping Policy

Slide 5-22

Network traffic shaping is a mechanism for limiting a virtual machine's consumption of available network bandwidth.

Average rate, peak rate, and burst size are configurable.



A virtual machine's network bandwidth can be controlled by enabling the network traffic shaper. The network traffic shaper, when used on a standard switch, shapes only outbound network traffic. To control inbound traffic, use a load-balancing system, or turn on rate-limiting features on your physical router.

Configuring Traffic Shaping

Slide 5-23

A traffic-shaping policy is defined by average bandwidth, peak bandwidth, and burst size. You can establish a traffic-shaping policy for each port group and each distributed port or distributed port group:

- Traffic shaping is disabled by default.
- Parameters apply to each virtual NIC in the standard switch.
- On a standard switch, traffic shaping controls only outbound traffic.

The screenshot shows the 'Production - Edit Settings' window with the 'Traffic shaping' tab selected. The settings are as follows:

Property	Value
Status	Enabled (Override checked)
Average bandwidth (kbit/s)	102400
Peak bandwidth (kbit/s)	204800
Burst size (KB)	102400

The ESXi host shapes only outbound traffic by establishing parameters for the following traffic characteristics:

- **Average bandwidth (Kbps):** Establishes the number of kilobits per second to allow across a port, averaged over time. The average bandwidth is the allowed average load.
- **Peak bandwidth (Kbps):** The maximum number of kilobits per second to allow across a port when it is sending a burst of traffic. This number tops the bandwidth that is used by a port whenever the port is using its burst bonus.
- **Burst size (KB):** The maximum number of kilobytes to allow in a burst. If this parameter is set, a port might gain a burst bonus if it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified in **Average bandwidth**, the port might be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of kilobytes that have accumulated in the burst bonus and thus transfers at a higher speed.

Network traffic shaping is off by default.

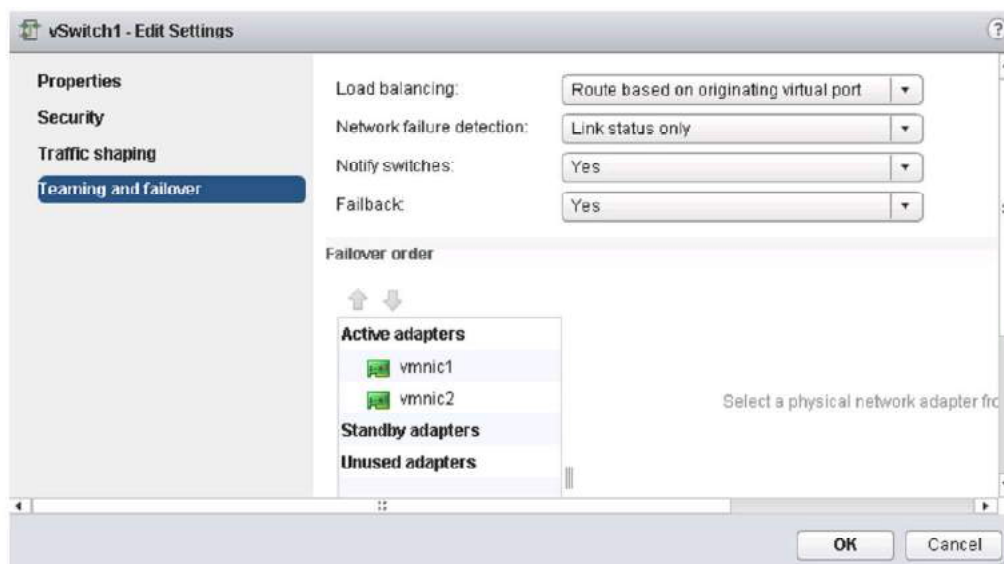
For information about configuring the traffic shaping policy in vSphere Web Client, see *VMware PowerCLI User's Guide* at <https://www.vmware.com/support/developer/PowerCLI/>.

Although you can establish a traffic-shaping policy at either the virtual switch level or the port group level, settings at the port group level override settings at the virtual switch level.

NIC Teaming and Failover Policies

Slide 5-24

Administrators can edit the NIC teaming and failover policy by configuring specific options.



NIC teaming enables you to increase the network capacity of a virtual switch by including two or more physical NICs in a team. To determine how the traffic is rerouted in case of adapter failure, you include physical NICs in a failover order. To determine how the virtual switch distributes the network traffic between the physical NICs in a team, you select load balancing algorithms depending on the needs and capabilities of your environment:

- **NIC teaming policy:** You can use NIC teaming to connect a virtual switch to multiple physical NICs on a host to increase the network bandwidth of the switch and to provide redundancy. A NIC team can distribute the traffic between its members and provide passive failover in case of adapter failure or network outage. You set NIC teaming policies at the virtual switch or port group level for a vSphere standard switch and at the port group or port level for a vSphere distributed switch.
- **Load balancing policy:** The load balancing policy determines how network traffic is distributed between the network adapters in a NIC team. vSphere virtual switches load balance only the outgoing traffic. Incoming traffic is controlled by the load balancing policy on the physical switch.

- **Failback policy:** By default, a failback policy is enabled on a NIC team. If a failed physical NIC returns online, the virtual switch sets the NIC back to active by replacing the standby NIC that took over its slot.

If the physical NIC that stands first in the failover order experiences intermittent failures, the failback policy might lead to frequent changes in the NIC that is used. The physical switch sees frequent changes in MAC addresses, and the physical switch port might not accept traffic immediately when an adapter becomes online. To minimize such delays, you might consider changing the following settings on the physical switch

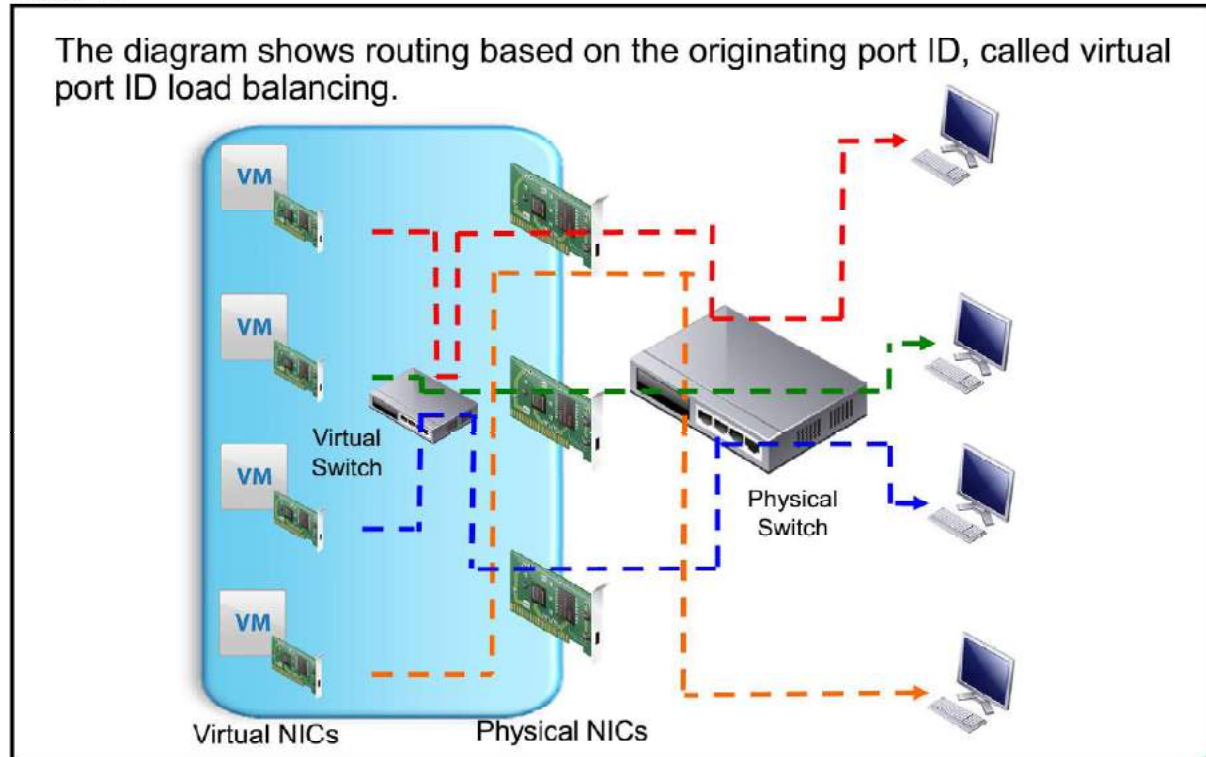
- **Notify switches policy:** By using the notify switches policy, you can determine how the ESXi host communicates failover events. When a physical NIC connects to the virtual switch or when traffic is rerouted to a different physical NIC in the team, the virtual switch sends notifications over the network to update the lookup tables on physical switches. Notifying the physical switch offers lowest latency when a failover or a migration with vSphere vMotion occurs.

NIC teaming and failover policies enable you to determine how network traffic is distributed between adapters and how to reroute traffic if an adapter fails. NIC teaming policies include load-balancing and failover settings. Default NIC teaming and failover policies are set for the entire standard switch. These default settings can be overridden at the port group level. The policies show what is inherited from the settings at the switch level.

For information about configuring NIC teaming, failover, and load balancing settings in vSphere Web Client, see *vSphere Networking* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Load-Balancing Method: Originating Virtual Port ID

Slide 5-25

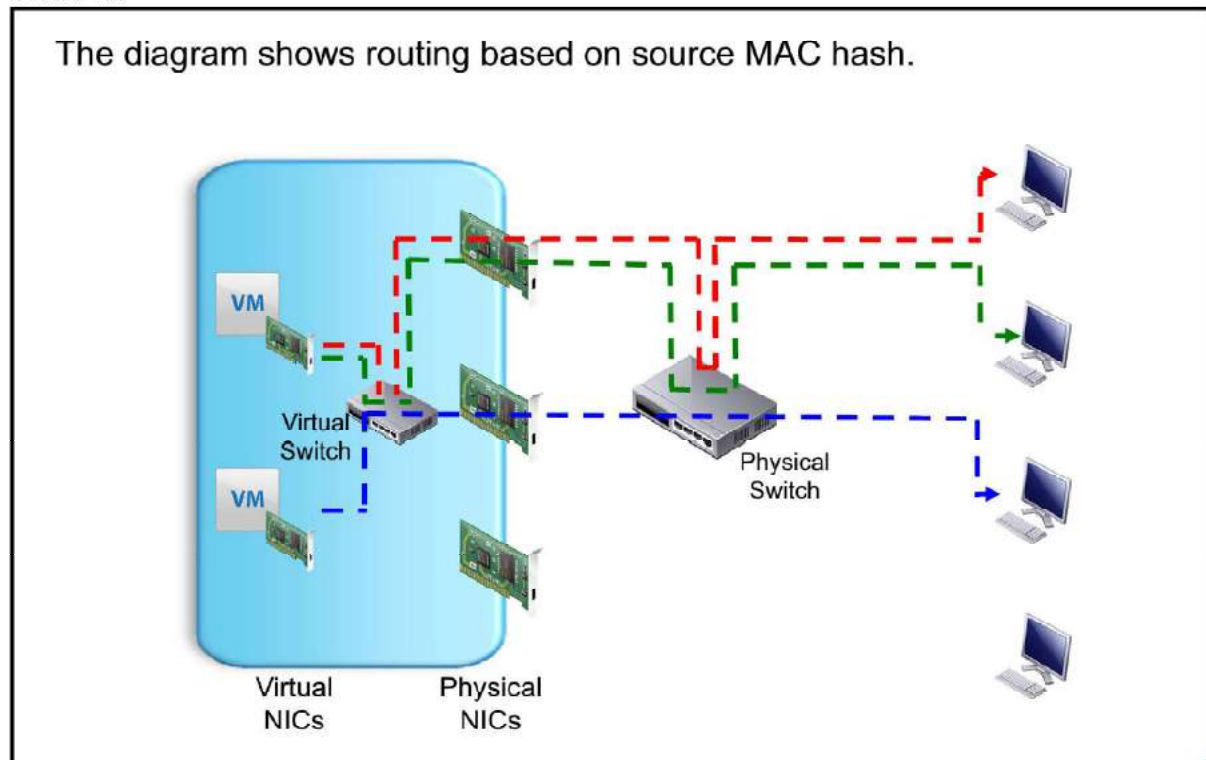


With this method, a virtual machine's outbound traffic is mapped to a specific physical NIC. The NIC is determined by the ID of the virtual port to which this virtual machine is connected. This method is simple and fast and does not require the VMkernel to examine the frame for necessary information.

When the load is distributed in the NIC team by using the port-based method, no single-NIC virtual machine gets more bandwidth than can be provided by a single physical adapter.

Load-Balancing Method: Source MAC Hash

Slide 5-26



In this load-balancing method, each virtual machine's outbound traffic is mapped to a specific physical NIC that is based on the virtual NIC's MAC address.

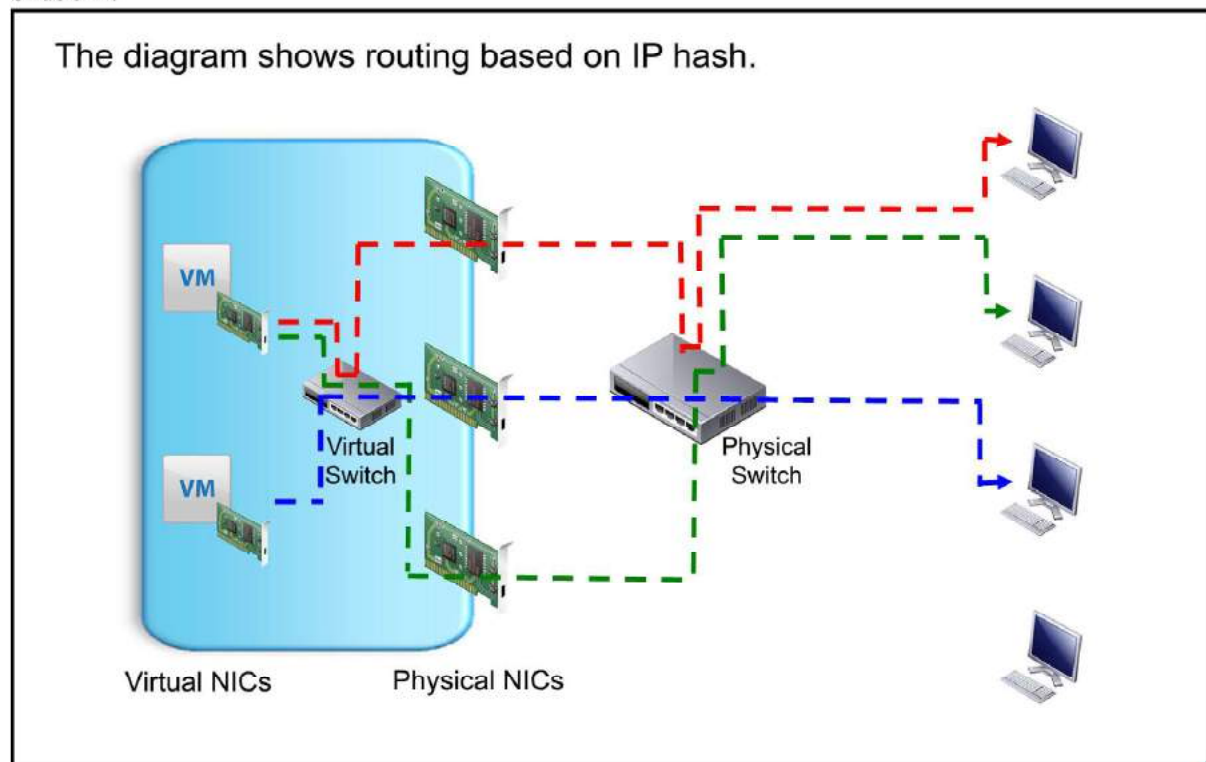
This method has low overhead and is compatible with all switches, but it might not spread traffic evenly across all the physical NICs.

When the load is distributed in the NIC team by using the MAC-based method, no single-NIC virtual machine gets more bandwidth than can be provided by a single physical adapter.

You can also balance your traffic based on the current traffic loads of the physical NICs. The NIC with less load is more likely to be chosen

Load-Balancing Method: Source and Destination IP Hash

Slide 5-27



In this load-balancing method, a NIC for each outbound packet is selected based on its source and destination IP addresses.

The IP-based method requires 802.3ad link aggregation support or EtherChannel on the switch. The Link Aggregation Control Protocol is a method to control the bundling of several physical ports to form a single logical channel. LACP is part of the IEEE 802.3ad specification.

EtherChannel is a port trunking technology that is used primarily on Cisco switches. This technology enables grouping several physical Ethernet links to create one logical Ethernet link for providing fault tolerance and high-speed links between switches, routers, and servers.

When the load is distributed in the NIC team using the IP-based method, a single-NIC virtual machine might use the bandwidth of multiple physical adapters.

The IP-based load-balancing method only affects outbound traffic. For example, a virtual machine might choose a particular NIC to communicate with a particular destination virtual machine. The return traffic might not arrive on the same NIC as the outbound traffic. The return traffic might arrive on another NIC in the same NIC team. For a list of ESXi host requirements for link aggregation, see VMware knowledge base article 1001938 at <http://kb.vmware.com/kb/1001938>.

Detecting and Handling Network Failure

Slide 5-28

The VMkernel can use link status or beaconing, or both, to detect a network failure.

Network failure is detected by the VMkernel, which monitors the link state and performs beacon probing.

The VMkernel notifies physical switches of changes in the physical location of a MAC address.

Failover is implemented by the VMkernel based on configurable parameters:

- **Failback:** How the physical adapter is returned to active duty after recovering from failure.
- **Load-balancing option:** Use explicit failover order. Always use the vmnic uplink at the top of the active adapter list.

Monitoring the link status provided by the network adapter detects failures like cable pulls and physical switch power failures. This monitoring does not detect configuration errors, such as a physical switch port being blocked by the Spanning Tree Protocol or misconfigured VLAN membership. This method cannot detect upstream, nondirectly connected physical switch or cable failures.

Beaconing introduces a load of a 62-byte packet approximately every 1 second per physical NIC. When beaconing is activated, the VMkernel sends out and listens for probe packets on all NICs in the team. This technique can detect failures that link-status monitoring alone cannot. Consult your switch manufacturer to confirm the support of beaconing in your environment. See VMware knowledge base article 1005577 at <http://kb.vmware.com/kb/1005577>.

A physical switch can be notified by the VMkernel whenever a virtual NIC is connected to a virtual switch. A physical switch can also be notified whenever a failover event causes a virtual NIC's traffic to be routed over a different physical NIC. The notification is sent out over the network to update the lookup tables on physical switches. In most cases, this notification process is desirable because otherwise virtual machines would experience greater latency after failovers and vSphere vMotion operation. But do not set this option when the virtual machines connected to the port group are running unicast-mode Microsoft Network Load Balancing (NLB). NLB in multicast mode is unaffected. For more about the NLB issue, see VMware knowledge base article 1556 at <http://kb.vmware.com/kb/1556>.

When using explicit failover order, always use the highest order uplink from the list of active adapters that pass failover-detection criteria.

The failback option determines how a physical adapter is returned to active duty after recovering from a failure. If **Failback** is set to **Yes**, the failed adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took its place at the time of failure. If **Failback** is set to **No**, a failed adapter is left inactive even after recovery, until another currently active adapter fails, requiring its replacement.

Physical Network Considerations

Slide 5-29

Your virtual networking environment relies on the physical network infrastructure. As a vSphere administrator, you should discuss your vSphere networking needs with your network administration team.

The following issues are topics for discussion:

- Number of physical switches
- Network bandwidth that is required
- Physical switch configuration support for 802.3ad, for NIC teaming
- Physical switch configuration support for 802.1Q, for VLAN tagging
- Physical switch configuration support for Link Aggregation Control Protocol (LACP)
- Network port security
- Link Layer Discovery Protocol (LLDP) and Cisco Discovery Protocol (CDP) and their operation modes, such as listen, broadcast, listen and broadcast, and disabled

Lab 7: Using Standard Switches

Slide 5-30

Create a standard switch and a port group

1. View the Standard Switch Configuration
2. Create a Standard Switch with a Virtual Machine Port Group
3. Attach Your Virtual Machines to the New Virtual Machine Port Group

Review of Learner Objectives

Slide 5-31

You should be able to meet the following objectives:

- Explain how to set the security policies for a standard switch port group
- Explain how to set the traffic-shaping policies for a standard switch port group
- Explain how to set the NIC teaming and failover policies for a standard switch port group

Key Points

Slide 5-32


- The following connection types can exist on a virtual switch: virtual machine port group, VMkernel, and physical uplinks.
- A standard switch is a virtual switch configuration for a single host.
- Network policies set at the standard switch level can be overridden at the port group level.
- A distributed switch provides centralized management and monitoring of the networking configuration of all hosts that are associated with the switch.

Questions?

MODULE 6

Configuring and Managing Virtual Storage

Slide 6-1



You Are Here

Slide 6-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
- 6. Configuring and Managing Virtual Storage**
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Importance

Slide 6-3

Storage options give you the flexibility to set up your storage based on your cost, performance, and manageability requirements.

Shared storage is useful for disaster recovery, high availability, and moving virtual machines between hosts.


Module Lessons

Slide 6-4

- | | |
|-----------|-----------------------|
| Lesson 1: | Storage Concepts |
| Lesson 2: | Fibre Channel Storage |
| Lesson 3: | iSCSI Storage |
| Lesson 4: | VMFS Datastores |
| Lesson 5: | NFS Datastores |
| Lesson 6: | vSAN Datastores |

Lesson 1: Storage Concepts

Slide 6-5



Learner Objectives

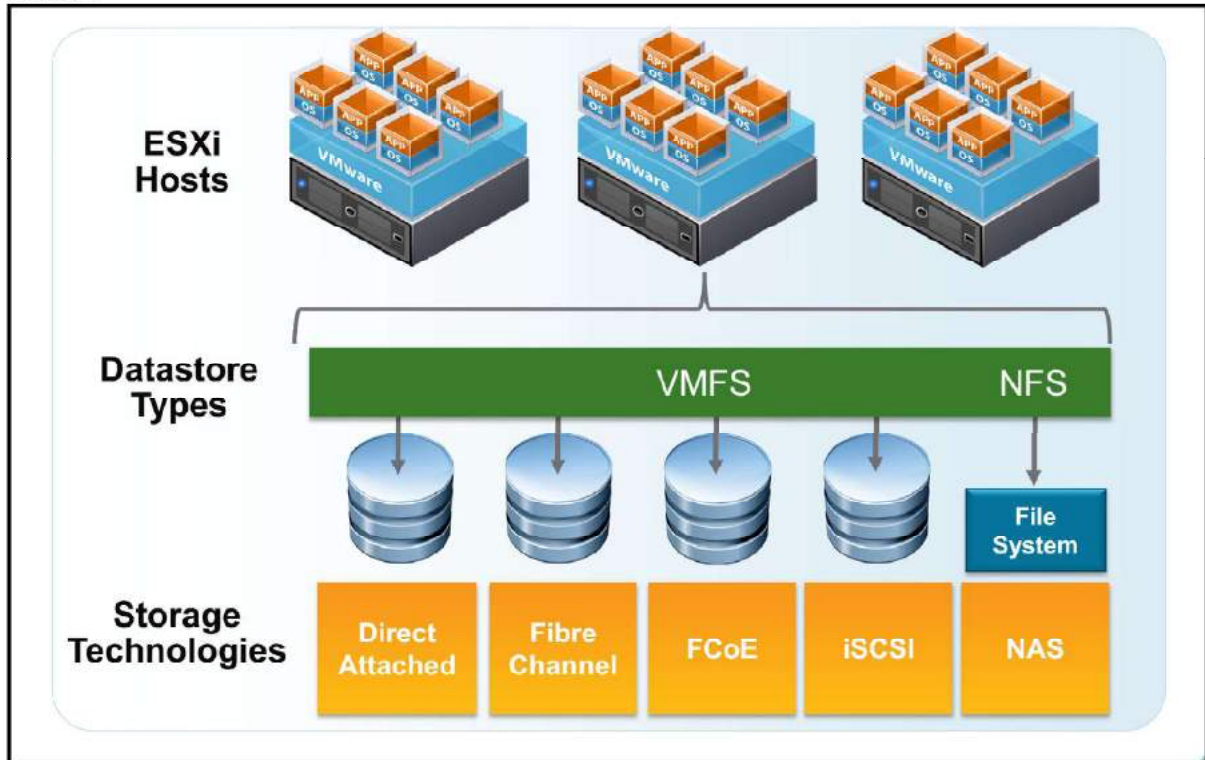
Slide 6-6

By the end of this lesson, you should be able to meet the following objective:

- Describe vSphere storage technologies and datastores

Basic Storage Overview

Slide 6-7



ESXi hosts should be configured so that they have shared access to datastores. Datastores are logical containers that hide specifics of each storage device and provide a uniform model for storing virtual machine files. Depending on the type of storage that you use, datastores can be formatted with VMFS, or with a file system native to a storage device that is shared using the NFS protocol.

Several storage technologies are supported by ESXi hosts in the vSphere environment:

- Direct-attached storage: Internal or external storage disks or arrays attached to the host through a direct connection instead of a network connection.
- Fibre Channel: A high-speed transport protocol used for SANs. Fibre Channel encapsulates SCSI commands, which are transmitted between Fibre Channel nodes. In general, a Fibre Channel node is a server, a storage system, or a tape drive. A Fibre Channel switch interconnects multiple nodes, forming the fabric in a Fibre Channel network.
- FCoE: The Fibre Channel traffic is encapsulated into Fibre Channel over Ethernet (FCoE) frames. These FCoE frames are converged with other types of traffic on the Ethernet network. By carrying both Ethernet and Fibre Channel traffic on the same Ethernet link, the utilization of the physical infrastructure has greatly increased. FCoE also reduces the total number of network ports and cabling.

- iSCSI: A SCSI transport protocol, enabling access to storage devices and cabling over standard TCP/IP networks. iSCSI maps SCSI block-oriented storage over TCP/IP. Initiators, such as an iSCSI host bus adapter (HBA) in an ESXi host, send SCSI commands to targets, located in iSCSI storage systems.
- NAS: Storage shared over standard TCP/IP networks at the file system level. NAS storage is used to hold NFS datastores. The NFS protocol does not support SCSI commands.

iSCSI, NAS, and FCoE can run over Gigabit Ethernet or 10 Gigabit Ethernet. 10GigE provides increased storage performance levels and sufficient bandwidth that permits multiple types of high-bandwidth protocol traffic to coexist on the same network.

Storage Protocol Overview

Slide 6-8

Storage Protocol	Boot from SAN Support	vSphere vMotion Support	vSphere HA Support	vSphere DRS Support	Raw Device Mapping Support
Fibre Channel	•	•	•	•	•
FCoE	•	•	•	•	•
iSCSI	•	•	•	•	•
NFS		•	•	•	
DAS		•			•
Virtual Volumes		•	•	•	
vSAN		•	•	•	

Direct-attached storage, as opposed to SAN storage, is where many administrators install ESXi. Local storage is also ideal for small environments due to the cost savings associated with purchasing and managing a SAN. The drawback is that you lose many of the features that make virtualization a worthwhile investment, for example, balancing the workload on a specific ESXi host. Direct-attached storage can also be used to store noncritical data, such as:

- CD-ROM ISO images
- Decommissioned virtual machines
- Virtual machine templates

In comparison, storage logical unit numbers (LUNs) must be pooled and shared so that all ESXi hosts can access them. Shared storage enables vSphere features like:

- vSphere vMotion
- vSphere HA
- vSphere DRS

Using shared SAN storage also enables robust features in vSphere, such as:

- Central repositories for virtual machine files and templates
- Clustering of virtual machines across ESXi hosts
- Allocation of large amounts (terabytes) of storage to your ESXi hosts

ESXi supports different methods of booting from the SAN to avoid handling maintenance of additional local storage or if you have diskless hardware configurations, such as blade systems. When you set up your host to boot from a SAN, your host's boot image is stored on one or more LUNs in the SAN storage system. When the host starts, it boots from the LUN on the SAN rather than from its local disk.

ESXi hosts allow booting from software iSCSI, a supported independent hardware SCSI adapter, and a supported dependent hardware iSCSI adapter. The network adapter must support only the iSCSI Boot Firmware Table (iBFT) format, which is a method of communicating parameters about the iSCSI boot device to an operating system. You will learn more about independent and dependent hardware iSCSI later.

About Datastores


Slide 6-9

A datastore is a logical storage unit that can use disk space on one physical device or span several physical devices.

Datastores are used to hold virtual machine files, templates, and ISO images.

Types of datastores:

- VMFS
- NFS
- vSAN
- vSphere Virtual Volumes



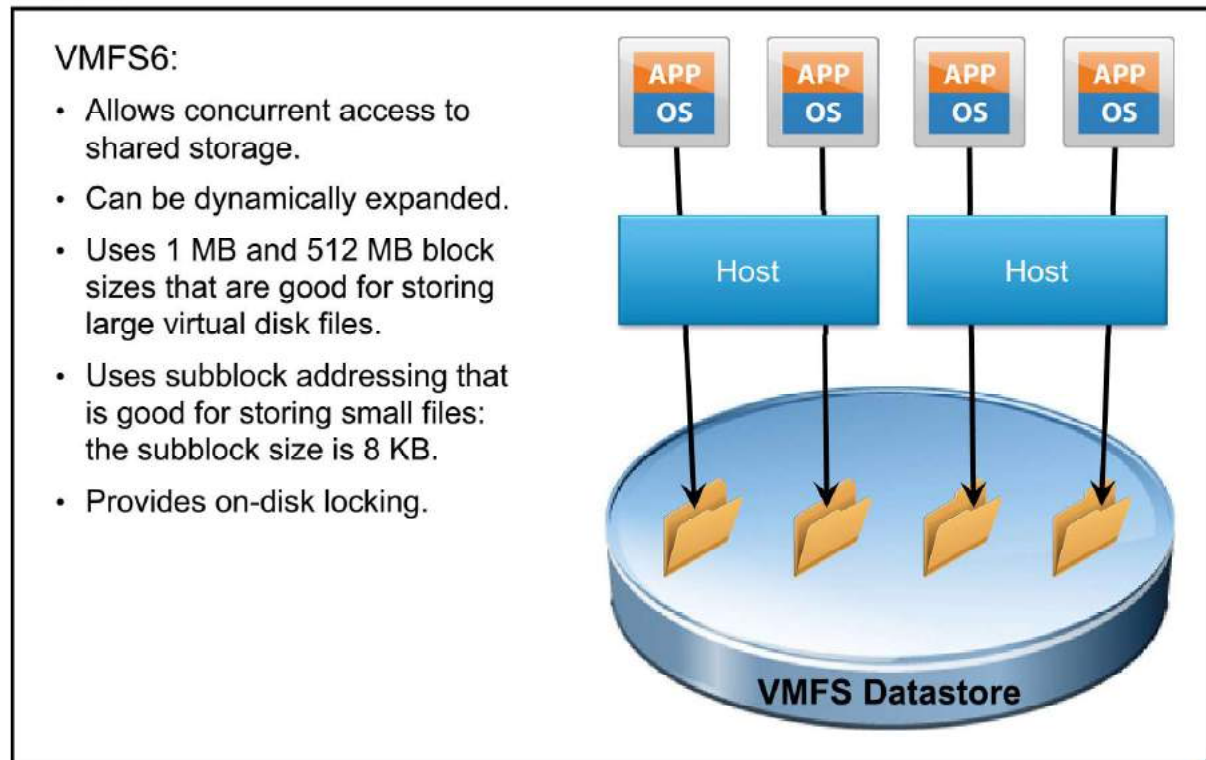
A datastore is a generic term for a container that holds files and objects. Datastores are logical containers, analogous to file systems, which hide specifics of each storage device and provide a uniform model for storing virtual machine files.

vSphere supports four types of file systems: VMFS, NFS, vSAN, and vSphere Virtual Volumes. You can display all datastores that are available to your hosts and analyze their properties.

A virtual machine is stored as a set of files in its own directory or a group of objects in a datastore. Datastores can also be used to store ISO images, floppy images, virtual machines templates, and so on.

About VMFS6

Slide 6-10



VMFS is a clustered file system that allows multiple ESXi hosts to read and write to the same storage device simultaneously. The clustered file system enables unique, virtualization-based services, including:

- Migration of running virtual machines from one ESXi host to another without downtime
- Automatic restarting of a failed virtual machine on a separate ESXi host
- Clustering of virtual machines across various physical servers

VMFS enables IT organizations to greatly simplify virtual machine provisioning by efficiently storing the entire machine state in a central location. VMFS enables multiple ESXi hosts to concurrently access shared virtual machine storage.

The size of a VMFS datastore can be increased dynamically while virtual machines residing on the VMFS datastore are powered on and running. A VMFS datastore efficiently stores both large and small files belonging to a virtual machine. A VMFS datastore can support virtual disk files. A virtual disk file has a maximum of 62 TB in size. A VMFS datastore uses subblock addressing to make efficient use of storage for small files.

VMFS provides block-level distributed locking to ensure that the same virtual machine is not powered on by multiple servers at the same time. If an ESXi host fails, the on-disk lock for each virtual machine will be released so that virtual machines can be restarted on other ESXi hosts.

On the slide, each ESXi host has two virtual machines running on it. The lines connecting the virtual machines to the virtual machine disks (VMDKs) are logical representations of the association and allocation of the larger VMFS datastore. The VMFS datastore includes one or more LUNs. The virtual machines see the assigned storage volume only as a SCSI target from within the guest operating system. The virtual machine contents are only files on the VMFS volume.

VMFS can be deployed on three kinds of SCSI-based storage devices:

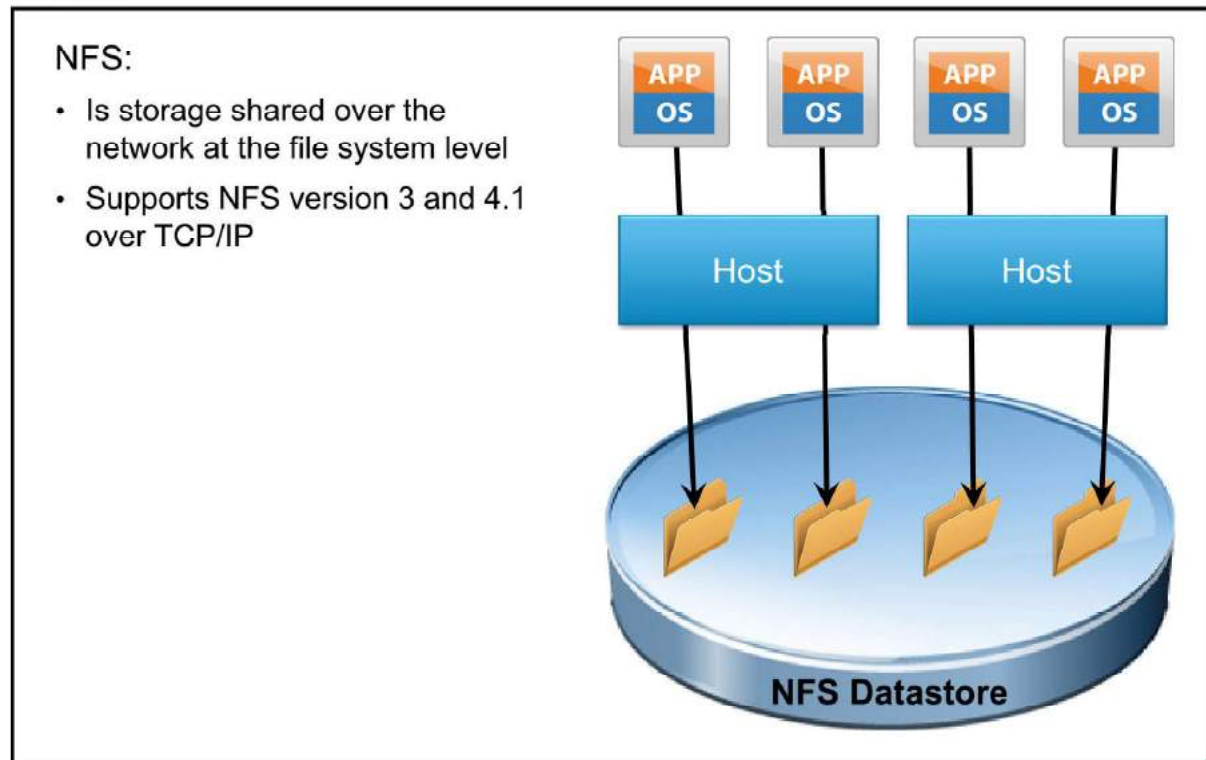
- Direct-attached storage
- Fibre Channel storage
- iSCSI storage

A virtual disk stored on a VMFS datastore always appears to the virtual machine as a mounted SCSI device. The virtual disk hides the physical storage layer from the virtual machine's operating system.

For the operating system in the virtual machine, VMFS preserves the internal file system semantics. Thus the operating system running in the virtual machine sees a native file system, not VMFS. These semantics ensure correct behavior and data integrity for applications running on the virtual machines.

About NFS

Slide 6-11



NFS is a file-sharing protocol that ESXi hosts use to communicate with a NAS device. NAS is a specialized storage device that connects to a network and can provide file access services to ESXi hosts.

NFS datastores are treated like VMFS datastores because they can be used to hold virtual machine files, templates, and ISO images. In addition, an NFS volume allows the vSphere vMotion migration of virtual machines whose files reside on an NFS datastore. The NFS client built into ESXi uses NFS protocol version 3 to communicate with the NAS/NFS servers.

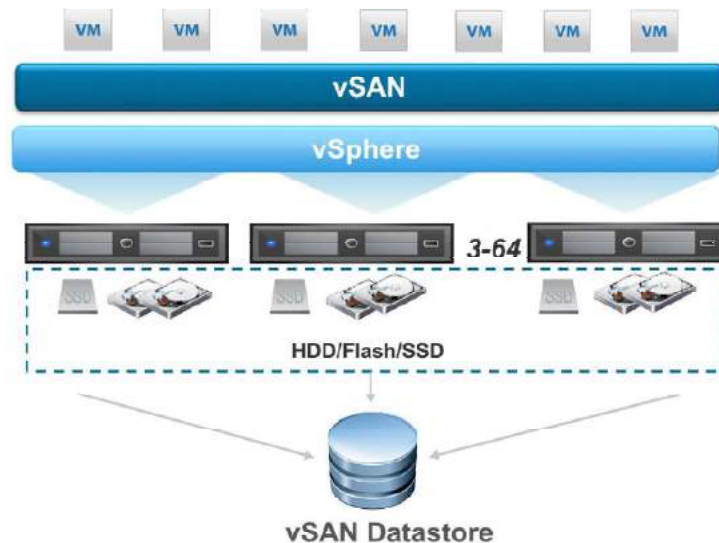
ESXi hosts do not use the Network Lock Manager protocol, which is a standard protocol used to support the file locking of NFS-mounted files. VMware has its own locking protocol. NFS locks are implemented by creating lock files on the NFS server. Lock files are named `.lck-fileid`, where `fileid` is the value of the `fileid` field. When a lock file is created, an update is periodically sent to the lock file to inform other ESXi hosts that the lock is still active. The lock file updates generate small (84-byte) WRITE requests to the NFS server.

vSAN Overview

Slide 6-12

vSAN is hypervisor-converged, software-defined storage for virtual environments.

By clustering host-attached hard disks (HDDs) or solid-state drives (SSDs), vSAN creates an aggregated datastore shared by virtual machines.



When vSAN is enabled on a cluster, a single vSAN datastore is created. This datastore uses the storage components of each host in the cluster.

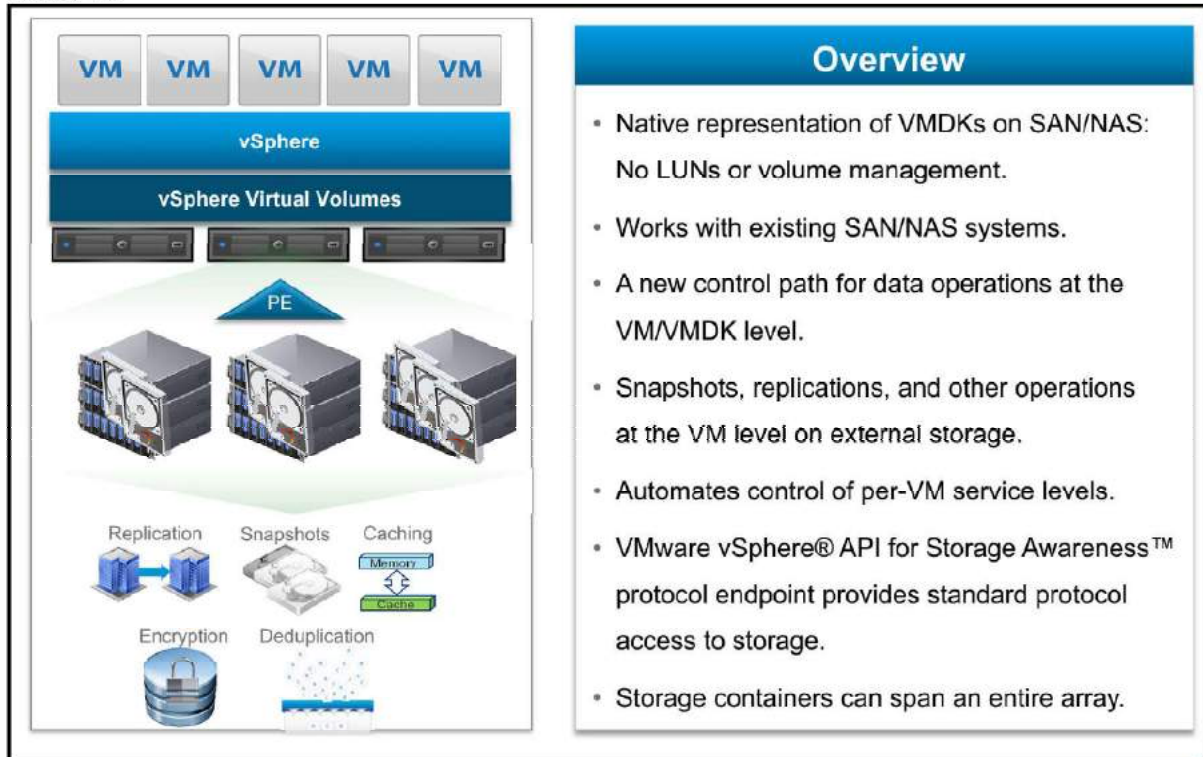
The storage is mounted by using Object Store File System (OSFS). vSAN stores and manages the data on the vSAN datastore as flexible data containers called objects. An object is a logical volume that has its data and metadata distributed and accessed across the entire cluster. In the ESXi storage stack, these objects appear as devices.

Although a single vSAN datastore is created for the entire vSAN cluster, the datastore can have multiple storage policies associated with it. These storage policies can be configured with different storage capabilities.

vSAN can be configured as hybrid or all-flash storage. In a hybrid storage architecture, vSAN pools server-attached HDDs and SSDs to create a distributed shared datastore that abstracts the storage hardware to provide a software-defined storage tier for virtual machines. Flash is used as a read cache/write buffer to accelerate performance and magnetic disks provide capacity and persistent data storage. Alternately, vSAN can be deployed as an all-flash storage architecture in which flash devices are used as a write cache while SSDs provide capacity, data persistence and consistent, fast response times. The all-flash architecture allows tiering of SSDs for a cost-effective implementation: a write-intensive, enterprise-grade SSD cache tier and a read-intensive, lower cost SSD capacity tier.

About vSphere Virtual Volumes

Slide 6-13



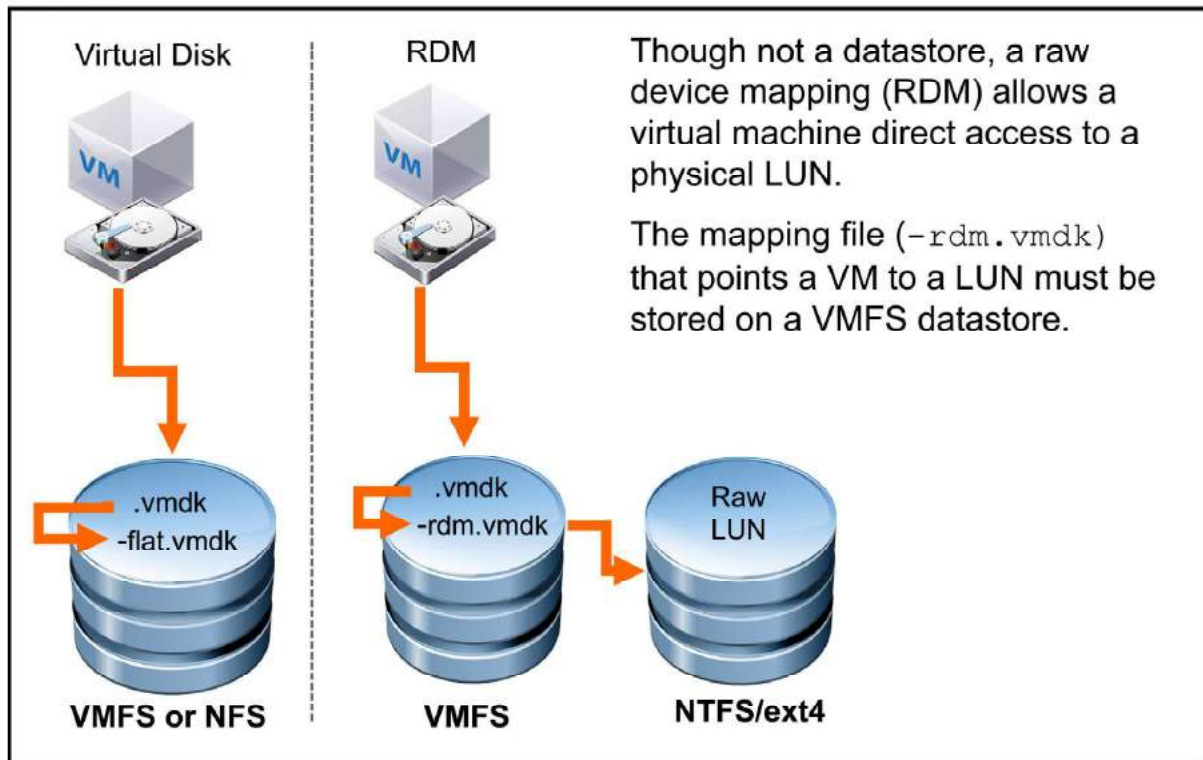
vSphere Virtual Volumes introduces a new storage paradigm that is designed to address the requirements of next-generation storage in the software-defined data center.

vSphere Virtual Volumes provides the following benefits:

- Lower cost of storage
- Reduced storage management overhead
- Greater scalability
- Better response to data access and analytical requirements

About Raw Device Mapping

Slide 6-14



An RDM is a file stored in a VMFS volume that acts as a proxy for a raw physical device.

Instead of storing virtual machine data in a virtual disk file stored on a VMFS datastore, you can store the guest operating system data directly on a raw LUN. Storing the data is useful if you are running applications in your virtual machines that must know the physical characteristics of the storage device. Mapping a raw LUN enables you to use existing SAN commands to manage storage for the disk.

An RDM is recommended when a virtual machine must interact with a real disk on the SAN. This condition exists when you make disk array snapshots or have a large amount of data that you do not want to move onto a virtual disk as part of a physical-to-virtual conversion.

Physical Storage Considerations

Slide 6-15

You should discuss vSphere storage needs with your storage administration team, including the following items:

- LUN sizes
- I/O bandwidth
- I/O requests per second that a LUN is capable of
- Disk cache parameters
- Zoning and masking
- Identical LUN presentation to each ESXi host
- Active-active or active-passive arrays
- Export properties for NFS datastores

As the vSphere administrator, before implementing your vSphere environment, you should discuss the storage needs with your storage administration team:

- LUN sizes
- I/O bandwidth that is required by your applications
- Disk cache parameters, zoning, and masking
- Identical LUN presentation to each ESXi host (if canonical names are not presented)
- Which multipathing setting to use (active-active or active-passive) for your storage arrays
- What NFS settings to use

For information to help you plan for your storage needs, see *vSphere Storage* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Review of Learner Objectives

Slide 6-16

You should be able to meet the following objective:

- Describe vSphere storage technologies and datastores

Lesson 2: Fibre Channel Storage

Slide 6-17



Learner Objectives

Slide 6-18

By the end of this lesson, you should be able to meet the following objectives:

- Describe uses of Fibre Channel with ESXi
- Describe Fibre Channel components and addressing
- Explain how multipathing with Fibre Channel work

About Fibre Channel

Slide 6-19

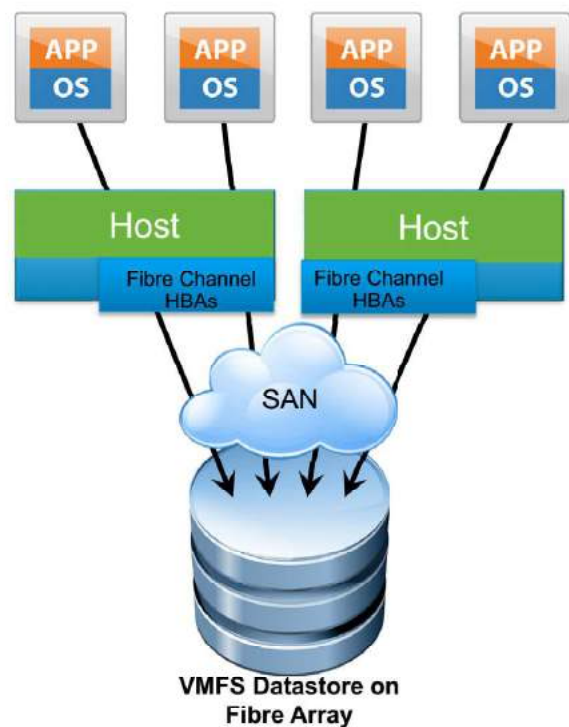
Fibre Channel stores virtual machine files remotely on a Fibre Channel SAN.

A Fibre Channel SAN is a specialized high-speed network that connects your hosts to high-performance storage devices.

The network uses the Fibre Channel protocol to transport SCSI traffic from virtual machines to the Fibre Channel SAN devices.

ESXi supports:

- 16 Gbps Fibre Channel
- Fibre Channel over Ethernet (FCoE)



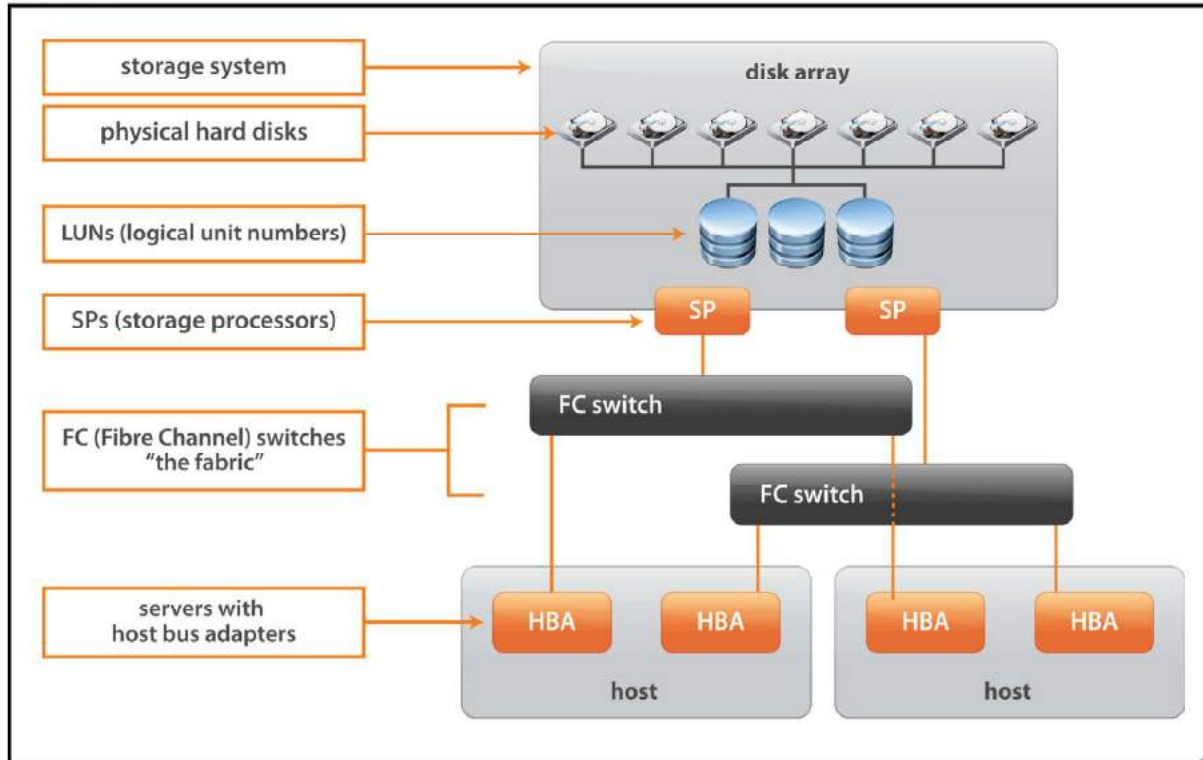
To connect to the Fibre Channel SAN, your host should be equipped with Fibre Channel host bus adapters (HBAs).

Unless you use Fibre Channel direct connect storage, you need Fibre Channel switches to route storage traffic. If your host contains Fibre Channel over Ethernet (FCoE) adapters, you can connect to your shared Fibre Channel devices by using an Ethernet network.

In this configuration, a host connects to a SAN fabric, which consists of Fibre Channel switches and storage arrays, using a Fibre Channel adapter. LUNs from a storage array become available to the host. You can access the LUNs and create datastores for your storage needs. The datastores use the VMFS format.

Fibre Channel SAN Components

Slide 6-20



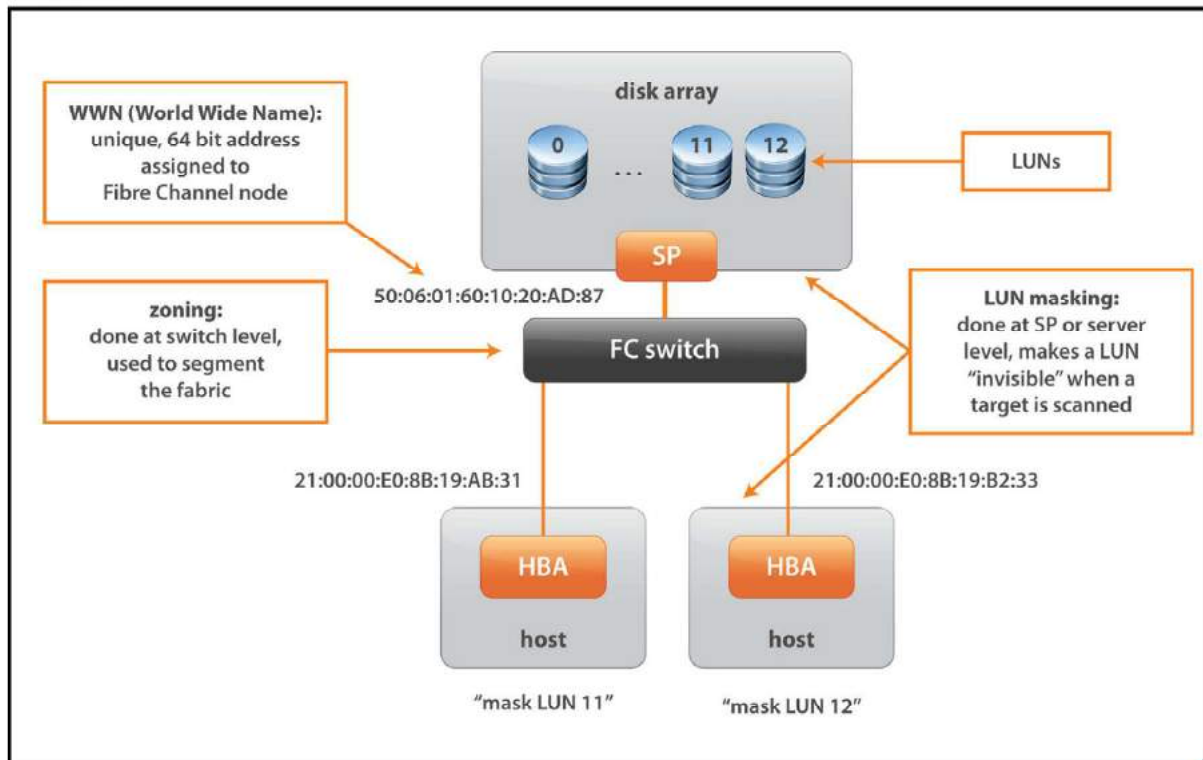
In its simplest form, a SAN consists of one or more servers attached to a storage array using one or more SAN switches. Each server might host numerous applications that require dedicated storage for applications processing.

The following components are involved:

- SAN switches: SAN switches connect various elements of the SAN. In particular, SAN switches might connect hosts to storage arrays. SAN switches also allow administrators to set up path redundancy if a path failure occurs from host server to switch or from storage array to switch.
- Fabric: The SAN fabric is the actual network portion of the SAN. When one or more SAN switches are connected, a fabric is created. The Fibre Channel protocol is used to communicate over the entire network. A SAN can consist of multiple interconnected fabrics. Even a simple SAN often consists of two fabrics for redundancy.
- Connections (HBAs and storage processors): Host servers and storage systems are connected to the SAN fabric through ports in the fabric:
 - A host connects to a fabric port through an HBA.
 - Storage devices connect to the fabric ports through their storage processors.

Fibre Channel Addressing and Access Control

Slide 6-21



A port connects from a device into the SAN. Each node in the SAN includes each host, storage device, and fabric component (router or switch). Each node in the SAN has one or more ports that connect it to the SAN. Ports can be identified in the following ways:

- **World Wide Port Name (WWPN):** A globally unique identifier for a port which allows certain applications to access the port. The Fibre Channel switches discover the WWPN of a device or host and assign a port address to the device.
- **Port_ID:** Within SAN, each port has a unique port ID that serves as the Fibre Channel address for that port. The FC switches assign the port ID when the device logs in to the fabric. The port ID is valid only while the device is logged on.

You can use zoning and LUN masking to segregate SAN activity and restrict access to storage devices.

You can protect access to storage in your vSphere environment by using zoning and LUN masking with your SAN resources. For example, you might manage zones defined for testing independently within the SAN so that they do not interfere with activity in the production zones. Similarly, you might set up different zones for different departments.

When you set up zones, take into account any host groups that are set up on the SAN device.

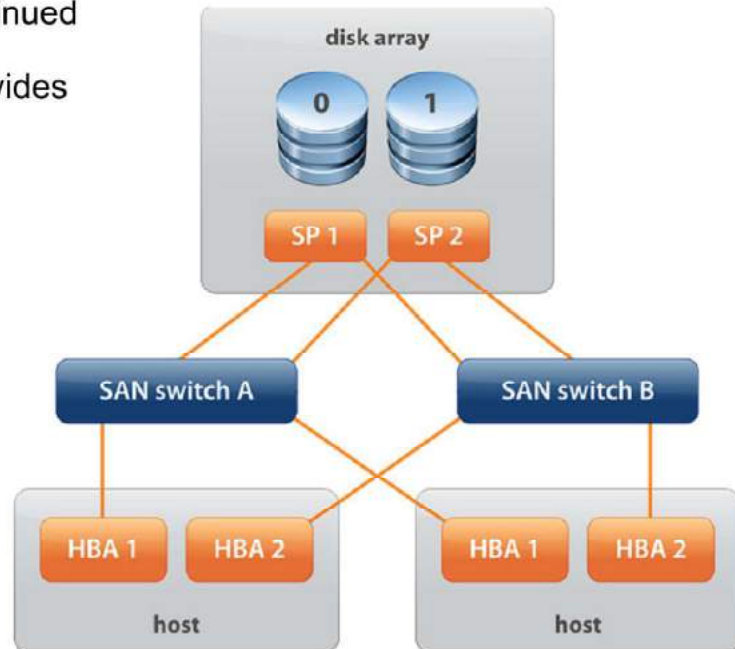
Zoning and masking capabilities for each SAN switch and disk array and the tools for managing LUN masking are vendor specific.

See your SAN vendor's documentation and *vSphere Storage* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Multipathing with Fibre Channel

Slide 6-22

Multipathing enables continued access to SAN LUNs if hardware fails. It also provides load balancing.



An Fibre Channel path describes a route:

- From a specific HBA port in the host
- Through the switches in the fabric
- Into a specific storage port on the storage array.

A given host might be able to access a LUN on a storage array through more than one path. Having more than one path from a host to a LUN is called multipathing.

By default, ESX/ESXi hosts use only one path from a host to a given LUN at any time. If the path actively being used by the VMware ESX/ESXi host fails, the server selects another of the available paths. The process of detecting a failed path and switching to another is called path failover. A path fails if any of the components along the path (HBA, cable, switch port, or storage processor) fails.

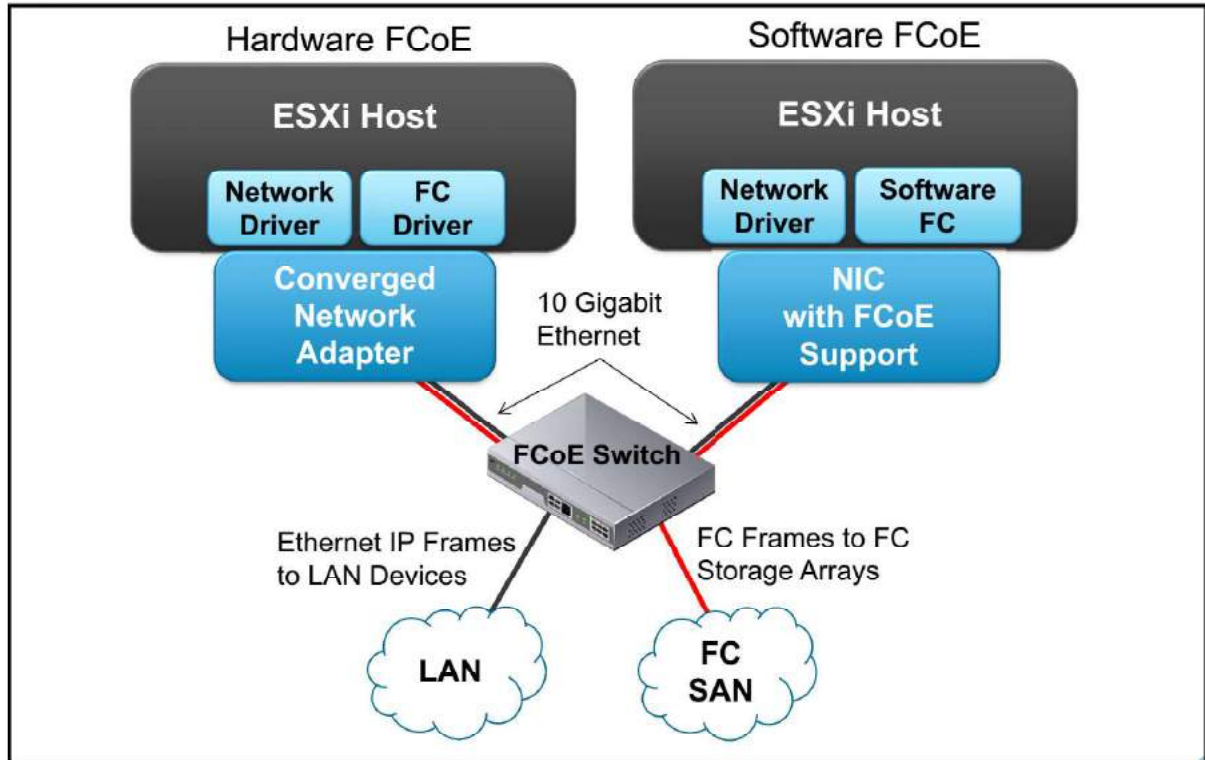
Active-active and active-passive disk arrays: Distinguishing between active-active and active-passive disk arrays can be useful:

- An active-active disk array allows access to the LUNs simultaneously through all the storage processors that are available without significant performance degradation. All the paths are active at all times (unless a path fails).

- In an active-passive disk array, one storage processor is actively servicing a given LUN. The other storage processor acts as a backup for the LUN and might be actively servicing other LUN I/O. I/O can be sent only to an active processor. If the primary storage processor fails, one of the secondary storage processors becomes active, either automatically or through administrative intervention.

FCoE Adapters

Slide 6-23



Configuring Software FCoE: Creating a VMkernel Port


Slide 6-24

Step 1: Connect the VMkernel to physical FCoE NICs that are installed on your host.

The VLAN ID and the priority class are discovered during FCoE initialization:

- The priority class is not configured in vSphere.

ESXi supports the maximum of four network adapter ports used for software FCoE.

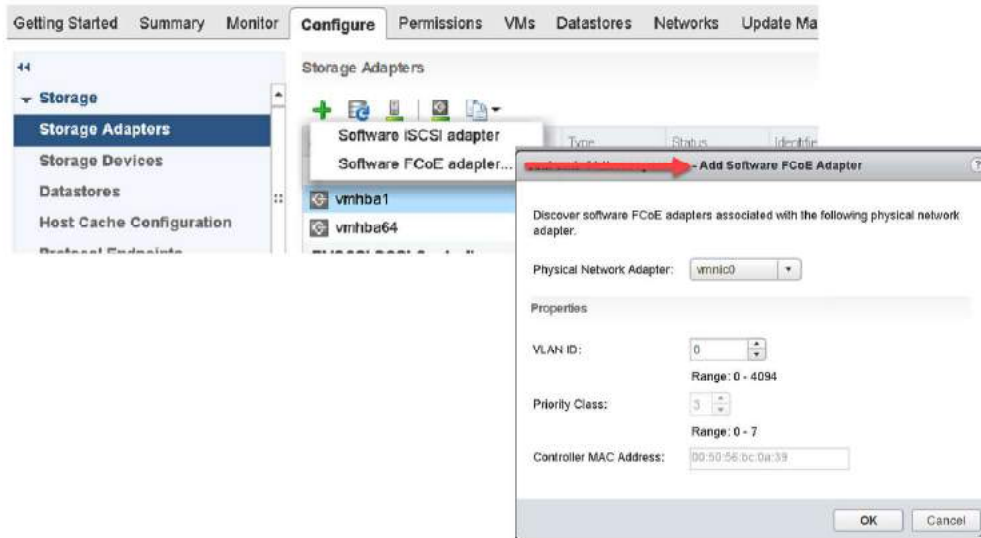


Configuring Software FCoE: Activating the Software FCoE Adapter

Slide 6-25


Step 2: Add the software FCoE adapter.

Select host > **Configuration** tab > **Storage Adapters** link > +



Multipathing with Software FCoE

Slide 6-26



Multipathing is a technique that lets you use more than one physical path that transfers data between the host and an external storage device.

To support path switching with Fibre Channel SAN, the ESXi host typically has two HBAs available from which the storage array can be reached through one or more switches. Alternatively, the setup can include one HBA and two storage processors so that the HBA can use a different path to reach the disk array.

In Fibre Channel multipathing, multiple paths connect each host with the storage device. For example, if HBA1 or the link between HBA1 and the switch fails, HBA2 takes over and provides the connection between the server and the switch. The process of one HBA taking over for another HBA is called HBA failover.

To configure a software FCoE adapter, you must have a dedicated VMkernel adapter. Software FCoE passes configuration information through the Data Center Bridging Exchange (DCBX) protocol by using the Cisco Discovery Protocol (CDP) VMkernel module.

Reviewing Learner Objectives


Slide 6-27

By the end of this lesson, you should be able to meet the following objectives:

- Describe uses of Fibre Channel with ESXi
- Describe Fibre Channel components and addressing
- Explain how multipathing with Fibre Channel work

Lesson 3: iSCSI Storage

Slide 6-28



Learner Objectives

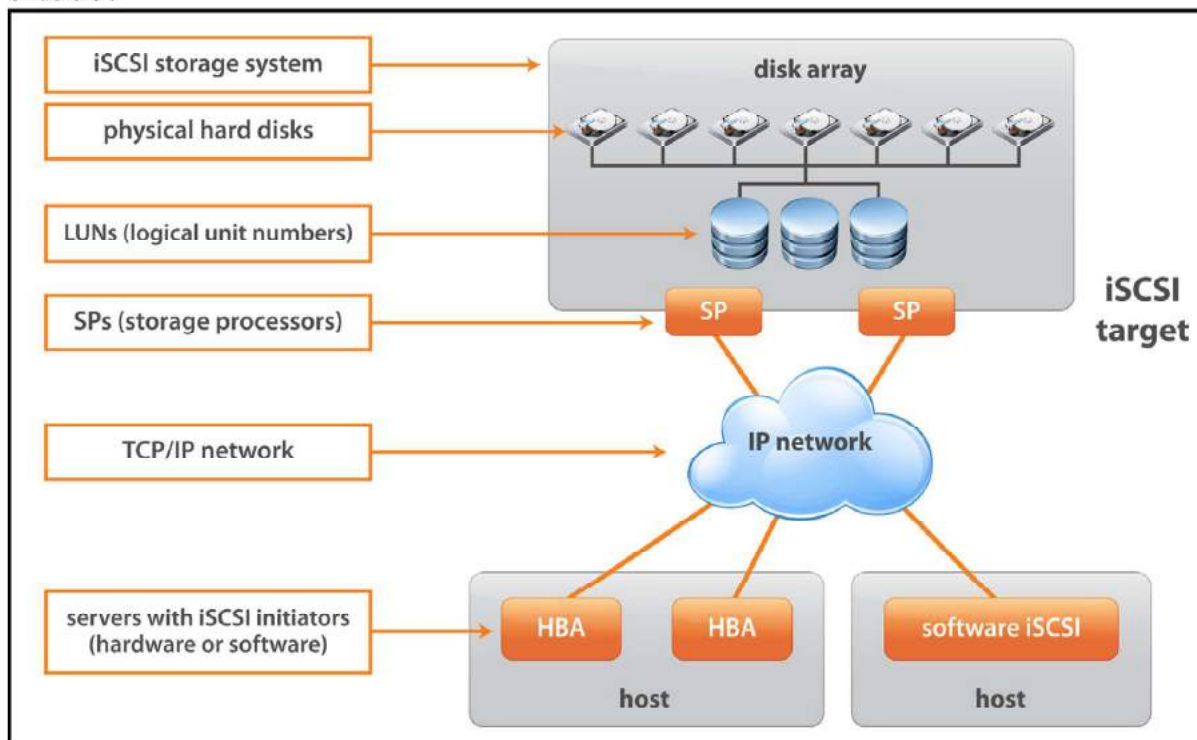
Slide 6-29

By the end of this lesson, you should be able to meet the following objectives:

- Describe uses of IP storage with ESXi
- Describe iSCSI components and addressing
- Configure iSCSI initiators
- Identify storage device naming conventions

iSCSI Components

Slide 6-30



An iSCSI SAN consists of an iSCSI storage system, which contains one or more LUNs and one or more storage processors. Communication between the host and the storage array occurs over a TCP/IP network.

The ESXi host is configured with an iSCSI initiator. An initiator can be hardware-based, in which case the initiator is an iSCSI host bus adapter (HBA). Or the initiator can be software-based, known as the iSCSI software initiator.

An initiator transmits SCSI commands over the IP network. A target receives SCSI commands from the IP network. You can have multiple initiators and targets in your iSCSI network. iSCSI is SAN-oriented because of the following reasons:

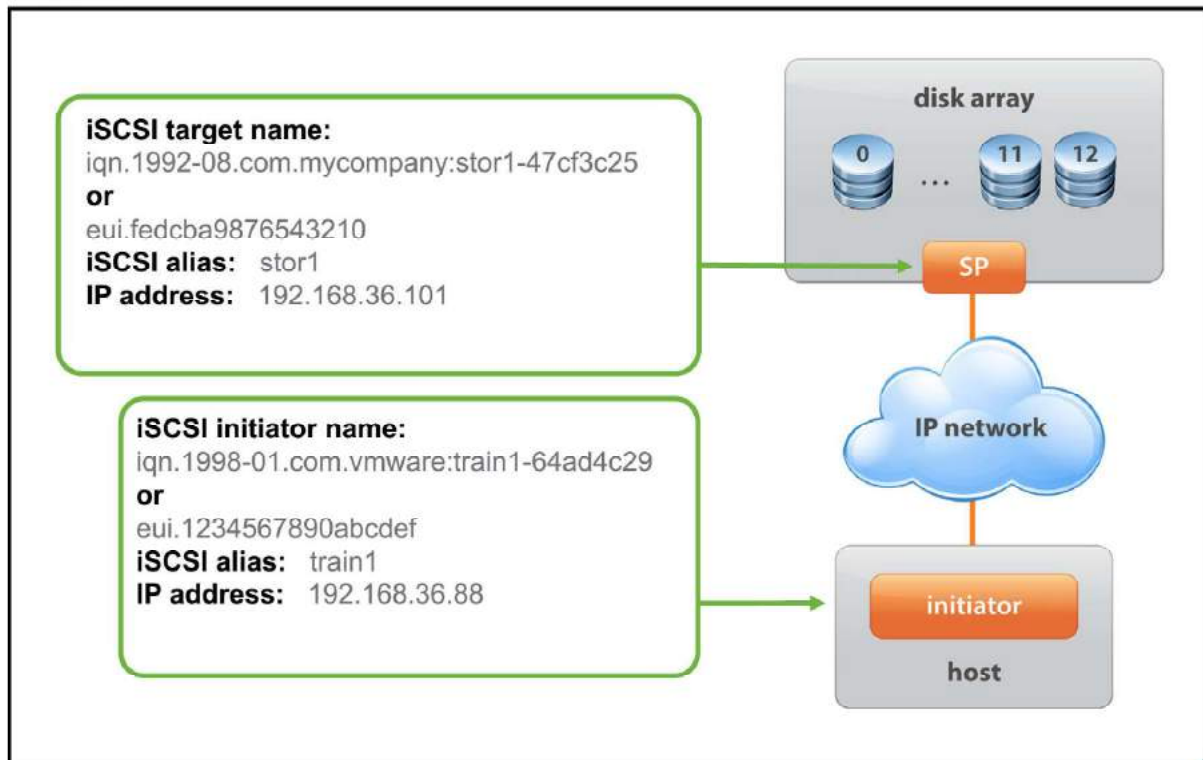
- The initiator finds one or more targets
- A target presents LUNs to the initiator
- The initiator sends SCSI commands to a target

An initiator resides in the ESXi host. Targets reside in the storage arrays that are supported by the ESXi host.

iSCSI arrays can use various mechanisms, including IP address, subnets, and authentication requirements, to restrict access to targets from hosts.

iSCSI Addressing

Slide 6-31



The main addressable, discoverable entity is an iSCSI node. An iSCSI node can be an initiator or a target. An iSCSI node requires a name so that storage can be managed regardless of address.

The iSCSI name can use one of the following formats: the iSCSI qualified name (IQN) or the extended unique identifier (EUI).

The IQN can be up to 255 characters long. The following naming convention is used:

- The prefix `iqn`
- A date code specifying the year and month in which the organization registered the domain or subdomain name used as the naming authority string
- The organizational naming authority string, which consists of a valid, reversed domain or subdomain name
- (Optional) A colon (:), followed by a string of the assigning organization's choosing, which must make each assigned iSCSI name unique

The following EUI naming convention is used:

- The prefix `eui`, followed by a 16-character name. The name includes 24 bits for a company name that is assigned by the IEEE and 40 bits for a unique ID, such as a serial number.

Storage Device Naming Conventions

Slide 6-32

Storage devices are identified in several ways:

- **Runtime name:** Uses the convention `vmhbaN:C:T:L`. This name is not persistent through reboots.
- **Target:** Identifies iSCSI target address and port.
- **LUN:** A unique identifier designated to individual or collections of hard disk devices. A logical unit is addressed by the SCSI protocol or SAN protocols that encapsulate SCSI, such as iSCSI or Fibre Channel.

Adapter Details

Runtime Name	Target	LUN	Status
vmhba33:C0:T1:L1	iqn.2008-08.com.starwindsoftware:shared:172.20.10.10:3260	1	Active (I/O)
vmhba33:C0:T1:L2	iqn.2008-08.com.starwindsoftware:shared:172.20.10.10:3260	2	Active (I/O)
vmhba33:C0:T1:L3	iqn.2008-08.com.starwindsoftware:shared:172.20.10.10:3260	3	Active (I/O)
vmhba33:C0:T1:L4	iqn.2008-08.com.starwindsoftware:shared:172.20.10.10:3260	4	Active (I/O)
vmhba33:C0:T1:L5	iqn.2008-08.com.starwindsoftware:shared:172.20.10.10:3260	5	Active (I/O)
vmhba33:C0:T1:L6	iqn.2008-08.com.starwindsoftware:shared:172.20.10.10:3260	6	Active (I/O)

On ESXi hosts, SCSI storage devices use various identifiers. Each identifier serves a specific purpose. For example, the VMkernel requires an identifier, generated by the storage device, which is guaranteed to be unique to each LUN. If a unique identifier cannot be provided by the storage device, the VMkernel must generate a unique identifier to represent each LUN or disk.

The disk identifiers referenced in the slide are not user friendly, so a third, more user-friendly naming convention is created after each reboot to reference each disk. This name can be used when you are using command-line utilities to interact with storage that is recognized by an ESXi host.

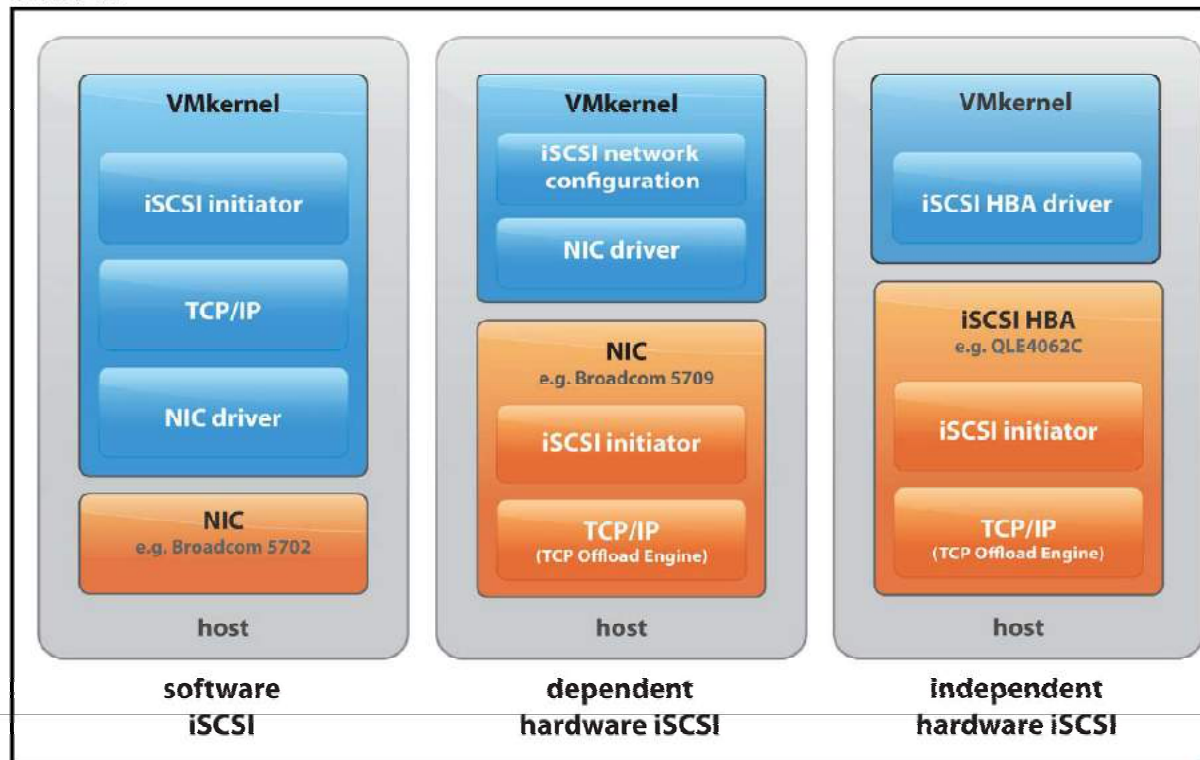
The following SCSI storage device identifiers are available:

- **Runtime name:** The name of the first path to the device. The runtime name is created by the host. It is not a reliable identifier for the device, because it is not persistent. The runtime name might change if you add HBAs to the ESXi host.
- **iSCSI name:** A worldwide unique name for identifying the node. iSCSI uses the iSCSI qualified name (IQN) and extended unique identifier (EUI). IQN uses the format `iqn.yyyy-mm.naming-authority:unique name`.

Storage device names appear in various panels in vSphere Web Client and vSphere Client.

iSCSI Adapters

Slide 6-33



To access iSCSI targets, your host uses iSCSI initiators. The initiators transport SCSI requests and responses, encapsulated into the iSCSI protocol, between the host and the iSCSI target. Your host supports two types of initiators: software iSCSI and hardware iSCSI.

A software iSCSI initiator is VMware code built in to the VMkernel. The initiator enables your host to connect to the iSCSI storage device through standard network adapters. The software iSCSI initiator handles iSCSI processing while communicating with the network adapter. With the software iSCSI initiator, you can use iSCSI technology without purchasing specialized hardware.

A hardware iSCSI initiator is a specialized third-party adapter capable of accessing iSCSI storage over TCP/IP. Hardware iSCSI initiators are divided into two categories: dependent hardware iSCSI and independent hardware iSCSI.

A dependent hardware iSCSI initiator, also known as an iSCSI host bus adapter, is a standard network adapter which includes the iSCSI offload function. To use this type of adapter, you must configure networking for the iSCSI traffic and bind the adapter to an appropriate VMkernel iSCSI port.

An independent hardware iSCSI adapter handles all iSCSI and network processing and management for your ESXi host.

Setting Up iSCSI Adapters

Slide 6-34

You set up software or hardware adapters before an ESXi host can work with a SAN.

Supported iSCSI adapter types (vmhba):

- Software adapter
- Hardware adapter:
 - Independent hardware adapter
 - Dependent hardware adapter

The iSCSI software adapter uses standard NICs to connect the ESXi host to a remote iSCSI target on the IP network. In this case, VMkernel networking configuration is required.

The third-party independent iSCSI hardware adapter offloads the iSCSI and network processing and management from the ESXi host. In this case, VMkernel networking configuration is not required.

The third-party dependent iSCSI hardware adapter depends on networking and iSCSI configuration management interfaces. In this case, VMkernel networking configuration is required.

For configuration information, see *vSphere Storage* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

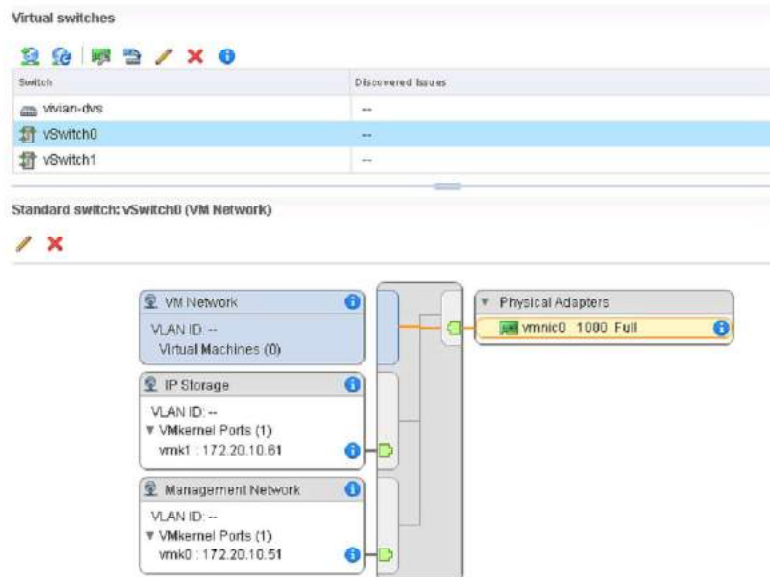
ESXi Network Configuration for IP Storage

Slide 6-35

A VMkernel port must be created for ESXi to access software iSCSI. The same port can be used to access NAS/NFS storage.

To optimize your vSphere networking setup, separate iSCSI networks from NAS/NFS networks:

- Physical separation is preferred.
- If physical separation is not possible, use VLANs.



Networking configuration for software iSCSI involves creating a VMkernel port on a virtual switch to handle your iSCSI traffic.

Depending on the number of physical adapters that you want to use for the iSCSI traffic, networking setup can be different:

- If you have one physical network adapter, you need a VMkernel port on a virtual switch.
- If you have two or more physical network adapters for iSCSI, you can use these adapters for host-based multipathing.

A best practice is to isolate your iSCSI network from other networks for performance and security reasons. Physically separate the networks. If physically separating the networks is impossible, logically separate the networks from one another on a single virtual switch by configuring a separate VLAN for each network.

For the configuration steps for adding a VMkernel port to a virtual switch, see *vSphere Storage* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Creating Datastores and Discovering iSCSI Targets

Slide 6-36

Based on the environment and storage needs, you can create VMFS, NFS, or virtual datastores as repositories for virtual machines.

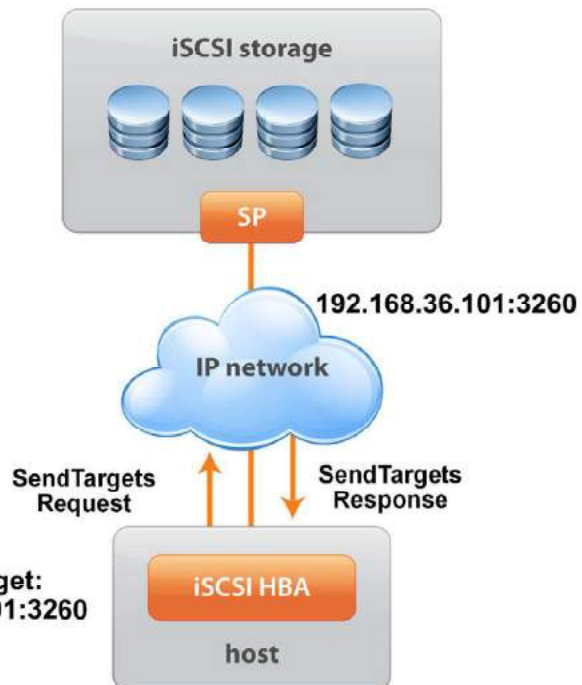
The iSCSI adapter discovers storage resources on the network and determines which resources are available for access.

An ESXi host supports the following discovery methods:

- Static
- Dynamic, also called SendTargets

The SendTargets response returns the IQN and all available IP addresses.

iSCSI Target:
192.168.36.101:3260



The ESXi host supports the following iSCSI target-discovery methods:

- **Static discovery:** The initiator does not have to perform discovery. The initiator knows in advance all the targets that it will contact and uses their IP addresses and domain names to communicate with them.
- **Dynamic discovery or SendTargets discovery:** Each time the initiator contacts a specified iSCSI server, it sends the SendTargets request to the server. The server responds by supplying a list of available targets to the initiator. The names and IP addresses of these targets appear as static targets in vSphere Client. You can remove a static target that was added by dynamic discovery. If you remove the target, the target might be returned to the list during the next rescan operation. The target might also be returned to the list if the HBA is reset or the host is rebooted.

iSCSI Security: CHAP

Slide 6-37

iSCSI initiators use CHAP for authentication purposes.

By default, CHAP is not configured.

ESXi supports two types of CHAP authentication:

- Unidirectional
- Bidirectional

ESXi also supports per-target CHAP authentication.

vmhba33 - Edit Authentication

The initiator uses these settings for authentication for all targets unless otherwise overridden by the specific target settings. Make sure that these parameters match on the storage side.

Authentication Method: None

Outgoing CHAP Credentials (target authenticates the initiator)

Name: Use initiator name

Secret:

Incoming CHAP Credentials (initiator authenticates the target)

Name: Use initiator name

Secret:

OK Cancel

You can implement CHAP to provide authentication between iSCSI initiators and targets.

ESXi supports the following CHAP authentication methods:

- Unidirectional or one-way CHAP: The target authenticates the initiator, but the initiator does not authenticate the target. You must specify the CHAP secret so that your initiators can access the target.
- Bidirectional or mutual CHAP: An additional level of security enables the initiator to authenticate the target. You must specify different target and initiator secrets.

CHAP uses a three-way handshake algorithm to verify the identity of your host and, if applicable, of the iSCSI target when the host and target establish a connection. The verification is based on a predefined private value, or CHAP secret, that the initiator and target share. ESXi implements CHAP as defined in RFC 1994.

ESXi supports CHAP authentication at the adapter level. In this case, all targets receive the same CHAP secret from the iSCSI initiator. For both software iSCSI and dependent hardware iSCSI initiators, ESXi also supports per-target CHAP authentication.

Before configuring CHAP, check whether CHAP is enabled at the iSCSI storage system and check the CHAP authentication method that the system supports. If CHAP is enabled, you must enable it

for your initiators, making sure that the CHAP authentication credentials match the credentials on the iSCSI storage. Although VMware recommends using CHAP in your iSCSI SAN implementation, consult with your storage vendor to ensure that best practices are followed.

You can protect your data in additional ways. For example, you may protect your iSCSI SAN by giving it a dedicated standard switch. You may also configure the iSCSI SAN on its own VLAN to improve performance and security. Some inline network devices may be implemented to provide encryption and further data protection.

For more security related information, see *vSphere Storage* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Multipathing with iSCSI Storage

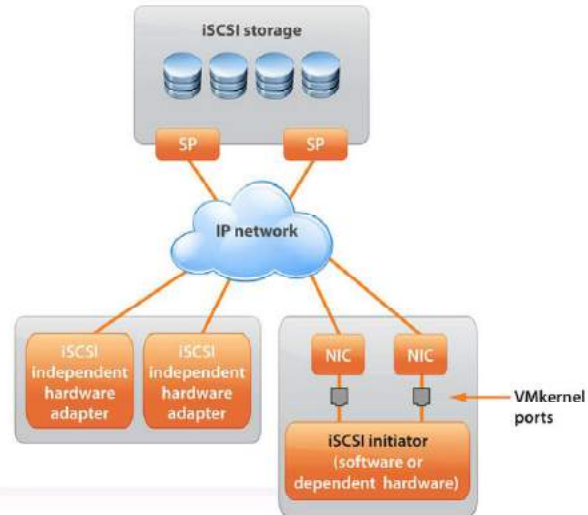
Slide 6-38

Software or dependent hardware iSCSI:

- Uses multiple NICs.
- Connects each NIC to a separate VMkernel port.
- Associates VMkernel ports with the iSCSI initiator.

Independent hardware iSCSI:

- Uses two or more hardware iSCSI adapters.



Adapter Details

Port Group	VMkernel Ad...	Port Group Policy	Path Status	Physical Network Adapter
IP storage (vSwitch0)	vmk1	Compliant	Not used	vmnic0 (1 Gbit/s, Full)

When setting up your ESXi host for multipathing and failover, you can use multiple hardware iSCSI adapters or multiple NICs. The choice depends on the type of iSCSI initiators on your host.

With software iSCSI and dependent hardware iSCSI, you can use multiple NICs that provide failover for iSCSI connections between your host and iSCSI storage systems. For this setup, because multipathing plug-ins do not have direct access to physical NICs on your host, you must first connect each physical NIC to a separate VMkernel port. Then you use a port-binding technique to associate all VMkernel ports with the iSCSI initiator. As a result, each VMkernel port connected to a separate NIC becomes a different path that the iSCSI storage stack and its multipathing plug-in can use. For dependent hardware iSCSI, ensure that the physical network card is correctly installed and appears on the host's **Manage** tab under **Networking**.

With independent hardware iSCSI, the host typically has two or more hardware iSCSI adapters available, from which the storage system can be reached using one or more switches. Alternatively, the setup might include one adapter and two storage processors so that the adapter can use a different path to reach the storage system.

After iSCSI multipathing is set up, each port on the ESXi system has its own IP address, but they all share the same iSCSI initiator IQN. When iSCSI multipathing is configured, the VMkernel routing table is not consulted when identifying the outbound NIC to use. Instead, iSCSI multipathing is

managed using vSphere multipathing modules. Due to the latency that can be incurred, VMware does not recommend routing iSCSI traffic. For more about configuring iSCSI multipathing, see *vSphere Storage* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Lab 8: Accessing iSCSI Storage

Slide 6-39

Configure access to an iSCSI datastore

1. Add a VMkernel Port Group to a Standard Switch
2. Configure the iSCSI Software Adapter and Connect It to the Storage

Review of Learner Objectives


Slide 6-40

You should be able to meet the following objectives:

- Describe uses of IP storage with ESXi
- Describe iSCSI components and addressing
- Configure iSCSI initiators
- Identify storage device naming conventions

Lesson 4: VMFS Datastores

Slide 6-41



Learner Objectives

Slide 6-42

By the end of this lesson, you should be able to meet the following objectives:

- Create a VMFS datastore
- Explain VMFS locking mechanisms
- Increase the size of a VMFS datastore
- Delete a VMFS datastore

Using VMFS Datastores with ESXi Hosts

Slide 6-43

Use VMFS datastores when accessing block-level storage:

- VMFS is optimized for storing and accessing large files.
- A VMFS datastore can have a maximum volume size of 64 TB.

Use RDMs if any of the following conditions are true of your virtual machine:

- It is taking storage array-level snapshots.
- It is clustered to a physical machine.
- It has large amounts of data that you do not want to convert into a virtual disk.

VMFS datastores primarily serve as repositories for the files of virtual machines. A VMFS datastore is optimized for storing and accessing large files like virtual disks and memory images of suspended virtual machines. The maximum size of a VMFS datastore is 64 TB.

You can use an NFS datastore to store your virtual machines, templates, and ISO images, but not all functions are supported. For example, you cannot store an RDM on an NFS datastore. A VMFS datastore is required for an RDM to store the RDM mapping file (*-rdm.vmdk).

As for RDMs, choose RDMs over VMFS datastores if:

- A virtual machine is using storage array-level snapshot applications
- A virtual machine is clustered with a physical machine
- You want to keep the virtual machine's data on a raw disk instead of converting it to a virtual disk because, for example, the data disk is very large

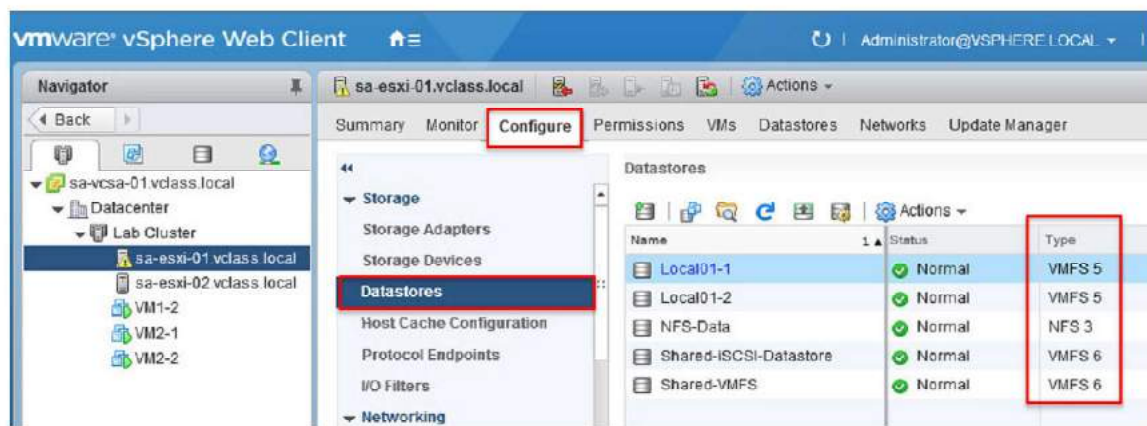
Otherwise, use a VMFS datastore to store the files of your virtual machines to use features like template deployment and for portability.

Creating and Viewing VMFS Datastores

Slide 6-44

VMFS datastores serve as repositories for virtual machines.

Using the New Datastore wizard, you can create VMFS datastores on any SCSI-based storage devices that the host discovers, including Fibre Channel, iSCSI, and local storage devices.



By viewing datastores, you can determine the type of datastores in use. For example, you can determine whether your datastore is an NFS datastore, a VMFS datastore, or any other type of datastores.

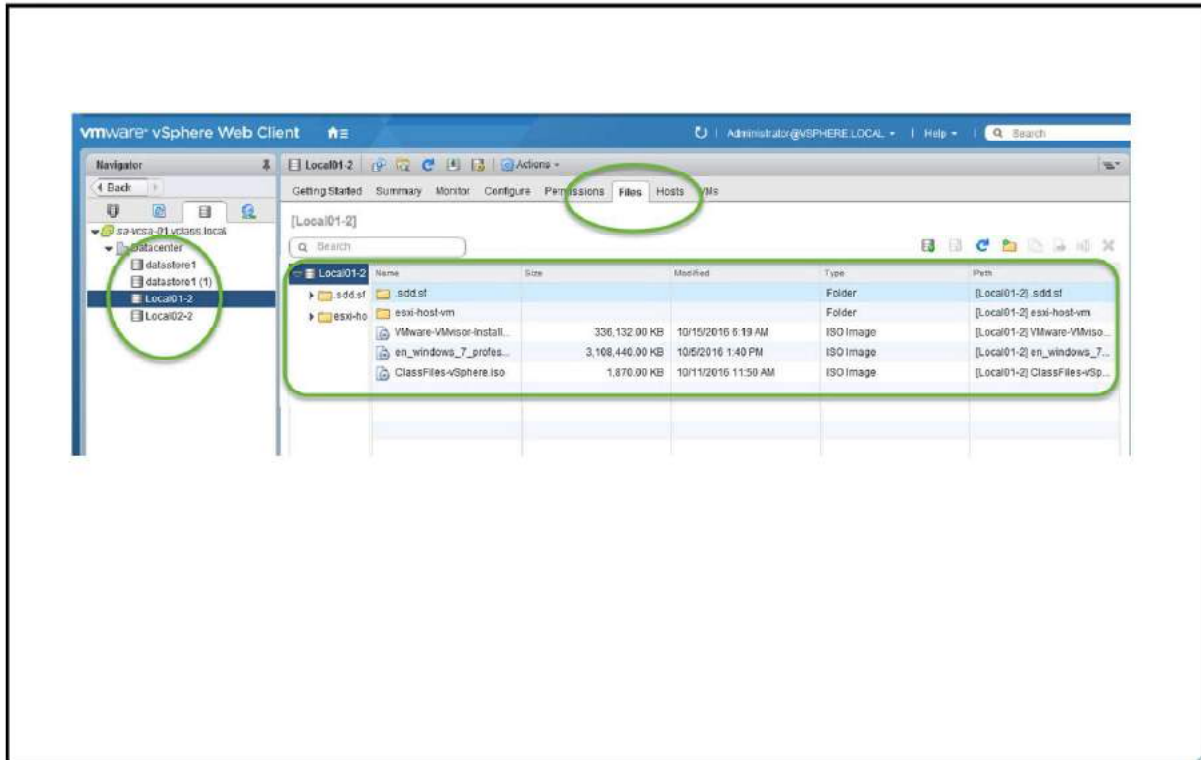
ATS+SCSI mechanism: A VMFS datastore that supports the ATS+SCSI mechanism is configured to use ATS and attempts to use it when possible. If ATS fails, the VMFS datastore reverts to SCSI reservations. In contrast with the ATS locking, the SCSI reservations lock an entire storage device while an operation that requires metadata protection is performed. After the operation completes, VMFS releases the reservation and other operations can continue.

Datastores that use the ATS+SCSI mechanism include VMFS5 datastores that were upgraded from VMFS3. In addition, new VMFS5 or VMFS6 datastores on storage devices that do not support ATS use the ATS+SCSI mechanism.

If your VMFS datastore reverts to SCSI reservations, you might notice performance degradation caused by excessive SCSI reservations. For information about how to reduce SCSI reservations, see *vSphere Troubleshooting* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Browsing Datastore Contents

Slide 6-45



The Datastores pane shows the datastores that are available in many views. The tab lists all datastores currently configured for the ESXi host.

The example shows the contents of the VMFS datastore named datastore1. The contents are the virtual machine folders. The files for each virtual machine are in the respective virtual machine folder.

Managing Overcommitted Datastores

Slide 6-46

A datastore becomes overcommitted when the total provisioned space of thin-provisioned disks is greater than the size of the datastore.

Actively monitor your datastore capacity:

- Alarms assist through notifications:
 - Datastore disk overallocation
 - Virtual machine disk usage
- Use reporting to view space usage.

Actively manage your datastore capacity:

- Increase the datastore capacity when necessary.
- Use vSphere Storage vMotion to mitigate space usage problems on a particular datastore.

Using thin-provisioned virtual disks for your virtual machines is a way to make the most of your datastore capacity. But if your datastore is not sized properly, it can become overcommitted. A datastore becomes overcommitted when its thin-provisioned virtual disk's full capacity is greater than the datastore's capacity. If disk capacity is managed correctly, then you, the vSphere administrator, should not run out of space (despite external issues like budgeting).

When a datastore reaches capacity, vSphere Web Client prompts you to provide more space on the underlying VMFS datastore and freezes the virtual machine until you do so. Monitor your datastore capacity by using alarms to alert you of how much a datastore's disks are overallocated or how much disk space a virtual machine is using. You can also use the storage reports to view disk space usage, for example, the Show all Datastores report in the **Storage** tab.

Manage your datastore capacity. Dynamically increase the size of your datastore when necessary. You can also use vSphere Storage vMotion to mitigate space usage issues. For example, with vSphere Storage vMotion, you can migrate a virtual machine off a datastore. The migration can be done with a change from virtual disks of thick format to thin format at the target datastore.

Increasing the Size of a VMFS Datastore

Slide 6-47

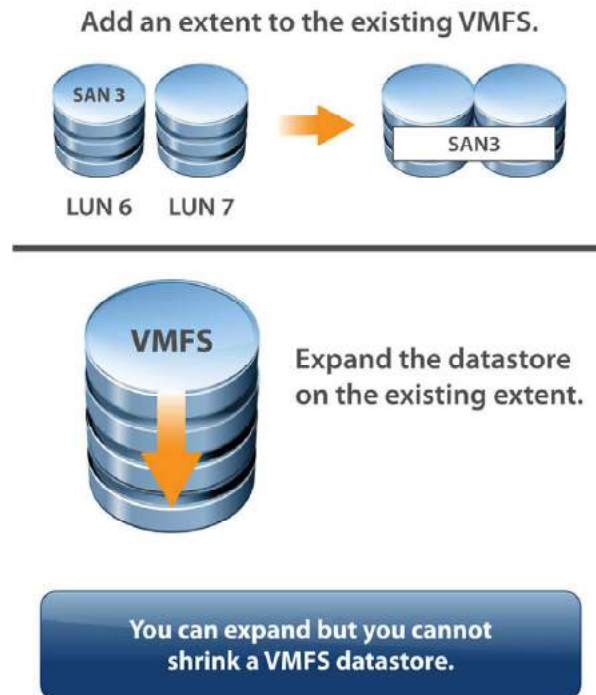
In general, before making any changes to your storage allocation:

- Perform a rescan to ensure that all hosts see the most current storage.
- Record the unique identifier.

Increase a VMFS datastore's size to give it more space or possibly to improve performance.

Ways to dynamically increase the size of a VMFS datastore:

- Add an extent (LUN).
- Expand the datastore within its extent.



In general, before you make storage allocation changes, a good practice is to rescan to ensure that all hosts see the current storage view. Also, record the unique identifier of the volume that you want to expand, for example, the NAA ID. You need this information to identify the VMFS datastore whose size you want to increase.

You can dynamically increase the capacity of a VMFS datastore if the datastore has insufficient disk space. Insufficient disk space is realized when you create a virtual machine or you try to add more disk space to a virtual machine.

Use one of the following methods:

- Add an extent to the VMFS datastore: An extent is a partition on a LUN. You can add an extent to any VMFS datastore. The datastore can stretch over multiple extents, up to 32.
- Expand the VMFS datastore: Increase the size of the VMFS datastore in its extent. Only extents with free space immediately after them are expandable. As a result, rather than adding the new extent, you can expand the existing extent so that it fills the available adjacent capacity.

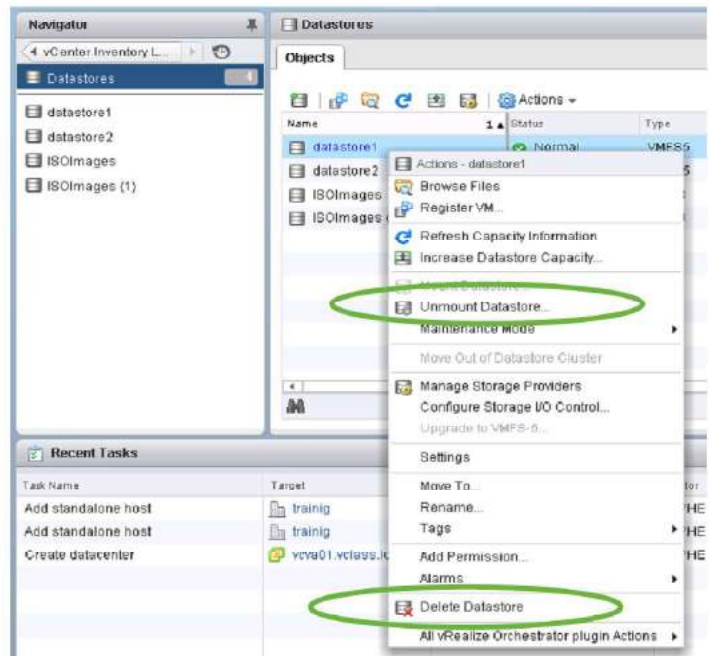
For more information about growing or expanding a VMFS volume or datastore, see VMware knowledge base article 1017662 at <http://kb.vmware.com/kb/1017662>.

Deleting or Unmounting a VMFS Datastore

Slide 6-48

An unmounted datastore remains intact, but cannot be seen from the hosts that you specify. The datastore continues to appear on other hosts, where it remains mounted.

A deleted datastore is destroyed and disappears from all hosts that have access to it. All virtual machine files on the datastore are permanently removed.



When you unmount a datastore, it remains intact, but can no longer be seen from the hosts that you specify. Unmounting a VMFS datastore preserves the files on the datastore but makes the datastore inaccessible to the ESXi host. The datastore is still visible and continues to serve other hosts, where it remains mounted. Do not perform any configuration operations that might result in I/O to the datastore while the unmount is in progress.

You can delete any type of VMFS datastore, including copies that you have mounted without resignaturing. When you delete a datastore, it is permanently destroyed and disappears from all hosts that have access to the datastore. This operation permanently removes all files that are associated with virtual machines on the datastore. Although you can delete the datastore without unmounting, it is preferable that you unmount the datastore first. Deleting a VMFS datastore destroys the pointers to the files on the datastore. So the files disappear from all hosts that have access to the datastore.

Before you delete or unmount a VMFS datastore, power off all virtual machines whose disks reside on the datastore. If you do not power off the virtual machines and you try to continue, an error message tells you that the resource is busy. Before you unmount a VMFS datastore, use vSphere Web Client to verify the following:

- No virtual machines reside on the datastore.

- The datastore is not part of a datastore cluster.
- The datastore is not managed by vSphere Storage DRS.
- VMware vSphere® Storage I/O Control is disabled.
- The datastore is not used for vSphere HA heartbeat.

To keep your data, back up the contents of your VMFS datastore before you delete the datastore.


Multipathing Algorithms

Slide 6-49

Arrays provide various features. Some offer active-active storage processors. Others offer active-passive storage processors.

vSphere offers native path selection, load-balancing, and failover mechanisms.

Third-party vendors can create their own software to be installed on ESXi hosts. The third-party software enables hosts to properly interact with the storage arrays.



The Pluggable Storage Architecture is a VMkernel layer responsible for managing multiple storage paths and providing load balancing. An ESXi host can be attached to storage arrays with the following storage processor configuration:

- Active-active
- Active-passive

VMware offers native load-balancing and failover mechanisms. Examples of VMware path selection policies include:

- Round Robin
- Most Recently Used (MRU)
- Fixed

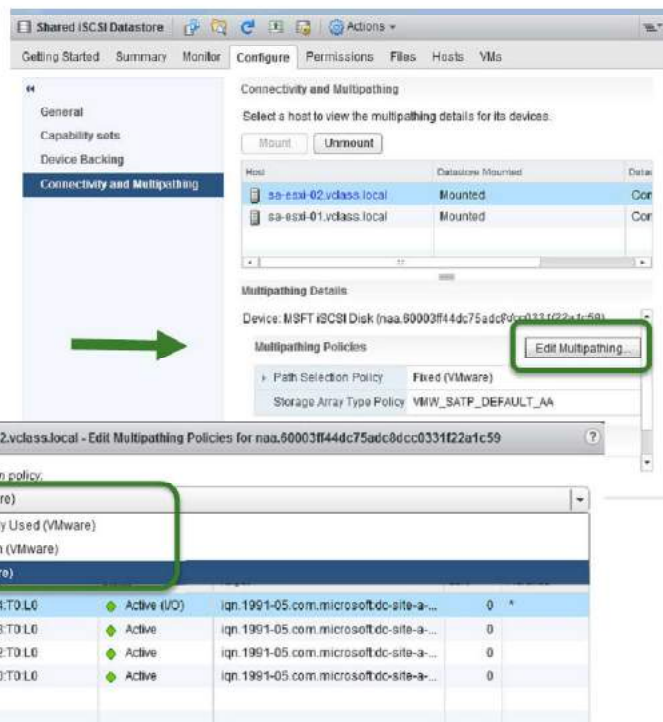
Third-party vendors are able to design their own load-balancing techniques and failover mechanisms for particular storage array types to add support for new arrays. Third-party vendors do not need to provide internal information or intellectual property about the array to VMware.

Configuring Storage Load Balancing

Slide 6-50

Path selection policies exist for:

- Scalability:
 - Round Robin: A multipathing policy that performs load balancing across paths
- Availability:
 - MRU
 - Fixed



Multiple paths can exist to a datastore from an ESXi host. You can view these paths in the datastore properties by clicking the **Manage** tab and the **Settings** tab.

The following path selection policies are supported for multipathing with Fibre Channel or iSCSI:

- Fixed path selection policy: The host always uses the preferred path to the disk when that path is available. If the host cannot access the disk through the preferred path, it tries the alternative paths. This is the default policy for active-active storage devices.
- Most Recent Used (MRU) path selection policy: The host uses the most recent path to the disk until this path becomes unavailable. That is, the host does not revert back until this path becomes unavailable. A failover to a new path is performed. If the original path becomes available again, the host does not fail back to the original path. MRU is the default policy for active-passive storage devices and is required for those devices.
- Round Robin path selection policy: The host uses a path selection algorithm that rotates through all available paths. In addition to path failover, the Round Robin policy supports load balancing across the paths. Before using this policy, check with storage vendors to find out whether a Round Robin configuration is supported on their storage.

Lab 9: Managing VMFS Datastores

Slide 6-51

Create and manage VMFS datastores

1. Rename a VMFS Datastore
2. Create VMFS Datastores for the ESXi Host
3. Expand a VMFS Datastore to Consume Unused Space on a LUN
4. Remove a VMFS Datastore
5. Extend a VMFS Datastore
6. Create a Second Shared VMFS Datastore Using iSCSI

Review of Learner Objectives

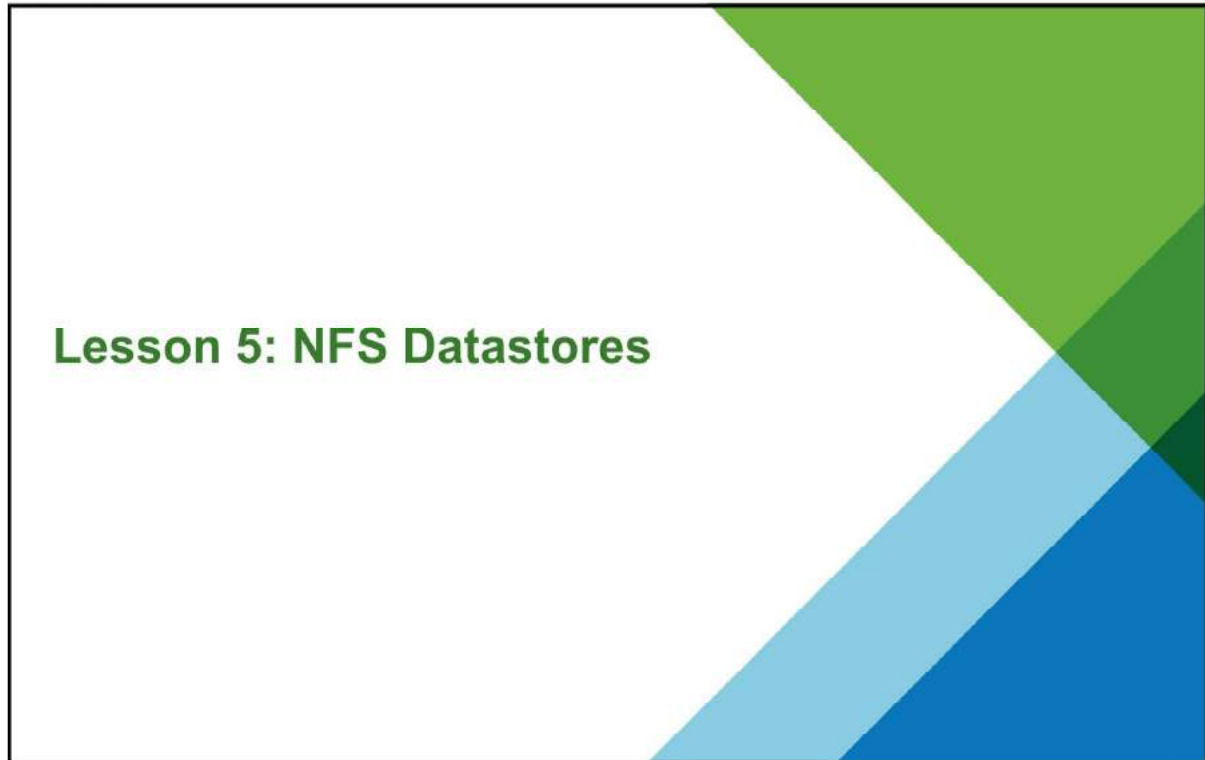
Slide 6-52

You should be able to meet the following objectives:

- Create a VMFS datastore
- Explain VMFS locking mechanisms
- Increase the size of a VMFS datastore
- Delete a VMFS datastore

Lesson 5: NFS Datastores

Slide 6-53



Learner Objectives

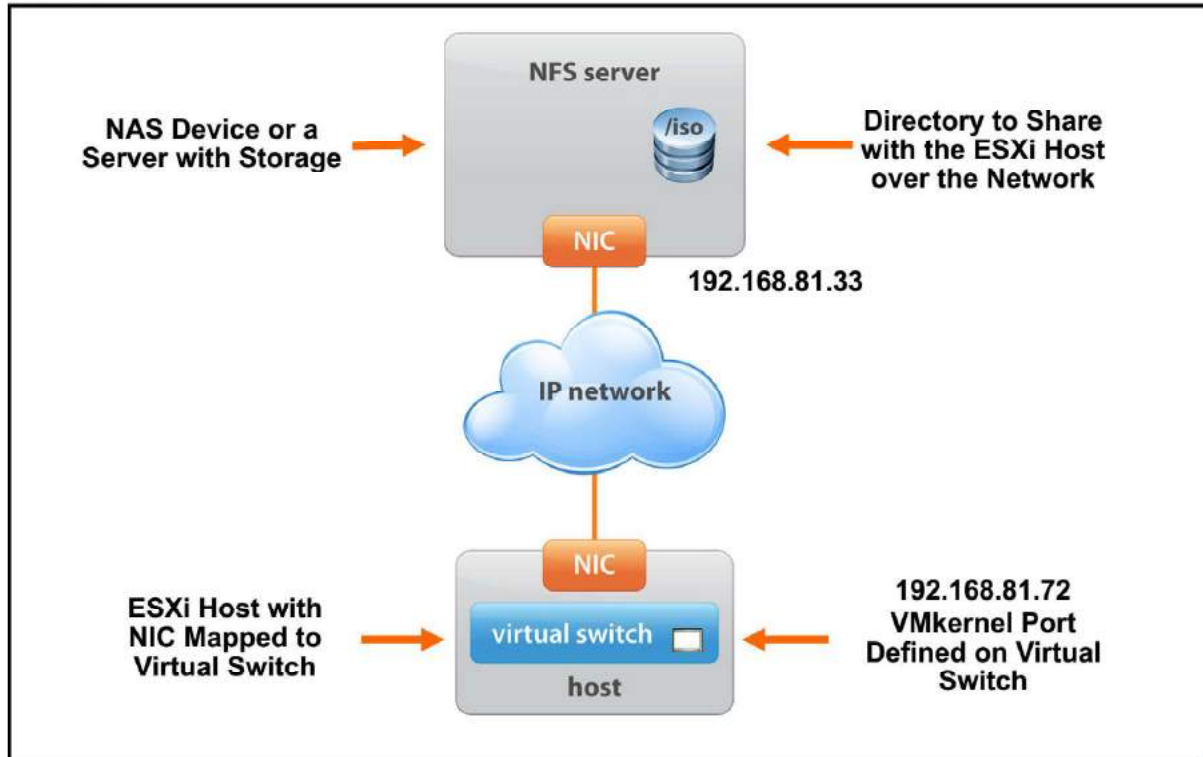
Slide 6-54

By the end of this lesson, you should be able to meet the following objectives:

- Describe NFS components
- Describe the differences between NFS v3 and NFS v4.1
- Configure and manage NFS datastores

NFS Components

Slide 6-55



An NFS file system is located on a NAS device that is called the NFS server. The NFS server contains one or more directories that are shared with the ESXi host over a TCP/IP network. An ESXi host accesses the NFS server through a VMkernel port that is defined on a virtual switch.

Configuring an NFS Datastore

Slide 6-56

Create a VMkernel port:

- For better performance and security, separate your NFS network from the iSCSI network.

Provide the following information:

- NFS version: v3 or v4.1
- Datastore name
- NFS server names or IP addresses
- Folder on the NFS server, for example, `/templates` and `/nfs_share`
- Select hosts that will mount the datastore
- Whether to mount the NFS file system read-only
- Authentication parameters

For each ESXi host to access an NFS datastore over the network, a VMkernel port must be configured on a virtual switch. The name of this port can be anything that you want. For performance and security reasons, a good practice is to isolate your NFS networks from the other networks, such as your iSCSI network and your virtual machine networks.

NFS v3 and NFS v4.1

Slide 6-57

NFS v3:

- ESXi managed multipathing
- AUTH_SYS (root) authentication
- VMware proprietary file locking
- Client-side error tracking

NFS v4.1:

- Native multipathing and session trunking
- Optional Kerberos authentication
- Built-in file locking
- Server-side error tracking



An NFS datastore can be created as either NFS v3 or NFS v4.1. Various compatibility issues between the two NFS versions preclude accessing datastores using both protocols at the same time from different hosts. If a datastore is configured as NFS v4.1, all hosts that access that datastore must mount the share as NFS 4.1. Data corruption can occur if hosts access a datastore with the wrong NFS version.

As of the release of vSphere 6, NFS v4.1 is not compatible with VMware vSphere® Storage DRS™, Storage I/O Control, VMware Site Recovery Manager™, and vSphere Virtual Volumes, because of server protocol locking.

NFS Version Compatibility with Other vSphere Technologies

Slide 6-58

	NFS v3	NFS v4.1
vSphere vMotion and vSphere Storage vMotion	Yes	Yes
vSphere HA	Yes	Yes
vSphere Fault Tolerance	Yes	Yes
vSphere DRS and vSphere DPM	Yes	Yes
Stateless ESXi and Host Profiles	Yes	Yes
vSphere Storage DRS and Storage I/O Control	Yes	No
Site Recovery Manager	Yes	No
vSphere Virtual Volumes	Yes	No

vSphere 6 supports NFS v4.1 to overcome many limitations when using NFS v3. In vSphere 6, both NFS v3 and NFS v4.1 shares can be used, although some important constraints must be taken into consideration when designing a vSphere environment in which both versions will be used.

NFS v4.1 provides the following enhancements:

- Native multipathing and session trunking: NFS v4.1 provides multipathing for servers that support session trunking. When trunking is available, you can use multiple IP addresses to access a single NFS volume. Client ID trunking is not supported.
- Kerberos authentication: NFS v4.1 introduces Kerberos authentication in addition to the traditional AUTH_SYS method used by NFS v3.
- Improved built-in file locking
- Enhanced error recovery using server-side tracking of open files and delegations
- Many general efficiency improvements including session leases and less protocol overhead

The NFS v4.1 client offers the following new features:

- Stateful locks with share reservation using a mandatory locking semantic

- Protocol integration, side-band (auxiliary) protocol no longer required to lock and mount
- Trunking (true NFS multipathing), where multiple paths (sessions) to the NAS array can be created and load distributed across those sessions.
- Enhanced error recovery to mitigate server crash and loss of connectivity

NFS Datastore Best Practices

Slide 6-59

Best practices:

- Configure an NFS array to allow only one NFS protocol.
- Use either NFS v3 or NFS v4.1 to mount the same NFS share across all ESXi hosts.
- Exercise caution when mounting an NFS share. Mounting an NFS share as NFS v3 on one ESXi host and as NFS v4.1 on another host can lead to data corruption.



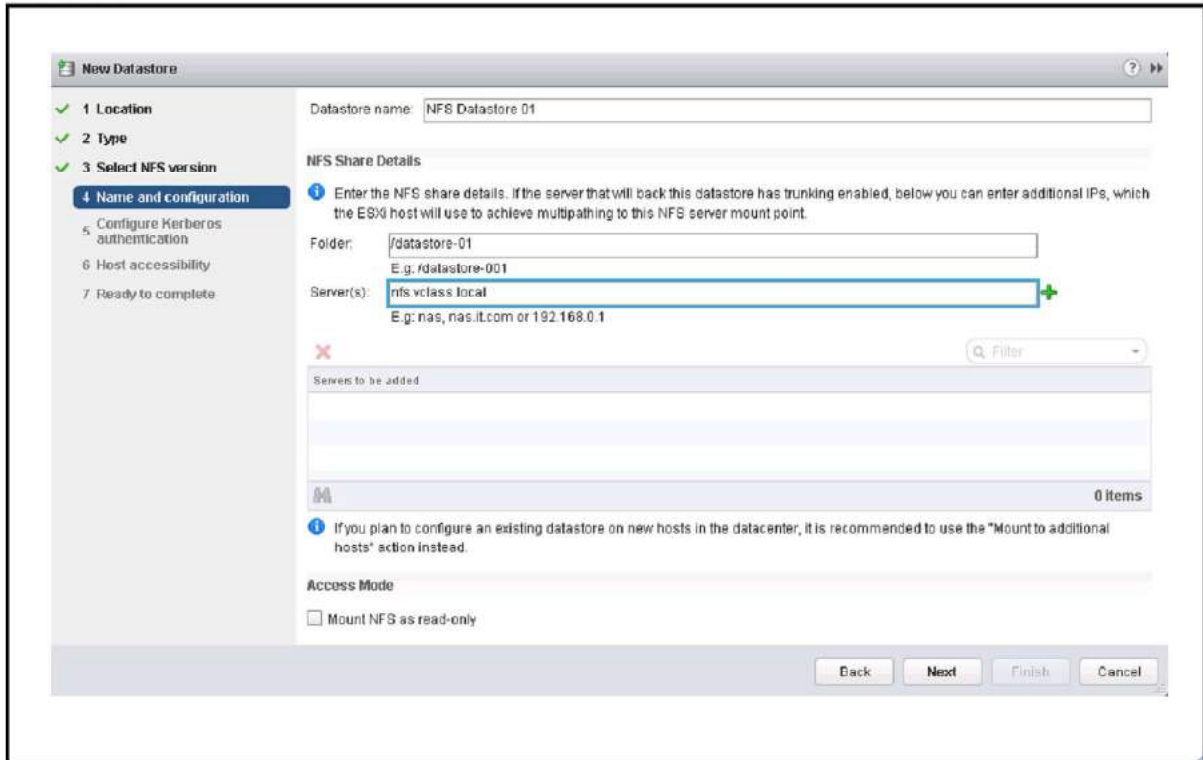
NFS v3 locking is not compatible with NFS v4.1:

- NFS v3 uses proprietary client-side cooperative locking. NFS v4.1 uses server-side locking.

An NFS datastore can be created as either NFS v3 or NFS v4.1. Various compatibility issues between the two NFS versions preclude accessing datastores using both protocols at the same time from different hosts. If a datastore is configured as NFS v4.1, all hosts that access that datastore must mount the share as NFS 4.1. Data corruption can occur if hosts access a datastore with the wrong NFS version.

NFS Datastore Name and Configuration

Slide 6-60

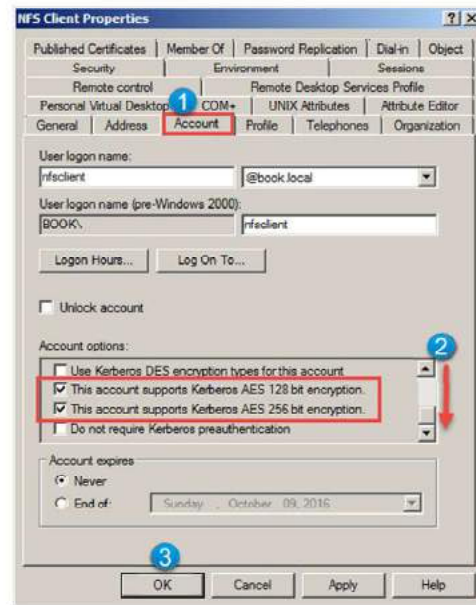


Configuring AD and NFS Servers to Use Kerberos

Slide 6-61

Before enabling Kerberos on ESXi hosts:

1. Create an account in AD for NFS v4.1 access.
 - Enable Kerberos AES encryption.
 - Set the account expire option to **Never**.
2. Configure NFS servers to use Kerberos.
3. Configure NFS server shares to grant full access to the AD account used.



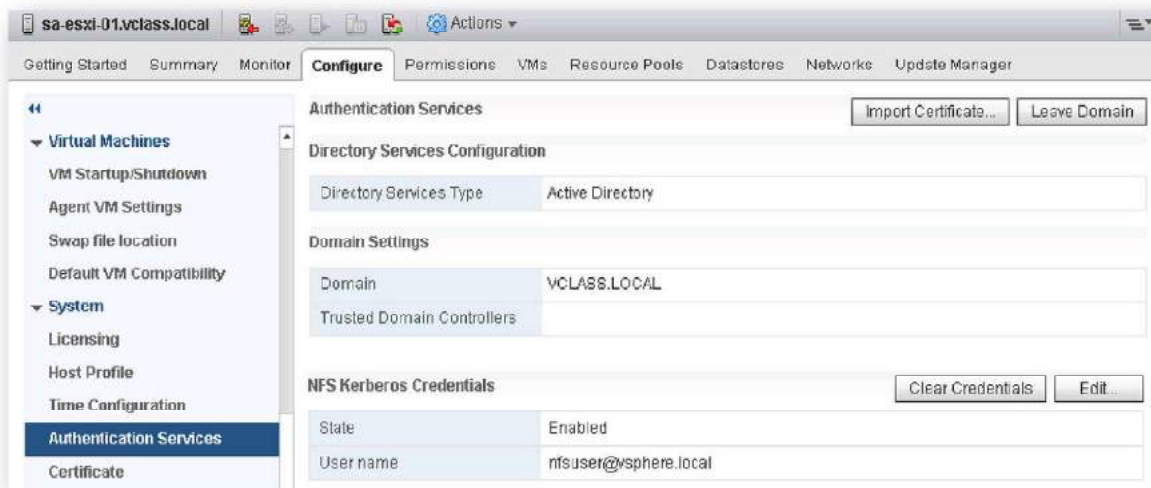
Each ESXi host must be added to the Active Directory domain.

A number of configuration steps must be taken to prepare each ESXi host to use Kerberos authentication. Kerberos authentication requires that all nodes involved: the Active Directory Server, the NFS servers, and the ESXi hosts, be synchronized so that little to no time drift exists. Kerberos authentication fails if any significant drift exists between the nodes. To prepare your ESXi host to use Kerberos authentication, configure the NTP client settings to reference a common NTP server (or the domain controller if applicable).

Configuring ESXi Host Authentication and NFS Kerberos Credentials

Slide 6-62

As a requirement of Kerberos authentication, you add each ESXi host to the AD domain. Then, you configure NFS Kerberos credentials.



Considerations When Using Kerberos for NFS

Slide 6-63

Be aware of the UID and GID on the files:

- For NFS v3, UID, and GID are root.
- Accessing files created with NFS v3 from the Kerberos client using NFS v4.1 results in permission-denied errors.

Use the same AD user on all ESXi hosts:

- vSphere vMotion and other features might fail if individual hosts use different user accounts.
- Use host profiles to avoid errors.

Time must be synchronized:

- Time synchronization is required for successful Kerberos authentication.
- Configure all components to synchronize to a common Network Time Protocol server.

Kerberos must be configured on the NFS servers and ESXi hosts before creating an NFS datastore to use Kerberos authentication.

Administrators must consider the following when planning on using NFS Kerberos:

- NFS v3 and v4.1 use different authentication credentials, resulting in incompatible UID/GID on files.
- Using different Active Directory users on different hosts that access the same NFS share can cause the vSphere vMotion migration to fail.
- NFS Kerberos configuration can be automated using host profiles to reduce configuration conflicts.

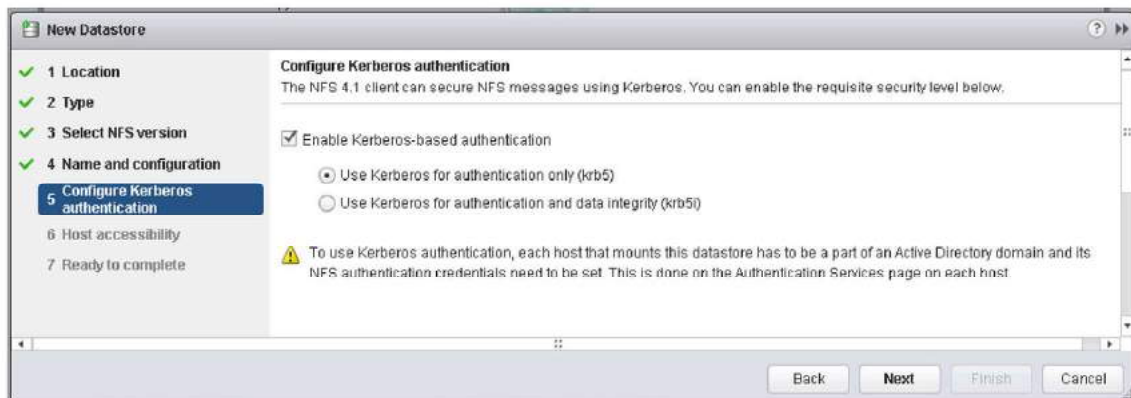
Time must be synchronized between all participating components.

Configuring a Datastore to Use Kerberos

Slide 6-64

Enable Kerberos authentication when creating each datastore.

- Kerberos5 authentication
- Kerberos5i authentication and data integrity

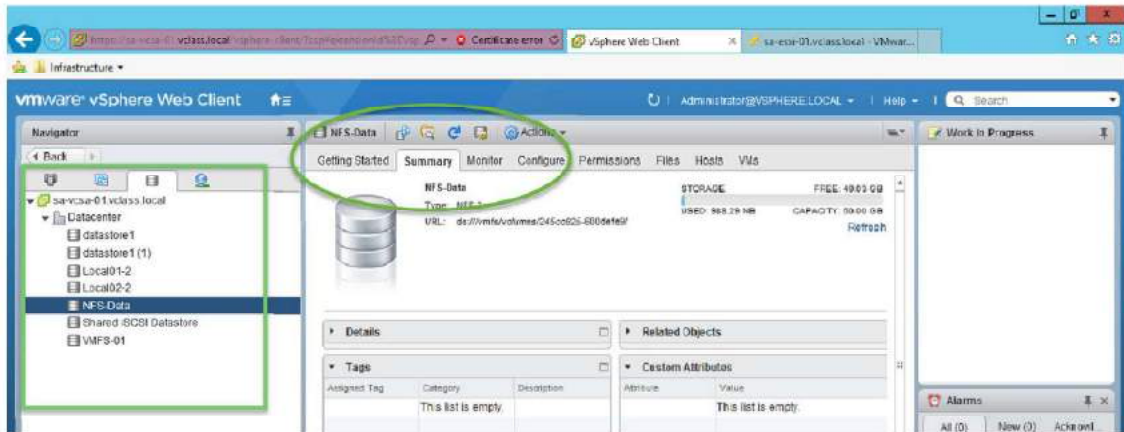


After the previously discussed configuration steps are implemented, you can configure the datastore to use Kerberos authentication.

Viewing IP Storage Information

Slide 6-65

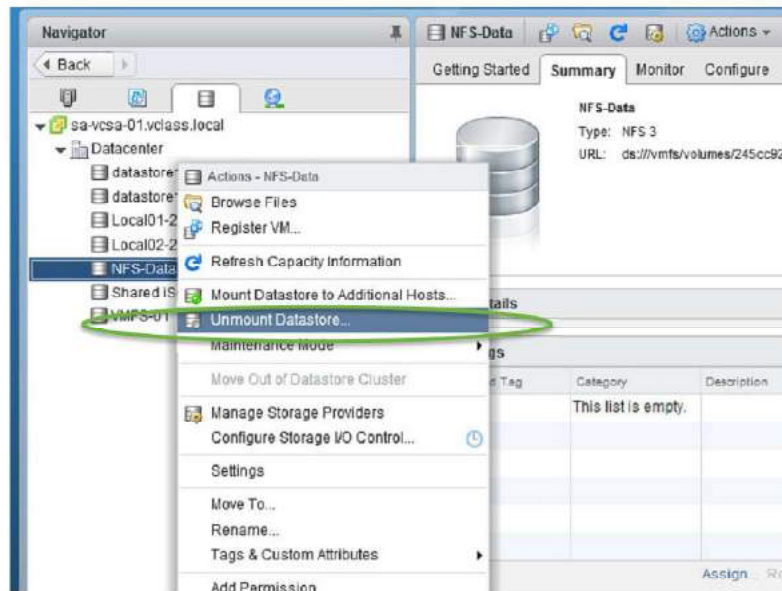
You can view the details of the VMFS or NFS datastores that you created.



Unmounting an NFS Datastore

Slide 6-66

Unmounting an NFS datastore causes the files on the datastore to become inaccessible to the ESXi host.



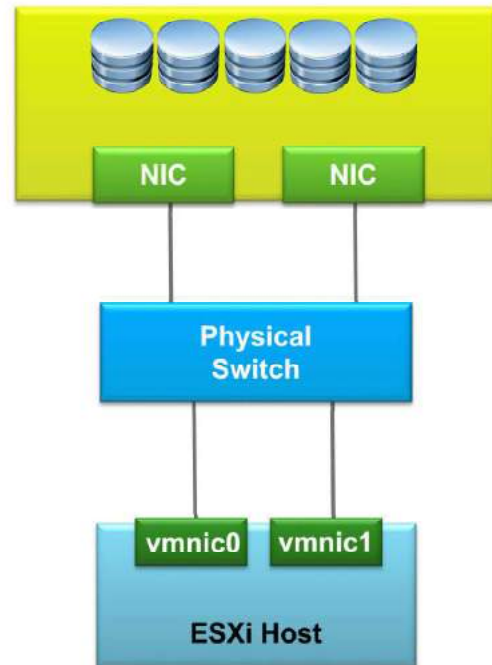
When an NFS datastore is unmounted, the files located on the NFS datastore become inaccessible to the ESXi host. Before unmounting an NFS datastore, you must stop all virtual machines whose disks reside on the datastore.

Multipathing and NFS v4.1 Storage

Slide 6-67

One recommended configuration for NFS version 4.1 multipathing:

- Configure one VMkernel port.
- Use adapters attached to the same physical switch to configure NIC teaming.
- Configure the NFS server with multiple IP addresses:
 - IP addresses can be on the same subnet.
- To better utilize multiple links, configure NIC teams with the IP hash load-balancing policy.



To create a highly available NAS architecture, you must avoid single points of failure. Examples of a single point of failure include the NIC card in an ESXi host, and the cable between the NIC card and the switch. One example is shown on the slide. To avoid single points of failure and to create a highly available NAS architecture, configure the ESXi host with redundant NIC cards and redundant physical switches.

The best approach is to install multiple NICs on an ESXi host and configure them in NIC teams. NIC teams should be configured on separate external switches, with each NIC pair configured as a team on the respective external switch. In addition, you may apply load-balancing algorithm, based on the link aggregation protocol type supported on the external switch, such as 802.3ad or EtherChannel.

As another example, an even higher level of performance and high availability can be achieved with cross-stack EtherChannel-capable switches. With certain network switches, you can team ports across two or more separate physical switches that are managed as one logical switch. NIC teaming across virtual switches provides additional resilience as well as some performance optimization. Having more paths available to the ESXi host can improve performance by enabling distributed load sharing.

Only one active path is available for the connection between the ESXi host and a single storage target (LUN or mount point). Although alternative connections might be available for failover, the bandwidth for a single datastore and the underlying storage is limited to what a single connection can provide. To leverage more available bandwidth, an ESXi host would require multiple connections from the ESXi host to the storage targets. You might need to configure multiple datastores, each using separate connections between the ESXi host and the storage.

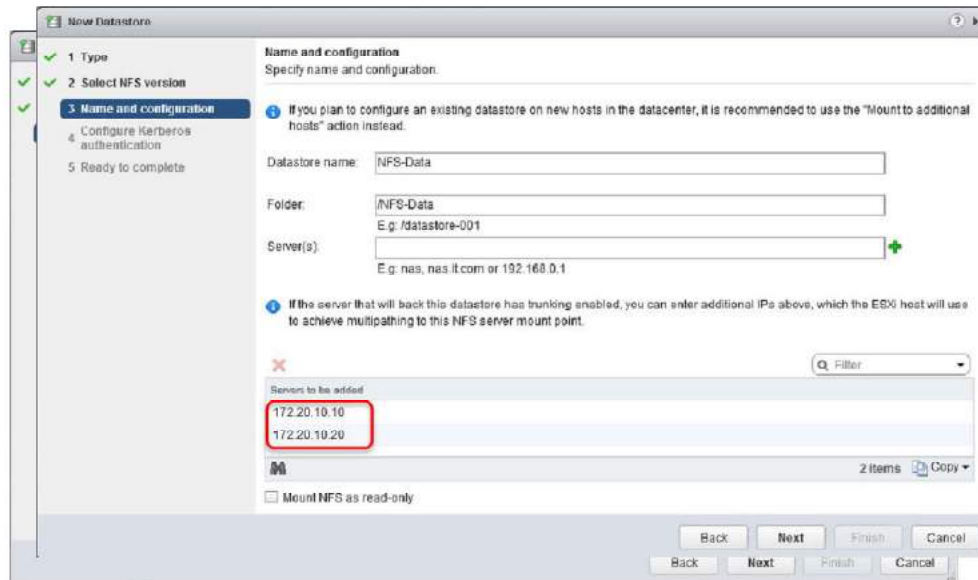
The table shows the recommended configuration for NFS multipathing.

External Switches Support Cross-Stack EtherChannel	External Switches Do Not Support Cross-Stack EtherChannel
Configure one VMkernel port.	Configure two or more VMkernel ports on different virtual switches on different subnets.
Configure NIC teaming by using adapters attached to separate physical switches.	Configure NIC teaming with adapters attached to the same physical switch.
Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.	Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.
To use multiple links, configure NIC teams with the IP hash load-balancing policy	To use multiple links, allow the VMkernel routing table to make decisions on which link to send packets (requires multiple datastores).

Enabling Session Trunking and Multipathing

Slide 6-68

Multiple IP addresses are configured for each NFS v4.1 datastore.



NFS v4.1 supports native multipathing and session trunking. To enable multipathing, the administrator enters multiple server IP addresses when configuring the datastore.

Lab 10: Accessing NFS Storage

Slide 6-69

Configure access to an NFS datastore

1. Configure Access to NFS Datastores
2. View NFS Storage Information

Review of Learner Objectives


Slide 6-70

You should be able to meet the following objectives:

- Describe NFS components
- Describe the differences between NFS v3 and NFS v4.1
- Configure and manage NFS datastores

Lesson 6: vSAN Datastores

Slide 6-71



Learner Objectives

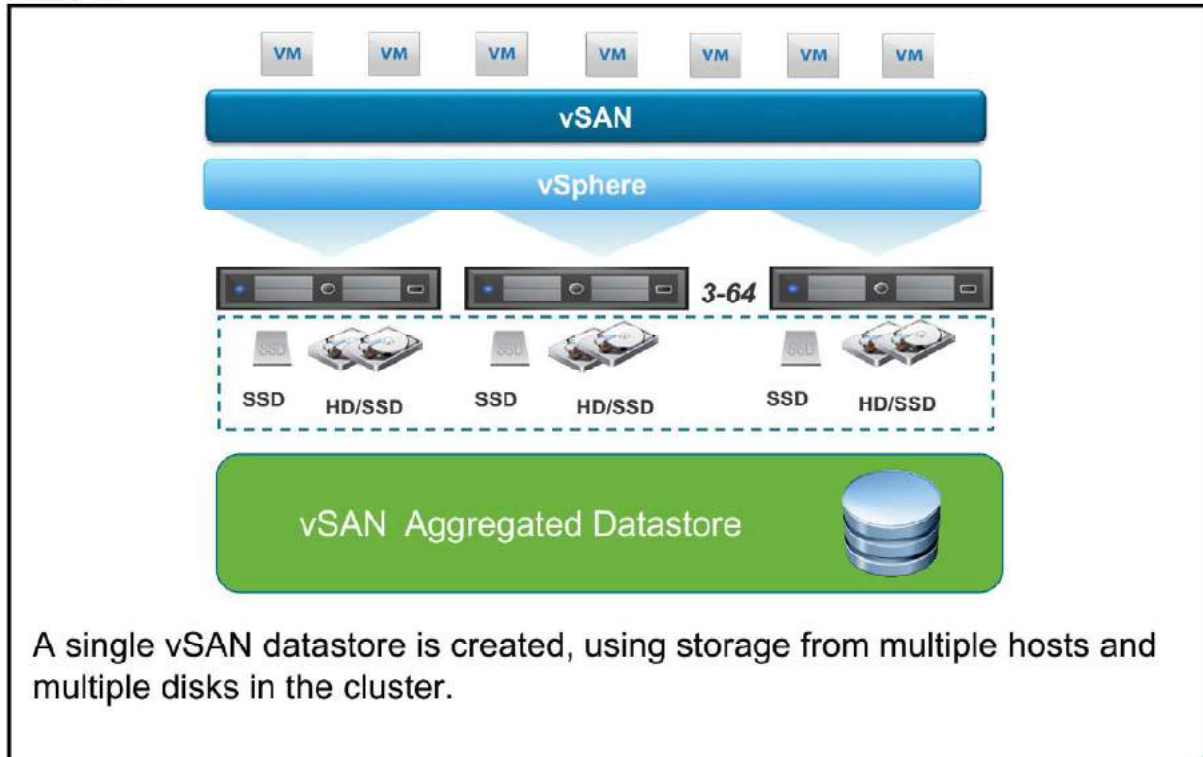
Slide 6-72

By the end of this lesson, you should be able to meet the following objectives:

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Describe the steps for configuring vSAN
- Explain how to create and use vSAN storage policies

About vSAN

Slide 6-73



A single vSAN datastore is created, using storage from multiple hosts and multiple disks in the cluster.

Starting from vSphere 5.5, VMware introduced software-defined storage with vSAN datastores. vSAN aggregates local storage devices (flash devices and magnetic disks) to provide a clustered datastore that can be utilized by virtual machines.

vSAN 6.x provides two different configurations: hybrid and all-flash. The hybrid configuration uses server-based flash devices to provide a cache layer for optimal performance while using magnetic disks to provide capacity and persistent data storage. The all-flash configuration uses flash devices for both the cache layer and capacity layer.

In a vSAN environment, a number of ESXi hosts are configured to form a vSAN cluster. All of the ESXi hosts communicate together through a dedicated vSAN network. Hosts without local storage can share their compute resources and take advantage of the clustered storage resources. Local storage is combined on each host to form up to five local disk groups. A disk group includes one flash cache device and up to seven capacity devices.

The disk groups of all of the ESXi hosts in the vSAN cluster are combined to create a vSAN datastore. Only a single datastore exists per vSAN cluster. Object Store File System (OSFS) enables the VMFS volumes from each host to be combined and mounted as a single datastore. This datastore contains all the virtual machine files, including the VMDK files. Each of the VMDK files can have a different virtual machine storage policy created that defines how data is stored on the disks of the

datastore. These virtual machine storage policies are configured to take advantage of the capabilities of vSAN.

vSAN datastores help administrators to use software-defined storage in the following ways:

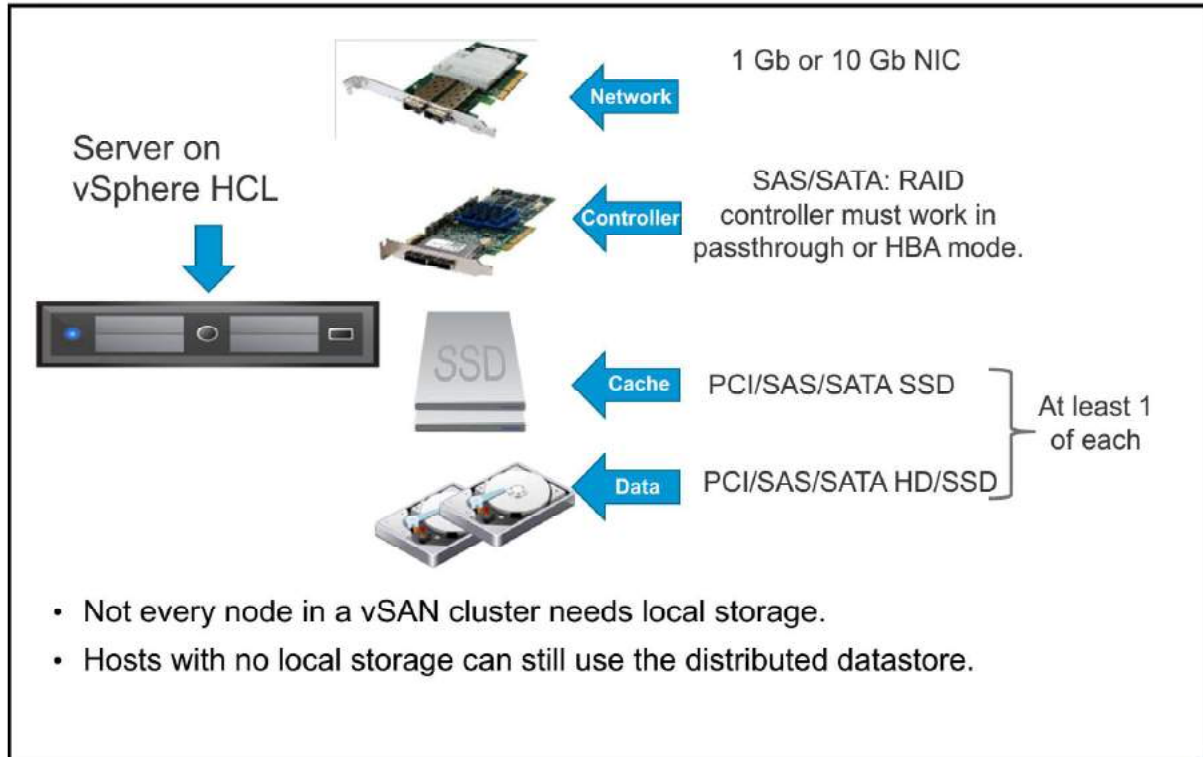
- Storage Policy per virtual machine architecture: Multiple policies per datastore allow for each virtual machine to have different storage.
- vSphere and vCenter Server integration: vSAN capability is built in and requires no appliance. You create a vSAN cluster, just like vSphere HA or vSphere DRS.
- Scale-out storage: Up to 64 ESXi hosts can be in a cluster. Scale out by populating new nodes in the cluster or by setting vSAN to scan for empty disks and add them automatically
- Built-in resiliency: A default policy mirrors all objects for virtual machines that are not configured for the vSAN.

vSAN 6 introduces a number of new features and enhancements:

- New on-disk format
- New performance snapshots
- Fault domains designed to withstand rack failure
- Default per-VM storage policy assignable at the datastore level
- Visualization of vSAN datastore utilization
- Projection of utilization based on possible storage policy configuration
- Resync status displayed in the vSphere Web Client interface
- New disk serviceability functions

vSAN Requirements

Slide 6-74



The requirements for vSAN start with a machine running vCenter Server. This requirement is necessary for management because vSAN is fully integrated into vSphere.

A vSAN cluster must have a minimum of three hosts that contribute capacity to the cluster. A host must not participate in other clusters beside the vSAN cluster.

Each vSAN cluster requires a dedicated network that each host can communicate with:

- Dedicated 1 Gbps network (10 Gbps recommended) for hybrid disk groups.
- Dedicated 10 Gbps network for all-flash disk groups.

VMware recommends a 10 Gb network, preferably with a NIC team of 2 x 10 Gb NICs for fault tolerance purposes. Testing has shown that a 1 Gb network can work, but this size of network is not supported. The vSAN network must be configured for IPv4 or IPv6 and support multicast.

All hosts in the cluster do not need to have local storage. A host without storage is used for its compute resources while leveraging the vSAN datastore. All hosts that have local storage must have the following devices:

- One Serial Attached SCSI (SAS) or SATA solid-state disk (SSD) or PCIe flash device and one or more magnetic drives for each hybrid disk group.

- One SAS or SATA SSD or PCIe flash device and one or more flash disks with flash capacity enabled for all-flash disk groups.

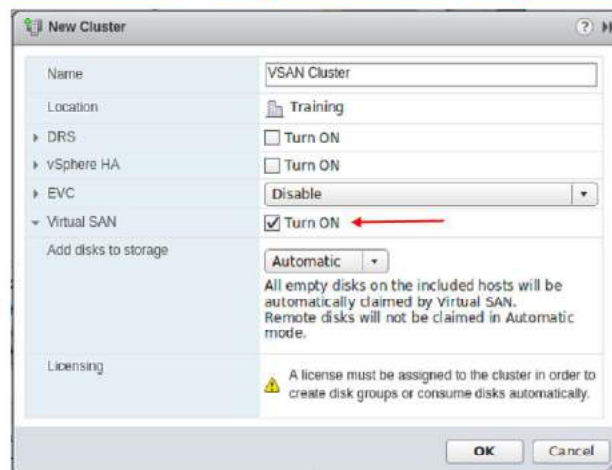
In addition, each host should have 32 GB or more of memory to accommodate a maximum number of five disk groups and a maximum number of seven capacity devices per disk group.

For more information, see *Administering VMware Virtual SAN* at <https://www.vmware.com/support/pubs/virtual-san-pubs.html>.

Configuring a vSAN Datastore

Slide 6-75

A vSAN datastore is configured in a few steps.



vSAN datastores are fully integrated with vSphere 6.

For information about enabling or disabling a vSAN cluster, see VMware knowledge base article 2058322 at <http://kb.vmware.com/kb/2058322>.

Configuration for vSAN can only be done through vSphere Web Client.

Disk Groups

Slide 6-76

vSAN disk groups composed of flash-based devices and magnetic disks require:

- One flash device:
 - Maximum of one flash device per disk group
- One HD/SSD:
 - Supports up to seven devices per disk group
- Maximum of five disk groups per host



On each ESXi host that contributes its local devices to a vSAN cluster, devices are organized into disk groups. A disk group is a unit of physical storage capacity on a host and a group of physical devices that provide performance and capacity to the vSAN cluster.

The disk groups of all ESXi hosts in a vSAN cluster are combined to create a vSAN datastore. The size of a vSAN datastore is governed by the number of capacity devices per ESXi host and the number of ESXi hosts in the cluster. The vSAN datastore is used to store virtual machine files, including virtual machine disks (VMDKs).

Viewing vSAN Cluster Summary

Slide 6-77

In vSphere Web Client, the **Summary** tab of the vSAN cluster displays the general vSAN configuration information.

The screenshot displays the vSphere Web Client interface for a vSAN cluster. The main interface shows the 'Virtual SAN' tab selected, with a green circle highlighting it. A green arrow points from this tab to a detailed summary panel on the right. The summary panel provides the following information:

Virtual SAN	
Add disks to storage	Manual
Hosts	8 hosts
SSD disks in use	8 of 8 eligible
Data disks in use	16 of 16 eligible
Total capacity of VSAN datastore	76.00 GB
Free capacity of VSAN datastore	66.13 GB
Network status	✓ Normal

At the bottom of the summary panel, there is a 'Disk management' link.

Using vSAN

Slide 6-78

Capabilities define the capacity, performance, and availability characteristics of the underlying physical storage. The vSAN cluster presents these capabilities to vCenter Server, where they can be consumed by virtual machines.

Requirements outline the needs of a virtual machine.

Virtual machine storage policies specify the virtual machine requirements so that the virtual machine can be placed appropriately on the vSAN datastore.



After you enable vSAN on a cluster, a single vSAN datastore is created. This datastore appears as another type of datastore on the list of datastores that might be available, including vSphere Virtual Volumes, VMFS, and NFS.

A single vSAN datastore can provide different service levels for each virtual machine or each virtual disk. In vCenter Server, storage characteristics of the vSAN datastore appear as a set of capabilities.

You can reference these capabilities when defining a storage policy for virtual machines. When you later deploy virtual machines, vSAN uses this policy to place virtual machines in the optimal manner based on the requirements of your virtual machine.

Objects in vSAN Datastores

Slide 6-79

In a vSAN datastore, files are grouped into four types of objects:

- Namespaces
- Virtual disks
- Snapshots
- Swap files

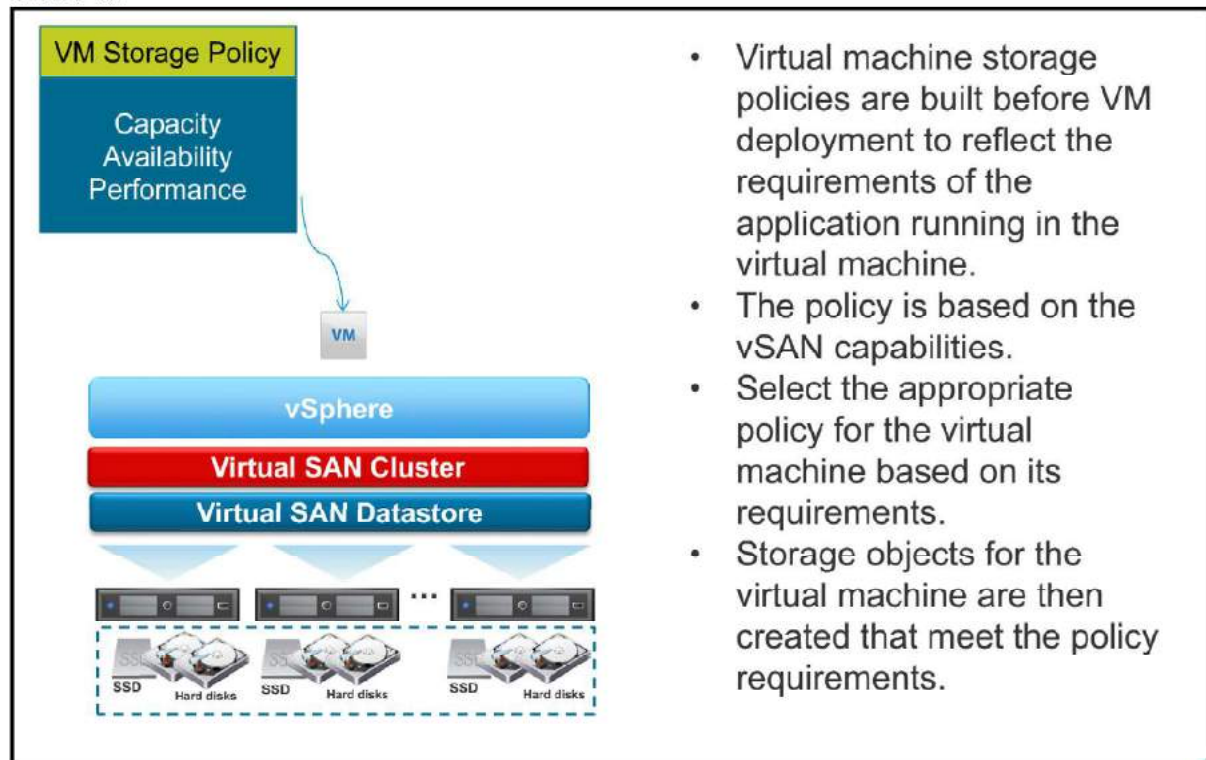


A vSAN cluster stores and manages data as flexible data containers called objects. When you provision a virtual machine on a vSAN datastore, a set of objects is created. These objects are of the following types:

- VM home namespace: Stores the virtual machine metadata (configuration files).
- VMDK: Virtual machine disk.
- Snapshot deltas: Created when snapshots of the virtual machine are taken.
- VM swap: Virtual machine swap file, which is created when the virtual machine is powered on.
- VM memory: Virtual machine's memory state when a virtual machine is suspended, or when a snapshot is taken of a virtual machine and its memory state is preserved.

Virtual Machine Storage Policies

Slide 6-80



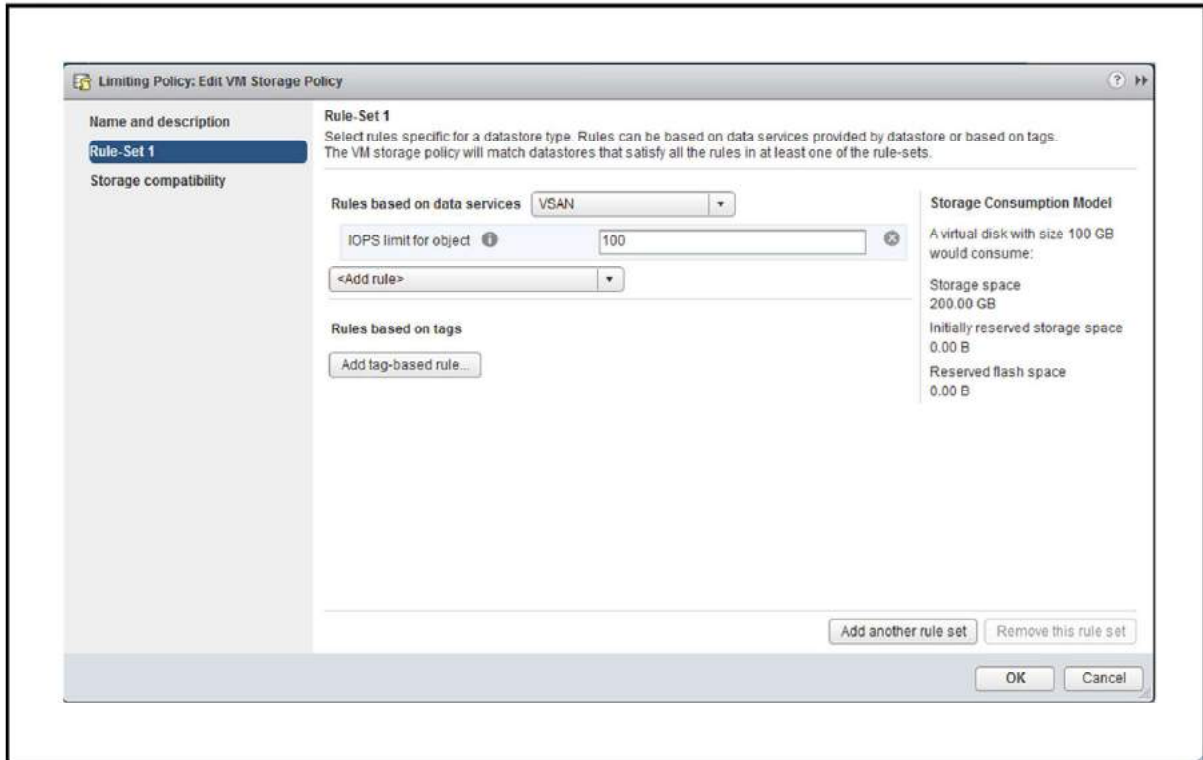
- Virtual machine storage policies are built before VM deployment to reflect the requirements of the application running in the virtual machine.
- The policy is based on the vSAN capabilities.
- Select the appropriate policy for the virtual machine based on its requirements.
- Storage objects for the virtual machine are then created that meet the policy requirements.

Virtual machine storage policies are a set of storage properties that are configured by an administrator to be used by virtual machines. Each of these storage policies is created to reflect a set of capabilities that meets the availability, performance, and storage requirements of some of the virtual machines in an environment. These storage policies are created in advance of the deployment of virtual machines.

When a virtual machines is deployed, the administrator chooses the VM Storage Policy that meets its requirements for the creation of its VMDK file.

Configuring Virtual Machine Storage Policies

Slide 6-81



The capabilities of VM Storage Policies reflect the level of Redundant Array of Independent Nodes (RAIN) support that they offer. VM Storage Policies define how objects are configured to offer availability.

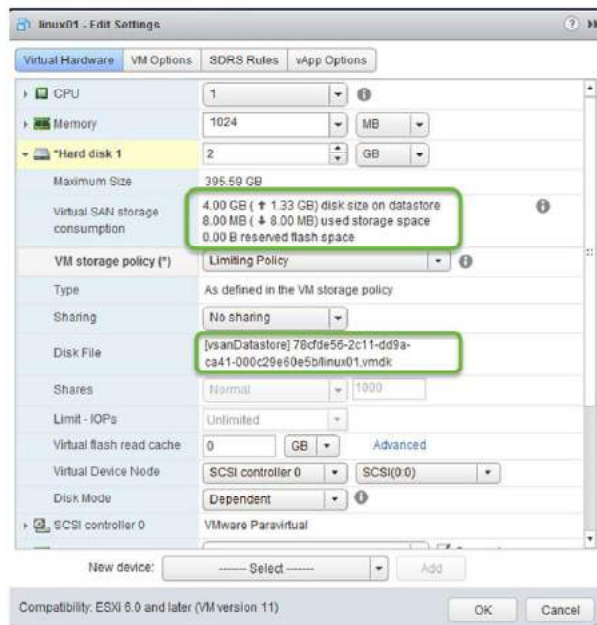
Viewing a Virtual Machine's vSAN Datastore

Slide 6-82

The consumption of vSAN storage is based on the virtual machine's storage policy.

The virtual machine's hard disk view provides the following information:

- Summarizes the total storage size and the used storage space
- Displays the virtual machine storage policy
- Shows the location of disk files on a vSAN datastore



vSAN Cluster Member Maintenance Mode Options

Slide 6-83

Before you shut down, reboot, or disconnect a host that is a member of a vSAN cluster, you must place the host in maintenance mode.

When you place a host in maintenance mode, you can select a specific evacuation mechanism.

When any member node of a vSAN cluster enters maintenance mode, the cluster capacity is automatically reduced because the member node no longer contributes storage to the cluster.

Option	Action
Ensure Accessibility	Moves enough components to ensure operational integrity of objects.
Full Data Migration	All components are evacuated from the host.
No Data Migration	No action is taken, which can result in degraded objects.

Ensuring accessibility guarantees that all objects are accessible but does not ensure that all underlying components of those objects are accessible. Depending on the storage policy, accessibility enables quicker entrance into maintenance mode but increases the risk of a second component failure leading to data loss. To ensure that no data loss occurs, use the Full Data Migration option. This option evacuates all components from the host. The evacuation can take a long time, depending on the configuration of the storage policies. The administrator must ensure that resources are available on the remaining hosts.

Removing a Host from a vSAN Cluster

Slide 6-84

To remove a host that is participating in a vSAN cluster:

1. Place the host in maintenance mode.
2. Delete the disk groups associated with the host.
3. Remove the vSAN VMkernel ports.
4. Remove the host from the cluster.

For more information about detaching an ESXi host from a vSAN cluster and removing the vSAN cluster, see VMware knowledge base article 2072347 at <http://kb.vmware.com/kb/2072347>.

Review of Learner Objectives

Slide 6-85

You should be able to meet the following objectives:

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Describe the steps for configuring vSAN
- Explain how to create and use vSAN storage policies

Key Points

Slide 6-86


- You use VMFS datastores to hold virtual machine files.
- Shared storage is integral to vSphere features such as vSphere vMotion, vSphere HA, and vSphere DRS.
- vSAN enables the use of vSphere HA, vSphere vMotion, and vSphere Storage vMotion without requiring external shared storage.
- vSAN clusters direct-attached server disks to create shared storage designed for virtual machines.

Questions?

MODULE 7

Virtual Machine Management

Slide 7-1



You Are Here

Slide 7-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
- 7. Virtual Machine Management**
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Importance

Slide 7-3

Because virtual machines are the basis of all IT operations in a virtual infrastructure, you should know how to properly deploy, configure, and manage them.

Failure to properly populate your virtual machine inventory might have unsatisfactory consequences.


Module Lessons

Slide 7-4

- | | |
|-----------|------------------------------------|
| Lesson 1: | Creating Templates and Clones |
| Lesson 2: | Working with Content Libraries |
| Lesson 3: | Modifying Virtual Machines |
| Lesson 4: | Migrating Virtual Machines |
| Lesson 5: | Creating Virtual Machine Snapshots |

Lesson 1: Creating Templates and Clones

Slide 7-5



Learner Objectives

Slide 7-6

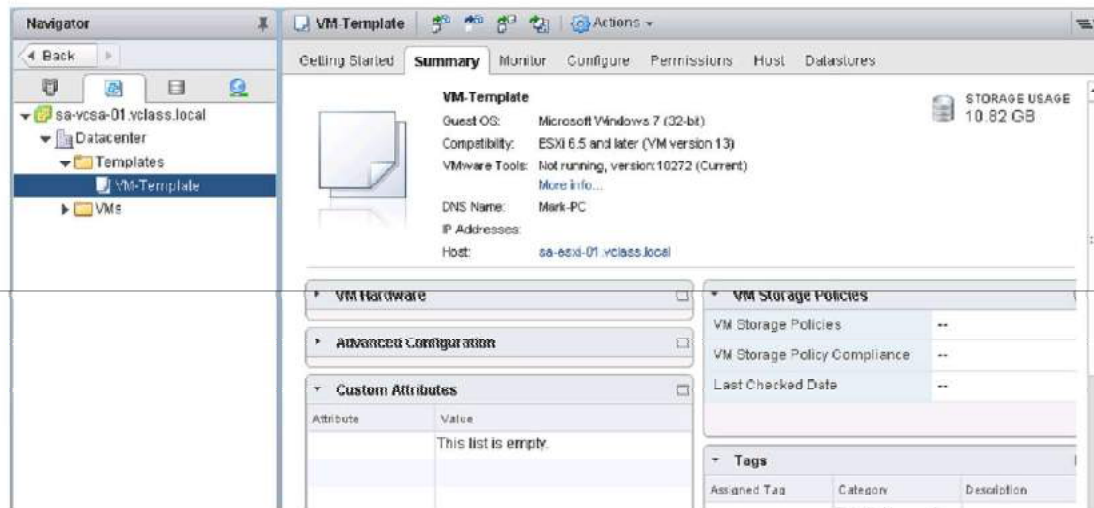
By the end of this lesson, you should be able to meet the following objectives:

- Create a template
- Deploy a virtual machine from a template
- Clone a virtual machine
- Enable guest operating system customization by vCenter Server
- Upgrade the virtual machine's hardware

Using a Template

Slide 7-7

A template is a master copy of a virtual machine. A template is used to create and provision new virtual machines.



A template is an image that typically includes:

- A guest operating system
- A set of applications
- A specific virtual machine configuration that provides virtual counterparts to hardware components

Creating templates makes provisioning of virtual machines much faster and less error prone than provisioning physical machines creating a new virtual machine by using the Create New Virtual machine wizard.

Templates coexist with virtual machines in the inventory. You can organize collections of virtual machines and templates into arbitrary folders and apply permissions to virtual machines and templates. Virtual machines can be changed into templates without the need to make a full copy of the virtual machine files and the creation of a new object.

You can deploy a virtual machine from a template. The deployed virtual machine is added to the folder that the user selected when the template was created.

Creating a Template

Slide 7-8

Clone the virtual machine to a template:

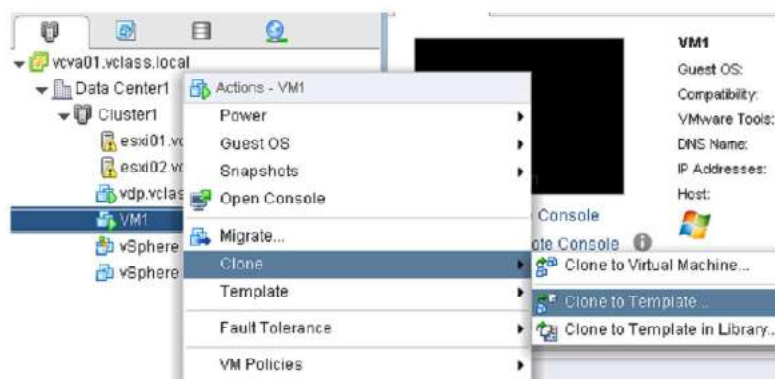
- The virtual machine can be powered on or powered off.

Convert the virtual machine to a template:

- The virtual machine must be powered off.

Clone a template:

- Used to create a new template based on a template that existed previously.



Clone to Template offers you the choice of format in which to store the virtual machine's virtual disks:

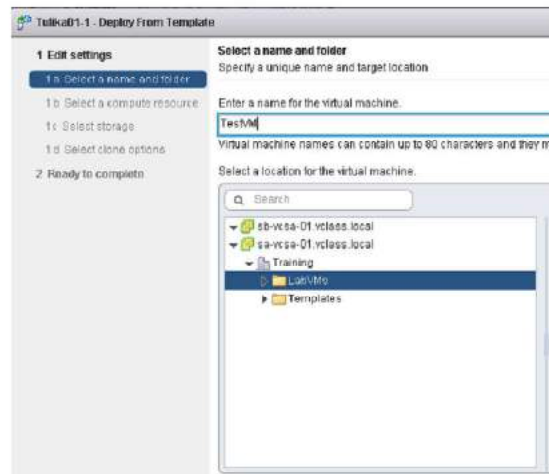
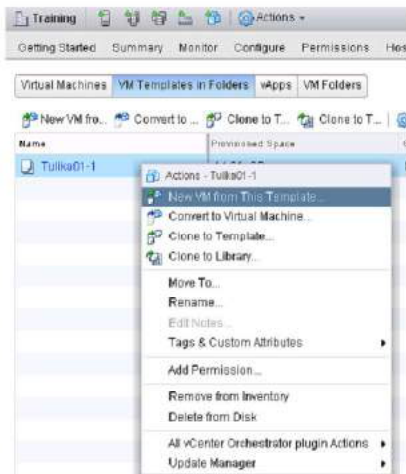
- Same format as source
- Thin-provisioned format
- Thick-provisioned lazy zeroed format
- Thick-provisioned eager zeroed format

Convert to Template does not offer a choice of format and leaves the virtual machine's disk file intact.

Deploying a Virtual Machine from a Template

Slide 7-9

To deploy a virtual machine, you must provide information, such as the virtual machine name, inventory location, host, datastore, and guest operating system customization data.

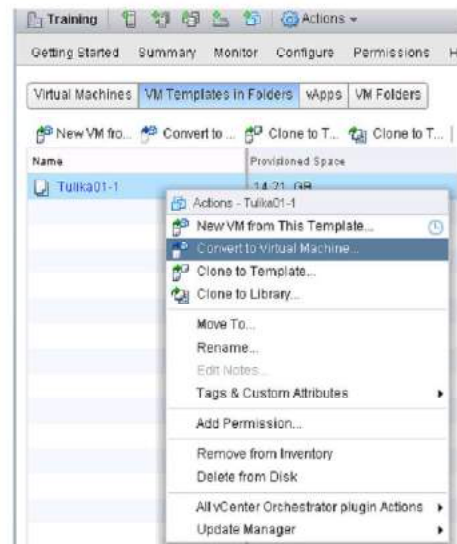


Updating a Template

Slide 7-10

Update a template to include new patches, make system changes, and install new applications:

1. Convert the template to a virtual machine.
2. Place the virtual machine on an isolated network to prevent user access.
3. Make appropriate changes to the virtual machine.
4. Convert the virtual machine to a template.



To update your template to include new patches or software, you do not have to create a new template. Instead, first convert the template to a virtual machine. This conversion enables you to power on the virtual machine.

For added security, prevent users from accessing the virtual machine while you are updating it. Either disconnect the virtual machine from the network or place it on an isolated network.

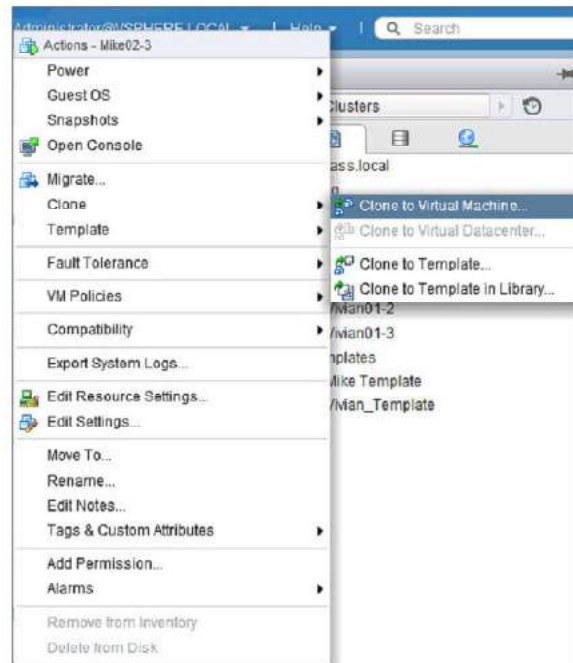
Log in to the virtual machine's guest operating system and apply the patch or install the software. When you have finished, power off the virtual machine and convert it to a template again.

Cloning a Virtual Machine

Slide 7-11

Cloning a virtual machine creates a virtual machine that is an exact copy of the original:

- Cloning is an alternative to deploying a virtual machine.
- The virtual machine being cloned can be powered on or powered off.



Cloning a virtual machine is an alternative to deploying a virtual machine from a template. Cloning a virtual machine creates a duplicate of the virtual machine with the same configuration and installed software as the original.

You can customize the guest operating system of the clone to change the virtual machine name, network settings, and other properties. Customizing the guest operating system prevents conflicts that might occur when a virtual machine and a clone with identical guest operating system settings are deployed simultaneously.

To clone a virtual machine, you must be connected to vCenter Server. You cannot clone virtual machines if you connect vSphere Client directly to an ESXi host.

The virtual machine that you clone can be powered on or powered off. When you clone a virtual machine that is powered on, services and applications are not automatically quiesced when the virtual machine is cloned.

Customizing the Guest Operating System

Slide 7-12

Use the Guest Operating System Customization wizard to make virtual machines created from the same template or clone unique.

- Customizing a guest operating system enables you to change:
 - Computer name
 - Network settings
 - License settings
 - Windows Security Identifier

During cloning or deploying virtual machines from a template:

- You can create a specification to prepare the guest operating systems of virtual machines.
- Specifications can be stored in the database.
- You can edit specifications in the Customization Specifications Manager.
- Windows and Linux operating systems are supported.

When you clone a virtual machine or deploy it from a template, you can customize its guest operating system to change properties, including:

- Computer name
- Network settings
- License settings

Customizing guest operating systems can help prevent conflicts that occur when virtual machines with identical settings are deployed, such as conflicts because of duplicate computer names or IP addresses.

You can specify the customization settings by using the Guest Customization wizard during cloning or deployment. Or, you can create customization specifications, which are customization settings stored in the vCenter Server database. During cloning or deployment, you can select a customization specification to apply to the new virtual machine. Use the Customization Specification Manager to manage customization specifications that you create with the Guest Customization wizard.

The guest operating system that is being customized must have VMware Tools installed. The guest operating system must also be installed on a disk attached to SCSI node 0:0 in the virtual machine configuration.

To enable guest operating system customization, vCenter Server must first be configured for this task. To customize virtual machines running operating systems that predate Windows 2008 and Windows Vista, you must install Microsoft Sysprep tools on the vCenter Server system. Sysprep tools are built in to the Windows 2008, Windows Vista, and later Windows operating systems.

Customization of Linux guest operating systems requires that Perl must be installed in the Linux guest operating system.

Guest operating system customization is supported on multiple Linux distributions. To verify customization support for Linux distributions and compatible ESXi hosts, see *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

For more about guest operating system customization, see *vSphere Virtual Machine Administration Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Review of Learner Objectives


Slide 7-13

You should be able to meet the following objectives:

- Create a template
- Deploy a virtual machine from a template
- Clone a virtual machine
- Enable guest operating system customization by vCenter Server
- Upgrade the virtual machine's hardware

Lesson 2: Working with Content Libraries

Slide 7-14



Learner Objectives

Slide 7-15

By the end of this lesson, you should be able to meet the following objectives:

- Identify benefits of a content library
- Describe types of content libraries
- Deploy a virtual machine from the content library

About the Content Library

Slide 7-16

A content library is a repository of OVF templates and other files that can be shared and synchronized across vCenter Server systems.



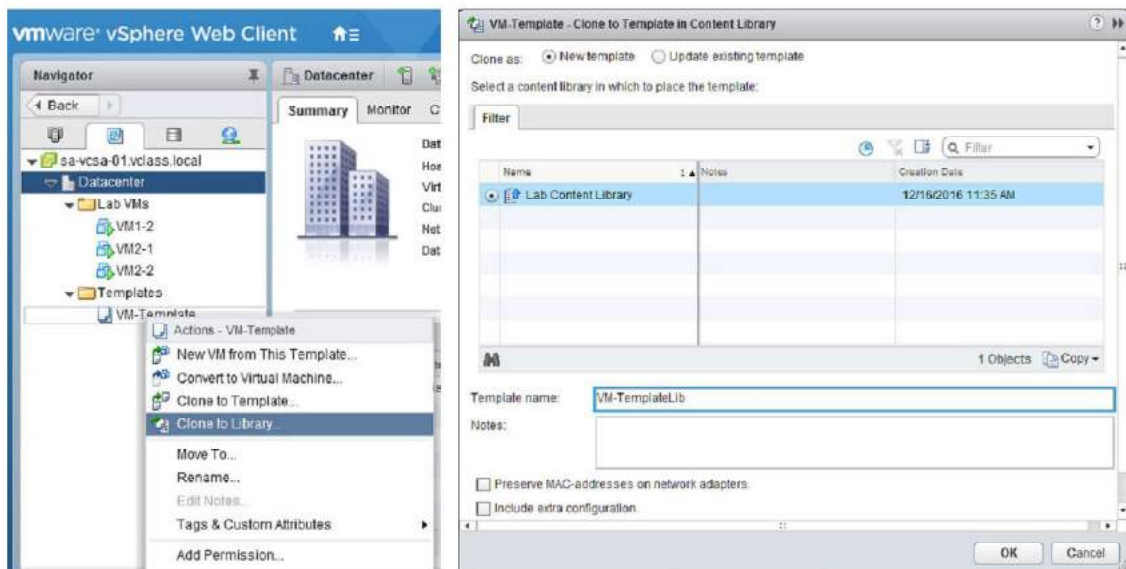
The content library helps manage templates and other file types used in vCenter Server instances globally. Organizations might have multiple vCenter Server instances in data centers around the globe. On these vCenter Server instances, organizations might have a collection of templates, ISO images, and so on. The challenge is that all of these items are independent of one another, with different versions of these files and templates on various vCenter Server instances.

The content library is the solution. IT can store OVF templates, ISO images, or any other file types in a central location. The templates, images, and files can be published, enabling other content libraries to subscribe to and download content. The content library keeps content up to date by periodically synchronizing with the publisher, ensuring that the latest version is available.

Adding Templates to a Content Library

Slide 7-17

Library items are VM templates, vApp templates, or other VMware objects that can be contained in a content library.



Library items are VM templates, vApp templates, or other VMware objects that can be contained in a content library. VMs and vApps have several files, such as log files, disk files, memory files, and snapshot files that are part of a single library item. You can create library items in a specific local library or remove items from a local library. You can also upload files to an item in a local library so that the libraries subscribed to it can download the files to their NFS or SMB server, or datastore.

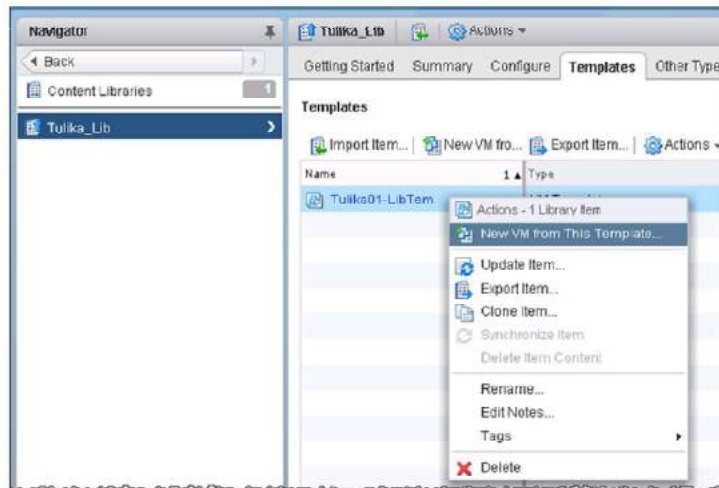
Deploying VMs from Templates in a Content Library

Slide 7-18

The templates in the library can be used to deploy virtual machines and vApps.

Each VM template, vApp template, or another type of file in a library is a library item.

You can mount an ISO file directly from a content library.



Content library was introduced with vSphere 6.0. This feature provides simple and effective way of managing content, such as VM templates, vApps, ISO images, and scripts. Content library also allows the content to be distributed across sites and vCenter Servers. This distribution can be achieved by publishing a library at one site and creating subscribed libraries at other sites, so that their content is synchronized with the publisher. To summarize, content library brings out consistency, compliance, and efficiency in deploying workloads at scale.

You can create and manage a content library from a single vCenter Server instance, but you can share the library items with other vCenter Server instances if HTTPS is allowed between them.

Benefits of Content Libraries

Slide 7-19



The Content Library API provides services that allow you to create and manage content libraries programmatically. You can create a local library and publish it for the entire virtual environment.

Using content libraries, administrators can perform the following functions:

- Store, version and share content: Storage and consistency is a key reason to install and use a content library. Sharing the content and ensuring that the content is kept up to date is a major task. For example, assume that you have a main vCenter Server instance and you create a central content library to store the master copies of OVF templates, ISO images, and other file types. You can publish this content library to allow other libraries, which can be located anywhere in the world, to subscribe and download an exact copy of the data. When an OVF template is added, modified, or deleted from the published catalog, the subscriber synchronizes with the publisher and the libraries are updated with the latest content.
- Perform distributed file management.
- Publish content to public or subscribed content.
- Synchronize content across sites and vCenter Server systems.
- Back up content libraries using vSphere datastores, vCenter directory, or NFS.
- Mount an ISO file directly from a content library

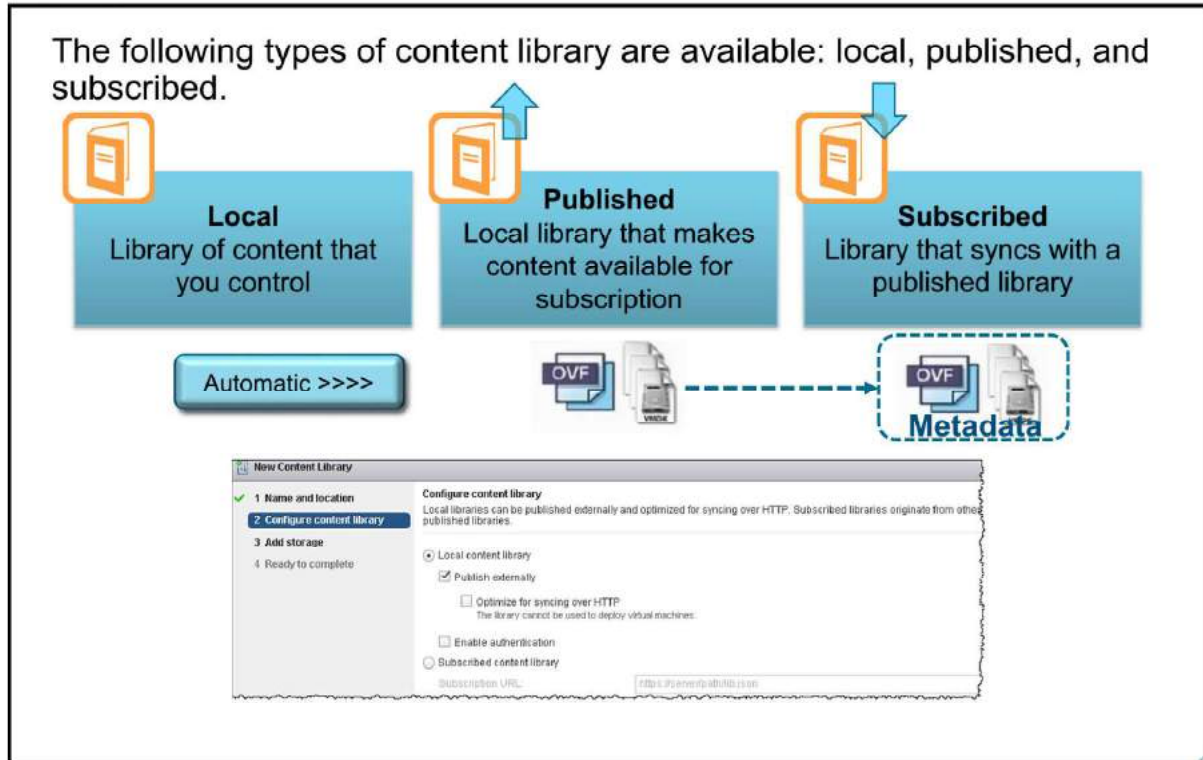
Subscribers might not always want to download and store all the content from the published library. They might not have the need or the space to store all of the content available to them. Content libraries can address such issues and implement storage efficiencies.

When a publisher updates its content, subscribers are notified of the changes. A subscriber can choose to receive either the full data or only the metadata. When the subscriber chooses to receive the full data, the subscriber can choose to download either the entire library or individual items.

Because content libraries are made globally available for subscription, concerns about security are expected. As a security measure, content libraries can be password-protected during publication.

Types of Content Library

Slide 7-20



You can create libraries local and subscribed libraries:

- You can create a local library as the source for content you want to save or share. Create the local library on a single vCenter Server instance. You can add items to a local library or remove them.
- You can publish a local library and this content library service endpoint can be accessed by other vCenter Server instances in your virtual environment. When you publish a library, you can configure the authentication method, which a subscribed library must use to authenticate to it.
- You can create a subscribed library and populate its content by synchronizing to a local library. A subscribed library contains copies of the local library files or just the metadata of the library items. The local library can be located on the same vCenter Server instance as the subscribed library, or the subscribed library can reference a local library on a different vCenter Server instance. You cannot add library items to a subscribed library. You can only add items to the source library. After synchronization, both libraries contain the same items.

Lab 11: Using Templates and Clones

Slide 7-21

Deploy a new virtual machine from a template and clone a virtual machine

1. Create a Virtual Machine Template
2. Create Customization Specifications
3. Deploy a Virtual Machine from a Template
4. Create a Content Library
5. Clone a VM Template to a Template in a Content Library
6. Deploy a Virtual Machine from a VM Template in the Content Library
7. Clone a Powered-On Virtual Machine

Review of Learner Objectives

Slide 7-22

You should be able to meet the following objectives:

- Identify benefits of a content library
- Describe types of content libraries
- Deploy a virtual machine from the content library

Lesson 3: Modifying Virtual Machines

Slide 7-23



Lesson 3: Modifying Virtual Machines

Learner Objectives

Slide 7-24

By the end of this lesson, you should be able to meet the following objectives:

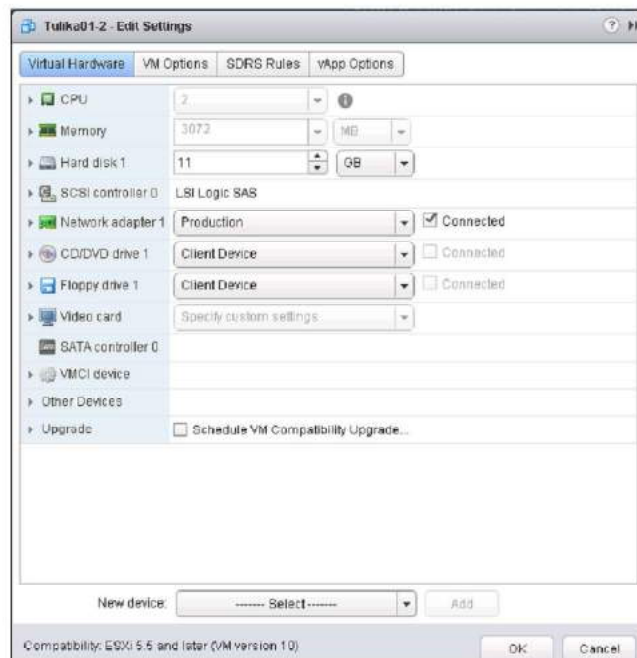
- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk
- Add a raw device mapping (RDM) to a virtual machine

Modifying Virtual Machine Settings

Slide 7-25

You can modify a virtual machine's configuration in its Edit Settings dialog box:

- Add virtual hardware:
 - Some hardware can be added while the virtual machine is powered on.
- Remove virtual hardware:
 - Some hardware can be removed only when the virtual machine is powered off
- Set virtual machine options.
- Control a virtual machine's CPU and memory resources.



You might have to modify a virtual machine's configuration, for example, to add a network adapter or to add a virtual disk. All virtual machine changes can be made while the virtual machine is powered off. Some hardware changes can be made to the virtual machine while it is powered on.

Besides adding virtual hardware, you can also remove virtual hardware and set various virtual machine options.

All virtual machine configuration is done in the virtual machine Edit Settings dialog box.

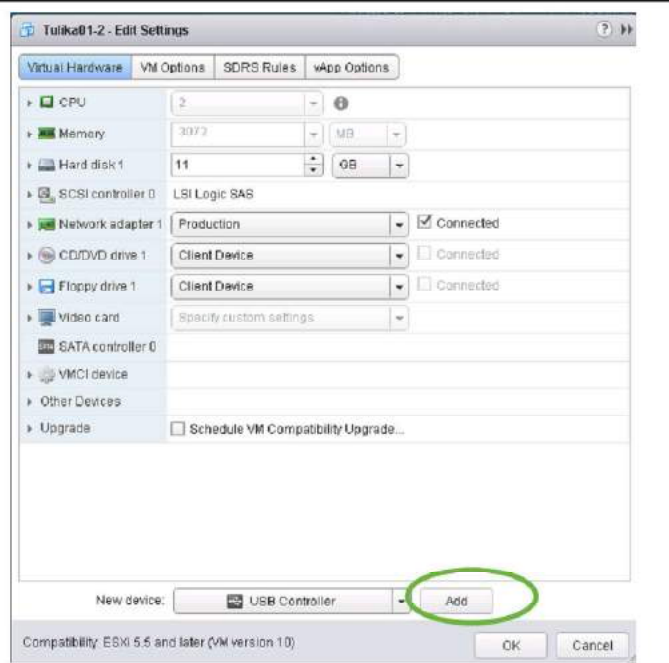
Hot-Pluggable Devices

Slide 7-26

The hot-plug option enables you to add resources to a running virtual machine

Examples of hot-pluggable devices: USB controllers, Ethernet adapters, and hard disk devices.

With supported guest operating systems, you can also add CPU and memory while the virtual machine is powered on.



Adding devices to a physical server or removing devices from a physical server requires you to physically interact with the server in the data center. When you use virtual machines, resources can be added dynamically without a disruption in service. You must shut down a virtual machine to remove hardware, but you can reconfigure the virtual machine without entering the data center.

CPU and memory can be added while the virtual machine is powered on. This feature is called the CPU hot-plug and memory hot-add option, which is supported only on guest operating systems that support this feature. These options are disabled by default. To use these hot-plug features, the following requirements must be satisfied:

- You must install VMware Tools.
- The virtual machine must use hardware version 7 or later.
- The guest operating system in the virtual machine must support CPU and memory hot-plug features.
- The hot-plug options must be enabled in the **CPU** or **Memory** options on the **Virtual Hardware** tab.

See *vSphere Virtual Machine Administration Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Creating an RDM

Slide 7-27

An RDM (a `-rdm.vmdk` file) enables a virtual machine to gain direct access to a physical LUN.

Encapsulating disk information in the RDM enables the VMkernel to lock the LUN so that only one virtual machine can write to the LUN.

You must define the following items when creating an RDM:

- Target LUN: LUN that the RDM will map to
- Mapped datastore: Stores the RDM file with the virtual machine or on a different datastore
- Compatibility mode
- Virtual device node



An RDM is a file that has a `.vmdk` extension, but the file contains only disk information describing the mapping to the LUN on the ESXi host. The data is stored on the LUN.

An RDM supports the following compatibility modes:

- Physical compatibility (passthrough) mode: Allows the guest operating system to access the hardware directly. Physical compatibility is useful if you are using SAN-aware applications in the virtual machine. But a LUN configured for physical compatibility cannot be cloned, made into a template, or migrated (using vSphere vMotion, cold migration can be performed) if the migration involves copying the disk. LUNs configured for passthrough mode can be a maximum of 64 TB in size. For physical compatibility mode RDM, the file name is `-rdmp.vmdk`.
- Virtual compatibility mode: Allows the virtual machine to use VMware snapshots and other advanced features. Virtual compatibility enables the LUN to behave as if it were a virtual disk. When you clone the disk, make a template out of it, or migrate it (if the migration involves copying the disk), the contents of the LUN are copied to a virtual disk (`-flat.vmdk`) file.

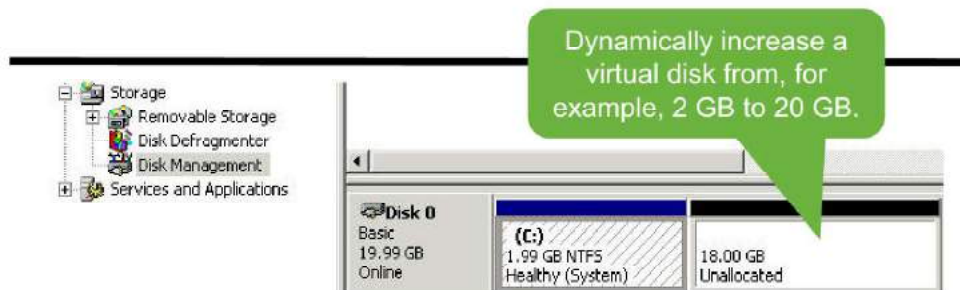
You can migrate virtual machines with RDMs with virtual machines powered on or powered off. Raw LUNs cannot be migrated because they are raw disks presented from the SAN. However, the RDM pointer files can be relocated if necessary.

Dynamically Increasing a Virtual Disk's Size

Slide 7-28

You can increase the size of a virtual disk that belongs to a powered-on virtual machine:

- It must not have snapshots attached.
- Might require system tools to make new space usable.



You can increase the size of a virtual disk that belongs to a powered-on virtual machine if the virtual machine lacks snapshots.

After you increase the size of a virtual disk, you might need to increase the size of the file system on this disk. Use the appropriate tool in the guest operating system to enable the file system to use the newly allocated disk space.

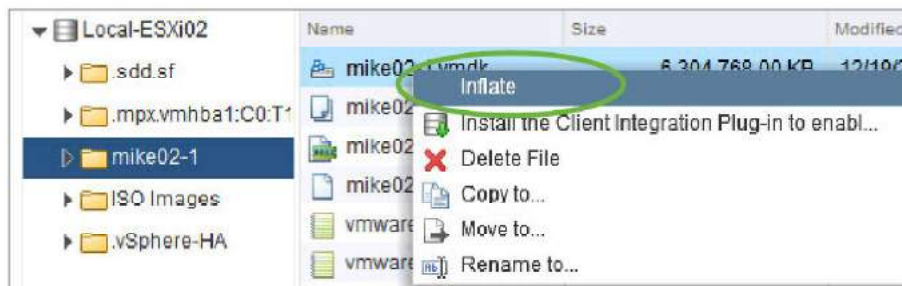
Inflating a Thin-Provisioned Disk

Slide 7-29

Thin-provisioned virtual disks can be converted to a thick, eager-zeroed format.

To inflate a thin-provisioned disk:

- The virtual machine must be powered off.
- Right-click the virtual machine's .vmdk file and select **Inflate**.



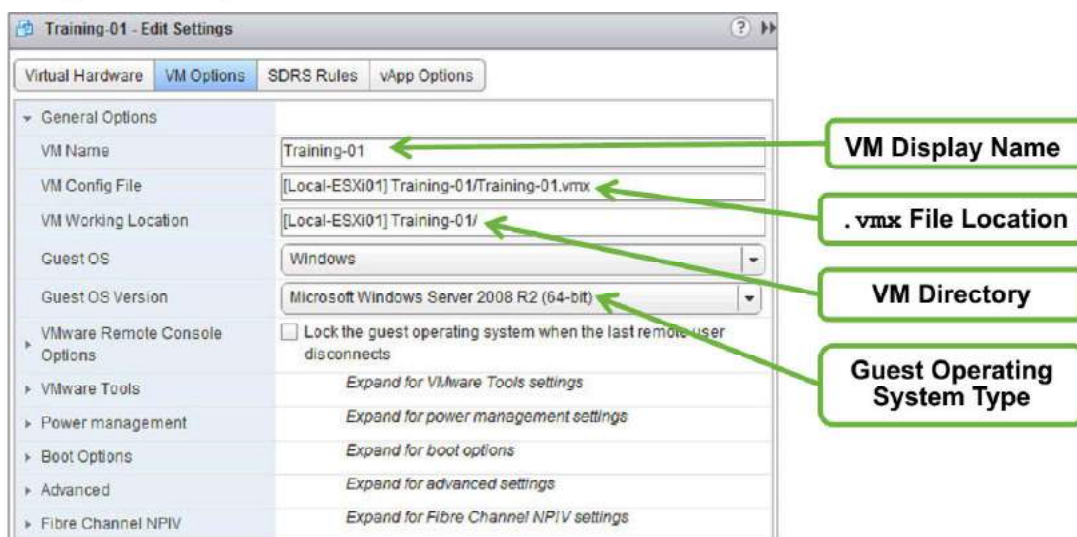
Or, you can use vSphere Storage vMotion and select a thick-provisioned disk as the destination.

When you inflate a thin-provisioned disk, the inflated virtual disk occupies the entire datastore space originally provisioned to it. Inflating a thin-provisioned disk converts a thin disk to a virtual disk in thick-provision format.

Virtual Machine Options

Slide 7-30

On the **VM Options** tab, you can set or change virtual machine options to run VMware Tools scripts, control user access to the remote console, configure startup behavior, and more.



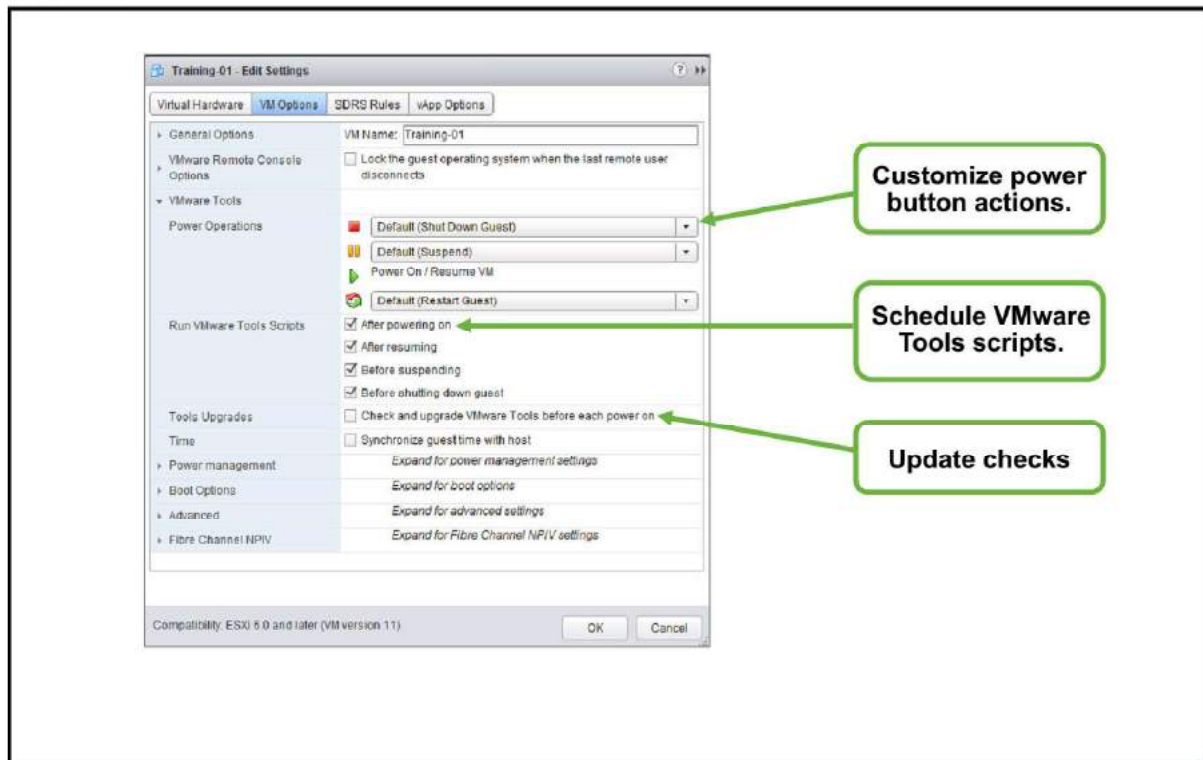
You can use the VM Options panel to modify things like the display name used for the virtual machine and the type of guest operating system installed. The location and name of the configuration file (.vmx file) and the location of the virtual machine's directory are also shown.

You can select the text for the configuration file and the working location to copy and paste them into a document. However, only the display name and the guest operating system type can be modified.

Changing the display name does not change the names of all of the virtual machine files or the directory that the virtual machine is stored in. When a virtual machine is created, the filenames and the directory name associated with the virtual machine are based on its display name. But changing the display name later does not modify the filename and the directory name.

VMware Tools Options

Slide 7-31



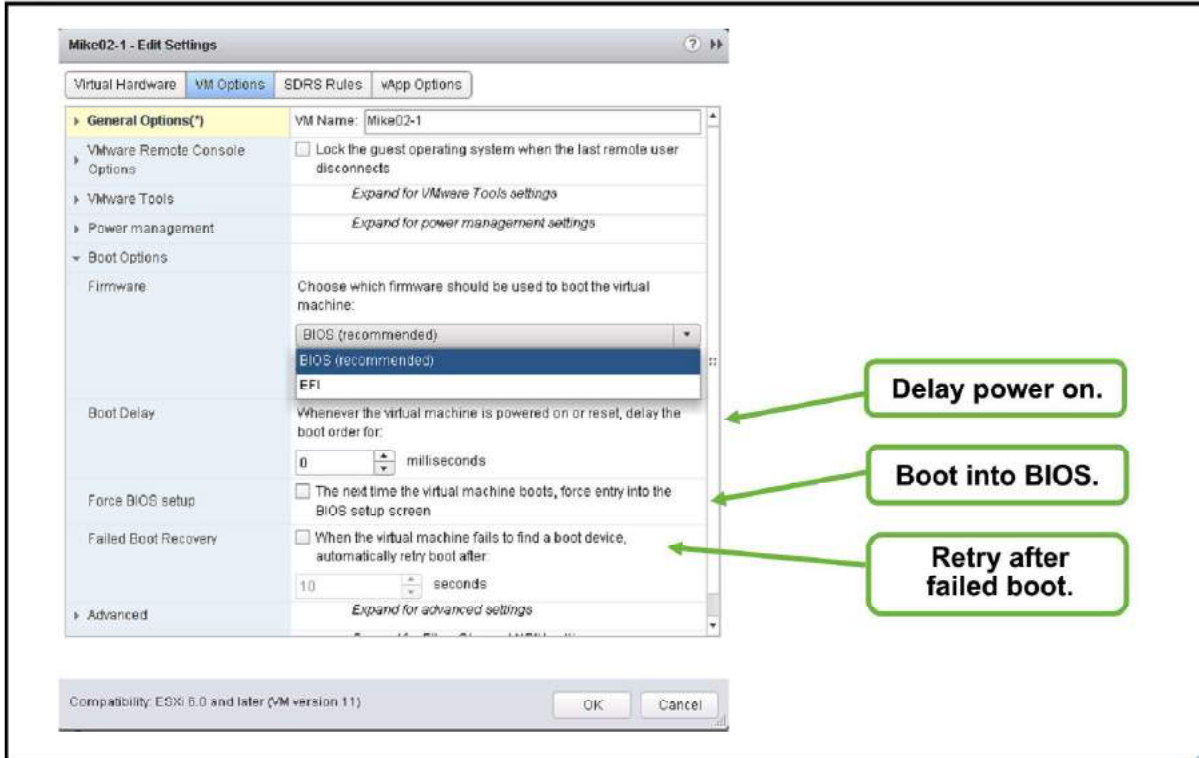
The VMware Tools panel controls how VMware Tools in the virtual machine responds to certain external events. You can use these controls to customize the power buttons on the virtual machine. The virtual machine must be powered off to change these settings.

The **Check and upgrade VMware Tools before each power on** check box can be configured to check whether a newer version of VMware Tools exists. If a newer version exists, VMware Tools is upgraded when the virtual machine is power cycled. The **Synchronize guest time with host** check box enables the guest operating system's clock to synchronize with the host.

For information about time keeping best practices for the guest operating systems that you are using, see VMware knowledge base articles 1318 at <http://kb.vmware.com/kb/1318> and 1006427 at <http://kb.vmware.com/kb/1006427>.

Boot Options

Slide 7-32



Occasionally, you might need to set the boot options.

- When you build a virtual machine and select a guest operating system, **BIOS** or **EFI** is selected by default, depending on the firmware supported by the operating system. Mac OS X Server guest operating systems support only Extensible Firmware Interface (EFI). If the operating system supports BIOS and EFI, you can change the default here.
- The Boot Delay panel enables you to set a delay between the time when a virtual machine is turned on and the guest operating system starts to boot. A delayed boot can help stagger virtual machine startups when several virtual machines are being powered on.
- You can use the Force BIOS Setup panel to change the BIOS settings like forcing a virtual machine to boot from a CD-ROM. The next time the virtual machine powers on, it goes straight into BIOS. A forced entry into BIOS is much easier than powering on the virtual machine, opening a console, and quickly trying to press the F2 key to go into BIOS.
- You can use the Failed Boot Recovery panel to have the virtual machine retry booting after 10 seconds (the default) if the virtual machine fails to find a boot device.

Lab 12: Modifying Virtual Machines

Slide 7-33

Modify a virtual machine's hardware and add a raw LUN to a virtual machine

1. Increase the Size of a VMDK File
2. Adjust Memory Allocation on a Virtual Machine
3. Rename a Virtual Machine in the vCenter Server Inventory
4. Add and Remove a Raw LUN on a Virtual Machine

Review of Learner Objectives


Slide 7-34

You should be able to meet the following objectives:

- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk
- Add an RDM to a virtual machine

Lesson 4: Migrating Virtual Machines

Slide 7-35



Learner Objectives

Slide 7-36

By the end of this lesson, you should be able to meet the following objectives:

- Verify vSphere vMotion requirements, including CPU constraints and guidelines
- Perform a vSphere vMotion migration
- Perform a vSphere Storage vMotion migration
- Perform a shared-nothing vSphere vMotion migration
- Describe the major enhancements to vSphere vMotion in vSphere 6

Migrating Virtual Machines

Slide 7-37

Migration means moving a virtual machine from one host, datastore, or vCenter Server system to another host, datastore, or vCenter Server system.

- Types of migrations:
 - Cold: Migrate a powered-off virtual machine to a new host or datastore.
 - Suspended: Migrate a suspended virtual machine to a new host or datastore
 - vSphere vMotion: Migrate a powered-on virtual machine to a new host.
 - vSphere Storage vMotion: Migrate a powered-on virtual machine's files to a new datastore.
 - Shared-nothing vSphere vMotion: Migrate a powered-on virtual machine to a new host and a new datastore simultaneously.

You can move virtual machines from one host or storage location to another location using hot or cold migration. For example, with vSphere vMotion you can move powered on virtual machines away from a host to perform maintenance, to balance loads, to collocate virtual machines that communicate with each other, to move virtual machines apart to minimize fault domain, to migrate to new server hardware, and so on.

You can perform several types of migration according to the virtual machine resource type:

- Change compute resource only: Moving a virtual machine but not its storage to another compute resource, such as a host, cluster, resource pool, or vApp.
- Change storage only: Moving a virtual machine and its storage, including virtual disks, configuration files, or a combination of these, to a new datastore on the same host.
- Change both compute resource and storage: Moving a virtual machine to another host and moving its disk or virtual machine folder to another datastore. When you move a virtual machine network between distributed switches, the network configuration and policies that are associated with the network adapters of the virtual machine are transferred to the target switch.

Comparison of Migration Types

Slide 7-38

Migration Type	Virtual Machine Power State	Change Host or Datastore?	Across vCenter Servers?	Shared Storage Required?	CPU Compatibility
Cold	Off	Host or datastore or both	Yes	No	Different CPU families allowed
Suspended	Suspended	Host or datastore or both	Yes	No	Must meet CPU compatibility requirements
vSphere vMotion	On	Host	Yes	Yes	Must meet CPU compatibility requirements
vSphere Storage vMotion	On	Datastore	Yes	No	N/A
Shared-nothing vSphere vMotion	On	Both	Yes	No	Must meet CPU compatibility requirements

A deciding factor behind using a particular migration technique is the purpose of performing the migration. For example, you might need to stop a host for maintenance but keep the virtual machines running. Use vSphere vMotion to migrate the virtual machines instead of performing a cold or suspended virtual machine migration. If you must move a virtual machine's files to another datastore to better balance the disk load or transition to another storage array, use vSphere Storage vMotion.

Some migration techniques, such as vSphere vMotion migration, have special hardware requirements that must be met to function properly. Other techniques, such as a cold migration or a suspended virtual machine migration, do not have special hardware requirements to function properly.

Migration of a suspended virtual machine and migration with vSphere vMotion can be called hot migration because it enables migration of a virtual machine without powering it off.

vSphere vMotion Migration

Slide 7-39

A vSphere vMotion migration moves a powered-on virtual machine from one host to another.

vSphere vMotion provides these capabilities:

- Improves overall hardware use
- Continuous virtual machine operation while accommodating scheduled hardware downtime
- vSphere DRS balancing virtual machines across hosts

vSphere vMotion migrates running virtual machines from one ESXi host to another ESXi host with no disruption or downtime. vSphere vMotion enables vSphere DRS to migrate running virtual machines from one host to another to balance the load.

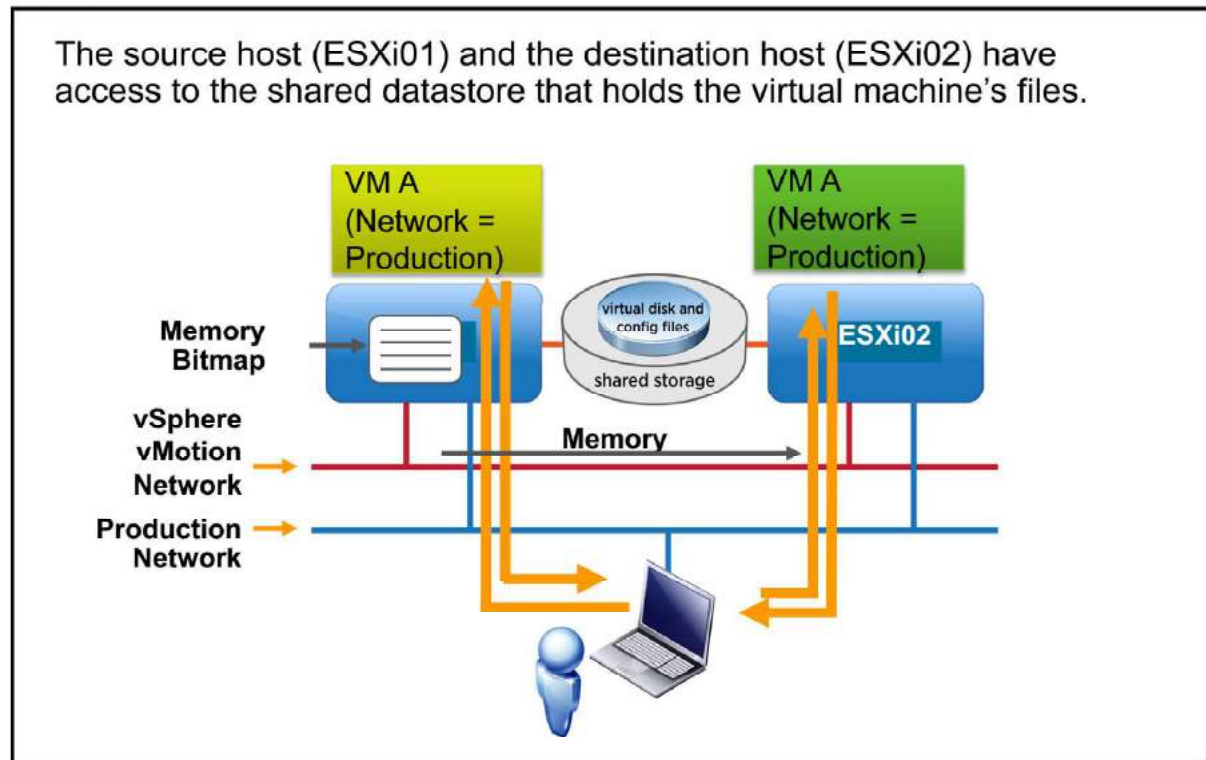
With vSphere vMotion, the entire state of the virtual machine is moved from one host to another while the data storage remains in the same datastore.

The state information includes the current memory content and all the information that defines and identifies the virtual machine. The memory content includes transaction data and whatever bits of the operating system and applications are in memory. The definition and identification information stored in the state includes all the data that maps to the virtual machine hardware elements, including:

- BIOS
- Devices
- CPU
- MAC addresses for the Ethernet cards

vSphere vMotion Migration Workflow

Slide 7-40



A vSphere vMotion migration consists of the following steps:

1. Shadow VM is created on the destination host.
2. The virtual machine's memory state is copied over the vSphere vMotion network from the source host to the target host through the vSphere vMotion network. Users continue to access the virtual machine and, potentially, update pages in memory. A list of modified pages in memory is kept in a memory bitmap on the source host.
3. After completing the first pass of memory state copy, another pass of memory copy is performed to copy any pages that changed during the last iteration. Continue this iterative memory copying until no changed pages remain.
4. After most of the virtual machine's memory is copied from the source host to the target host, the virtual machine is quiesced. No additional activity occurs on the virtual machine. In the quiesce period, vSphere vMotion transfers the virtual machine device state and memory bitmap to the destination host.

5. Immediately after the virtual machine is quiesced on the source host, the virtual machine is initialized and starts running on the target host. A Gratuitous Address Resolution Protocol (GARP) request notifies the subnet that virtual machine A's MAC address is now on a new switch port.
6. Users access the virtual machine on the target host instead of the source host.
7. The memory pages that the virtual machine was using on the source host are marked as free.

vSphere vMotion Migration Requirements

Slide 7-41

A virtual machine must meet the following requirements:

- By default, migrating a virtual machine with vSphere vMotion produces an error if the virtual machine has an active connection to an internal virtual switch.
- The virtual machine must not have a connection to a virtual device, such as a CD/DVD or floppy drive, with a local image mounted.
- The virtual machine must not have CPU affinity configured.
- If the virtual machine's swap file is not accessible to the destination host, vSphere vMotion must be able to create a swap file that is accessible to the destination host before migration can begin.
- If a virtual machine uses an RDM, the RDM and the physical disk to which it maps must be accessible by the destination host.

The vSphere vMotion migration produces an error in certain conditions. When an error is encountered, the migration does not proceed until you fix the error.

By default, migrating a virtual machine with vSphere vMotion produces an error if the virtual machine has an active connection to an internal virtual switch (that is, a switch that is not directly connected to a physical NIC). The intent of this restriction was to ensure server availability through a live migration, and avoid problems where the target ESX host might be misconfigured without an active connection to the physical network. However, this restriction does not take into account when a bridged-mode virtual machine is being used to filter traffic between the workload and the physical network. You may enable vSphere vMotion when using network security or service virtual machines in bridged mode. For information, see VMware knowledge base article 1006701 at <http://kb.vmware.com/kb/1006701>.

You cannot use vSphere vMotion to migrate a virtual machine that uses a virtual device that is backed by a device on the ESXi host. Disconnect these devices before migrating the virtual machine. vSphere vMotion also produces warnings in certain conditions, for example, when a virtual machine is configured to access a local CD-ROM drive or floppy image but is not connected to it. The vSphere vMotion migration still proceeds even if warnings have not been addressed.

For the complete list of vSphere vMotion migration requirements, see *vCenter Server and Host Management Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Host Requirements for vSphere vMotion Migration

Slide 7-42

Source and destination hosts must have these characteristics:

- Accessibility to all storage (Fibre Channel, iSCSI, or NAS) that is used by the virtual machine:
 - 128 concurrent vSphere vMotion migrations per VMFS or NFS datastore
- At least a 1 Gigabit Ethernet (1GigE) network:
 - Four concurrent vSphere vMotion migrations on a 1 Gbps network
 - Eight concurrent vSphere vMotion migrations on a 10 Gbps (or faster) network
- Compatible CPUs:
 - CPU feature sets of both the source and destination hosts must be compatible.
 - Some features can be hidden by using Enhanced vMotion Compatibility (EVC) or compatibility masks.

Each operation, such as a migration with vSphere vMotion or cloning a virtual machine, is assigned a resource cost. Each host, datastore, or network resource, has a maximum cost that it can support at any one time. Any new migration or provisioning operation that causes a resource to exceed its maximum cost does not proceed immediately, but is queued until other operations complete and release resources. Each of the network, datastore, and host limits must be satisfied for the operation to proceed.

For the complete list of limits on simultaneous migrations, see *vCenter Server and Host Management Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

The source and destination host must meet certain requirements for a vSphere vMotion migration to be successful:

- SAN visibility of virtual disks
- Gigabit Ethernet (or greater) interconnection
- Consistent network configuration, both physical and virtual
- Source and destination server CPUs from the same compatibility group

CPU Constraints on vSphere vMotion Migration

Slide 7-43

CPU Characteristics	Exact Match Required	Reason
Clock speeds, cache sizes, hyperthreading, and number of cores	N/A	Virtualized away by the VMkernel.
Manufacturer (Intel or AMD) family and generation (Opteron4, Intel Westmere)	Applicable	Instruction sets contain many small differences.
Presence or absence of SSE3, SSSE3, or SSE4.1 instructions	Applicable	Multimedia instructions usable directly by applications.
Virtualization hardware assist	For 32-bit VMs: N/A	Virtualized away by the VMkernel.
	For 64-bit VMs on Intel: Applicable	Intel 64-bit with VMware implementation uses Intel VT.
Execution-disable (NX/XD bit)	Applicable but customizable	Guest operating system relies on NX/XD bit if detected.

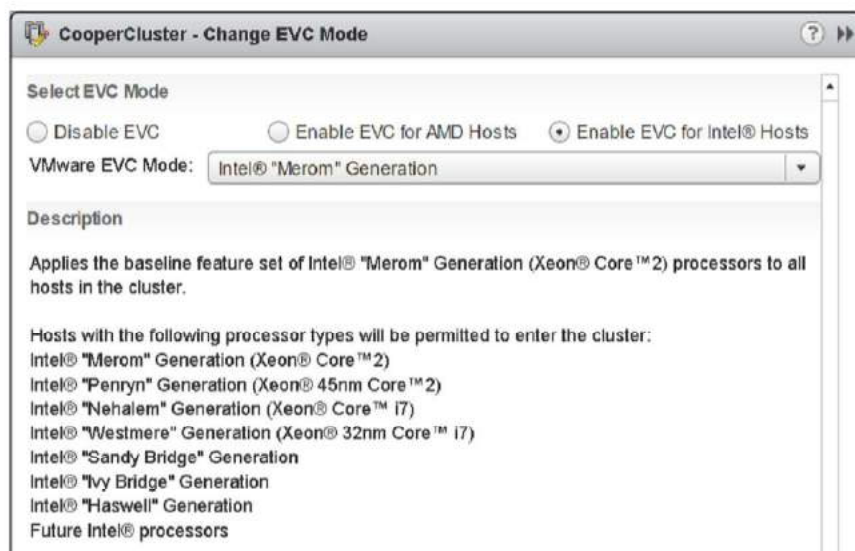
CPU compatibility between the source host and the target host is a vSphere vMotion requirement that must be met.

For example, if hyperthreading is enabled on the source host and disabled on the destination host, the vSphere vMotion migration continues because the VMkernel handles this difference in characteristics. But if the source host processor supports SSE4.1 instructions and the destination host processor does not support them, the hosts are considered incompatible and the vSphere vMotion migration fails. SSE4.1 instructions are application-level instructions that bypass the virtualization layer and might cause application instability if mismatched after a migration with vSphere vMotion.

Other Cluster Settings: EVC for vSphere vMotion

Slide 7-44

EVC is a cluster feature that prevents vSphere vMotion migrations from failing because of incompatible CPUs.

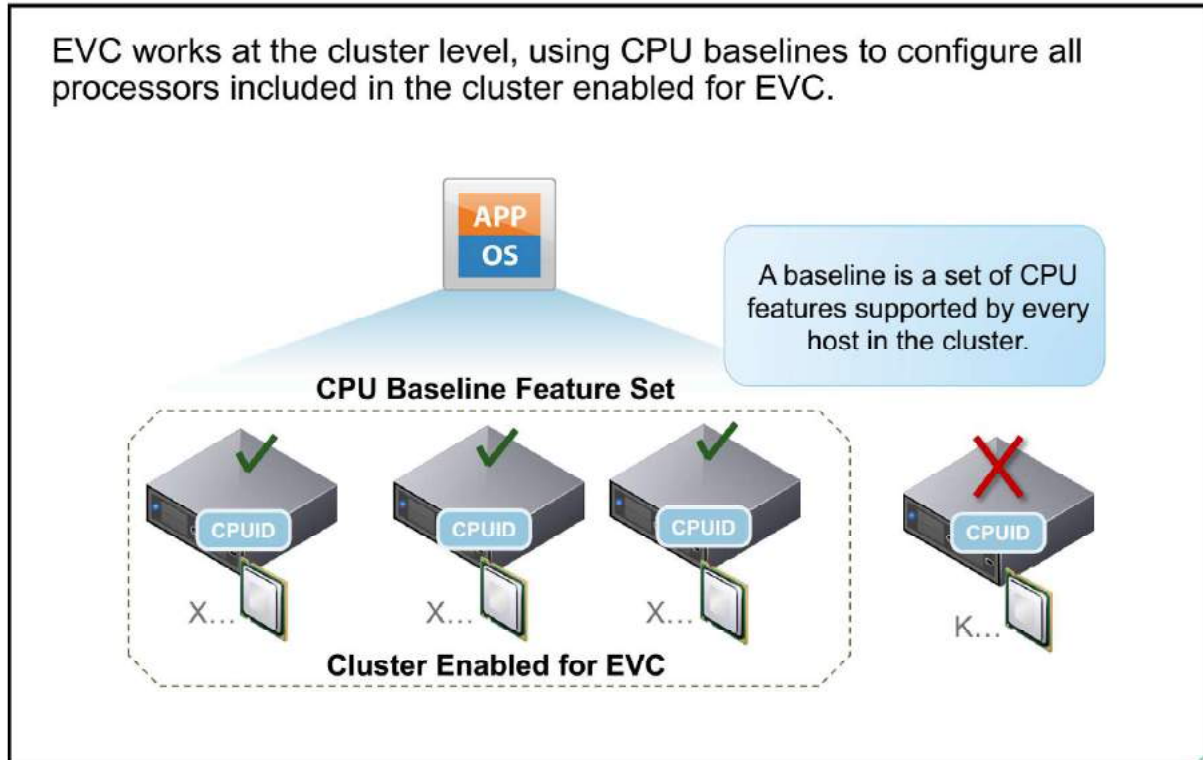


EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. Presenting the same CPU feature set prevents migrations with vSphere vMotion from failing because of incompatible CPUs.

For information about EVC processor support, see VMware knowledge base article 1003212 at <http://kb.vmware.com/kb/1003212>.

CPU Baselines for an EVC Cluster

Slide 7-45



EVC facilitates safe vSphere Motion migration across a range of CPU generations. With EVC, you can use vSphere vMotion to migrate virtual machines among CPUs that would otherwise be considered incompatible.

EVC allows vCenter Server to enforce vSphere vMotion compatibility among all hosts in a cluster by forcing hosts to expose a common set of CPU features (baseline) to virtual machines. A baseline is a set of CPU features that are supported by every host in the cluster. When you configure EVC, you set all host processors in the cluster to present the features of a baseline processor. After they have been enabled for a cluster, hosts that are added to the cluster are automatically configured to the CPU baseline.

Hosts that cannot be configured to the baseline are not permitted to join the cluster. Virtual machines in the cluster always see an identical CPU feature set, no matter which host they happen to run on. Because this process is automatic, EVC is easy to use and requires no specialized knowledge of CPU features and masks.

EVC Cluster Requirements

Slide 7-46

All hosts in the cluster must meet the following requirements:

- Use CPUs from a single vendor, either Intel or AMD:
 - Use Intel CPUs with Merom microarchitecture and later.
 - Use AMD first-generation Opteron CPUs and later.
- Be enabled for hardware virtualization: AMD-V or Intel VT
- Be enabled for execution-disable technology: AMD No eXecute (NX) or Intel eXecute Disable (XD)
- Be configured for vSphere vMotion migration

Applications in virtual machines must be CPU ID compatible.

Before you create an EVC cluster, ensure that the hosts you intend to add to the cluster meet the requirements.

EVC automatically configures hosts whose CPUs feature Intel FlexMigration and AMD-V Extended Migration technologies to be compatible with vSphere vMotion with hosts that use older CPUs. EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the real CPUs on the hosts differ.

You can use one of the following methods to create an EVC cluster:

- Create an empty cluster with EVC enabled and then move hosts into the cluster.
- Enable EVC on an existing cluster.

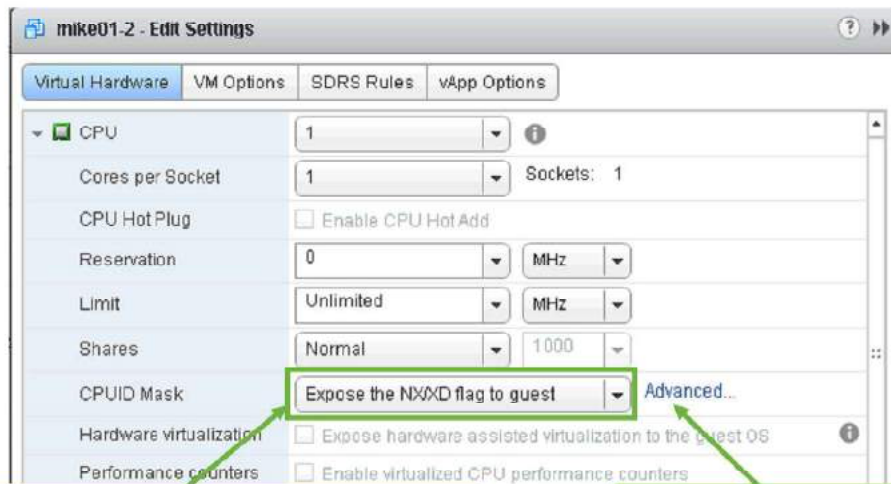
For EVC to function properly, the applications on the virtual machines must be written to use the CPUID machine instruction to discover CPU features as recommended by the CPU vendors. vSphere cannot support EVC with applications that do not follow the CPU vendor recommendations to discover CPU features.

To determine the EVC modes compatible with your CPU, search the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php?deviceCategory=server>. Search for the server model or CPU family, and click the entry in the CPU Series column to display the compatible EVC modes.

Hiding or Exposing NX/XD

Slide 7-47

AMD No Execute (NX) and Intel Execute Disable (XD) technologies mark memory pages as data-only to prevent malicious software exploits and buffer overflow attacks.



Choose between NX/XD security features and broadest vSphere vMotion compatibility.

For future CPU features, edit mask at the bit level.

If NX/XD technology is exposed on the source host, then it must be exposed on the destination host. NX/XD technology is exposed by default for all guest operating systems that can use it (trading off some compatibility for security by default).

Hiding the NX/XD flag increases vSphere vMotion compatibility between hosts, at the cost of disabling certain CPU security features for some guest operating systems and applications.

Identifying CPU Characteristics

Slide 7-48

To identify CPU characteristics, use the server and CPU specifications or use the VMware CPU identification utility.

```
Random_Init: Using random seed: 2044292605 (0x79d96dfd)
Reporting CPUID for 2 logical CPUs...

All CPUs are identical

Family: 06 Model: 17 Stepping: 6

ID1ECX      ID1EDX      ID81ECX     ID81EDX
0x00082201  0x0febfbff  0x00000001  0x20100000

Vendor      : Intel
Brand String : "Intel(R) Xeon(R) CPU           X5482  @ 3.20GHz"
SSE Support : SSE1, SSE2, SSE3, SSSE3, SSE4.1
Supports NX / XD : Yes
Supports CMPXCHG16B : Yes
Supports RDTSCP : No
Hyperthreading : No
Supports Flex Migration : Yes
Supports 64-bit Longmode : Yes
Supports 64-bit VMware : No
Supported EVC modes : None

PASS: Test 56983: CPUID
Press any key to reboot.
```

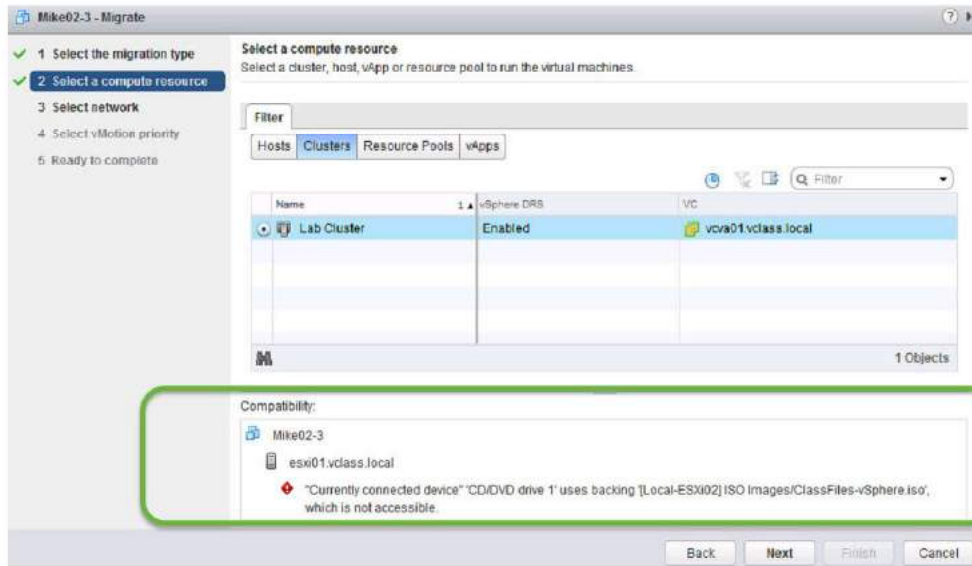
The server hardware's CPU specifications usually indicate whether the CPUs contain the features that affect vSphere vMotion compatibility. If the specifications of a server or its CPU features are unknown, you can display these features with the VMware CPU identification utility. You use this utility to boot a server and determine whether its CPUs contain features like SSE3, SSSE3, and NX/XD.

To download the VMware CPU Identification Utility for ESXi hosts, go to https://my.vmware.com/web/vmware/details/cpu_identification_utility/ZHcqYmR0dGhidGR3.

Checking vSphere vMotion Errors

Slide 7-49

When you select the host and cluster, a validation check is performed to verify that most vSphere vMotion requirements were met.



If validation succeeds, then you can continue in the wizard. If validation does not succeed, a list of vSphere vMotion errors and warnings is displayed in the Compatibility pane.

Warnings have yellow triangles. Errors have red diamonds. Warnings allow you to perform a vSphere vMotion migration. Errors do not allow you to continue. You must exit the wizard and fix all errors before retrying the migration.


If a failure occurs during the vSphere vMotion migration, the virtual machine is not migrated and continues to run on the source host.

vSphere Storage vMotion in Action

Slide 7-50

vSphere Storage vMotion uses an I/O mirroring architecture to copy disk blocks between source and destination:

1. Initiate storage migration.
2. Use the VMkernel data mover or vSphere Storage APIs - Array Integration to copy data.
3. Start a new virtual machine process.
4. Mirror I/O calls to file blocks that are already copied to virtual disk on the destination datastore.
5. Cut over to the destination virtual machine process to begin accessing the virtual disk copy.



The storage migration process does a single pass of the disk, copying all the blocks to the destination disk. If blocks are changed after they are copied, the blocks are synchronized from the source to the destination through the mirror driver, with no need for recursive passes. This capability results in a much shorter vSphere Storage vMotion operation because it can complete a migration in a single pass.

vSphere Storage vMotion uses a mirroring architecture, which mirrors the changed disk blocks after they have been copied to the destination. Comparing to older methods, the mirroring method produces more predictable results, shorter migration times, and fewer I/O operations. This method also guarantees migration success even when using a slow disk on the destination.

This approach guarantees complete transactional integrity and is fast enough to be unnoticeable to the end user. The mirror driver uses the VMkernel data mover to copy blocks of data from the source disk to the destination disk. The mirror driver synchronously mirrors writes to both disks during the vSphere Storage vMotion operation.

Finally, vSphere Storage vMotion operations are performed either internally on a single ESXi host or offloaded to the storage array. Operations performed internally on the ESXi host use a data mover built in to the VMkernel. Operations are offloaded to the storage array if the array supports vSphere Storage APIs - Array Integration, also called hardware acceleration.

vSphere Storage vMotion Guidelines and Limitations

Slide 7-51

Guidelines:

- Plan the migration and coordinate with administrators.
- Perform migrations during off-peak hours.
- Ensure that the host has access to source datastores and target datastores.

Limitations:

- Virtual machine disks must be in persistent mode or be RDMs.

A virtual machine and its host must meet certain resource and configuration requirements for the virtual machine disks (VMDKs) to be migrated with vSphere Storage vMotion. One of the requirements is that the host on which the virtual machine is running must have access both to the source datastore and to the target datastore.

During a migration with vSphere Storage vMotion, you can change the disk provisioning type. Migration with vSphere 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 does not succeed.

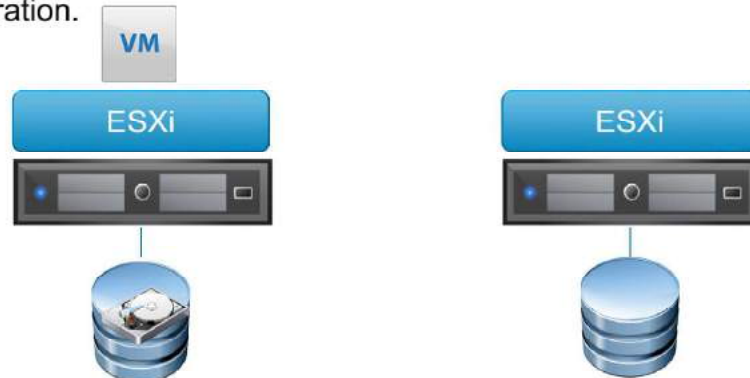
vSphere Storage vMotion is subject to the following limitations: VMDKs must be in persistent mode or be RDMs. For virtual compatibility mode RDMs, you can migrate the mapping file or convert to thick-provisioned or thin-provisioned disks during migration if the destination is not an NFS datastore. For physical compatibility mode RDMs, you can migrate only the mapping file.

Shared-Nothing vSphere vMotion Migration

Slide 7-52

Shared-nothing vSphere vMotion migration enables a virtual machine to change its host, datastores, networks, and vCenter Server instances simultaneously, even if the two hosts do not have a shared storage.

- This technique combines vSphere vMotion and vSphere Storage vMotion into a single operation.



You can migrate virtual machines between hosts, within and across clusters, data centers, and vCenter Server instances, beyond storage accessibility boundaries.

You can use vSphere vMotion to migrate virtual machines to a different compute resource and storage simultaneously. You can migrate virtual machines across vCenter Server instances. Unlike vSphere Storage vMotion, which requires a single host to have access to both the source and destination datastore, you can migrate virtual machines across storage accessibility boundaries.

vSphere vMotion does not require environments with shared storage. This feature is useful for performing cross-cluster migrations, when the target cluster machines might not have access to the source cluster's storage. Processes on the virtual machine continue to run during the migration with vSphere vMotion.

Shared-nothing vSphere vMotion migration is useful for virtual infrastructure administration tasks, including:

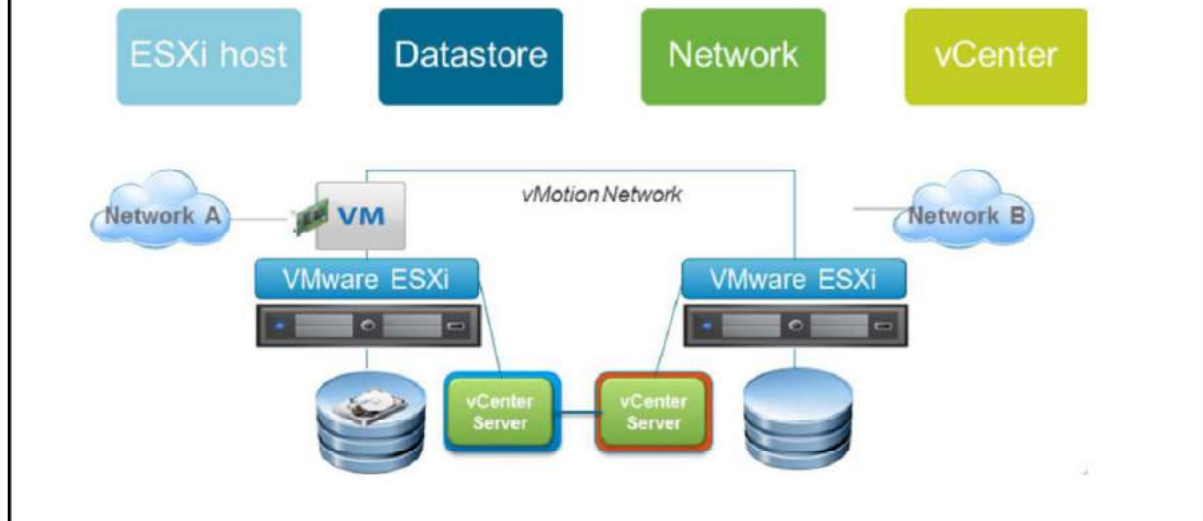
- **Host maintenance:** You can move virtual machines from a host to allow maintenance of the host.
- **Storage maintenance and reconfiguration:** You can move virtual machines from a storage device to allow maintenance or reconfiguration of the storage device without virtual machine downtime.
- **Storage load redistribution:** You can manually redistribute virtual machines or virtual disks to different storage volumes to balance capacity or improve performance.

For information about requirements and limitations for shared-nothing vSphere vMotion migration, see *vCenter Server and Host Management Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Shared-Nothing vSphere vMotion Migration Considerations

Slide 7-53

In vSphere 6, multiple changes can occur simultaneously with shared-nothing vSphere vMotion migrations.



Shared-nothing vSphere vMotion migration counts against the concurrent limitations of both vSphere vMotion and vSphere Storage vMotion. No more than two concurrent shared-nothing vSphere vMotion migrations are allowed. Because these instances count against vSphere Storage vMotion limits, running two concurrent shared-nothing vSphere vMotion migrations causes all attempted vSphere Storage vMotion instances to remain queued until one of the active shared-nothing vSphere vMotion migrations is completed. Similarly, shared-nothing vSphere vMotion migration instances also count against vSphere vMotion limits, at most eight concurrent vSphere vMotion instances per host. If two shared-nothing vSphere vMotion migration instances are active, then at most only six concurrent vSphere vMotion instances are allowed at the same time. If eight vSphere vMotion instances are active, new shared-nothing vSphere vMotion migration attempts are queued until one of the active vSphere vMotion instances is completed.

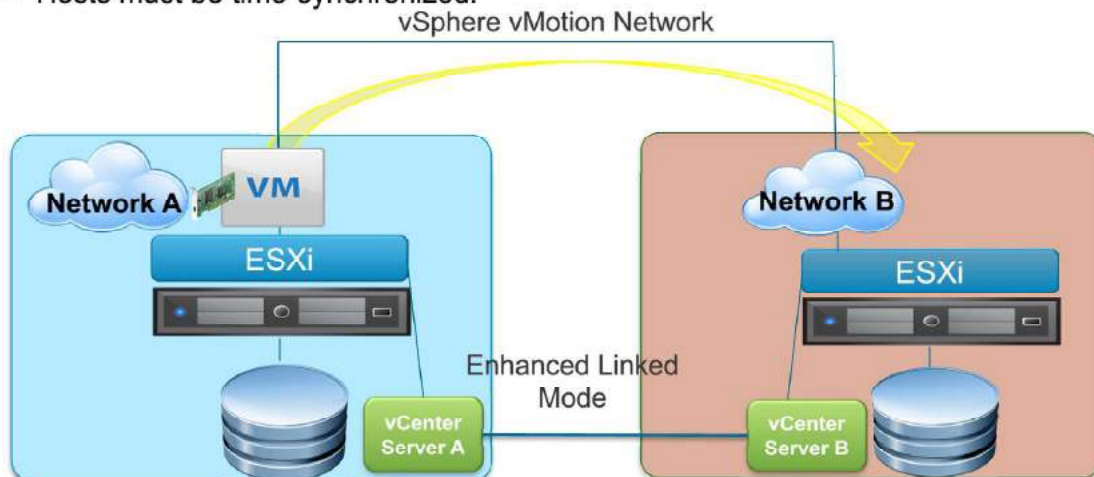
Shared-nothing vSphere vMotion migration behaves exactly like vSphere vMotion with respect to supporting multi-NICs. Likewise, shared-nothing vSphere vMotion migration supports either shared swap or unshared swap migrations just as vSphere vMotion does, with virtual machine home directory movement becoming an unshared swap migration. Shared-nothing vSphere vMotion migration instances are more expensive, which must be factored in when making migration decisions. Neither vSphere DRS nor vSphere Storage DRS uses shared-nothing vSphere vMotion migration technology. Although neither vSphere DRS nor vSphere Storage DRS recommends shared-nothing vSphere vMotion migrations, users can perform manual shared-nothing vSphere vMotion migrations in or across vSphere Storage DRS or vSphere DRS clusters.

Migration Between vCenter Server Instances

Slide 7-54

In vSphere 6, vSphere vMotion can migrate virtual machines between linked vCenter Server instances. This type of migration requires:

- ESXi hosts and vCenter Server systems must be upgraded to vSphere 6.
- vCenter Server instances must be in Enhanced Linked Mode.
- Hosts must be time-synchronized.



VMkernel Networking Layer and TCP/IP Stacks

Slide 7-55

The VMkernel networking layer provides connectivity to hosts and handles the standard system traffic of vSphere vMotion, IP storage, vSphere Fault Tolerance, vSAN, and others.

You can also create VMkernel adapters on the source and target vSphere Replication hosts to isolate the replication data traffic.

TCP/IP stacks at the VMkernel level:

- Default TCP/IP stack
- vSphere vMotion TCP/IP stack
- Provisioning TCP/IP stack
- Custom TCP/IP stacks

Consider the following key points about TCP/IP stacks at the VMkernel level:

- **Default TCP/IP stack:** Provides networking support for the management traffic between vCenter Server and ESXi hosts and for system traffic such as vSphere vMotion, IP storage, and vSphere Fault Tolerance.
- **vMotion TCP/IP stack:** Supports the traffic for live migration of virtual machines. Use the vMotion TCP/IP stack to provide better isolation for the vSphere vMotion traffic. After you create a VMkernel adapter on the vMotion TCP/IP stack, you can use only this stack for vSphere vMotion migration on this host. The VMkernel adapters on the default TCP/IP stack are disabled for the vSphere vMotion service. If a live migration uses the default TCP/IP stack while you configure VMkernel adapters with the vMotion TCP/IP stack, the migration completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are disabled for future vSphere vMotion sessions.
- **Provisioning TCP/IP stack:** Supports the traffic for virtual machine cold migration, cloning, and snapshot creation. You can use the provisioning TCP/IP stack to handle NFC traffic during long-distance vSphere vMotion migration. VMkernel adapters configured with the provisioning TCP/IP stack handle the traffic from cloning the virtual disks of the migrated virtual machines in long-distance vSphere vMotion. By using the provisioning TCP/IP stack, you can isolate the

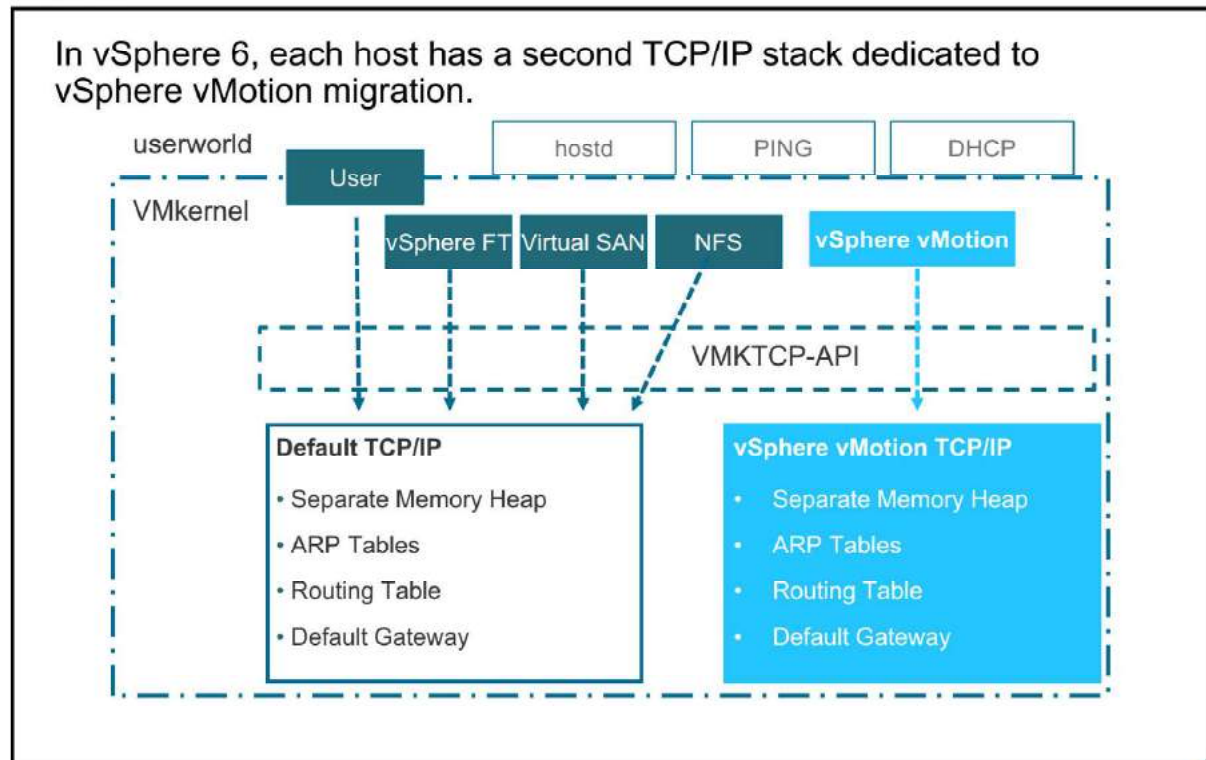
traffic from the cloning operations on a separate gateway. After you configure a VMkernel adapter with the provisioning TCP/IP stack, all adapters on the default TCP/IP stack are disabled for the provisioning traffic.

- Custom TCP/IP stacks: You can add custom TCP/IP stacks at the VMkernel level to handle networking traffic of custom applications.

Take appropriate security measures to prevent unauthorized access to the management and system traffic in your vSphere environment. For example, isolate the vSphere vMotion traffic in a separate network that includes only the ESXi hosts that participate in the migration. Isolate the management traffic in a network that only network and security administrators can access.

vSphere vMotion TCP/IP Stacks

Slide 7-56



vSphere vMotion TCP/IP stacks support the traffic for live migration of virtual machines. Use the vSphere vMotion TCP/IP stack to provide better isolation for the vSphere vMotion traffic. After you create a VMkernel adapter on the vSphere vMotion TCP/IP stack, you can use only this stack for vSphere vMotion migration on this host. The VMkernel adapters on the default TCP/IP stack are disabled for the vSphere vMotion service. If a live migration uses the default TCP/IP stack while you configure VMkernel adapters with the vMotion TCP/IP stack, the migration completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are disabled for future vSphere vMotion sessions.


Long-Distance vSphere vMotion Migration

Slide 7-57

Long-distance vSphere vMotion migration is an extension of vSphere vMotion migration across vCenter Server instances. This migration is targeted at environments where vCenter Server systems are spread across large geographic distances and where the latency across sites is high.

Use cases for long-distance vSphere vMotion migration:

- Permanent migrations
- Disaster avoidance
- VMware Site Recovery Manager™ disaster avoidance testing
- Multisite load balancing
- Follow-the-Sun scenario support



Networking Requirements for Long-Distance vSphere vMotion Migration

Slide 7-58

vSphere vMotion migrations between vCenter Server instances must connect over layer 3 connections:

- Virtual machine network:
 - L2 connection.
 - Same virtual machine IP address available at destination.
 - The round-trip time between the hosts can be up to 150 milliseconds.
- vSphere vMotion network:
 - L3 connection.
 - Secure (dedicated or encrypted).
 - 250 Mbps per vSphere vMotion operation.

Network Checks for Migrations Between vCenter Server Instances

Slide 7-59

vCenter Server performs several network compatibility checks to prevent the following configuration problems:

- MAC address compatibility on the destination host
- vSphere vMotion migration from a distributed switch to a standard switch
- vSphere vMotion migration between distributed switches of different versions
- vSphere vMotion migration to an internal network, for example, a network without a physical NIC

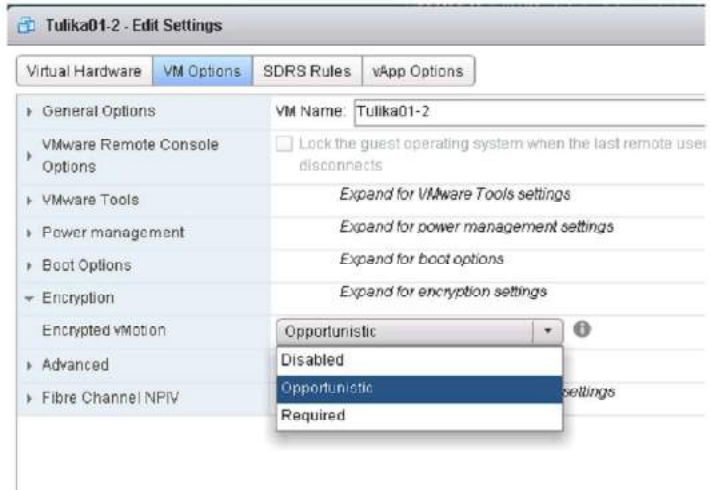
Encrypted vSphere vMotion

Slide 7-60

vSphere vMotion always uses encryption when migrating encrypted virtual machines.

- For virtual machines that are not encrypted, select one of the following encrypted vSphere vMotion options:

- Disabled
- Opportunistic: Encrypted vSphere vMotion is used if source and destination hosts support it.
- Required: If the source or destination host does not support encrypted vSphere vMotion, migration with vSphere vMotion fails.



Encrypted vSphere vMotion secures confidentiality, integrity, and authenticity of data that is transferred with vSphere vMotion. Encrypted vSphere vMotion supports all variants of vSphere vMotion, including migration across vCenter Server systems. Encrypted vSphere Storage vMotion is not supported.

You cannot turn off encrypted vSphere vMotion for encrypted virtual machines.

The Opportunistic state is default is for virtual machines that are not encrypted.

Lab 13: Migrating Virtual Machines

Slide 7-61

Use vSphere vMotion and vSphere Storage vMotion to migrate virtual machines

1. Migrate Virtual Machine Files from the Local Storage to the Shared Storage
2. Create a Virtual Switch and a VMkernel Port Group for vSphere vMotion Migration
3. Perform a vSphere vMotion Migration of a Virtual Machine on a Shared Datastore
4. Perform a Compute Resource and Storage Migration

Review of Learner Objectives


Slide 7-62

You should be able to meet the following objectives:

- Verify vSphere vMotion requirements, including CPU constraints and guidelines
- Perform a vSphere vMotion migration
- Perform a vSphere Storage vMotion migration
- Perform a shared-nothing vSphere vMotion migration
- Describe the major enhancements to vSphere vMotion in vSphere 6

Lesson 5: Creating Virtual Machine Snapshots

Slide 7-63



Lesson 5: Creating Virtual Machine Snapshots

Learner Objectives

Slide 7-64

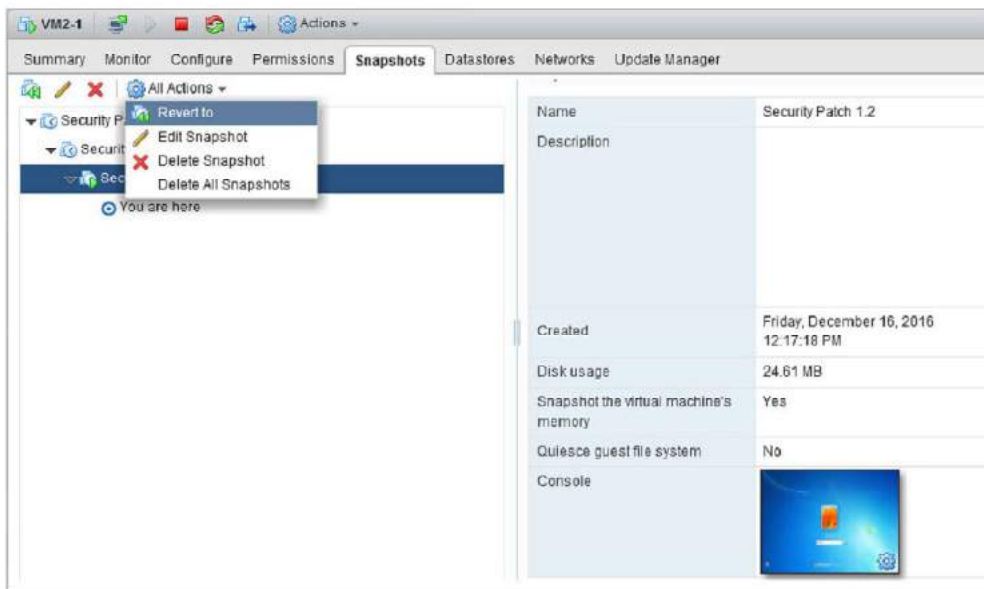
By the end of this lesson, you should be able to meet the following objectives:

- Take a snapshot of a virtual machine and manage multiple snapshots
- Delete virtual machine snapshots
- Consolidate snapshots

Virtual Machine Snapshots

Slide 7-65

Snapshots enable you to preserve the state of the virtual machine so that you can repeatedly return to the same state.



Snapshots are useful when you want to revert repeatedly to the same state but do not want to create multiple virtual machines. Examples include patching or upgrading the guest operating system in a virtual machine. Snapshots give you the ability to back out of the patch or upgrade process if problems occur during patching or upgrading.

The relationship between snapshots is like the relationship between a parent and a child. Snapshots are organized in a snapshot tree. In a snapshot tree, each snapshot has one parent and one or more children, except for the last snapshot, which has no children.

On the slide, the snapshot named Security patch 1.0 has two child snapshots, each named Security Patch 1.1 and 1.2.

A virtual machine snapshot mainly includes the following:


- Settings state: The virtual machine's settings (.nvram and .vmx) and power state
- Disk state: State of the virtual machine's associated disks
- Memory state: Contents of the virtual machine's memory (optional)

Virtual Machine Snapshot Files

Slide 7-66

A snapshot consists of a set of files: the memory state file (`.vmsn`), the description file (`-00000#.vmdk`), and the delta file (`-00000#-delta.vmdk`).

The snapshot list file (`.vmsd`) keeps track of the virtual machine's snapshots.



A virtual machine can have one or more snapshots. Each snapshot consists of the following:

- Delta disk: When you take a virtual machine snapshot, the state of the virtual disk at the time the snapshot is taken is preserved. When this preservation occurs, the guest operating system cannot write to its `.vmdk` file. Instead, changes are captured in an alternate file named `vmname-number-delta.vmdk`.
- Memory state file: `VM_name-Snapshot#.vmsn`, where # is the next number in the sequence, starting with 1. This file holds the active memory state of the virtual machine when the snapshot was taken. If memory is captured, the size of this file is the size of the virtual machine's maximum memory. If memory is not captured, the file is much smaller.
- Disk descriptor file: `VM_name-00000#.vmdk`. This file is a small text file that contains information about the snapshot.
- Snapshot delta file: `VM_name-00000#-delta.vmdk`. This file contains the changes to the virtual disk's data since the snapshot was taken.
- Snapshot active memory file: `VM_name-Snapshot#.vmem`. This file contains the contents of the virtual machine memory if the option to include memory is selected when creating the snapshot.

- *VM_name.vmsd* is the snapshot list file and is created at the time that the virtual machine is created. It maintains snapshot information for a virtual machine so that it can create a snapshot list in vSphere Web Client. This information includes the name of the snapshot *.vmsn* file and the name of the virtual disk file.
- The snapshot state file has a *.vmsn* extension and is used to store the state of a virtual machine when a snapshot is taken. A new *.vmsn* file is created for every snapshot that is created on a virtual machine and is deleted when the snapshot is deleted. The size of this file varies, based on the options selected when the snapshot is created. For example, including the memory state of the virtual machine in the snapshot increases the size of the *.vmsn* file.

You can exclude one or more of the VMDKs from a snapshot by designating a virtual disk in the virtual machine as an independent disk. Placing a virtual disk in independent mode is typically done when the virtual disk is created. If the virtual disk was created without enabling independent mode, you must power off the virtual machine to enable it.

Other files might also exist, depending on the virtual machine hardware version. For example, each snapshot of a virtual machine that is powered on has an associated *snapshot_name_number.vmem* file, which contains the guest operating system main memory, saved as part of the snapshot.


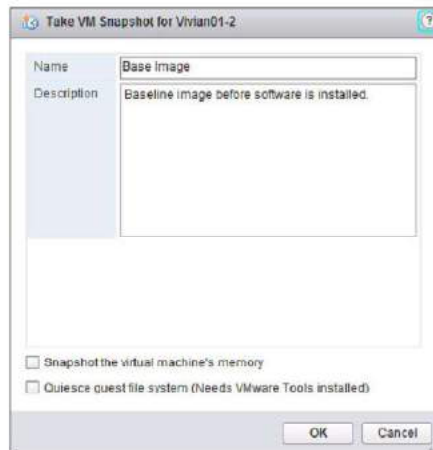
Taking a Snapshot

Slide 7-67

You can take a snapshot while a virtual machine is powered on, powered off, or suspended.

A snapshot captures the state of the virtual machine: memory state, settings state, and disk state.

Virtual machine snapshots are not recommended as a virtual machine backup strategy.



A snapshot captures the entire state of the virtual machine at the time that you take the snapshot, including:

- **Memory state:** The contents of the virtual machine's memory. The memory state is captured only if the virtual machine is powered on and if you select the **Snapshot the virtual machine's memory** check box (selected by default).
- **Settings state:** The virtual machine settings.
- **Disk state:** The state of all the virtual machine's virtual disks.

At the time that you take the snapshot, you can also quiesce the guest operating system. This action quiesces the file system of the guest operating system.

Snapshots of physical compatibility mode RDM disks are not supported.

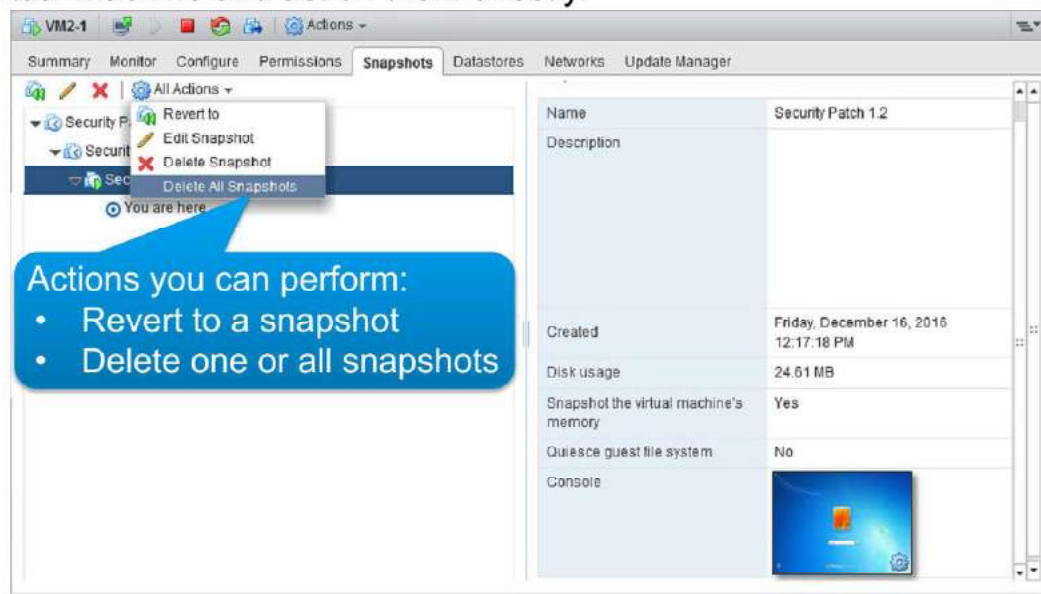
CAUTION

Virtual machine snapshots are not recommended as a virtual machine backup strategy.

Managing Snapshots

Slide 7-68

The **Snapshot** tab enables you to review all snapshots for the active virtual machine and act on them directly.



In the **Snapshot** tab, you can perform the following actions:

- **Edit:** Edit the snapshot name and description.
- **Delete:** Removes the snapshot from the Snapshot Manager and consolidates the snapshot files to the parent snapshot disk and merge with the virtual machine base disk.
- **Delete All:** Commits all the intermediate snapshots before the current-state icon (You Are Here) to the Tulika01-03 VMDK file and removes all snapshots for that virtual machine.
- **Revert to:** Enables you to restore, or revert to, a particular snapshot. The snapshot that you restore becomes the current snapshot.

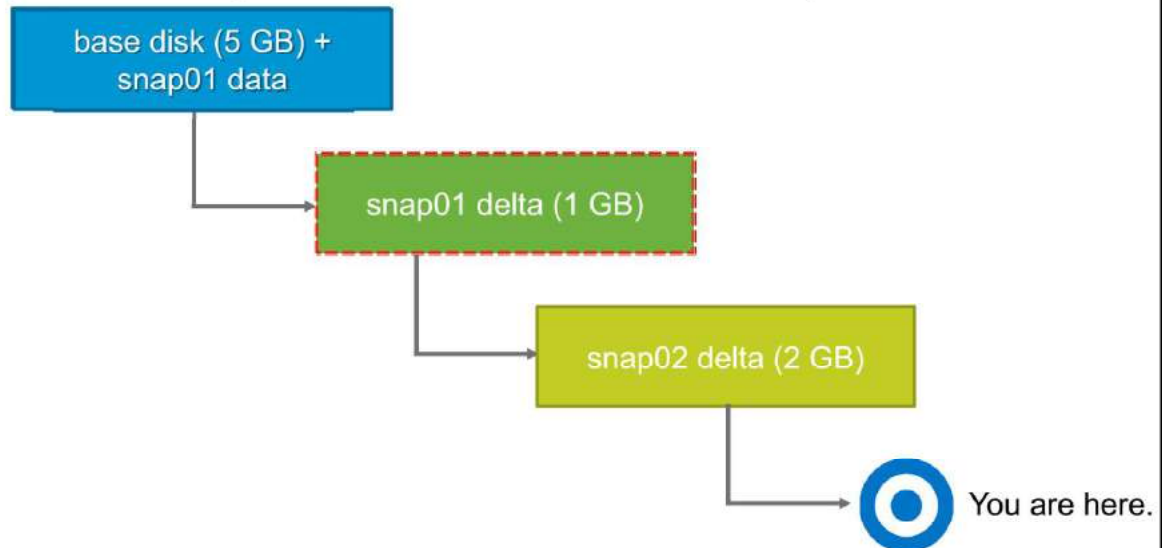
When you revert to a snapshot, you return all these items to the state that they were in at the time that you took the snapshot. If you want the virtual machine to be suspended, powered on, or powered off when you start it, ensure that the virtual machine is in the correct state when you take the snapshot.

Deleting a snapshot (with **Delete** or **Delete All**) consolidates the changes between snapshots and previous disk states. Deleting a snapshot also writes to the parent disk all data from the delta disk that contains the information about the deleted snapshot. When you delete the base parent snapshot, all changes merge with the base VMDK.

Deleting a Virtual Machine Snapshot (1)

Slide 7-69

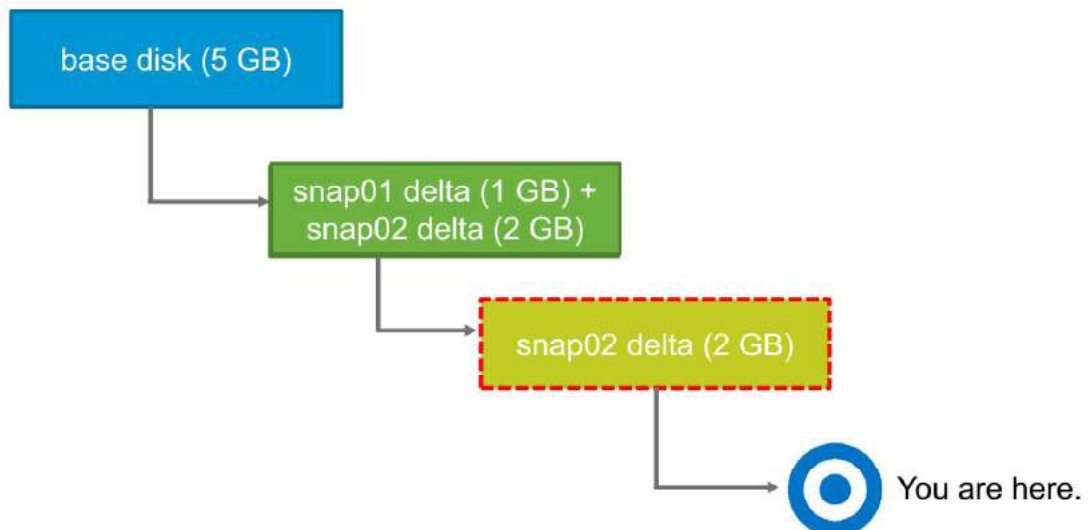
If you delete a snapshot one or more levels above You Are Here, the snapshot state is deleted. The snap01 data is committed into the previous state (base disk) and the foundation for snap02 is retained.



Deleting a Virtual Machine Snapshot (2)

Slide 7-70

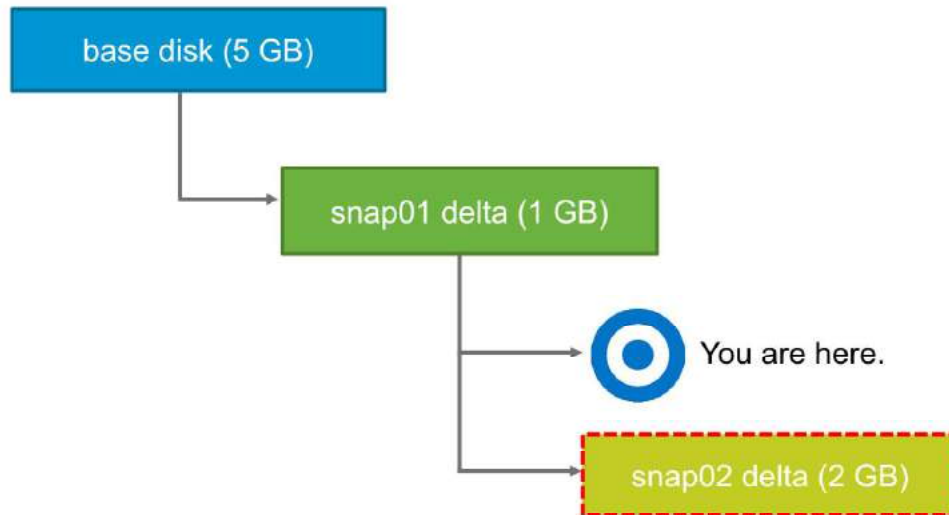
If you delete the current snapshot, the changes are committed to its parent. The snap02 data is committed into snap01 data, and the snap02 -delta.vmdk file is deleted.



Deleting a Virtual Machine Snapshot (3)

Slide 7-71

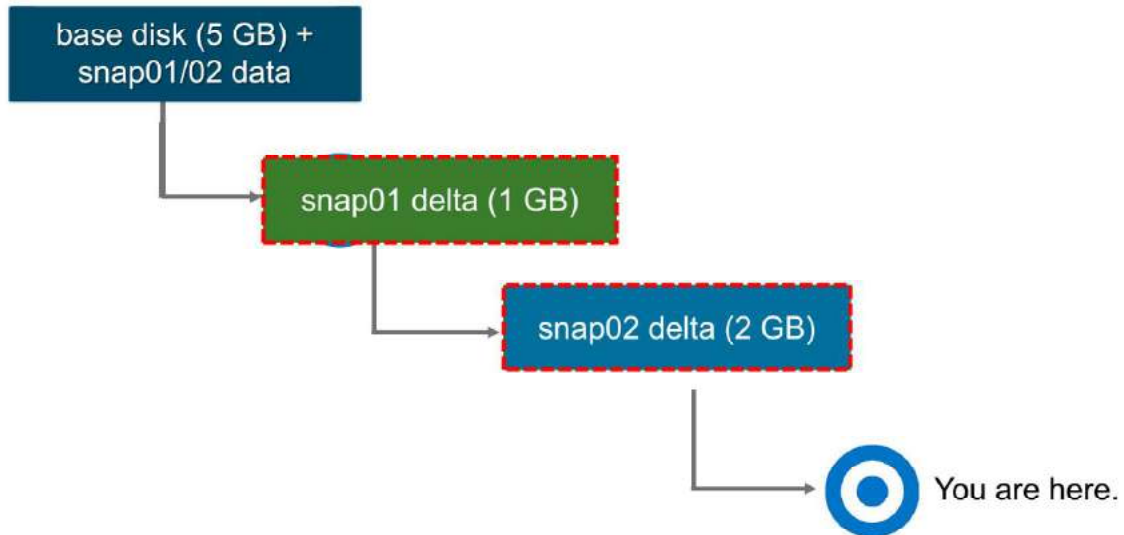
If you delete a snapshot one or more levels below You Are Here, subsequent snapshots are deleted and you can no longer return to those states. The snap02 data is deleted.



Deleting All Virtual Machine Snapshots

Slide 7-72

The delete-all-snapshots mechanism uses storage space efficiently. The size of the base disk does not increase. Just like a single snapshot deletion, changed blocks in the snapshot overwrite their counterparts in the base disk.



In the slide, snap01 is committed to the base disk before snap02 is committed.

All snapshots before You Are Here are committed all the way up to the base disk. All snapshots after You Are Here are discarded.

About Snapshot Consolidation

Slide 7-73

Snapshot consolidation is a method to commit a chain of snapshots to the base disks, when the Snapshot Manager shows that no snapshots exist, but the delta files still remain on the datastore.

Snapshot consolidation is intended to resolve problems that might occur with snapshots:

- The snapshot descriptor file is committed correctly. Hence the **Snapshot** tab shows that all the snapshots are deleted.
- The snapshot files (`-delta.vmdk`) are still part of the virtual machine.
- Snapshot files continue to expand until the virtual machine runs out of datastore space.

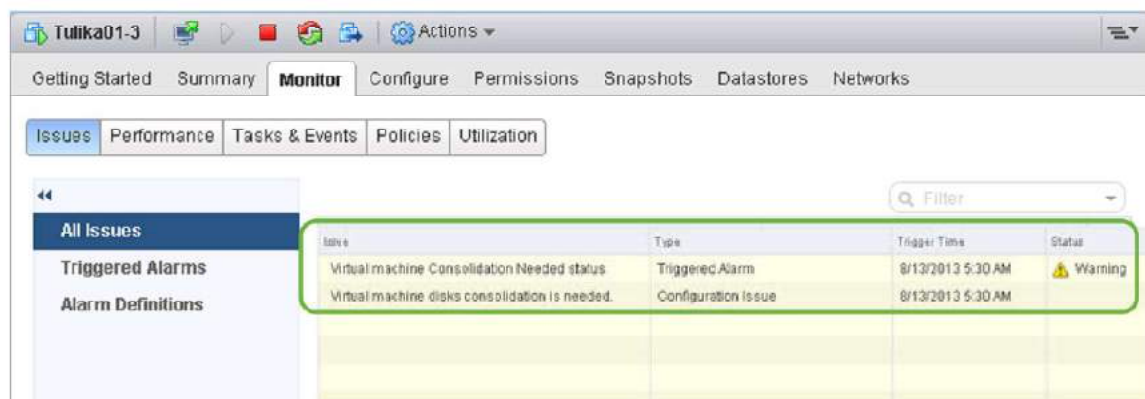
Snapshot consolidation is a way to clean unneeded snapshot delta files from a datastore. If no snapshots are registered for a virtual machine, but snapshot delta files exist, snapshot consolidation commits the chain of the snapshots indicated by the delta files and removes them.

If consolidation is not performed, the snapshot files might expand to the point of consuming all the remaining space on the virtual machine's datastore.

Discovering When to Consolidate

Slide 7-74

A warning on the **Monitor** > **Issues** tab of the virtual machine notifies the user that a consolidation is required.



The screenshot shows the vSphere Web Client interface for a virtual machine named 'Tulika01-3'. The 'Monitor' tab is selected, and the 'Issues' sub-tab is active. A table displays two issues related to consolidation, with the second row highlighted in yellow and a green box around it:

Issue	Type	Trigger Time	Status
Virtual machine Consolidation Needed status	Triggered Alarm	8/13/2013 5:30 AM	Warning
Virtual machine disks consolidation is needed.	Configuration Issue	8/13/2013 5:30 AM	

With snapshot consolidation, vCenter Server displays a warning when the descriptor and the snapshot files do not match. After the warning is displayed, the user can use vSphere Web Client to commit the snapshots, rather than committing snapshots from a command-line session.

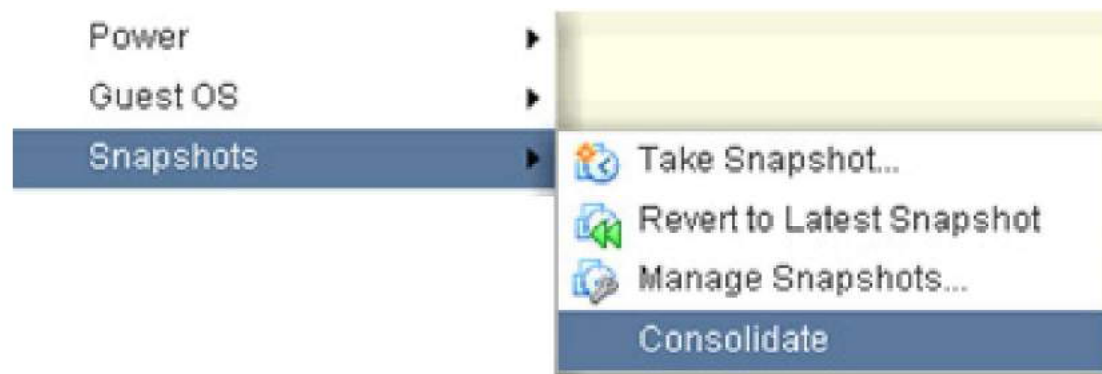
In the example, the **Monitor** tab displays a consolidation issue.

Performing Snapshot Consolidation

Slide 7-75

After the snapshot consolidation warning appears, the user can use vSphere Web Client to consolidate the snapshots:

- Select **Snapshots > Consolidate** to reconcile snapshots.
- All snapshot delta disks are committed to the base disks.



Lab 14: Managing Virtual Machines

Slide 7-76

Perform virtual machine management tasks

1. Unregister a Virtual Machine from the vCenter Server Appliance Inventory
2. Register a Virtual Machine in the vCenter Server Appliance Inventory
3. Unregister and Delete Virtual Machines from the Disk
4. Take Snapshots of a Virtual Machine
5. Revert the Virtual Machine to a Snapshot
6. Delete an Individual Snapshot
7. Delete All Snapshots

Review of Learner Objectives

Slide 7-77

You should be able to meet the following objectives:

- Take a snapshot of a virtual machine and manage multiple snapshots
- Delete virtual machine snapshots
- Consolidate snapshots

Key Points

Slide 7-78


- vCenter Server provides features for provisioning virtual machines, such as templates and cloning.
- By deploying virtual machines from a template, you can create many virtual machines easily and quickly.
- You can use vSphere vMotion to move virtual machines while they are powered on.
- You can use vSphere Storage vMotion to move virtual machines from one datastore to another datastore.
- You can use virtual machine snapshots to preserve the state of the virtual machine so that you can return to the same state repeatedly.

Questions?

MODULE 8

Resource Management and Monitoring

Slide 8-1



You Are Here

Slide 8-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. **Resource Management and Monitoring**
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
11. vSphere Update Manager

Importance

Slide 8-3

Although the VMkernel works proactively to avoid resource contention, maximizing performance requires both analysis and ongoing monitoring.

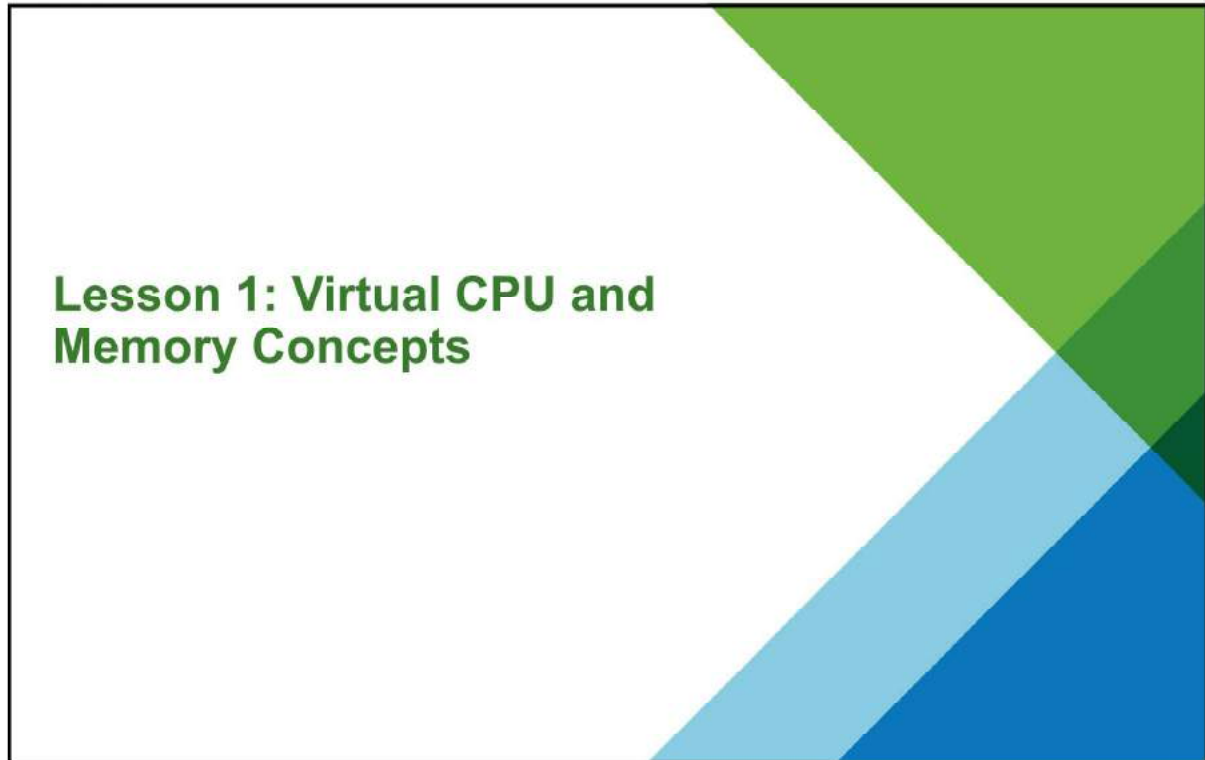
Module Lessons

Slide 8-4

- | | |
|-----------|--------------------------------------|
| Lesson 1: | Virtual CPU and Memory Concepts |
| Lesson 2: | Resource Controls and Resource Pools |
| Lesson 3: | Creating vApps |
| Lesson 4: | Monitoring Resource Use |
| Lesson 5: | Using Alarms |

Lesson 1: Virtual CPU and Memory Concepts

Slide 8-5



Learner Objectives

Slide 8-6

By the end of this lesson, you should be able to meet the following objectives:

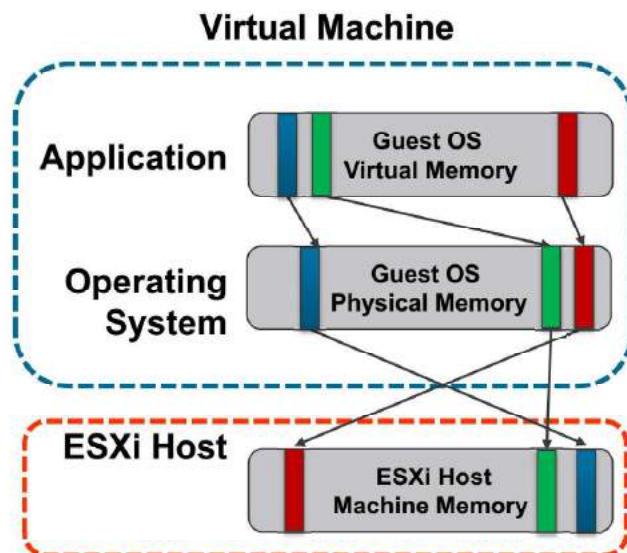
- Discuss CPU and memory concepts in a virtualized environment
- Describe what overcommitment of a resource means
- Identify additional technologies that improve memory utilization
- Describe how VMware vSphere® Virtual Symmetric Multiprocessing works and how hyperthreading is used by the VMkernel

Memory Virtualization Basics

Slide 8-7

vSphere has the following layers of memory:

- Guest operating system virtual memory is presented to applications by the operating system.
- Guest operating system physical memory is presented to the virtual machine by the VMkernel.
- Host machine memory that is managed by the VMkernel provides a contiguous, addressable memory space that is used by the virtual machine.



When running a virtual machine, the VMkernel creates a contiguous addressable memory space for the virtual machine. This memory space has the same properties as the virtual memory address space presented to applications by the guest operating system. This memory space allows the VMkernel to run multiple virtual machines simultaneously while protecting the memory of each virtual machine from being accessed by others. From the view of the application running in the virtual machine, the VMkernel adds an extra level of address translation that maps the guest physical address to the host physical address.

Virtual Machine Memory Overcommitment

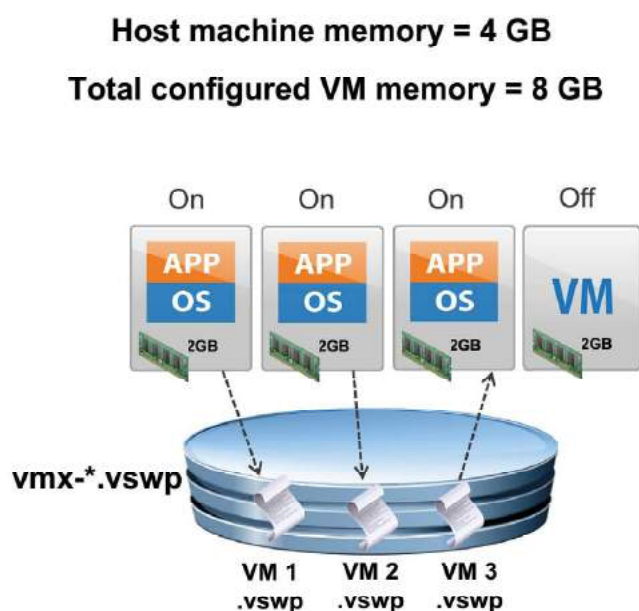
Slide 8-8

Memory is overcommitted when the combined working memory footprint of all virtual machines exceeds that of the host memory sizes.

Virtual machines do not always use their full allocated memory.

An ESXi host transfers memory from idle virtual machines to virtual machines that need more memory to improve memory utilization.

Memory overhead is stored in a swap file (`.vswp`).



The total configured memory sizes of all virtual machines might exceed the amount of available physical memory on the host. However, this condition does not necessarily mean memory is overcommitted. Memory is overcommitted when the working memory size of all virtual machines exceeds that of the ESXi host's physical memory size.

Because of the memory management techniques used by the ESXi host, your virtual machines can use more virtual RAM than the available physical RAM on the host. For example, you can have a host with 4 GB memory and run four virtual machines with 2 GB memory each. In that case, the memory is overcommitted. For example, if all four virtual machines are idle, the combined consumed memory might be below 4 GB. However, if all virtual machines are actively consuming memory, then their memory footprint might exceed 4 GB and the ESXi host might become overcommitted. An ESXi host can run out of memory if virtual machines consume all reservable memory in an overcommitted-memory environment. Although the powered-on virtual machines are not affected, a new virtual machine might fail to power on due to lack of memory. Overcommitment makes sense because, typically, some virtual machines are lightly loaded while others are more heavily loaded, and relative activity levels vary over time.

Extra memory from a virtual machine is gathered into a `.vswp` swap file. The memory overcommitment process on the host uses the `vmx-*.vswp` swap file to gather and track memory overhead. Memory from this file is swapped out to disk when host machine memory is overcommitted.

Memory Reclamation Techniques

Slide 8-9

Economize use of physical memory pages:

- Transparent page sharing allows pages with identical contents to be stored only once.

Deallocate memory from one virtual machine for another:

- Ballooning mechanism, active when memory is scarce, forces virtual machines to use their own paging areas.

Memory compression:

- Attempts to reclaim some memory performance when memory contention is high.

Host-level SSD swapping:

- Use of a solid-state drive (SSD) on the host for a host cache swap file might increase performance.

Page virtual machine memory out to disk:

- Use of VMkernel swap space is the last resort, because of poor performance.

The VMkernel uses various techniques when reclaiming memory on an ESXi host. Each technique is described in the order that the VMkernel uses it.

The first technique economizes the use of physical memory pages and is known as transparent page sharing (TPS). TPS allows pages with identical contents to be stored only once. Workloads consisting of multiple virtual machines often consume less memory than they would when running on physical machines. The effect of TPS is that it allows hosts to efficiently support higher levels of memory overcommitment.

During times of memory contention, the VMkernel looks for opportunities to reclaim idle or allocated but unused memory from virtual machines. The VMkernel transfers memory from idle virtual machines to virtual machines that need more memory, using the `vmmemctl` driver that is typically installed with VMware Tools. The `vmmemctl` driver is also called the memory balloon driver.

Memory compression is another technique that the VMkernel uses to reclaim host physical memory. This technique attempts to reclaim some memory by compressing pages when contention is high to avoid swapping out to a virtual machine swap file.

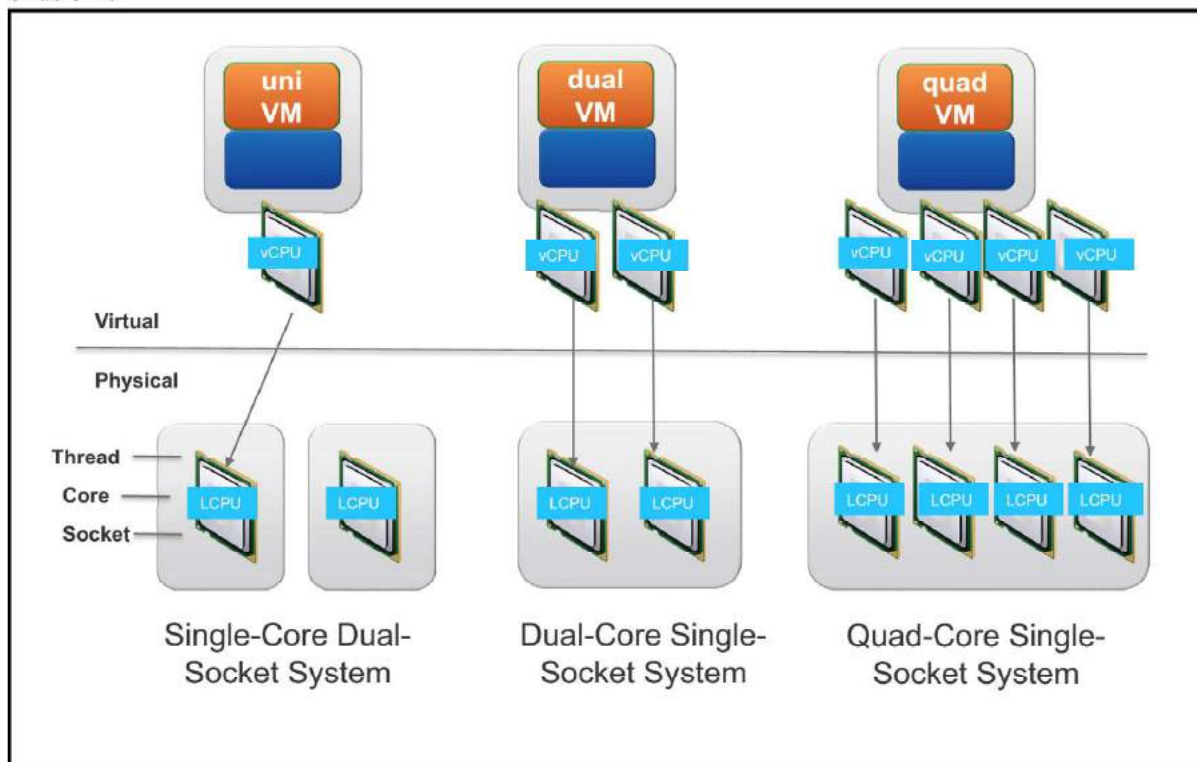
Another option for reclaiming memory is to create a host cache on a solid-state drive (SSD) on the host machine. This host cache will be used after TPS, ballooning, and memory compression have

been tried, and before using a swap file on disk. Having the host cache file on an SSD drive can improve the performance of this technique.

The last and least desirable technique employed is to page virtual machine memory out to disk. Swapping to disk is used when other techniques are temporarily unable to reclaim memory quickly enough to satisfy current system demand. A significant performance penalty is observed when this technique is used.

Virtual SMP

Slide 8-10



You can configure a virtual machine with up to 128 virtual CPUs (vCPUs). The VMkernel includes a CPU scheduler that dynamically schedules vCPUs on the physical CPUs of the host system.

The VMkernel scheduler considers socket-core-thread topology when making scheduling decisions. Intel and AMD have developed processors that combine multiple processor cores into a single integrated circuit, called a socket in this discussion. A socket is a single package that can have one or more physical CPUs, with each core having one or more logical CPUs (LCPU's), or threads. LCPU's provide the core with the ability to schedule one thread of execution. Each LCPU provides the core with the ability to schedule one thread of execution. On the slide, the first system is a single-core, dual-socket system that has two cores and so two LCPU's.

When a vCPU, of a single-vCPU or multi-vCPU virtual machine, must be scheduled, the VMkernel maps the vCPU to an available logical processor.

Hyperthreading


Slide 8-11

Hyperthreading enables a core to execute two threads, or sets of instructions, at the same time.

Hyperthreading provides more scheduler throughput.

To enable hyperthreading:

1. Verify that the host system supports hyperthreading.
2. Enable hyperthreading in the system BIOS.
3. Ensure that hyperthreading for the ESXi host is turned on.
4. Hyperthreading is enabled by default.



If hyperthreading is enabled on the host system, ESXi can execute two threads at the same time on each processor core (physical CPU). Hyperthreading provides more scheduler throughput. That is, hyperthreading provides more logical CPUs on which vCPUs can be scheduled. The drawback of hyperthreading is that it does not double the power of a core. So, if both threads of execution need the same on-chip resources at the same time, one thread has to wait. Still, ESXi uses recent advances in hyperthreading technology and, on systems that use these new technologies, performance is improved.

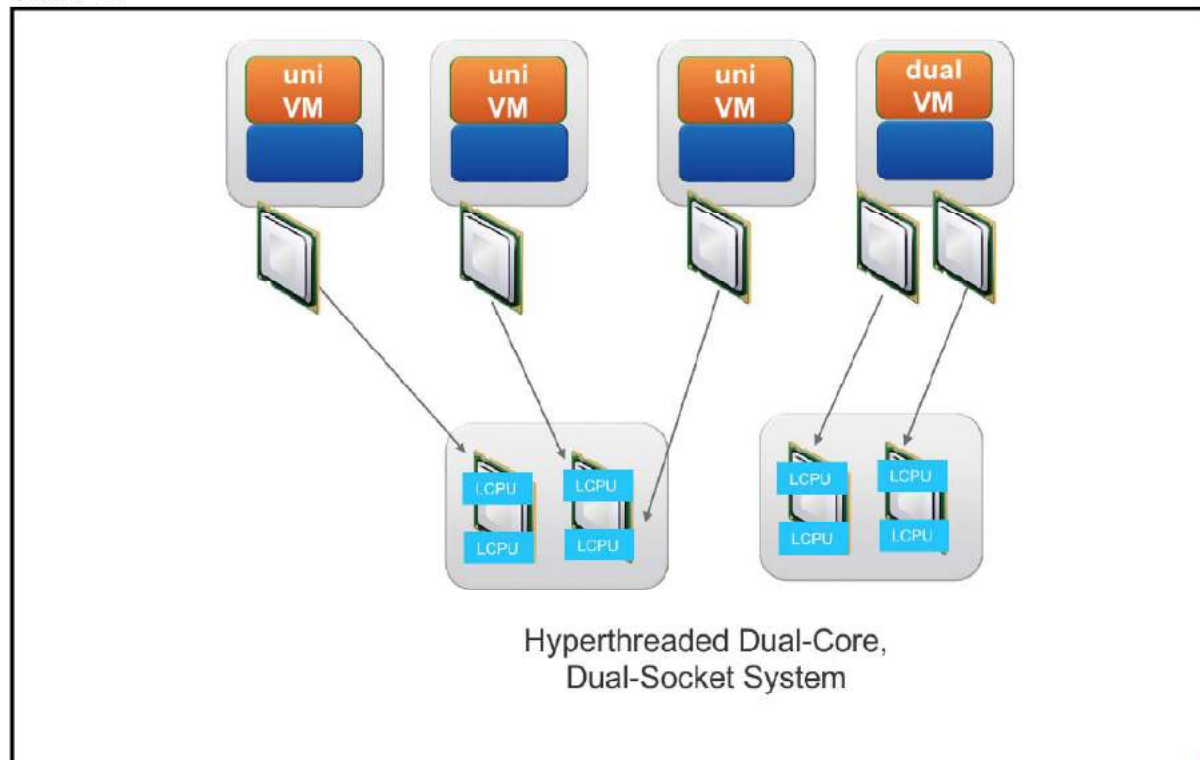
An ESXi host enabled for hyperthreading should behave almost exactly like a standard system. Logical processors on the same core have adjacent CPU numbers. So logical processors 0 and 1 are on the first core, logical processors 2 and 3 are on the second core, and so on.

To ensure that hyperthreading is functioning, consult the hardware documentation to see whether the BIOS includes support for hyperthreading. Then enable hyperthreading in the system BIOS. Some manufacturers call this option Logical Processor. Others call it Enable Hyperthreading.

Use vSphere Web Client to ensure that hyperthreading for your host is turned on. To access the hyperthreading option, go to the host's **Summary** tab and select **CPUs** under Hardware.

CPU Load Balancing

Slide 8-12



The CPU scheduler can use each logical processor independently to execute virtual machines, providing capabilities similar to traditional symmetric multiprocessing (SMP) systems. The VMkernel intelligently manages processor time to guarantee that the load is spread smoothly across processor cores in the system. Every 2 to 40 milliseconds (depending on the socket-core-thread topology), the VMkernel looks to migrate vCPUs from one logical processor to another to keep the load balanced. The VMkernel does its best to schedule virtual machines with multiple vCPUs on two different cores rather than on two logical processors on the same core. But if necessary, the VMkernel can map two vCPUs from the same virtual machine to threads on the same core.

If a logical processor has no work, it is put into a halted state. This action frees its execution resources and allows the virtual machine running on the other logical processor on the same core to use the full execution resources of the core. The VMkernel scheduler properly accounts for this halt time. So a virtual machine running with the full resources of a core is charged more than a virtual machine running on a half core. This approach to processor management ensures that the server does not violate the ESXi resource allocation rules.

Review of Learner Objectives

Slide 8-13

You should be able to meet the following objectives:

- Discuss CPU and memory concepts in a virtualized environment
- Describe what overcommitment of a resource means
- Identify additional technologies that improve memory utilization
- Describe how VMware Virtual SMP works and how hyperthreading is used by the VMkernel

Lesson 2: Resource Controls and Resource Pools

Slide 8-14



Lesson 2: Resource Controls and Resource Pools

Learner Objectives

Slide 8-15


By the end of this lesson, you should be able to meet the following objectives:

- Assign share values for CPU, memory, and disk resources
- Describe how virtual machines compete for resources
- Create a resource pool
- Set resource pool attributes
- Establish CPU and memory reservations and limits
- Describe expandable reservations

Shares, Limits, and Reservations

Slide 8-16

A virtual machine powers on only if its reservation can be guaranteed.



Because virtual machines simultaneously use the resources of an ESXi host, resource contention can often occur.

For proper resource management, vSphere has mechanisms to enable less, more, or an equal amount of access to a defined resource. vSphere also prevents a virtual machine from consuming large amounts of a resource and grants a guaranteed amount of a resource to a virtual machine whose performance is not adequate or requires a certain amount of a resource to run properly.

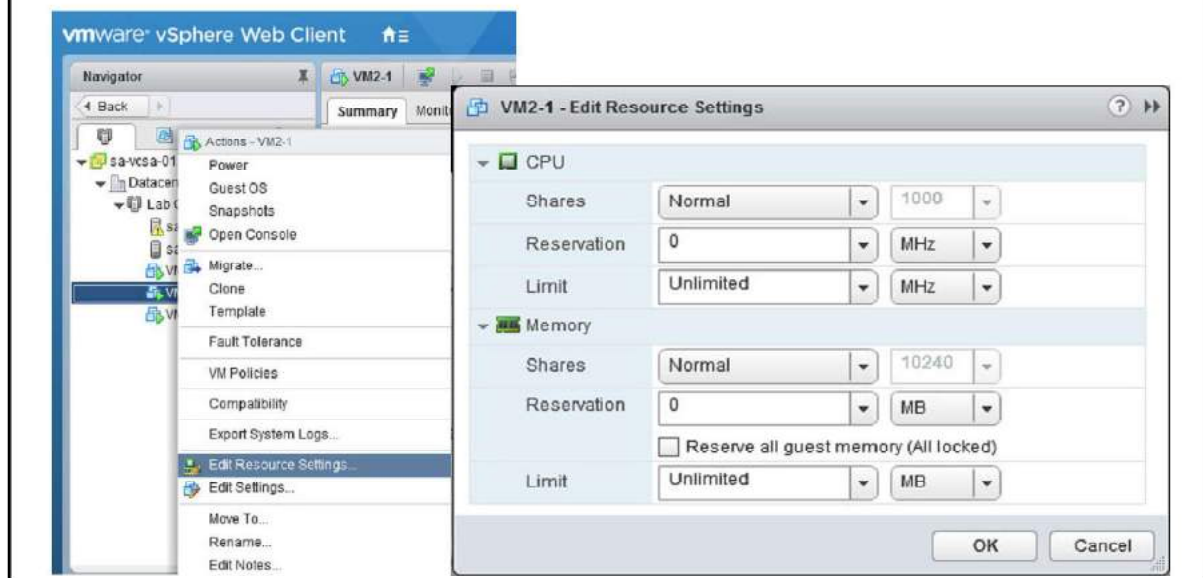
When host memory or CPU is overcommitted, a virtual machine's allocation target is somewhere between its specified reservation and specified limit, depending on the virtual machine's shares and the system load. vSphere uses a share-based allocation algorithm to achieve efficient resource use for all virtual machines and to guarantee a given resource to the virtual machines that need it most. Three configurable parameters control a virtual machine's access to a given resource:

- **Shares:** A value that specifies the relative priority or importance of a virtual machine's access to a given resource
- **Limit:** Limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a virtual machine.
- **Reservation:** A reservation specifies the guaranteed minimum allocation for a virtual machine.

Defining Resource Settings for Individual VMs

Slide 8-17

When available resource capacity does not meet the demands of the resource consumers, you might need to customize the amounts of resources that are allocated to VMs.



You must understand how Shares, Limits, and Reservations work. If you overprovision your ESXi host on memory and CPU, you need a tool to make sure that the right machines get the correct amount of resources. Shares, Limits, and Reservations serve this purpose.

A reservation is a guarantee on either memory or CPU for a virtual machine. You define the reservation in MB or MHz.

Shares specify the relative importance of a VM or resource pool. If a VM has twice as many shares of a resource as another VM, it is entitled to consume twice that resource when these two VMs are competing.

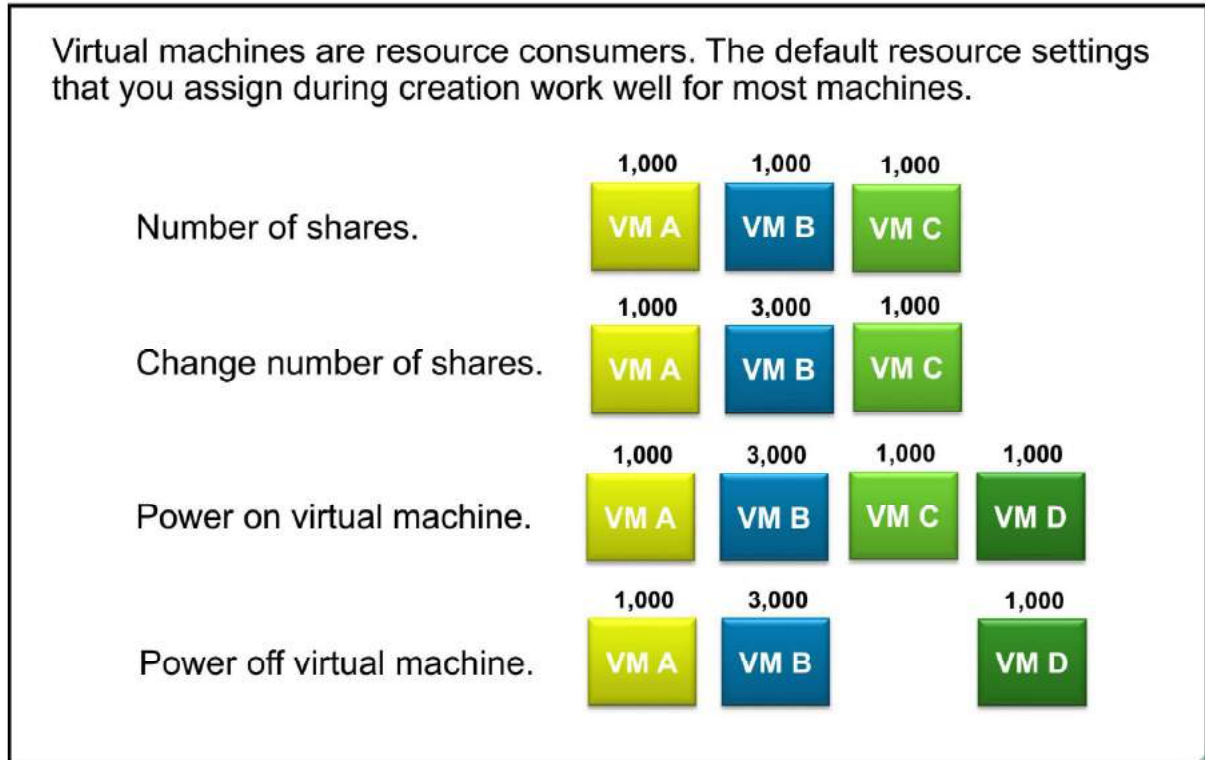
Limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a VM.

In most cases, it is not necessary to specify a limit. Specifying limits has the following benefits and drawbacks:

- **Benefits:** Assigning a limit is useful if you start with a small number of virtual machines and want to manage user expectations. The performance deteriorates as you add more virtual machines. You can simulate having fewer resources available by specifying a limit.
- **Drawbacks:** You might waste idle resources if you specify a limit. The system does not allow virtual machines to use more resources than the limit, even when the system is underutilized and idle resources are available. Specify the limit only if you have good reasons for doing so.

Resource Sharing by Virtual Machines

Slide 8-18



The proportional share mechanism applies to CPU, memory, storage I/O, and network I/O allocation. The mechanism operates only when virtual machines are contending for the same resource.

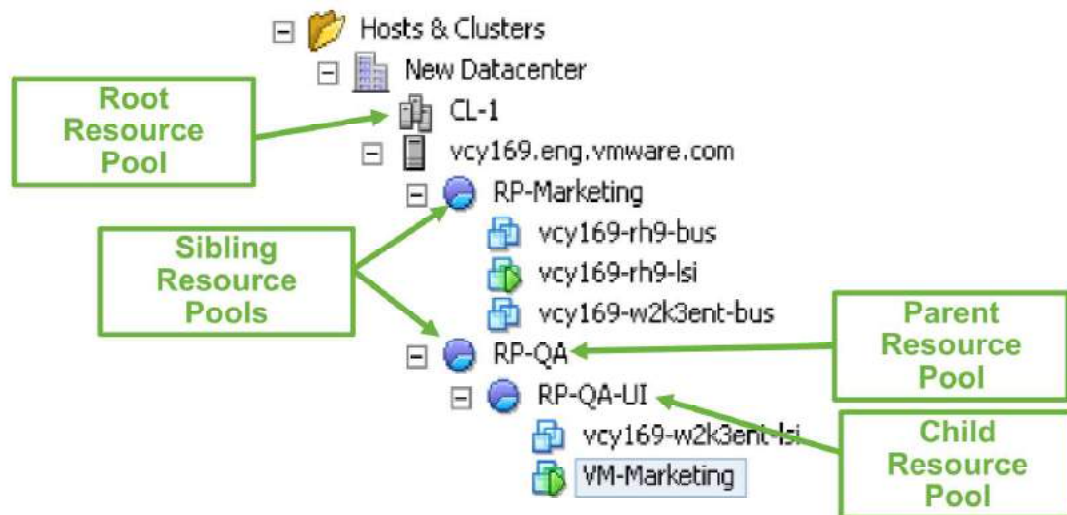
Shares guarantee that a virtual machine is given a certain amount of a resource (CPU, RAM, storage I/O, or network I/O). For example, consider the third row of virtual machines on the slide, where virtual machine D has been powered on with 1,000 shares. Before it was powered on, a total of 5,000 shares were available, but virtual machine D's addition increases the total shares to 6,000. The result is that the other virtual machines decline in value. But each virtual machine's share value still represents a minimum guarantee. Virtual machine A is still guaranteed one-sixth of the resource because it owns one-sixth of the shares.

You can add shares to a virtual machine while it is running, and it will get more access to that resource (assuming competition for the resource). When you add a virtual machine, it gets shares, too. The virtual machine's share amount factors into the total number of shares, but the existing virtual machines are guaranteed not to be starved for the resource. When you delete or power off a virtual machine, fewer total shares remain, so the surviving virtual machines get more access.

About Resource Pools

Slide 8-19

A resource pool is a logical abstraction of hierarchically managed CPU and memory resources.



A resource pool is a logical abstraction for managing resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources. In the example, on the slide RP-QA is the parent resource pool for RP-QA-UI. RP-Marketing and RP-QA are siblings.

A resource pool allows you as the administrator to divide and allocate resources to virtual machines and other resource pools. A resource pool allows you to control the aggregate CPU and memory resources of the compute resource. The compute resource can be either a standalone host or a vSphere DRS cluster. Resource pools are also used to delegate privileges to other users and groups.

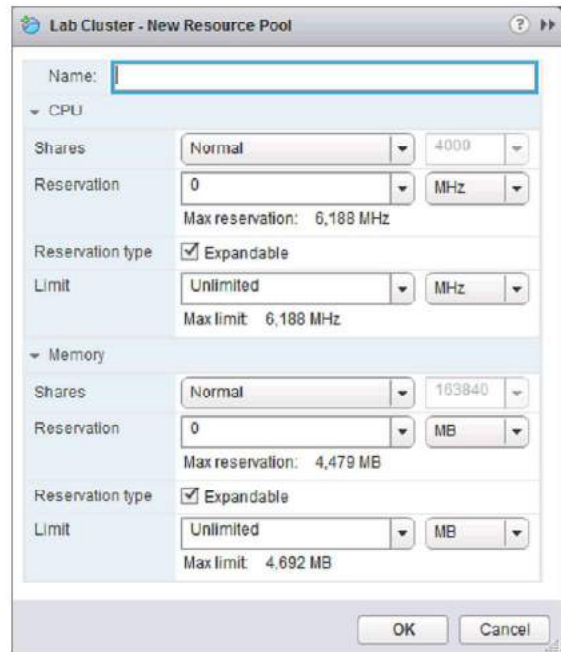
The topmost resource pool is called the root resource pool. Each standalone host and each vSphere DRS cluster has an (invisible) root resource pool that groups the resources of that host or cluster. The root resource pool does not appear, because the resources of the host (or cluster) and the root resource pool are always the same.

Resource Pool Attributes

Slide 8-20

You can create a child resource pool of any ESXi host, resource pool, or vSphere DRS cluster.

- Shares: Low, Normal, High, Custom
- Reservations: In MHz or GHz, MB or GB
- Limits:
 - In MHz or GHz, MB or GB.
 - Unlimited access, by default, up to maximum amount of resource accessible.
- Reservation type:
 - Expandable selected: Virtual machines and subpools can draw from this pool's parent.
 - Expandable deselected: Virtual machines and subpools can draw only from this pool, even if its parent has free resources.



Like virtual machines, a resource pool has reservation, limit, and share values for CPU and memory resources:

- **Shares:** A guarantee that the resource pool is given a certain proportion of CPU and memory resources. Resource pool shares work like virtual machine shares.
- **Reservation:** The minimum amount of resources that are required by the resource pool. For example, you can set a CPU reservation, which is the minimum amount of CPU that this pool must have.
- **Limit:** The maximum amount of resources that are given to this resource pool. By default, the resource pool is given unlimited access to the maximum amount of resource (specified by the limit).
- **Expandable reservation:** An attribute that is specific to a resource pool. This attribute allows a resource pool that cannot satisfy a reservation request to search through its hierarchy to find unreserved capacity to satisfy the reservation request.

Reasons to Use Resource Pools

Slide 8-21

Use of resource pools can provide these benefits:

- Flexible hierarchical organization
- Isolation between pools and sharing in pools
- Access control and delegation
- Separation of resources from hardware
- Management of sets of virtual machines running a multitier service
- Ability to prioritize virtual machine workloads

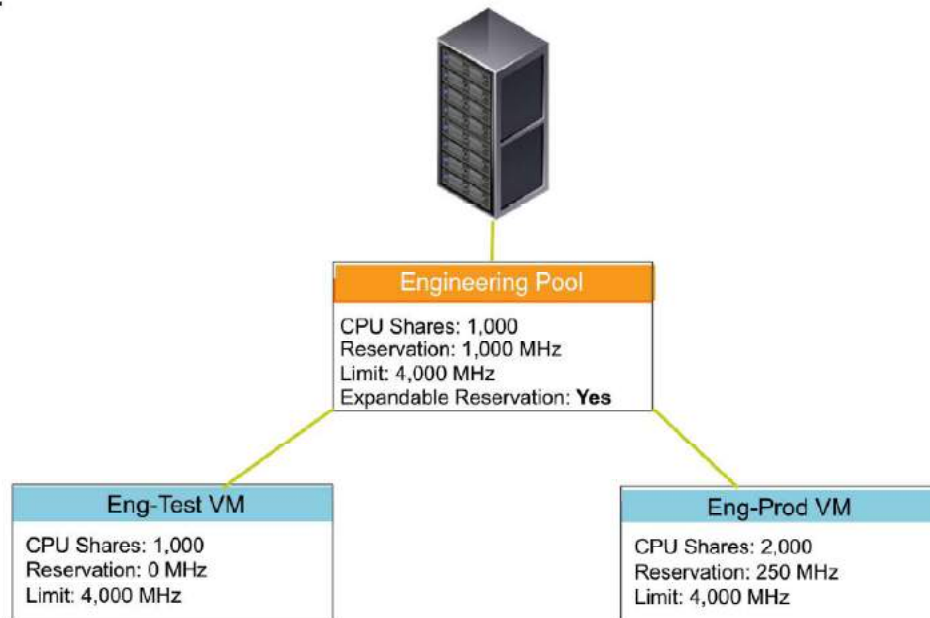
With resource pools, you can delegate control over resources of a standalone host or a vSphere DRS cluster. Using resource pools can result in the following benefits:

- **Flexible hierarchical organization:** Add, remove, or reorganize resource pools or change resource allocations as needed.
- **Isolation between pools, sharing in pools:** Top-level administrators can make a pool of resources available to a department-level administrator.
- **Access control and delegation:** Virtual machine creation and management are performed in the boundaries of the resources to which the resource pool is entitled. Delegation is usually done with permissions settings.
- **Separation of resources from hardware:** If you are using vSphere DRS clusters, the resources of all hosts are always assigned to the cluster.
- **Management of sets of virtual machines running a multitier service:** Group virtual machines for a multitier service in a resource pool.

Resource Pool Example

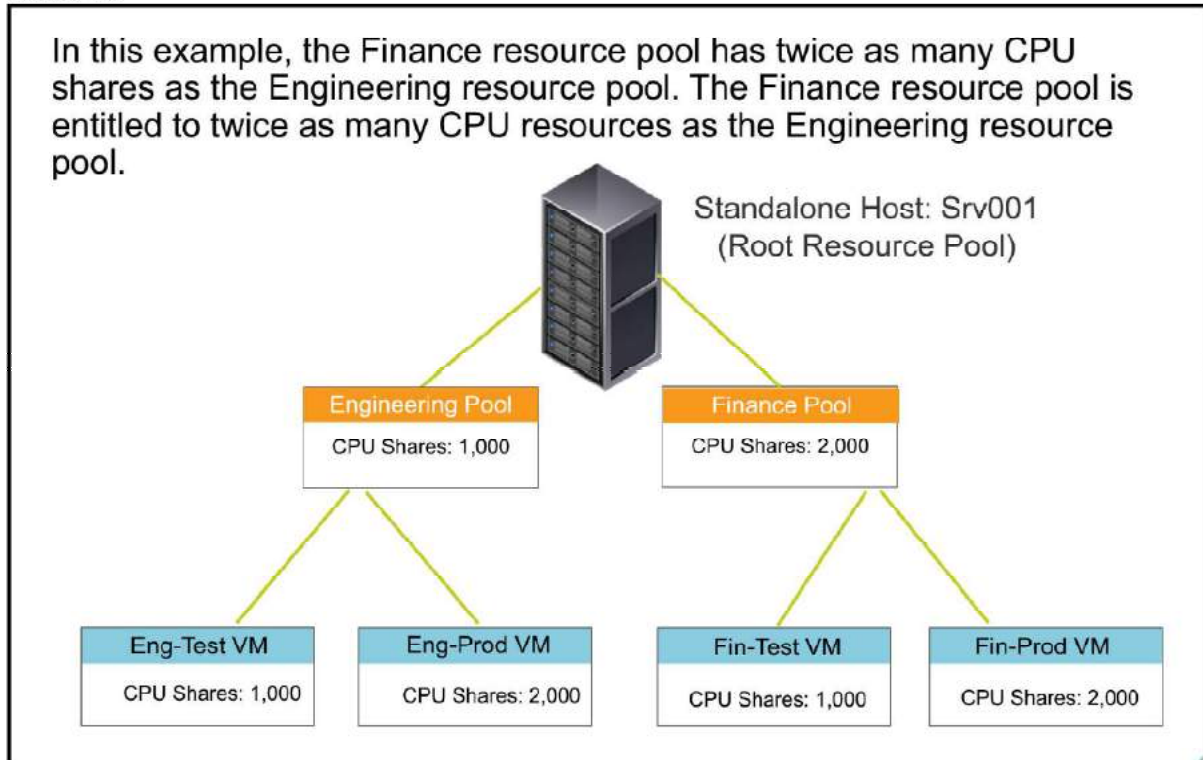
Slide 8-22

This example shows where resource attributes are set on a resource pool.



Resource Pools Example: CPU Shares

Slide 8-23



Resource pools can be organized hierarchically. The root resource pool is the topmost resource pool. The root resource pool includes the sum of all megahertz for all CPUs and the sum of all the installed RAM (in megabytes) available in the compute environment (standalone host or cluster).

On the slide, the root resource pool is a standalone host named Srv001. It has 12,000 MHz of CPU and 4 GB of RAM, available for use by other resource pools or virtual machines.

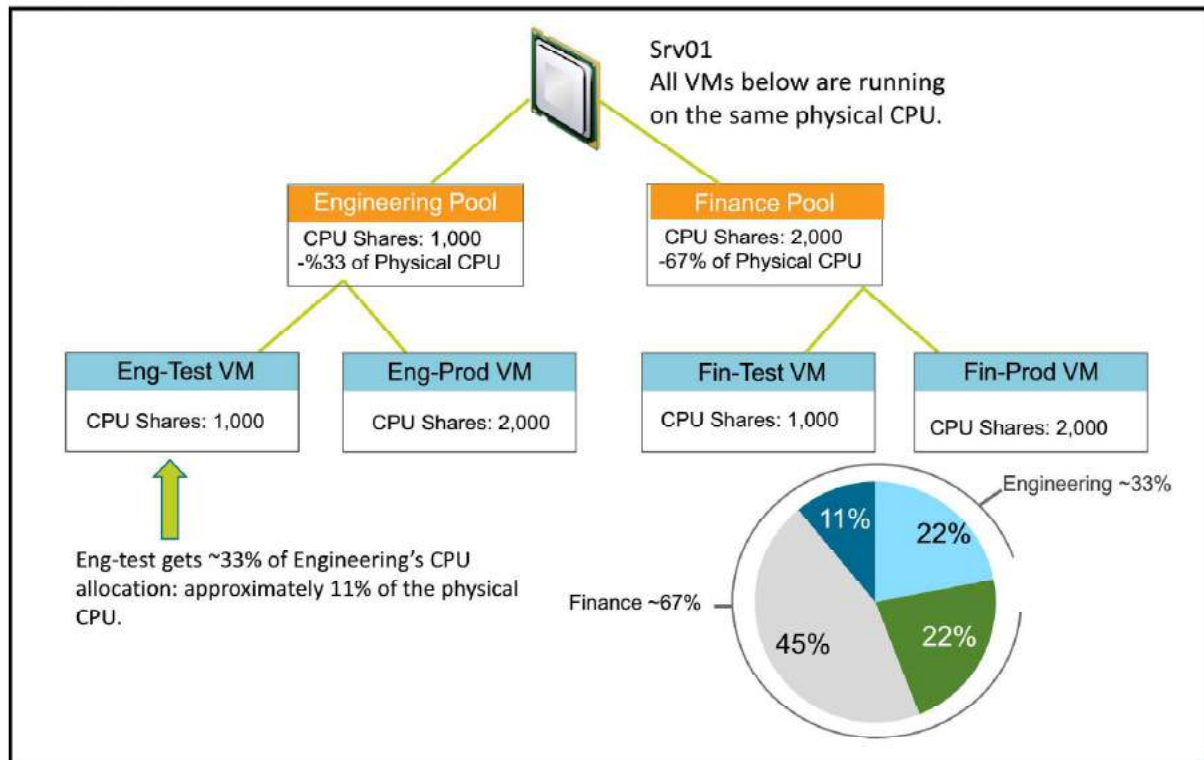
Except for the root resource pool, every resource pool has a parent resource pool. A resource pool might contain child resource pools or only virtual machines that are powered on in it.

A child resource pool is used to allocate resources from the parent resource pool for the child's consumers. Administrative control can also be delegated to individuals or organizations. A child resource pool cannot exceed the capacity of the parent resource pool. Creating a child pool reserves resources from the parent pool, whether or not virtual machines in the child pool are powered on.

Shares specify the relative priority or importance of either a resource pool or virtual machine. If a resource pool has twice as many shares of a resource as another resource pool, it is entitled to consume twice as much of that resource. The same thing can be applied to virtual machines.

Resource Pools Example: CPU Contention

Slide 8-24



Assume that all four virtual machines have been scheduled by the VMkernel onto the same physical CPU. The virtual machines are in direct competition with one another.

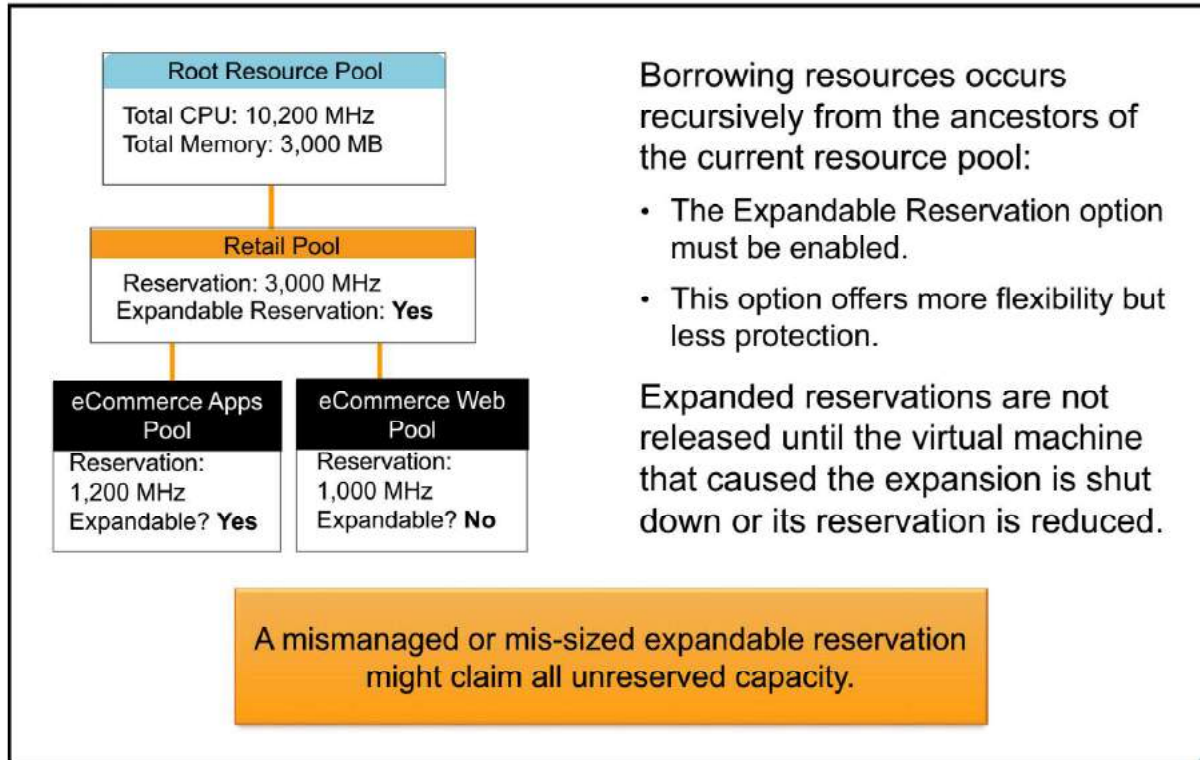
The Engineering pool gets 33 percent of that CPU and splits its allotment between virtual machines Eng-Test and Eng-Prod. Likewise, the Finance pool gets 67 percent of that CPU and splits its 67 percent allotment between virtual machines Fin-Test and Fin-Prod. A virtual machine's resource settings are constrained by the resources of the resource pool to which the virtual machine belongs.

The virtual machine Eng-Test gets approximately 33 percent of the CPU allocation of the Engineering resource pool $[1,000/(1,000+2,000)]$. This figure is equal to about 11 percent of the physical CPU (33 percent of 33 percent equals about 11 percent). Each of the virtual machines gets a percentage of the physical CPU allocated to its resource pool that is based on its individual share allocation.

The example on the slide uses approximations to explain how the number of shares affects the amount of CPU allocated to a virtual machine.

Expandable Reservation

Slide 8-25



Expandable reservation allows a resource pool that cannot satisfy a reservation request to search through its hierarchy to find unreserved capacity to satisfy the reservation request.

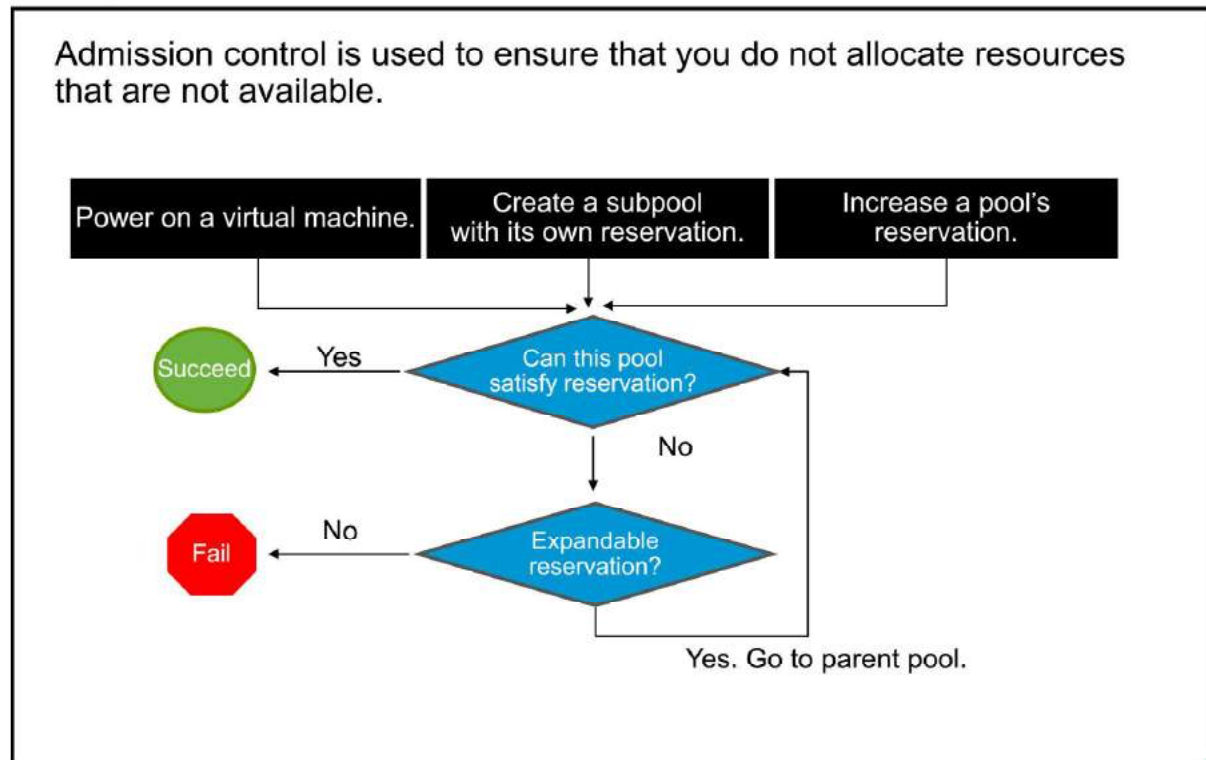
On the slide, the child resource pool eCommerce Apps has the **Expandable Reservation** option enabled. The reservation of a child resource pool cannot exceed that of its parent. The search for unused resources goes through the ancestry of the root resource pool or to the first resource pool that lacks the **Expandable Reservation** option enabled.

Use expandable reservation carefully. A single child resource pool can use all of its parent's available resources, leaving nothing directly available for other child resource pools. This can happen if virtual machines are improperly moved around to get them to start based upon resource reservations.

You might want to disable the **Expandable Reservation** option when you are giving a fixed amount of resources to a group. For example, an IT administrator whose customers are different organizations in the company that have paid for a fixed amount of CPU and memory resources might want to disable the **Expandable Reservation** option.

Admission Control for CPU and Memory Reservations

Slide 8-26



Certain operations must satisfy admission control:

- Powering on a virtual machine
- Creating a resource pool with its own reservation
- Increasing a resource pool's reservation

If the object (virtual machine or resource pool) resides in a resource pool with an expandable reservation, the parent of the current resource pool is consulted if necessary to satisfy the reservation.

Reservations cannot be overcommitted.

Resource Pool Summary Tab

Slide 8-27

The resource pool **Summary** tab displays information that applies to the host machine and its resources.

The screenshot shows the VMware vSphere Summary tab for a resource pool named 'Production'. The interface includes a Navigator on the left, a main content area with a 'Summary' tab selected, and several summary panels.

Summary Tab Data:

Category	Value
VMs and Templates	3 / 3
Powered on VMs	3 / 3
Child Resource Pools	0 / 0
Child vApps	0 / 0

Resource Settings:

Setting	Value
CPU	Normal (4000)
Memory	Normal (163840)

Resource Consumers:

Consumer Type	Count
Virtual Machines and Templates	3
Powered On Virtual Machines	3
Child Resource Pools	0

Related Objects:

Attribute	Value
This list is empty.	

Consider these important points about the **Summary** tab:

- The Resource Settings pane displays CPU and Memory share settings.
- The Resource Consumers pane displays the number of virtual machines, number of powered-on virtual machines, and child resource pools that are in the selected resource pool.
- The Tags pane shows tags assigned to objects that reside in the resource pool.

Resource Reservation Tab

Slide 8-28

On the **Resource Reservation** tab, you can view information about a resource pool's CPU, memory, and storage resources.

The screenshot shows the vSphere interface for a resource pool named 'Production'. The 'Resource Reservation' tab is active, displaying the following information:

- CPU:** 0 GHz (Total) / 6.19 GHz (Available)
- Configured Reservation:** 0.00 GHz
- Used Reservation:** 0.00 GHz
- Available Reservation:** 6.19 GHz
- Reservation Type:** Expandable

Name	Reservation (MHz)	Limit (MHz)
Mike02-2	0	Unlimited
Mike02-3	0	Unlimited
Vivian01-2	0	Unlimited

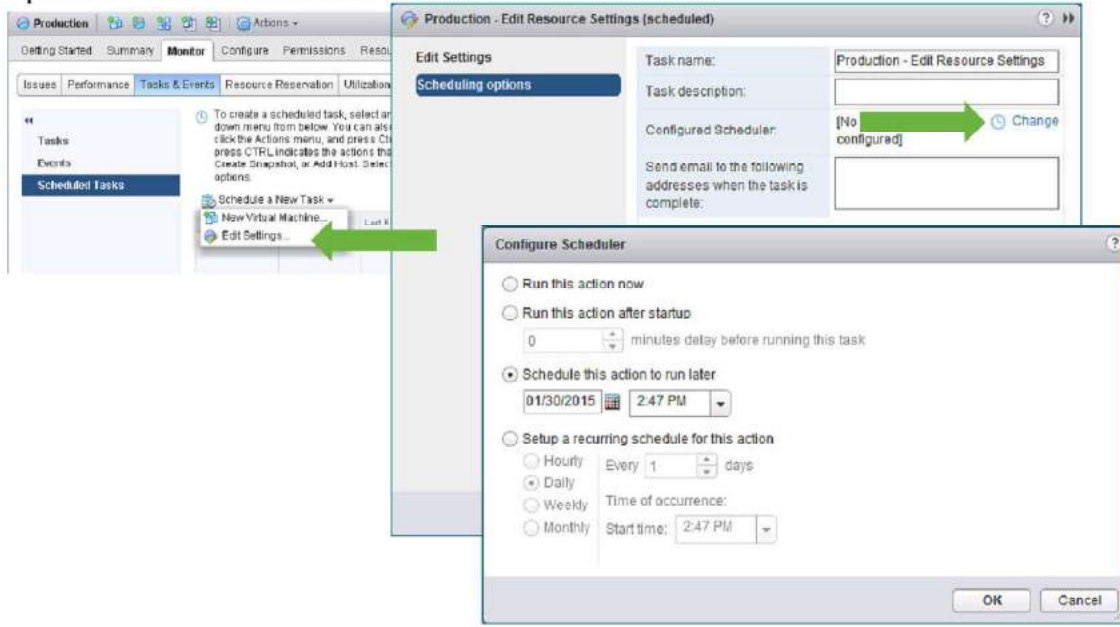
The following information is displayed for the pool's CPU and memory resources:

- The amount of CPU and memory reservation that is configured
- The type of reservation used (expandable or fixed)
- The amount of reservation in use by virtual machines and child pools
- The amount of CPU and memory that is available to be reserved

Scheduling Changes to Resource Settings

Slide 8-29

You can schedule a task to change the resource settings of a resource pool or virtual machine.



You can configure the task to change the shares, reservation, and limit for CPU or for memory or for both so that you can accommodate changing business priorities.

For example, at the end of each quarter you can give financial applications more CPU and memory resources than internal applications. In a retail organization, you can double the CPU and memory resource reservations for the virtual machines running the online store applications during the month of December.

Lab 15: Managing Resource Pools

Slide 8-30

Create and use resource pools on an ESXi host using vCenter Server

1. Create CPU Contention
2. Create Resource Pools
3. Verify Resource Pool Functionality

Review of Learner Objectives

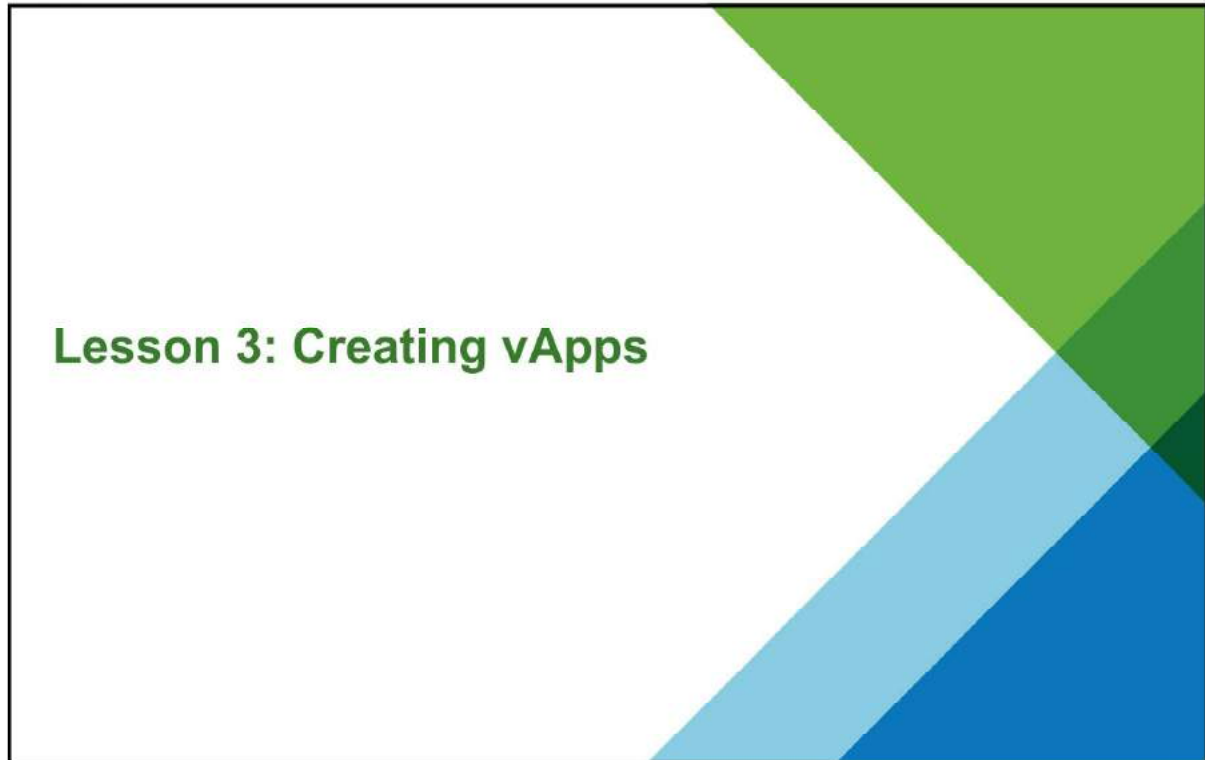
Slide 8-31

You should be able to meet the following objectives:

- Assign share values for CPU, memory, and disk resources
- Describe how virtual machines compete for resources
- Create a resource pool
- Set resource pool attributes
- Establish CPU and memory reservations and limits
- Describe expandable reservations

Lesson 3: Creating vApps

Slide 8-32



Learner Objectives

Slide 8-33

By the end of this lesson, you should be able to meet the following objectives:

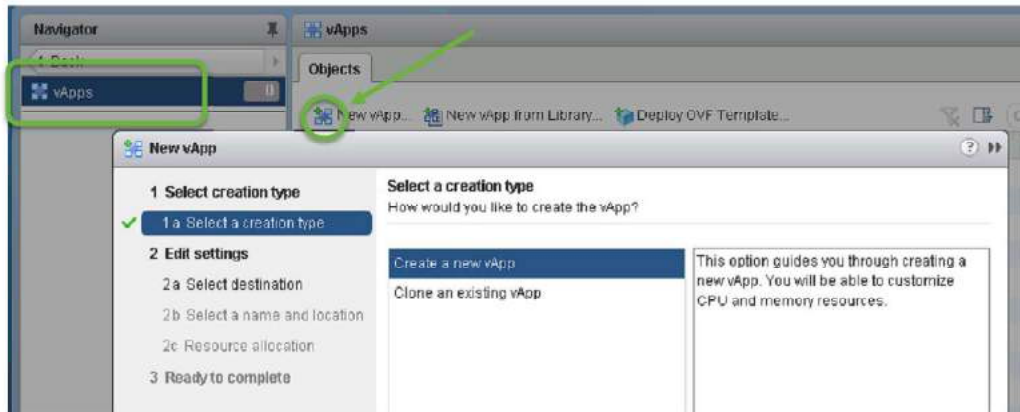
- Describe a vApp
- Build a vApp
- Use a vApp to manage virtual machines
- Deploy and export a vApp

Managing Virtual Machines with a vApp

Slide 8-34

A vApp is an object in the vCenter Server inventory:

- A vApp is a container for one or more virtual machines.
- A vApp can be used to package and manage multitiered applications.



You can use vSphere as a platform for running applications, such as multitiered applications. The applications can be packaged to run directly on top of vSphere.

In vSphere Web Client, a vApp is represented in the Hosts and Clusters view and in the VM and Templates view.

A vApp is a container for one or more virtual machines or vApps. A vApp shares functionality with virtual machines. A vApp can power on and power off and it can be cloned. The distribution format for a vApp can be either Open Virtualization Format (OVF) or Open Virtualization Appliance (OVA). The differences between these formats are:

- An OVF file is a collection of virtual machine files. The OVF file is an XML file that has information about the virtual disk files in the directory. When you export a virtual machine as an OVF file, a directory is created that has an OVF file and the VMDKs.
- OVA is the portable virtual machine format from XenSource, a third-party product. The OVA file is a single file that can be considered an archive, like a ZIP file, of all the files that belong to the OVF directory.

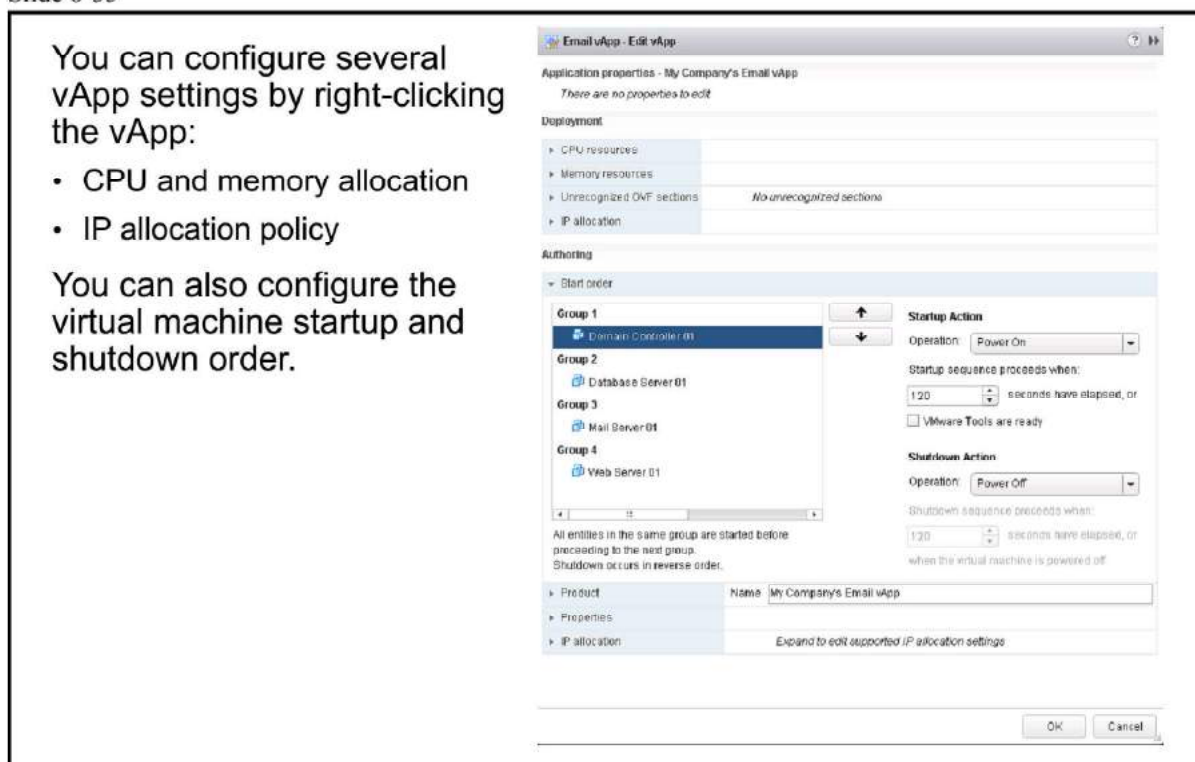
vApp Characteristics

Slide 8-35

You can configure several vApp settings by right-clicking the vApp:

- CPU and memory allocation
- IP allocation policy

You can also configure the virtual machine startup and shutdown order.



After creating the vApp, you can modify certain vApp settings:

- Resource allocation: Determines how CPU and memory should be allocated for the vApp.
- IP allocation policy: Determines how IP addresses are allocated for the vApp.
- Advanced settings: Product and vendor information, and custom properties.

You can change the order in which virtual machines (and nested vApps) in a vApp start up and shut down. You change this order by assigning virtual machines to groups. All entities in the same group are started before those in the next group. Shutdown is done in the reverse order. You can also specify delays and actions performed at startup and shutdown.

On the slide, the virtual machine Domain Controller 01 is started first. The virtual machine Database Server 01 starts after 120 seconds.

For more about vApps, see *vSphere Virtual Machine Administration Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Exporting and Deploying vApps

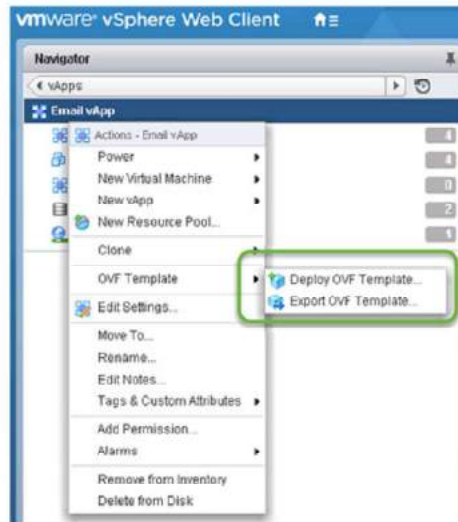
Slide 8-36

Exporting the vApp as an OVF or OVA template:

- Share with others.
- Use for archive purposes.

Deploying the OVF template:

- Deploy multitier vApps.
- Deploy OVF from VMware Virtual Appliance Marketplace.



Lab 16: Managing vApps

Slide 8-37

Perform vApp management tasks

1. Create a vApp
2. Power On a vApp
3. Remove a vApp

Review of Learner Objectives

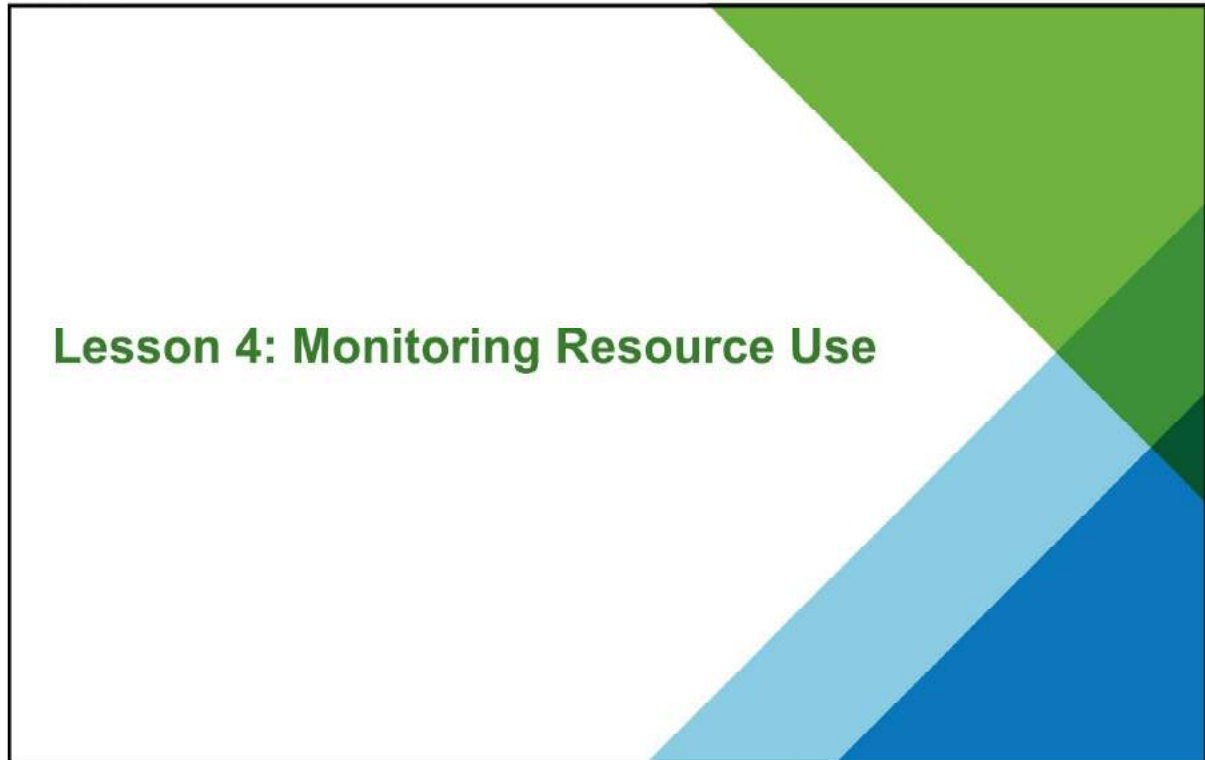
Slide 8-38

You should be able to meet the following objectives:

- Describe a vApp
- Build a vApp
- Use a vApp to manage virtual machines
- Deploy and export a vApp

Lesson 4: Monitoring Resource Use

Slide 8-39



Learner Objectives

Slide 8-40

By the end of this lesson, you should be able to meet the following objectives:

- Use the performance-tuning methodology and resource monitoring tools
- Use performance charts to view and improve performance
- Monitor the key factors that can affect the virtual machine's performance: CPU, memory, disk, and network bandwidth use

Performance-Tuning Methodology

Slide 8-41

Follow these best practices for performance-tuning your vSphere infrastructure:

- Assess performance:
 - Use appropriate monitoring tools.
 - Record a numerical benchmark before changes.
- Identify the limiting resource.
- Make more resources available:
 - Allocate more.
 - Reduce competition.
 - Log your changes.
- Benchmark again.

Do not make casual changes to production systems.

The best practice for performance tuning is to take a logical, step-by-step approach:

1. For a complete view of the performance situation of a virtual machine, use monitoring tools in the guest operating system and in vCenter Server. Record benchmarks before you make changes.
2. Identify the resource that the virtual machine relies on the most. That resource is most likely to affect the virtual machine's performance if the virtual machine is constrained by it.
3. Give a virtual machine more resources. Or decrease the resources of other virtual machines.
4. After making more of the limiting resource available to the virtual machine, take another benchmark and record changes.

Be extra cautious when making changes to production systems because a change might have a negative effect on the performance of the virtual machines.

Resource-Monitoring Tools

Slide 8-42

Many resource and performance monitoring tools are available to administrators to use with vSphere.

Inside the Guest OS

Perfmon DLL
Task Manager

Outside the Guest OS

vCenter Server
performance charts
vRealize Operations
Hyperic
vSphere/ESXi system logs
resxtop and esxtop

Tools in the guest operating system are available from sources external to VMware, and are utilized in various VMware applications. Many tools used outside of the guest OS are made available by VMware for use with vSphere and other applications.


A partial list of some of these resource monitoring tools is shown.

Guest Operating System Monitoring Tools

Slide 8-43

To monitor performance in the guest operating system, use tools that you are familiar with, such as Windows Task Manager.

Windows Task Manager



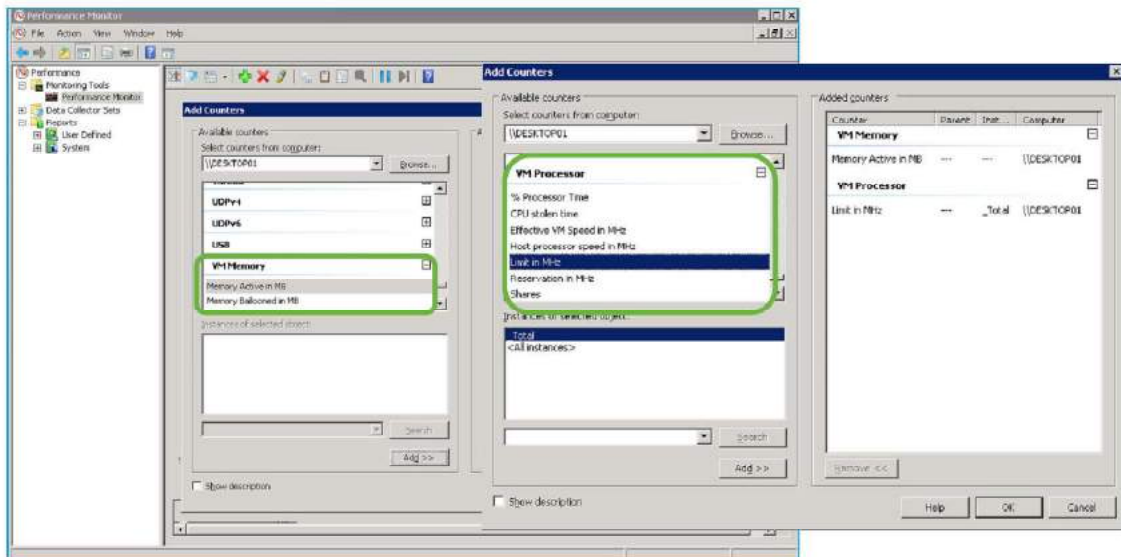
Windows Task Manager helps you measure CPU and memory use in the guest operating system.

The measurements that you take with tools in the guest operating system reflect resource use of the guest operating system, not necessarily of the virtual machine itself.

Using Perfmon to Monitor Virtual Machine Resources

Slide 8-44

The Perfmon DLL in VMware Tools provides virtual machine processor and memory objects to access host statistics inside a virtual machine.



VMware Tools includes a library of functions called the Perfmon DLL. Perfmon allows you to access key host statistics in a guest virtual machine. The Perfmon performance objects (VM Processor and VM Memory) allow you to view actual CPU and memory use alongside observed CPU and memory use of the guest operating system.

For example, you can use the VM Processor object to view the % Processor Time counter, which monitors the current load of the virtual machine's virtual processor. Likewise, you can use the Processor object and view the % Processor Time counter (not shown), which monitors the total use of the processor by all running processes.

Using esxtop to Monitor Virtual Machine Resources

Slide 8-45

`esxtop` is the primary real-time performance monitoring tool for vSphere:

- This tool can be run from an ESXi host local command line as `esxtop`.
- This tool can also be run remotely from VMware vSphere® Command-Line Interface as `resxtop`.
- `esxtop` works like the `top` performance utility found in Linux operating systems.

In this example, you enter `c` and uppercase `V` to view CPU metrics for virtual machines.

```
2:48:37pm up 7 days 13:36, 493 worlds, 3 VMs, 3 vCPUs; CPU load average: 0.03, 0.03, 0.02
PCPU USED(%): 1.8 3.9 AVG: 2.8
PCPU UTIL(%): 2.2 4.2 AVG: 3.2
```

ID	GID NAME	NWLD	%USED	%RUN	%SYS	%WAIT	%VMWAIT	%RDY	%IDLE	%OVLDP	%CSTP
52127	52127 VM2-1	8	0.15	0.14	0.00	800.00	0.72	0.03	99.67	0.00	0.00
43316	43316 VM1-2	8	0.11	0.10	0.01	800.00	0.74	0.07	99.67	0.00	0.00
55009	55009 VM2-2	8	0.11	0.10	0.01	800.00	0.73	0.12	99.63	0.00	0.00

You can run the `esxtop` utility by using vSphere ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

The first eight lines contain lowercase and uppercase letters to specify which fields appear in which order on the CPU, memory, storage adapter, storage device, virtual machine storage, network, interrupt, and CPU power panels. The letters correspond to the letters in the Fields or Order panels for the respective `esxtop` panel.

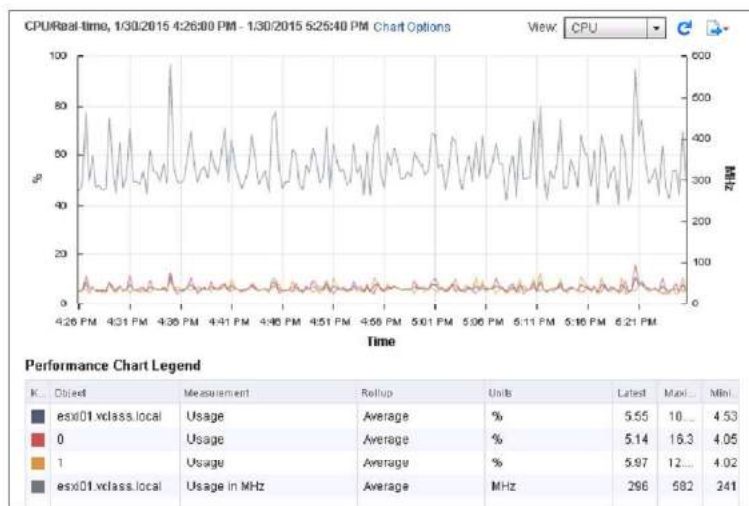
The ninth line contains information on the other options. If you saved a configuration in secure mode, you do not get an insecure `esxtop` without removing the `s` from the seventh line of your `esxtop50rc` file. A number specifies the delay time between updates. As in interactive mode, typing `c`, `m`, `d`, `u`, `v`, `n`, `I`, or `p` determines the panel with which `esxtop` starts.

About Monitoring Inventory Objects with Performance Charts

Slide 8-46

The vSphere statistics subsystem collects data on the resource usage of inventory objects:

- Counters and metric groups
- Collection levels and collection intervals
- Data availability



Data on a wide range of metrics is collected at frequent intervals, processed, and archived in the vCenter Server database.

You can access statistical information through command-line monitoring utilities or by viewing performance charts in vSphere Web Client and vSphere Client.

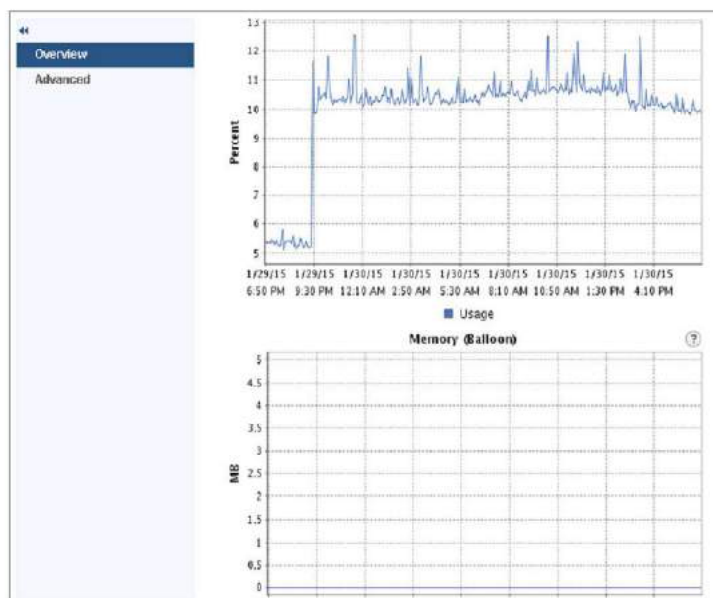
- Counters and metric groups: vCenter Server systems and hosts use data counters to query for statistics. A data counter is a unit of information relevant to a given inventory object or device. Each counter collects data for a different statistic in a metric group. For example, the disk metric group includes separate data counters to collect data for disk read rate, disk write rate, and disk usage. Statistics for each counter are rolled up after a specified collection interval. Each data counter consists of several attributes that are used to determine the statistical value collected.
- Collection levels and collection intervals: 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 your vCenter Server database.
- Data availability: Real-time data appears in the performance charts only for hosts and virtual machines that are powered on. Historical data appears for all supported inventory objects, but might be unavailable during certain circumstances.

Working with Overview Performance Charts

Slide 8-47

The overview performance charts display the most common metrics for an object in the inventory.

Host's Performance Charts
Partial Overview Panel



Two types of VMware performance charts are available in vSphere Web Client: Overview charts and Advanced charts.

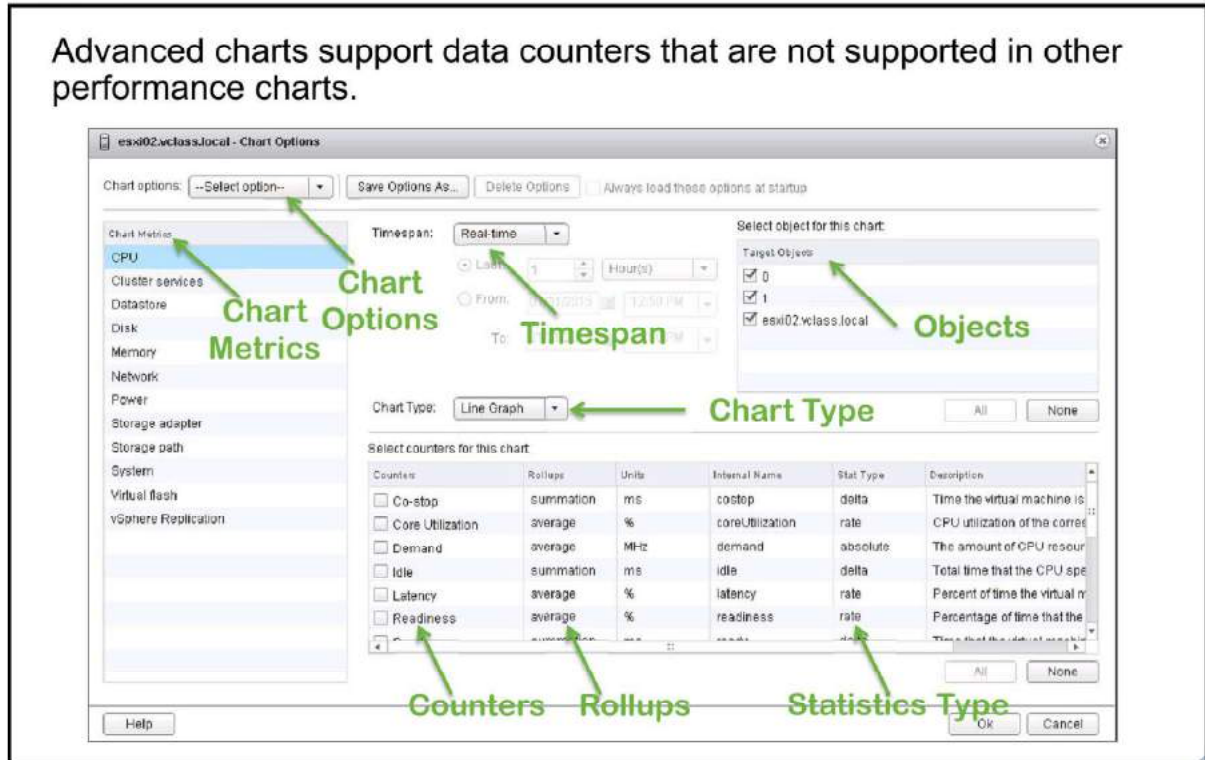
The Overview performance charts show the performance statistics that VMware considers most useful for monitoring performance and diagnosing problems.

Depending on the object that you select in the inventory, the performance charts in the Overview panel provide you with a quick visual representation of how your host or virtual machine is doing. The example shows a partial view of the overview performance charts for an ESXi host.

Working with Advanced Performance Charts

Slide 8-48

Advanced charts support data counters that are not supported in other performance charts.



Both vSphere Client and vSphere Web Client enable you to customize the appearance of the advanced performance charts.

Advanced charts include the following features:

- More information: Hover over a data point in a chart and details about that specific data point are displayed.
- Customizable charts: Change chart settings. Save custom settings to create your own charts.
- Export to spreadsheet.
- Save to image file or spreadsheet.

Chart Options: Real-Time and Historical

Slide 8-49

vCenter Server stores statistics at different specificities.

Time Interval	Data Frequency	Number of Samples
Real-Time (past hour)	20 seconds	180
Past Day	5 minutes	288
Past Week	30 minutes	336
Past Month	2 hours	360
Past Year	1 day	365

Real-time information is information generated for the past hour at a 20-second specificity. Historical information is information generated for the past day, week, month, or year, at varying specificities.

By default, vCenter Server has four collection intervals: day, week, month, and year. Each interval specifies a length of time that statistics are archived in the vCenter Server database. You can configure which intervals are enabled and for what period of time. You can also configure the number of data counters used during a collection interval by setting the collection level. Together, the collection interval and the collection level determine how much statistical data is collected and stored in your vCenter Server database. For example, using the table, past-day statistics show one data point every 5 minutes, for a total of 288 samples. Past-year statistics show 1 data point per day, or 365 samples.

Real-time statistics are not stored in the database. They are stored in a flat file on ESXi hosts and in memory on vCenter Server systems. ESXi hosts collect real-time statistics only for the host or the virtual machines available on the host. Real-time statistics are collected directly on an ESXi host every 20 seconds. If you query for real-time statistics, vCenter Server queries each host directly for the data. vCenter Server does not process the data at this point. vCenter Server only passes the data to vSphere Web Client and vSphere Client.

On ESXi hosts, the statistics are kept for 30 minutes, after which 90 data points are collected. The data points are aggregated, processed, and returned to vCenter Server. Then vCenter Server archives the data in the database as a data point for the day collection interval.

To ensure that performance is not impaired when collecting and writing the data to the database, cyclical queries are used to collect data counter statistics. The queries occur for a specified collection interval. At the end of each interval, the data calculation occurs.

Chart Types

Slide 8-50

Depending on the metric type and object, performance metrics are displayed in different types of charts.

Line chart:

- Each instance is shown separately.

Bar chart:


- Each instance is a bar in the chart.

Pie chart:

- Each instance is a slice in a circular pie.

Stacked chart:

- Graphs are stacked on top of one another.



You can use different chart types to change how metrics appear:

- Line chart: Displays metrics for a single inventory object. The data for each performance counter is plotted on a separate line in the chart. For example, a network chart for a host can contain two lines: one showing the number of packets received, and one showing the number of packets transmitted.
- Bar chart: Displays storage metrics for datastores in a selected data center. Each datastore is represented as a bar in the chart. Each bar displays metrics based on the file type: virtual disks, snapshots, swap files, and other files.
- Pie chart: Displays storage metrics for a single object, based on the file types or virtual machines. For example, a pie chart for a datastore can display the amount of storage space occupied by the virtual machines taking up the largest space.
- Stacked chart: Displays metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term Other. For example, a host's stacked CPU usage chart displays CPU usage metrics for the five virtual machines on the host that are consuming the most CPU. The Other amount contains the total CPU usage of the remaining virtual machines. The metrics for the host itself are displayed in separate line charts. Stacked charts are useful in comparing resource allocation and usage across multiple hosts or virtual machines. By default, the 10 child objects with the highest data counter values appear.

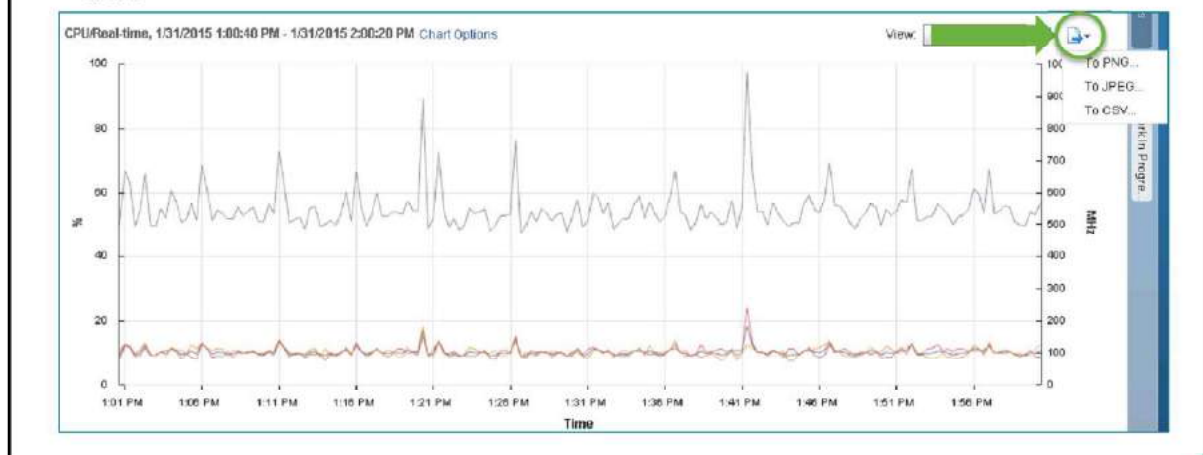
Saving Charts

Slide 8-51

You click the **Save Chart** icon above the graph to save performance chart information.

You can save information in these formats:

- PNG
- JPEG
- CSV



In vSphere Web Client, you can save data from the advanced performance charts to a file in various graphics formats or in Microsoft Excel format. When you save a chart, you select the file type and save the chart to the location of your choice.

Objects and Counters

Slide 8-52

The performance charts graphically display CPU, memory, disk, network, and storage metrics for devices and entities managed by vCenter Server.

Objects are instances or aggregations of devices:

- Examples: vCPU0, vCPU1, vmhba1:1:2, aggregate over all NICs

Counters identify which statistics to collect:

- Examples:
 - CPU: Used time, ready time, usage (%)
 - NIC: Network packets received
 - Memory: Memory swapped

vCenter Server allows the user to determine how much or how little information about a specific device type is displayed. You can control the amount of information a chart displays by selecting one or more objects and counters.

An object refers to an instance for which a statistic is collected. For example, you might collect statistics for an individual CPU (for example, vCPU0, vCPU1), all CPUs, a host, or a specific network device.

A counter represents the actual statistic that you are collecting. An example is the amount of CPU used or the number of network packets per second for a given device.

Statistics Type

Slide 8-53

The statistics type is the unit of measurement used during the statistics interval.

Statistics Type	Description	Example
Rate	Value over the current interval	CPU use (MHz)
Delta	Change from previous interval	CPU ready time
Absolute	Absolute value, independent of interval	Memory active

The statistics type refers to the measurement used during the statistics interval and is related to the unit of measurement.

The statistics type is one of the following:

- Rate: Value over the current statistics interval
- Delta: Change from previous statistics interval
- Absolute: Absolute value (independent of the statistics interval)

For example, CPU usage is a rate, CPU ready is a delta, and Memory active is an absolute value.

Rollup

Slide 8-54

Rollup is the conversion function between statistics intervals:

- 5 minutes of past-hour statistics are converted to 1 past-day value:
 - Fifteen 20-second statistics are rolled up into a single value.
- 30 minutes of past-day statistics are converted to 1 past-week value:
 - Six 5-minute statistics are rolled up into a single value.
- Other rollup types: Minimum, Maximum

Rollup Type	Conversion Function	Sample Statistic
Average	Average of data points	CPU use (average)
Summation	Sum of data points	CPU ready time (milliseconds)
Latest	Last data point	Uptime (days)

When looking at different historical intervals, data is displayed at different specificities. Past-hour statistics are shown at a 20-second specificity, and past-day statistics are shown at a 5-minute specificity. The averaging that is done to convert from one time interval to another is called rollup.

Different rollup types are available. The rollup type determines the type of statistical values returned for the counter:

- Average: The data collected during the interval is aggregated and averaged.
- Minimum: The minimum value is rolled up.
- Maximum: The maximum value is rolled up.

The minimum and maximum values are collected and displayed only in collection level 4. Minimum and maximum rollup types are used to capture peaks in data during the interval. For real-time data, the value is the current minimum or current maximum. For historical data, the value is the average minimum or average maximum.

For example, the following information for the CPU usage chart shows that the average is collected at collection level 1 and the minimum and maximum values are collected at collection level 4.

- Counter: usage
- Unit: Percentage (%)
- Rollup Type: Average (Minimum/Maximum)
- Collection Level: 1 (4)

Statistics levels include:

- Summation: The data collected is summed. The measurement displayed in the performance chart represents the sum of data collected during the interval.
- Latest: The data collected during the interval is a set value. The value displayed in the performance chart represents the current value.

For example, if you look at the CPU Used counter in a CPU performance chart, the rollup type is summation. This means that for a given 5-minute interval, the sum of all of the 20-second samples in that interval is represented.

Setting Log Levels

Slide 8-55

Setting log levels enables the user to control the quantity and type of information logged.

Examples of when to set log levels:

- When troubleshooting complex issues, set the log level to verbose or trivia. Troubleshoot and set it back to info.
- To control the amount of information being stored in the log files.


Option	Description
None	Turns off logging
Error (errors only)	Displays only error log entries
Warning (errors and warnings)	Displays warning and error log entries
Info (normal logging)	Displays information, error, and warning log entries
Verbose	Displays information, error, warning, and verbose log entries
Trivia (extended verbose)	Displays information, error, warning, verbose, and trivia log entries

Changes to the logging settings take effect immediately. You do not have to restart the vCenter Server system.

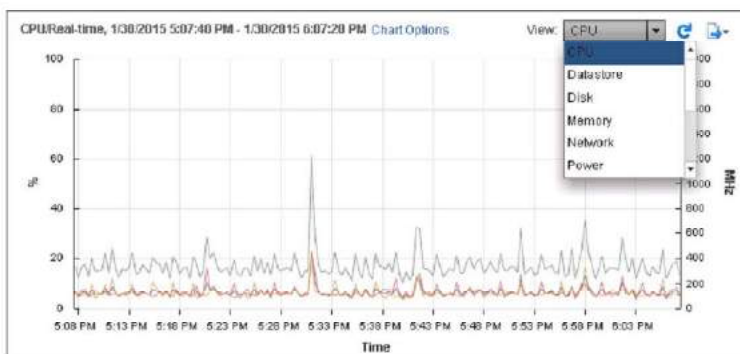
Interpreting Data from the Tools

Slide 8-56

vCenter Server monitoring tools and guest operating system monitoring tools provide different points of view.



Task Manager in Guest Operating System



CPU Usage Chart for Host

The key to interpreting performance data is to observe the range of data from the guest operating system, the virtual machine, and the host's perspective.

The CPU usage statistics in Task Manager, for example, do not give you the complete picture. You should also view CPU usage for the virtual machine and the host on which the virtual machine is located. Use the performance charts in vCenter Server to view this data.



CPU-Constrained Virtual Machine

Slide 8-57

If CPU usage is continuously high, the virtual machine is constrained by CPU. However, the host might have enough CPU for other virtual machines to run.

Multiple virtual machines are constrained by CPU if the following conditions are present:

- High CPU usage in the guest operating system
- Relatively high CPU ready values for the virtual machines



Single Virtual Machine CPU Usage

To check if a virtual machine is being constrained by CPU resources, check CPU usage in the guest operating system using, for example, Task Manager.

If CPU usage is high, check the virtual machine's CPU use. In vSphere Web Client and vSphere Client, select the virtual machine in the inventory, click the **Monitor** tab, and click the **Performance** button. Use either the overview charts or the advanced charts to view CPU usage. On the slide, an advanced chart tracking a virtual machine's CPU usage is displayed.

If a virtual machine's CPU usage remains high over a period of time, the virtual machine is constrained by CPU. Other virtual machines on the host might have enough CPU resources to satisfy their needs.

If more than one virtual machine is constrained by CPU, the key indicator is CPU ready time. Ready time refers to the interval when a virtual machine is ready to execute instructions but cannot, because it cannot get scheduled onto a CPU. Several factors affect the amount of ready time:

- Overall CPU use: You are more likely to see ready time when use is high because the CPU is more likely to be busy when another virtual machine becomes ready to run.
- Number of resource consumers (in this case, guest operating systems): When a host is running a larger number of virtual machines, the scheduler is more likely to need to queue a virtual machine behind virtual machines that are already running or queued.

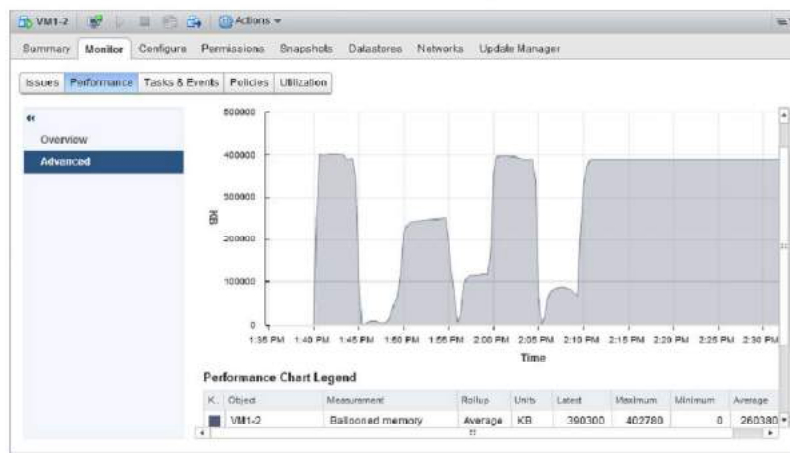
A good ready time value varies from workload to workload. To find a good ready time value for your workload, collect ready time data over time for each virtual machine. When you have this ready time data for each virtual machine, estimate how much of the observed response time is ready time. If the shortfalls in meeting response-time targets for the applications appear largely because of the ready time, take steps to address the excessive ready time.

Memory-Constrained Virtual Machine

Slide 8-58

Check the virtual machine's ballooning activity to determine if the virtual machine is constrained for memory:

- If ballooning activity is high, this state might not be a problem if all virtual machines have sufficient memory.
- If ballooning activity is high and the guest operating system is swapping, then the virtual machine is constrained for memory.




When a virtual machine experiences ballooning activity, some of the guest operating system's physical memory is being reclaimed from the virtual machine by the balloon driver. If a virtual machine experiences high ballooning values, this might not be a problem if the virtual machine continues to have the memory that it needs. But if a virtual machine experiences high ballooning activity over time and its guest operating system starts to page, the virtual machine might be constrained for memory.

Memory-Constrained Host

Slide 8-59

If active host-level swapping is occurring, then host memory is overcommitted.




If multiple virtual machines are being constrained for memory, you see high ballooning activity and the guest operating systems paging. You also see the virtual machine itself being swapped in and out by the VMkernel. This serious situation indicates that the host memory is overcommitted. The amount of memory on the host needs to be increased.

Monitoring Active Memory of a Virtual Machine

Slide 8-60

Monitor for increases in active memory on the host:

- Host active memory refers to active physical memory used by virtual machines and the VMkernel.
- If amount of active memory is high, this situation might lead to virtual machines that are memory-constrained.



A general memory counter to monitor over time is a host's active memory counter. Host active memory refers to the amount of physical memory that is actively being used by virtual machines and the VMkernel. If the active memory of certain virtual machines is continuously high, this state might lead to those virtual machines being constrained by memory.

Disk-Constrained Virtual Machines

Slide 8-61

Disk-intensive applications can saturate the storage or the path.

If you suspect that a virtual machine is constrained by disk access:

- Measure the throughput and latency between the virtual machine and storage.
- Use the advanced performance charts to monitor:
 - Read rate and write rate
 - Read latency and write latency

Disk performance problems are commonly caused by saturating the underlying physical storage hardware. You can use the vCenter Server advanced performance charts to measure storage performance at different levels. These charts can provide you with insight into a virtual machine's performance by allowing you to monitor everything from the virtual machine's datastore to a specific storage path.

If you select a host object, you can view throughput and latency for a datastore, a storage adapter, or a storage path. The storage adapter charts are available only for Fibre Channel storage. The storage path charts are available for Fibre Channel and iSCSI storage, not NFS.

If you select a virtual machine object, you can view throughput and latency for the virtual machine's datastore or specific virtual disk.

To monitor throughput, view the Read rate and Write rate counters. To monitor latency, view the Read latency and Write latency counters.

Monitoring Disk Latency

Slide 8-62

To determine disk performance problems, monitor two disk latency data counters:

- Kernel command latency:
 - The average time spent in the VMkernel per SCSI command.
 - High numbers (greater than 2 or 3 ms) represent either an overworked array or an overworked host.
- Physical device command latency:
 - The average time the physical device takes to complete a SCSI command.
 - High numbers (greater than 15 or 20 ms) represent a slow or overworked array.

A reliable way to determine whether your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. Use the advanced performance charts to view these statistics. In particular, monitor the following counters:

- Kernel command latency: This data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0 through 1 milliseconds. If the value is greater than 4 milliseconds, the virtual machines on the ESXi host are trying to send more throughput to the storage system than the configuration supports.
- Physical device command latency: This data counter measures the average amount of time, in milliseconds, for the physical device to complete a SCSI command. Depending on your hardware, a number greater than 15 milliseconds indicates that the storage array might be slow or overworked.

Network-Constrained Virtual Machines

Slide 8-63

Network-intensive applications often bottleneck on path segments outside the ESXi host:

- Example: WAN links between server and client

If you suspect that a virtual machine is constrained by the network:

- Verify that VMware Tools is installed.
 - Enhanced network drivers are available.
- Measure the effective bandwidth between the virtual machine and its peer system.
- Check for dropped receive packets and dropped transmit packets.

Like disk performance problems, network performance problems are commonly caused by saturating a network link between client and server. Use a tool like Iometer, or a large file transfer, to measure the effective bandwidth.

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use the advanced performance charts to examine the droppedTx and droppedRx network counter values of a virtual machine.

In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. In some instances, large packets can result in high network latency. When network packets are small, more packets are transferred, but the network speed is slower because more CPU is required to process the data.

Lab 17: Monitoring Virtual Machine Performance

Slide 8-64

Use the system monitoring tools to reflect the CPU workload

1. Create CPU Workload
2. Use Performance Charts to Monitor CPU Utilization
3. Undo Changes Made to the Virtual Machines

Review of Learner Objectives


Slide 8-65

You should be able to meet the following objectives:

- Use the performance-tuning methodology and resource monitoring tools
- Use performance charts to view and improve performance
- Monitor the key factors that can affect the virtual machine's performance: CPU, memory, disk, and network bandwidth use

Lesson 5: Using Alarms

Slide 8-66



Lesson 5: Using Alarms

Learner Objectives

Slide 8-67

By the end of this lesson, you should be able to meet the following objectives:

- Create alarms with condition-based triggers
- Create alarms with event-based triggers
- View and acknowledge triggered alarms

About Alarms

Slide 8-68

An alarm is a notification that occurs in response to selected events or conditions that occur with an object in the inventory.

Default alarms exist for various inventory objects:

- Many default alarms for hosts and virtual machines

You can create custom alarms for a wide range of inventory objects:

- Virtual machines, hosts, clusters, data centers, datastores, datastore clusters, networks, distributed switches, and distributed port groups

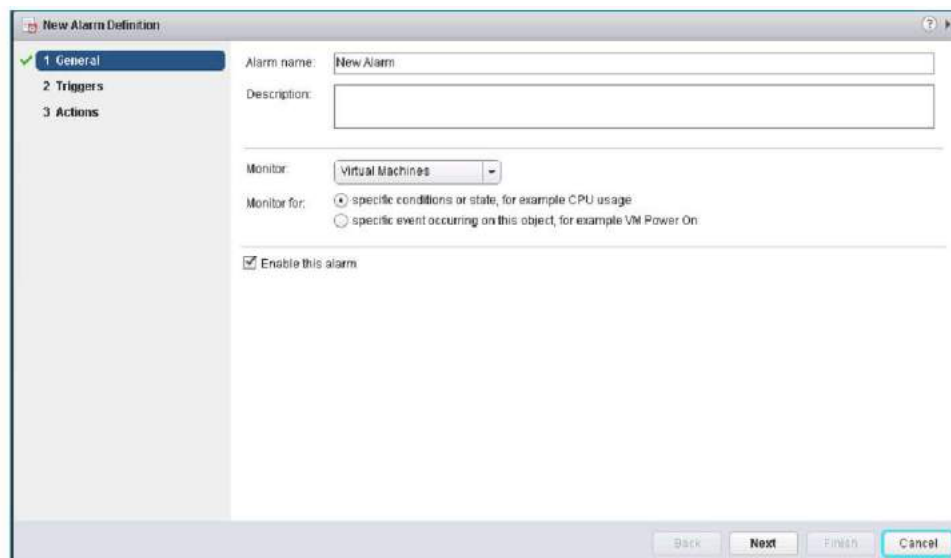
Alarms are notifications in response to selected events or conditions that occur with an object in the inventory. VMware provides a set of predefined alarms for most objects in the vCenter Server inventory. For example, alarms exist for host, virtual machine, and resource pool memory and CPU usage. You can also define custom alarms for virtual machines, hosts, clusters, data centers, datastores, datastore clusters, networks, distributed switches, and distributed port groups.

The predefined alarms are configurable. If the predefined alarms do not address the condition, state, or event that you want to monitor, define custom alarms instead of modifying predefined alarms.

Alarm Settings

Slide 8-69

To monitor your environment, you can create and modify alarm definitions in vSphere Web Client and vSphere Client.



The screenshot shows the 'New Alarm Definition' dialog box with the following fields and options:

- Alarm name:** New Alarm
- Description:** (empty text box)
- Monitor:** Virtual Machines (dropdown menu)
- Monitor for:**
 - specific conditions or state, for example CPU usage
 - specific event occurring on this object, for example VM Power On
- Enable this alarm

Navigation buttons at the bottom: Back, Next, Finish, Cancel.

On the General page, you name the alarm, give it a description, and give it an alarm type.

You also select what to monitor:

- **Monitor for specific conditions or state:** A condition-based alarm. You can create condition-based alarms for virtual machines, hosts, and datastores.
- **Monitor for specific events occurring on this object:** An event-based alarm. You can create event-based alarms for virtual machines, hosts, clusters, data centers, datastores, datastore clusters, networks, distributed virtual switches, and distributed virtual port groups.

You can also use the **General** page to enable or disable the alarm by selecting or deselecting the **Enable this alarm** check box.

Alarm Triggers

Slide 8-70

An alarm requires a trigger.

Types of triggers:

- Condition or state trigger: Monitors the current condition or state. Examples:
 - A virtual machine's current snapshot is above 2 GB in size.
 - A host is using 90 percent of its total memory.
 - A datastore has been disconnected from all hosts.
- Event: Monitors events. Examples:
 - The health of a host's hardware has changed.
 - A license has expired in the data center.
 - A host has left the distributed switch.


You configure alarm triggers to generate warnings and alerts when the specified criteria are met. Alarms have condition or state triggers and event triggers.

- Condition or state triggers: Monitor the current condition or state of virtual machines, hosts, and datastores. Conditions or states include power states, connection states, and performance metrics such as CPU and disk usage.
- Event triggers: Monitor events that occur in response to operations occurring with a managed object in the inventory or the vCenter Server system. For example, an event is recorded each time a virtual machine (which is a managed object) is cloned, created, deleted, deployed, and migrated.

Configuring Condition Triggers

Slide 8-71

Condition or state triggers monitor metrics for a host, virtual machine, or datastore.



On the slide, you can configure a condition trigger so that:

- A virtual machine’s CPU usage must be above 75 percent for more than 5 minutes to generate a warning
- Above 90 percent for more than 5 minutes to generate an alert

Time periods are used to ensure that the metric conditions are valid and not caused by incidental spikes.

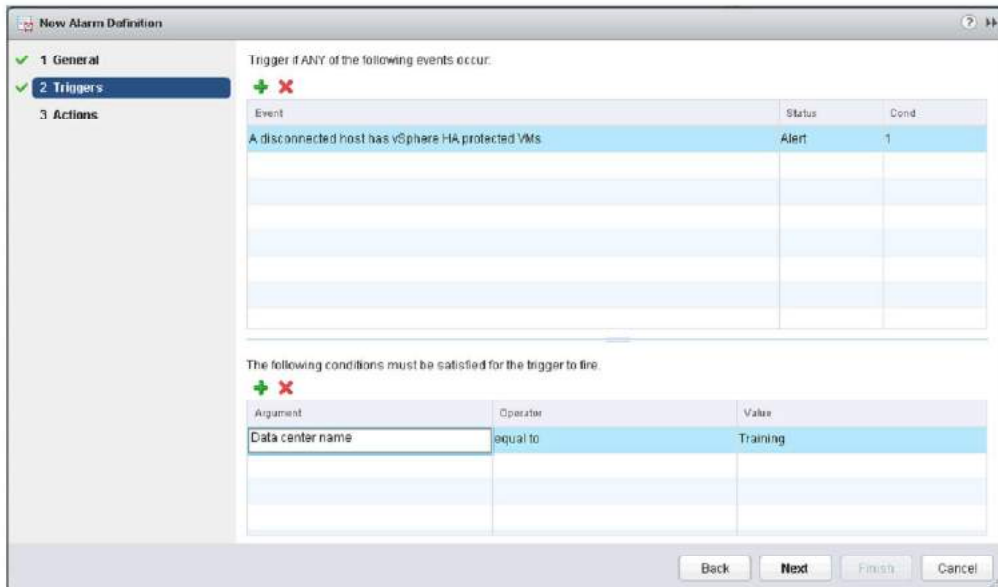
Also on the slide, you can configure a state trigger to generate an alert if a virtual machine has no heartbeat. When the triggering conditions are no longer true, a triggered alarm resets itself and no longer triggers.

If you add multiple triggers, you can choose to trigger the alarm if any one of the conditions is satisfied or if all of the conditions are satisfied.

Configuring Event Triggers

Slide 8-72

Event triggers monitor the current state of a host, virtual machine, or datastore.



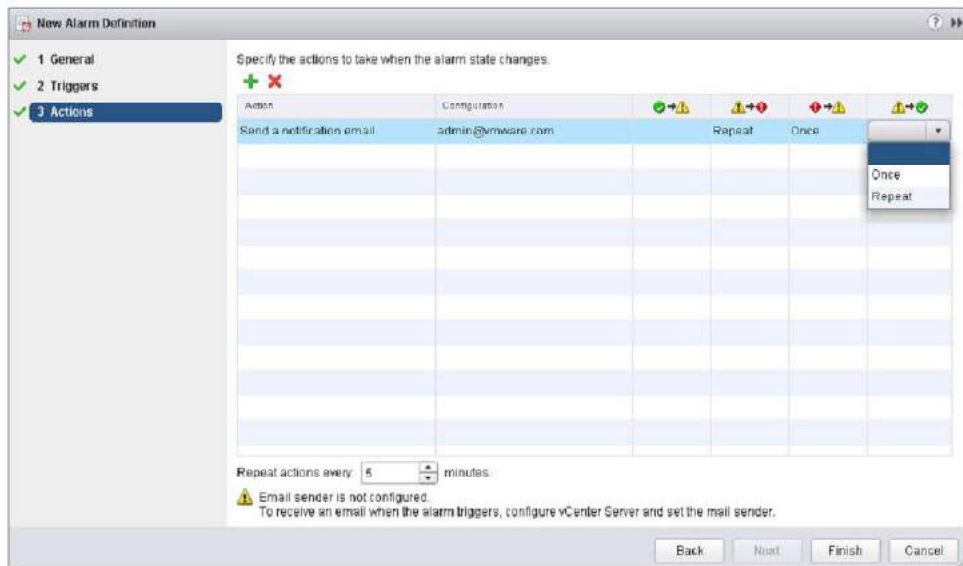
Event triggers do not rely on thresholds or durations. They use arguments, operators, and values to identify the triggering condition. In the example, the event trigger monitors the **A disconnected host has vSphere HA protected VMs** event.

Whenever a host that is housing virtual machines that are to be protected by HA is disconnected, it triggers the alert. A condition has also been configured to trigger the alert only if the guest is in a data center named Training.

Configuring Actions

Slide 8-73

You can define actions that the system performs when the alarm is triggered or changes status.



Alarms are composed of a trigger and an action. An action is the operation that occurs in response to the trigger, for example, sending an email notification to one or more administrators.

In the New Alarm Definition dialog box, use the **Actions** tab to specify actions to take when the alarm is triggered. Colors and shapes are used to denote the alarm's severity: a green circle is normal, a yellow triangle is a warning, and a red diamond is an alert.

You can set alarms to trigger when the state changes:

- From a green circle to a yellow triangle
- From a yellow triangle to a red diamond
- From a red diamond to a yellow triangle
- From a yellow triangle to a green circle

For every action, you can specify an option for each color transition:

- **Empty** indicates no interest in the transition.
- **Once** tells vCenter Server to do the action only one time.

- **Repeat** tells vCenter Server to repeat the action until another color change occurs. The default is 5 minutes. The maximum is 2 days. The mentioned default time value is related to **Repeat action every** at the bottom of the dialog box.

Every alarm type has the following actions:

- Send a notification email
- Send a notification trap
- Run a command.

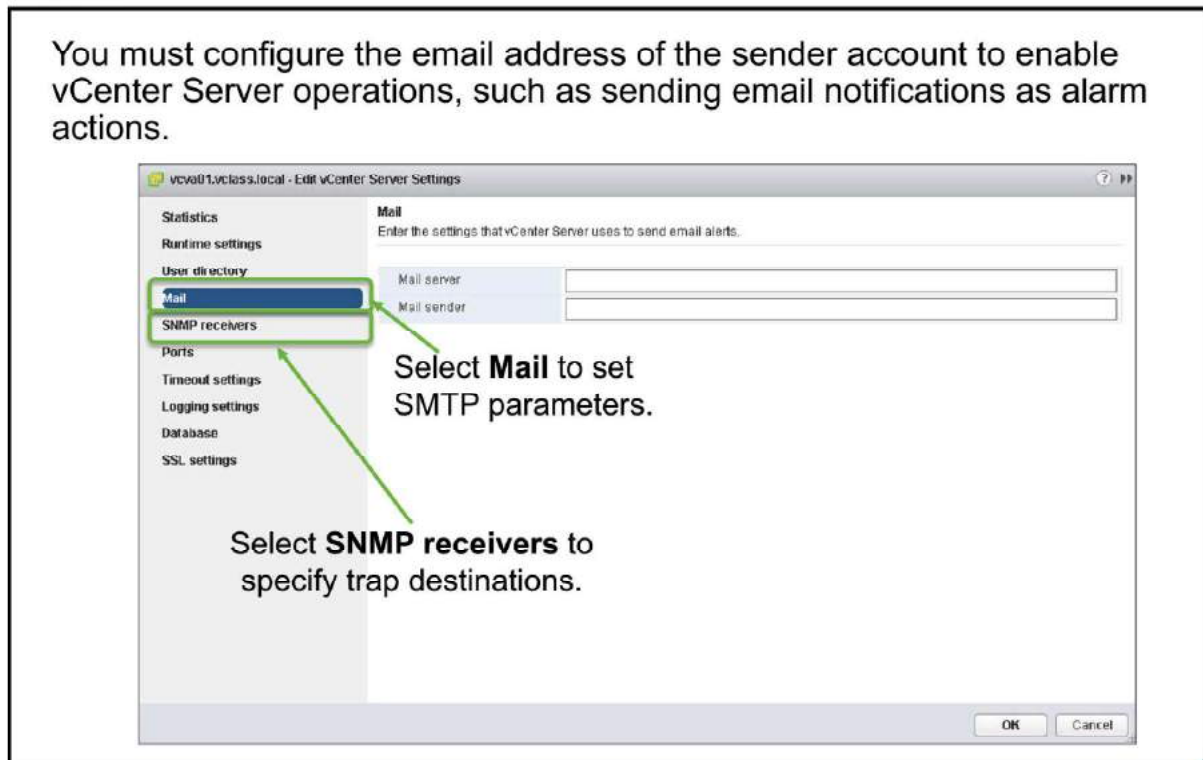
Virtual machine alarms and host alarms have more actions, such as:

- Migrate virtual machine
- Power on a virtual machine
- Power off a virtual machine
- Reboot guest on virtual machine
- Suspend a virtual machine
- Reboot host
- Shut down host

Configuring vCenter Server Notifications

Slide 8-74

You must configure the email address of the sender account to enable vCenter Server operations, such as sending email notifications as alarm actions.

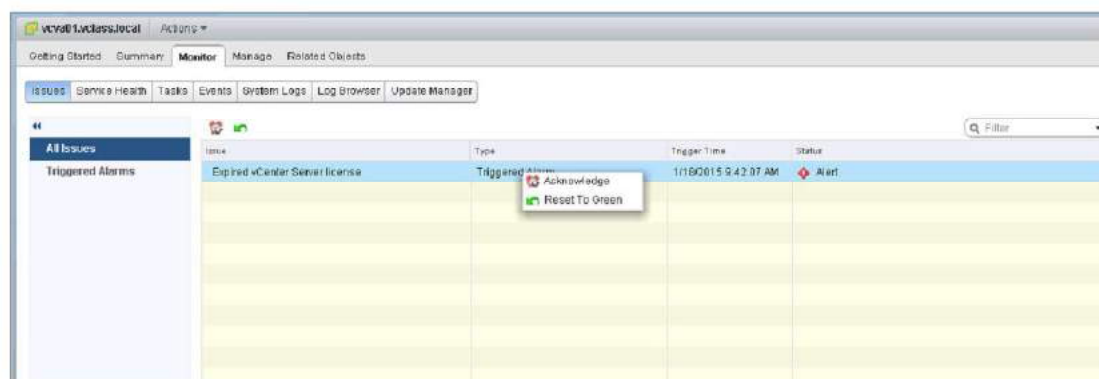


You can configure up to four receivers of SNMP traps. They must be configured in numerical order. Each SNMP trap requires a corresponding host name, port, and community.

Viewing and Acknowledging Triggered Alarms

Slide 8-75

The Acknowledge Alarm feature is used to track when triggered alarms are addressed.



After you acknowledge an alarm in vSphere Web Client or vSphere Client, its alarm actions are discontinued. Alarms are neither cleared nor reset when acknowledged.

Acknowledging an alarm lets other users know that you are taking ownership of the issue. For example, a host has an alarm set on it that monitors CPU usage and that sends an email to an administrator when the alarm is triggered. The host CPU usage spikes, triggering the alarm which sends an email to the host's administrator. The administrator acknowledges the triggered alarm to let other administrators know that the problem is being addressed, and to prevent the alarm from sending more email messages. However, the alarm is still visible in the system.

Lab 18: Using Alarms

Slide 8-76

Use the vCenter Server Appliance alarm feature

1. Create a Virtual Machine Alarm to Monitor a Condition
2. Create a Virtual Machine Alarm to Monitor an Event
3. Trigger Virtual Machine Alarms and Acknowledge the Alarms
4. Disable Virtual Machine Alarms

Review of Learner Objectives

Slide 8-77

You should be able to meet the following objectives:

- Create alarms with condition-based triggers
- Create alarms with event-based triggers
- View and acknowledge triggered alarms

Key Points

Slide 8-78

- For proper resource management, vSphere has mechanisms to enable less, more, or an equal amount of access to a defined resource.
- vSphere prevents a virtual machine from consuming large amounts of a resource and grants a guaranteed amount of a resource to a virtual machine whose performance is not adequate or requires a certain amount of a resource to run properly.
- A resource pool enables you to divide and allocate resources to virtual machines and other resource pools.
- A vApp is a container for one or more virtual machines. The vApp can be used to package and manage related applications.
- The **Performance** tab enables you to monitor the performance of a host or a virtual machine in real time or over a period of time.
- You use alarms to monitor the vCenter Server inventory and send notifications when selected events or conditions occur.

Questions?

MODULE 9

vSphere HA, vSphere Fault Tolerance, and Protecting Data

Slide 9-1



You Are Here

Slide 9-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. **vSphere HA, vSphere Fault Tolerance, and Protecting Data**
10. vSphere DRS
10. vSphere Update Manager

Importance

Slide 9-3

Most organizations rely on computer-based services like email, databases, and Web-based applications. The failure of any of these services can mean lost productivity and revenue.

Configuring highly available, computer-based services is extremely important for an organization to remain competitive in contemporary business environments.

Module Lessons

Slide 9-4

- | | |
|-----------|---|
| Lesson 1: | Introduction to vSphere HA |
| Lesson 2: | vSphere HA Architecture |
| Lesson 3: | Configuring vSphere HA |
| Lesson 4: | Introduction to vSphere Fault Tolerance |
| Lesson 5: | vSphere Replication and vSphere Data Protection |

Lesson 1: Introduction to vSphere HA

Slide 9-5



Lesson 1: Introduction to vSphere HA

Learner Objectives

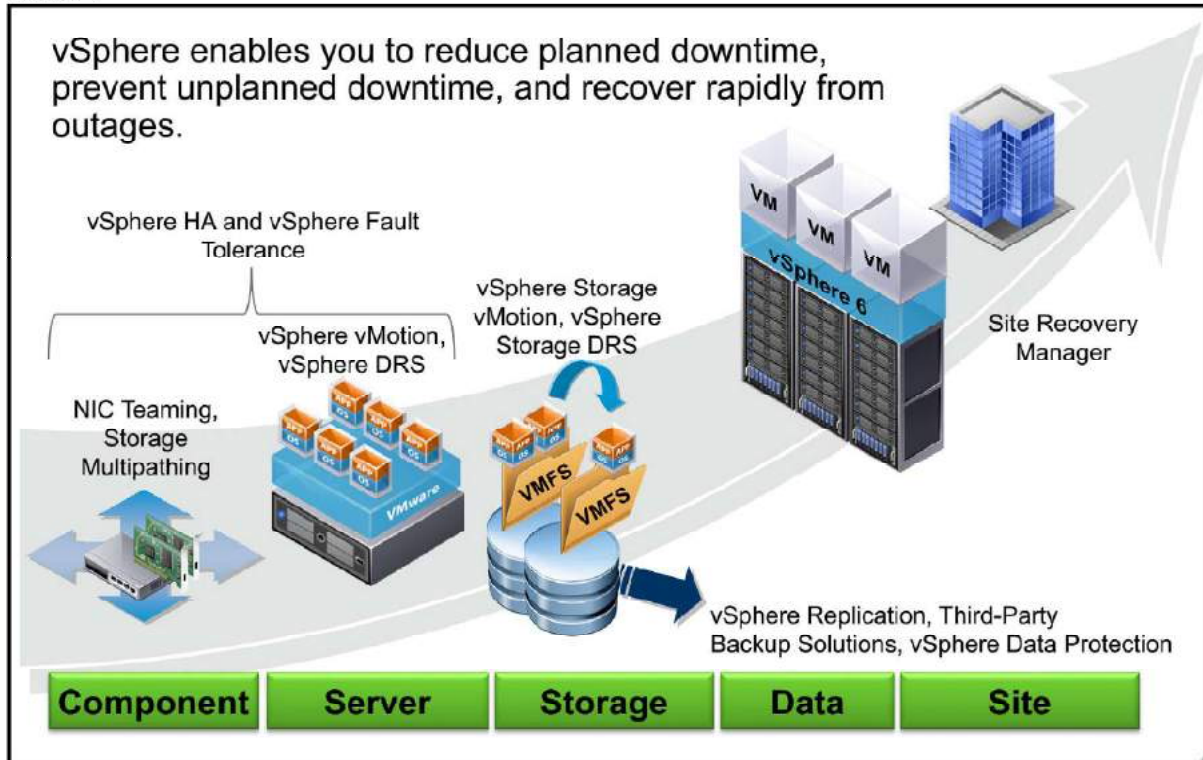
Slide 9-6

By the end of this lesson, you should be able to meet the following objectives:

- Describe options that you can configure to make your vSphere environment highly available
- Discuss the response of vSphere HA when an ESXi host, a virtual machine, or an application fails

Protection at Every Level

Slide 9-7



Downtime, whether planned or unplanned, brings with it considerable costs. However, solutions to ensure higher levels of availability have traditionally been costly, hard to implement, and difficult to manage.

VMware software makes it simpler and less expensive to provide higher levels of availability for important applications. With vSphere, organizations can easily increase the baseline level of availability provided for all applications and provide higher levels of availability more easily and cost effectively. With vSphere, you can:

- Provide higher availability independent of hardware, operating system, and applications.
- Reduce planned downtime for common maintenance operations.
- Provide automatic recovery in cases of failure.

Many methods ensure high availability in a virtualized environment. vSphere uses the following technologies to ensure that virtual machines running in the environment remain available:

- Virtual machine migration
- Multiple I/O adapter paths
- Virtual machine load balancing

- Fault tolerance
- Disaster recovery tools

vSphere HA provides a base level of protection for your virtual machines by restarting virtual machines in the event of a host failure. vSphere Fault Tolerance provides a higher level of availability, allowing users to protect any virtual machine from a host failure with no loss of data, transactions, or connections. vSphere Fault Tolerance provides continuous availability by ensuring that the states of the primary and secondary VMs are identical at any point in the instruction execution of the virtual machine

vSphere vMotion and vSphere Storage vMotion keep virtual machines available during a planned outage, for example, when hosts or storage must be taken offline for maintenance. System recovery from unexpected storage failures is simple, quick, and reliable with the encapsulation property of virtual machines. vSphere Storage vMotion can be used to support planned storage outages resulting from upgrades to storage arrays to newer firmware or technology and VMFS upgrades.

VMware vSphere® Replication™ enables a vSphere platform to protect virtual machines natively by copying their disk files to another location where they are ready to be recovered.

Virtual machine encapsulation is leveraged by backup applications such as vSphere Data Protection and third-party backup applications that support file and image-level backups to protect data. Backup solutions play prominent roles in recovering from deleted files or disks and corrupt or infected guest operating systems or file systems.

Site Recovery Manager allows you to quickly restore your organization's IT infrastructure, shortening the time that you experience a business outage. Site Recovery Manager automates setup, failover, and testing of disaster recovery plans. Site Recovery Manager requires that vCenter Server be installed at the protected site and at the recovery site. Site Recovery Manager also requires either host-based replication through vSphere Replication or preconfigured array-based replication between the protected site and the recovery site.

vCenter Server Availability: Recommendations

Slide 9-8

Make vCenter Server and the components that it relies on highly available.

vCenter Server relies on these major components:

- vCenter Server database:
 - Create a cluster for the database.
- Authentication identity source:
 - For example, vCenter Single Sign-On (a service within Platform Services Controller) and Active Directory.
 - Set up with multiple redundant servers.

Methods for making vCenter Server available:

- Use vSphere HA and vCenter Server High Availability to protect the vCenter Server virtual machine.

To provide high availability to vCenter Server, provide high availability for the components that it uses:

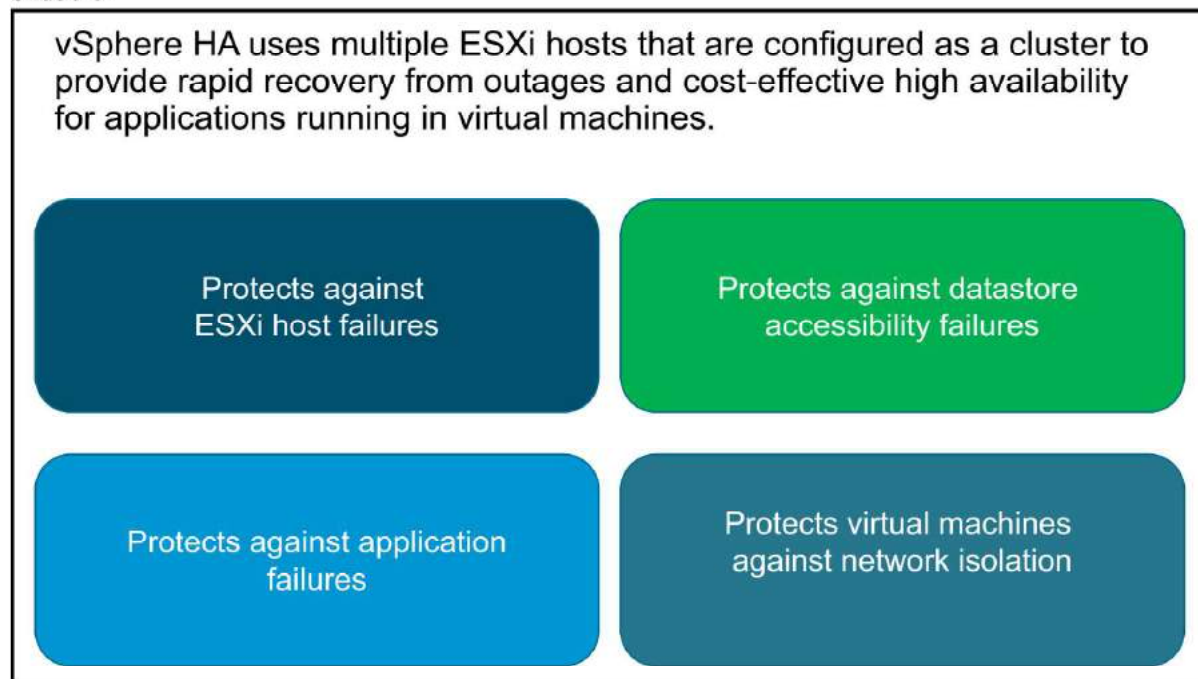
- vCenter Server database
- Authentication identity source, such as vCenter Single Sign-On or Active Directory (AD). vCenter Single Sign-On is one of the common infrastructure services within Platform Services Controller. If your environment uses AD, ensure that it is set up with multiple redundant servers.

High availability for vCenter Server can be implemented by using vSphere HA. vSphere HA protects against hardware and operating system failures.

To protect vCenter Server Appliance, you can use vCenter Server High Availability.

About vSphere HA

Slide 9-9



vSphere HA protects application availability in the following ways:

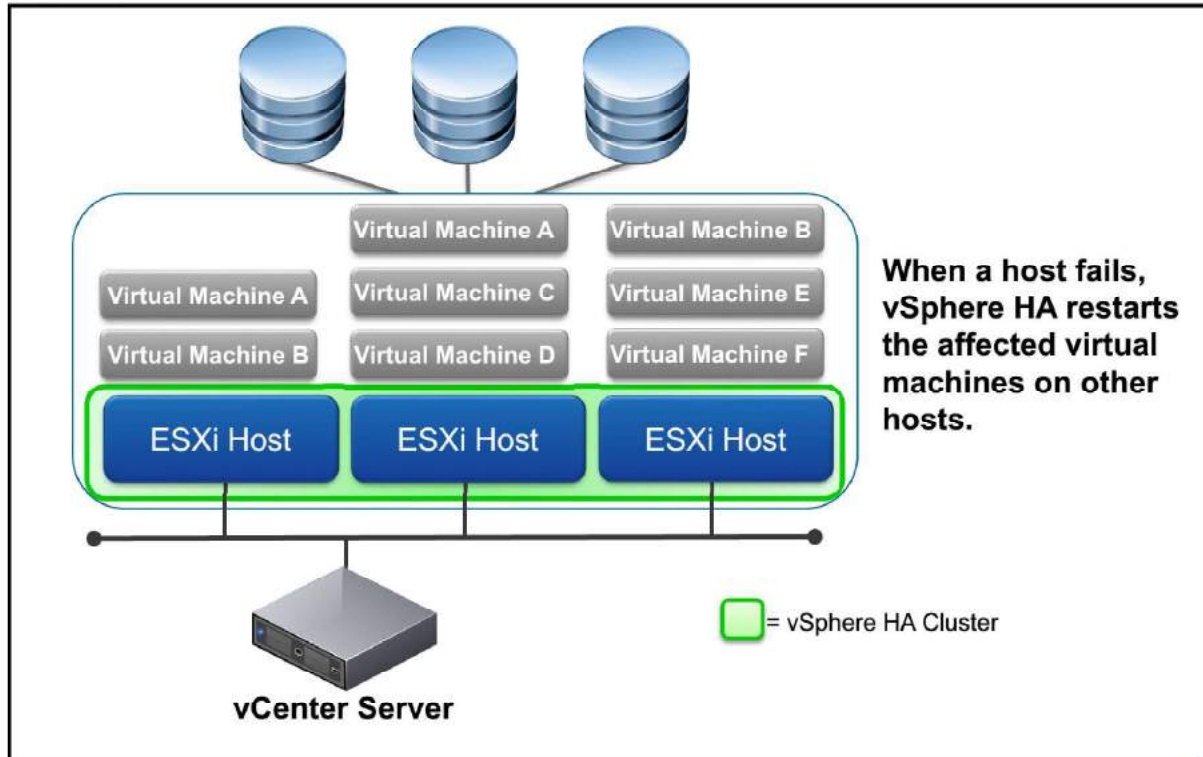
- It protects against a server failure by restarting the virtual machines on other hosts within the cluster.
- It protects against application failure by continuously monitoring a virtual machine and resetting it in the event that a failure is detected.
- It protects against datastore accessibility failures and provides automated recovery for affected virtual machines. With Virtual Machine Component Protection (VMCP), the affected virtual machines are restarted on other hosts that still have access to the 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.

Unlike other clustering solutions, vSphere HA provides the infrastructure to protect all workloads with the infrastructure:

- You do not need to install special software within the application or virtual machine. All workloads are protected by vSphere HA. After vSphere HA is configured, no actions are required to protect new virtual machines. They are automatically protected.
- You can combine vSphere HA with vSphere DRS to protect against failures and to provide load balancing across the hosts within a cluster.

vSphere HA Scenario: ESXi Host Failure

Slide 9-10

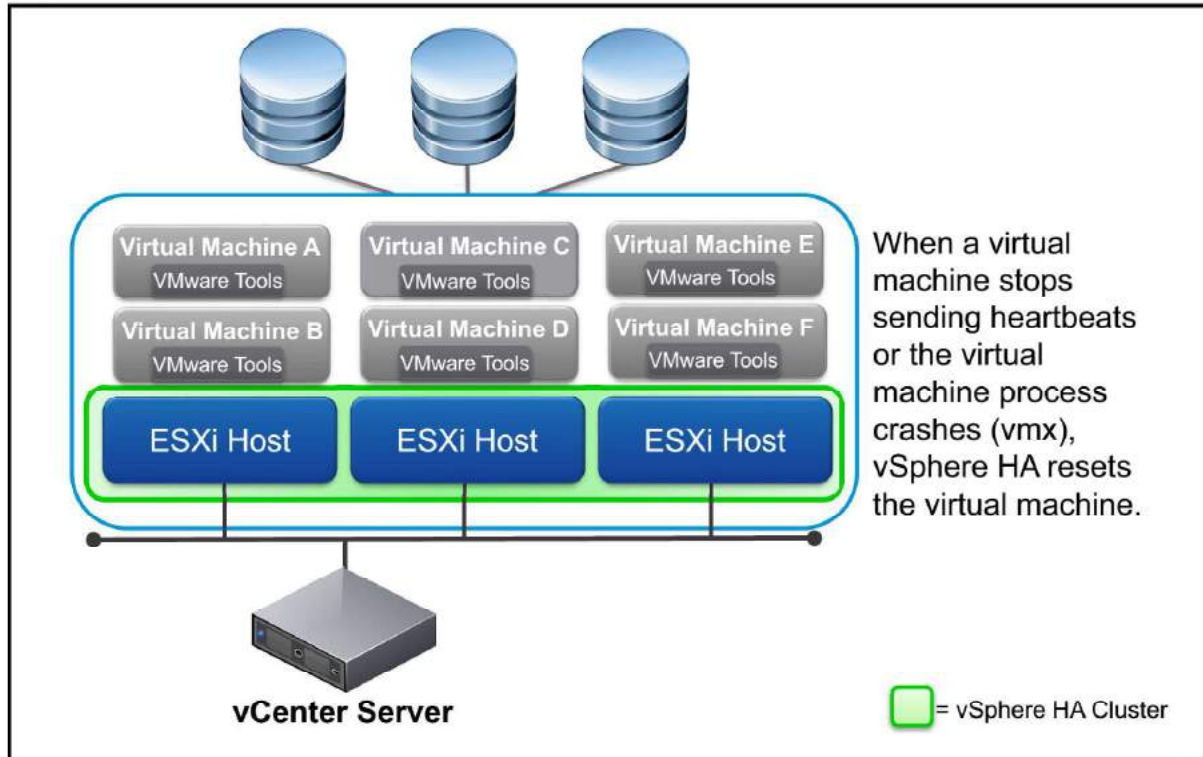


vSphere HA is also able to determine whether an ESXi host is isolated or has crashed. If an ESXi host crashes, vSphere HA has the responsibility of restarting virtual machines that were running on the failed host on the remaining hosts in the cluster.

In every cluster, downtime depends on how long it takes your guest operating systems and applications to restart when the virtual machine is failed over.

vSphere HA Scenario: Guest Operating System Failure

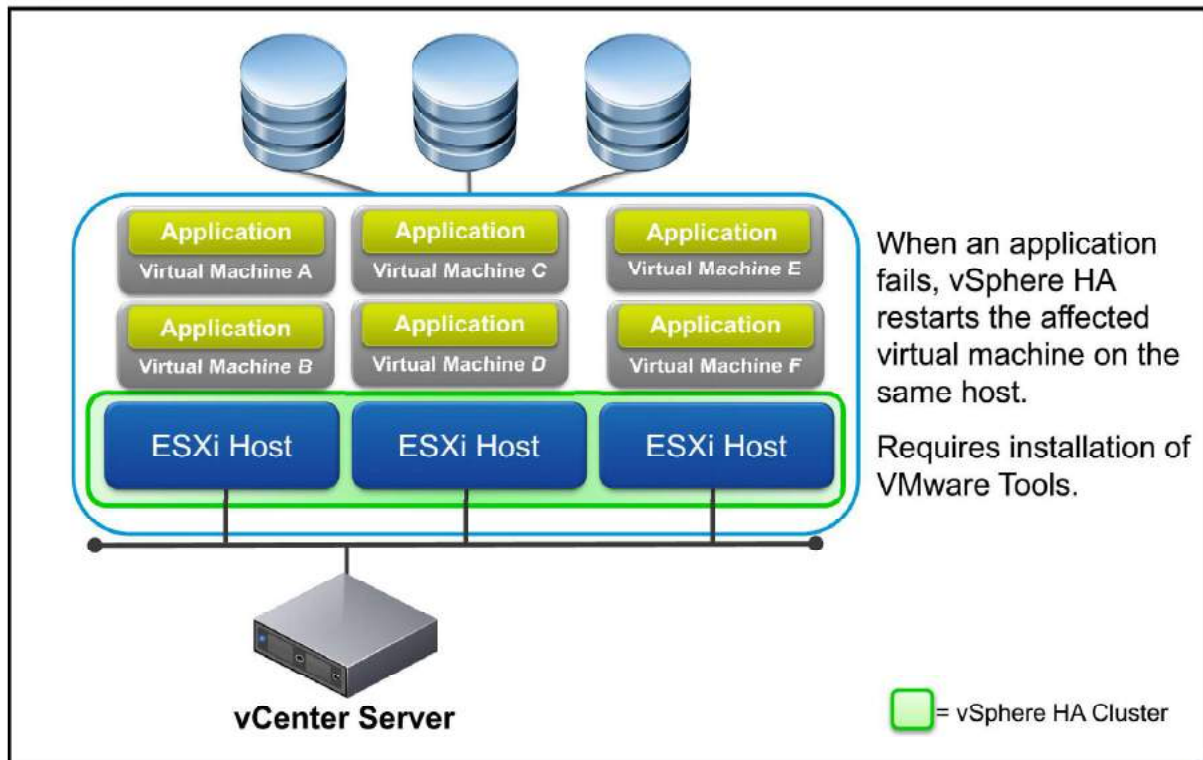
Slide 9-11



If Virtual Machine Monitoring is enabled, the vSphere HA agent on each individual host monitors VMware Tools in each virtual machine running on the host. When a virtual machine stops sending heartbeats, the guest operating system is reset. The virtual machine stays on the same host.

vSphere HA Scenario: Application Failure

Slide 9-12



The agent on each host can optionally monitor heartbeats of applications running in each virtual machine. When an application fails, the virtual machine on which the application was running is restarted on the same host. Application monitoring requires a third-party application monitoring agent designed to work with virtual machine application monitoring.

Importance of Redundant Heartbeat Networks

Slide 9-13

In a vSphere HA cluster, heartbeats have these characteristics:

- Heartbeats are sent between the master host and the slave hosts.
- They are used to determine whether a master host or slave host has failed.
- They are sent over a heartbeat network.

Redundant heartbeat networks ensure reliable failure detection, and minimize the chance of host isolation scenarios.

Heartbeat network implementation:

- Implemented by using a VMkernel port that is marked for management.
- Implemented by using a VMkernel port that is marked for virtual SAN traffic, when vSAN is in use.

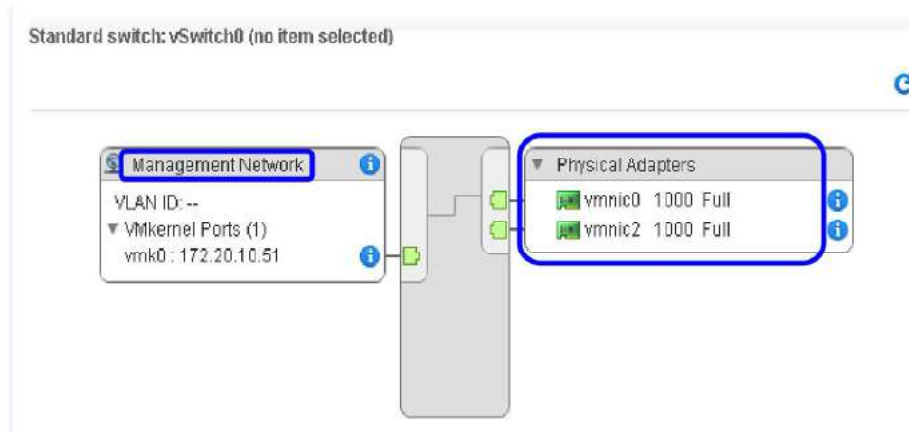
VMware recommends redundant heartbeat networking for your vSphere HA cluster. If you do not provide redundancy, your failover setup has a single point of failure. When a master host's connection fails, a second connection is still available to send heartbeats to other hosts.

Redundancy Using NIC Teaming

Slide 9-14

You can use NIC teaming to create a redundant heartbeat network on ESXi hosts.

Ports or port groups used must be VMkernel ports.

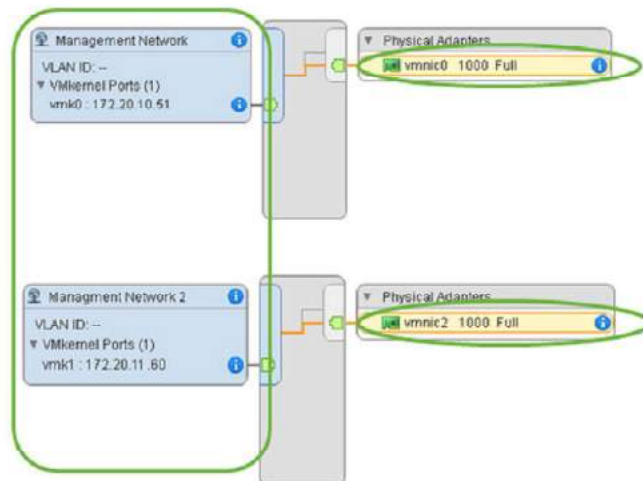


NIC Teaming on an ESXi Host

Redundancy Using Additional Networks

Slide 9-15

You can also create redundancy by configuring more heartbeat networks: On each ESXi host, create a second VMkernel port on a separate virtual switch with its own physical adapter.



Review of Learner Objectives

Slide 9-16

You should be able to meet the following objectives:

- Describe options that you can configure to make your vSphere environment highly available
- Discuss the response of vSphere HA when an ESXi host, a virtual machine, or an application fails

Lesson 2: vSphere HA Architecture

Slide 9-17



Lesson 2: vSphere HA Architecture

Learner Objectives


Slide 9-18

By the end of this lesson, you should be able to meet the following objectives:

- Describe the heartbeat mechanisms used by vSphere HA
- Identify and discuss other failure scenarios
- Recognize vSphere HA design considerations

vSphere HA Architecture: Agent Communication

Slide 9-19



When vSphere HA is enabled, the Fault Domain Manager (FDM) service starts on the member hosts. After the FDM agents have started, the cluster hosts are said to be in a fault domain. Hosts cannot participate in a fault domain if they are in maintenance mode, standby mode, or disconnected from vCenter Server. A host can be in only one fault domain at a time.

The fault domain is managed by a master host. All other hosts are called slave hosts. FDM agents on slave hosts all communicate with FDM on the master host.

To determine which host will be the master, an election process takes place. The system that can access the greatest number of datastores is elected the master. If more than one cluster hosts see the same number of datastores, the election process determines the master host by using the host managed object ID (MOID) assigned by vCenter Server.

The election process for a new master completes in approximately 15 seconds and occurs under these circumstances:

- vSphere HA is enabled
- The master encounters a system failure because of one of the following factors:
 - The master host is placed in maintenance mode

- The master host is placed in standby mode
- vSphere HA is reconfigured
- When the slave hosts cannot communicate with the master host due to a network problem

During the election process, the candidate vSphere HA agents communicate with each other over the management network or the vSAN network in a vSAN cluster by using User Datagram Protocol (UDP). All network connections are point-to-point. After the master agent has been determined, master and slave hosts communicate using the secure Transmission Control Protocol (TCP). When vSphere HA is started, vCenter Server contacts the master host agent and sends a list of hosts with membership in the cluster along with the cluster configuration. That information is saved to local storage on the master host and then pushed out to the slave hosts in the cluster. If additional hosts are added to the cluster during normal operation, the master agent sends an update to all hosts in the cluster.

The master host provides an interface for vCenter Server to query the state of and report on the health of the fault domain and virtual machine availability. vCenter Server tells the vSphere HA agent which virtual machines to protect along with their virtual machine-to-host compatibility list. The agent learns about state changes through `hostd` and vCenter Server learns of these through `vpix`. The master host monitors the health of the slave hosts and takes responsibility for virtual machines that were running on a failed slave host.

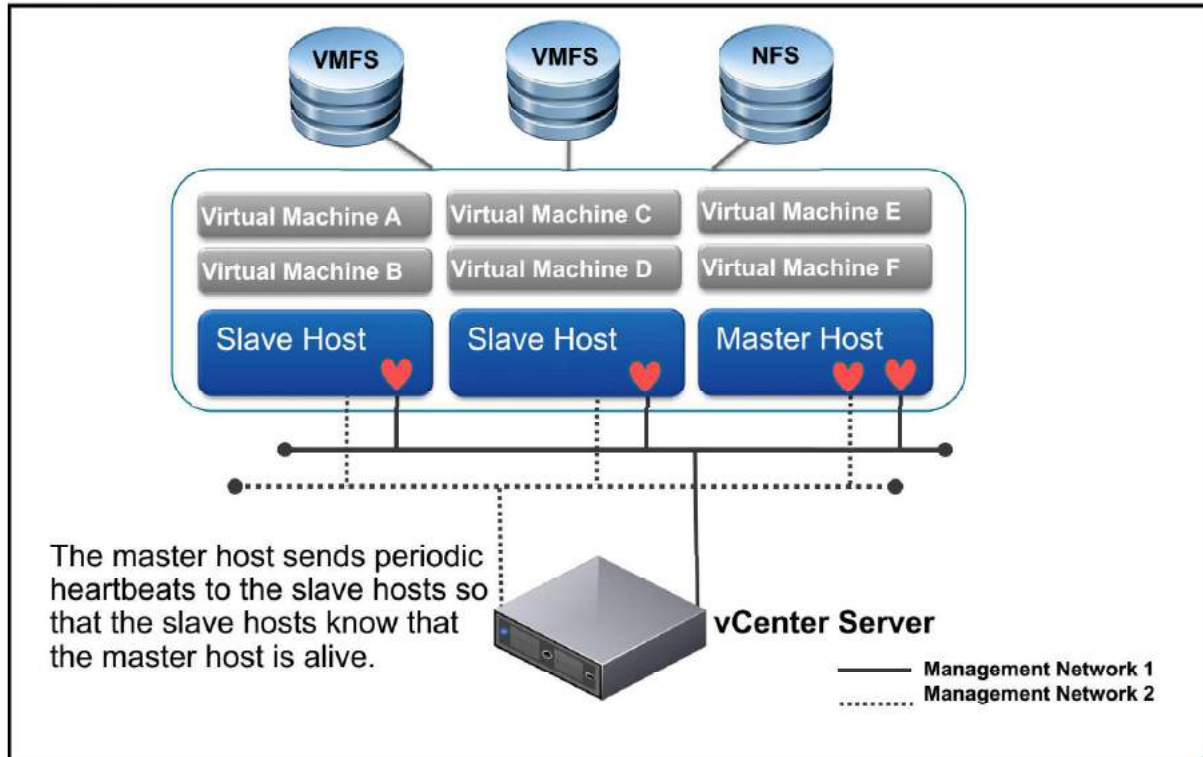
A slave host monitors the health of virtual machines running locally and sends state changes to the master host. A slave host also monitors the health of the master host.

vSphere HA is configured, managed, and monitored through vCenter Server. The cluster configuration data is maintained by the `vpixd` process which runs on the vCenter Server system. Cluster configuration changes are reported by the `vpixd` process to the master agent. The master agent advertises a new copy of the cluster configuration information and each slave fetches an updated copy. Each slave writes the updated configuration information to local storage. A list of protected virtual machines is stored on each datastore. The virtual machine list is updated after each user-initiated power-on (protected) and power off (unprotected) operation. The virtual machine list is updated after vCenter Server observes these operations.

A virtual machine becomes protected when an operation results in a power on. Reverting a virtual machine to a snapshot with memory state causes the virtual machine to power on and become protected. Similarly, a user action that causes the virtual machine to power off, for example, reverting to a snapshot without memory state or a standby operation performed in the guest, causes the virtual machine to become unprotected.

vSphere HA Architecture: Network Heartbeats

Slide 9-20

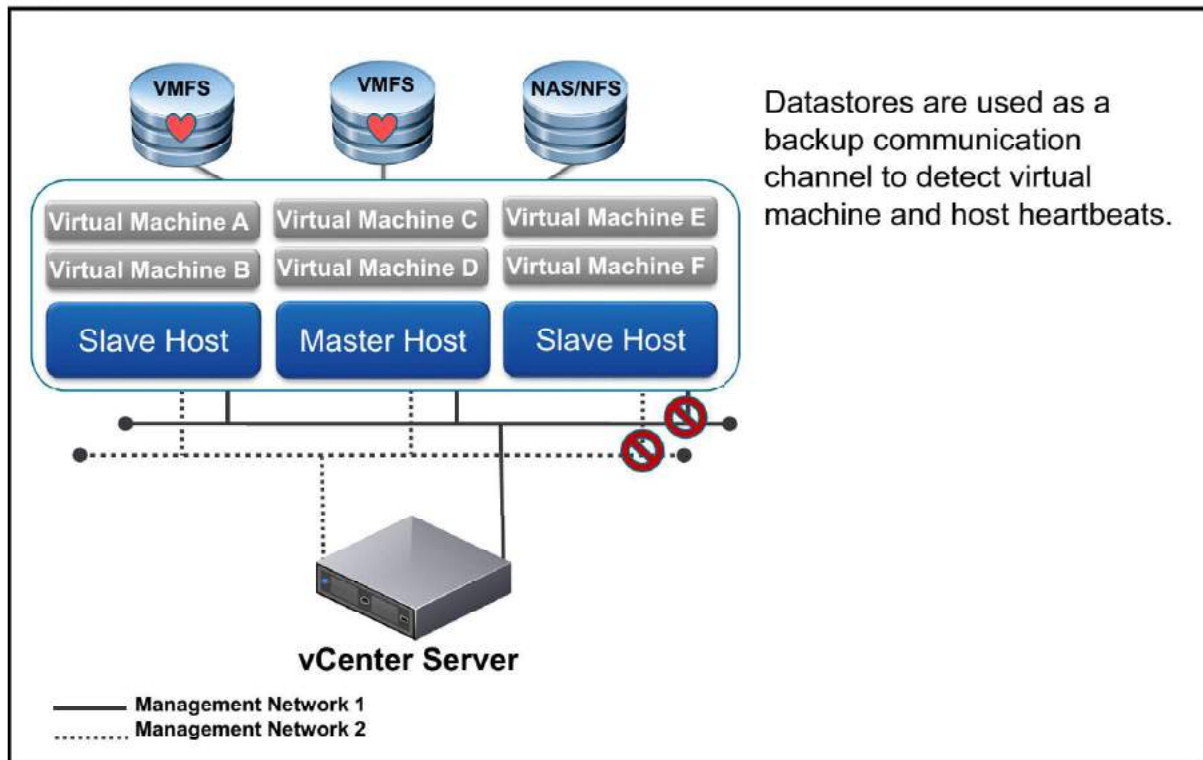


Heartbeats are sent to each slave host from the master host over all configured management networks. However, slave hosts use only one management network to communicate with the master. If the management network used to communicate with the master fails, the slave switches to another management interface to communicate with the master host.

If the slave host does not respond within predefined timeout period, the master host declares the slave host as agent unreachable. When a slave host is not responding, the master host attempts to determine the cause of the slave host's inability to respond. The master host must determine whether the slave host crashed or is not responding because of a network failure or the vSphere HA agent is in an unreachable state.

vSphere HA Architecture: Datastore Heartbeats

Slide 9-21



The datastore heartbeats are used to make the distinction between a failed and isolated or partitioned host. vSphere HA tries to restart virtual machines only in one of these situations:

- A host has failed (no network heartbeats, no ping, no datastore heartbeats).
- A host becomes isolated and the cluster's configured host isolation response is Power off or Shut down.

vSphere HA Failure Scenarios

Slide 9-22

Slave host failure

Master host failure

Network failure (Host isolation)

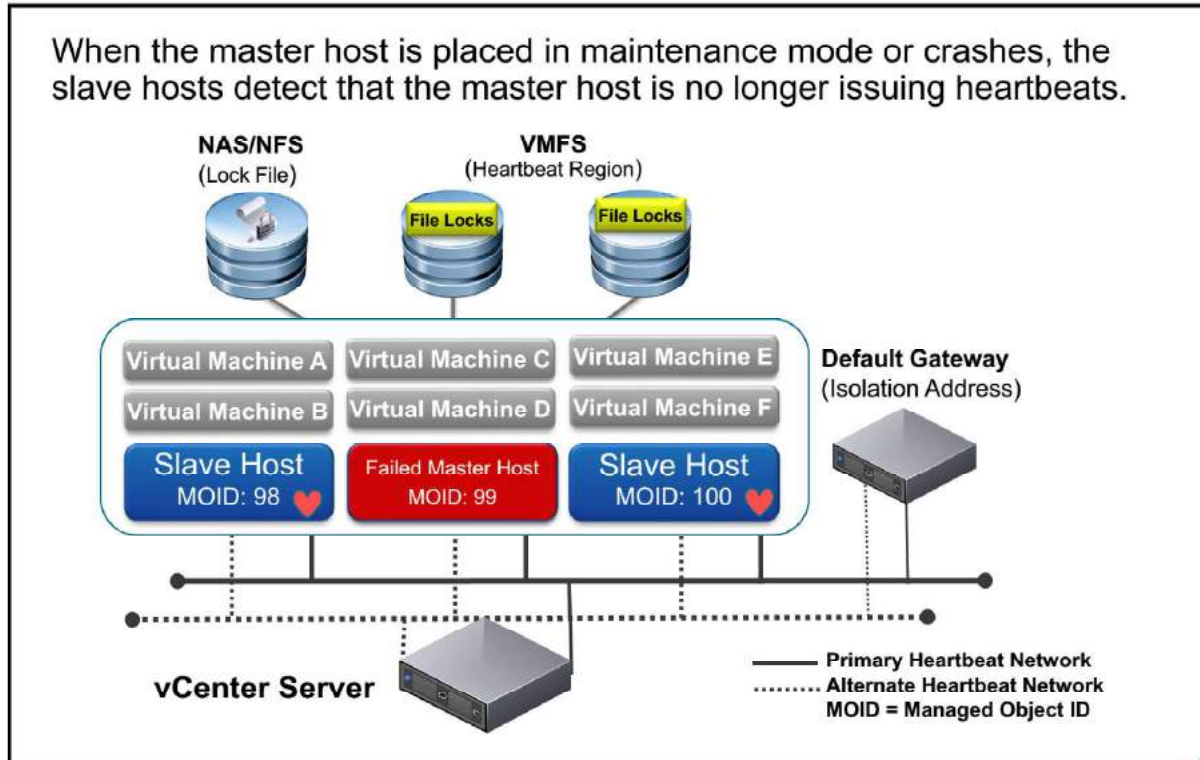
Datastore access failure:

- Virtual Machine Component Protection (VMCP)
 - All Paths Down
 - Permanent Device Loss

vSphere HA can also determine whether an ESXi host is isolated or has crashed. Isolation refers to when an ESXi host cannot see traffic coming from the other hosts in the cluster and cannot ping the default gateway. If an ESXi host crashes, vSphere HA must restart the virtual machines that were running on the failed host on the remaining hosts in the cluster. If the ESXi host is isolated because it cannot ping the default gateway (or isolation address) and sees no management network traffic, the host executes the Host Isolation Response.

Failed Master Host

Slide 9-24



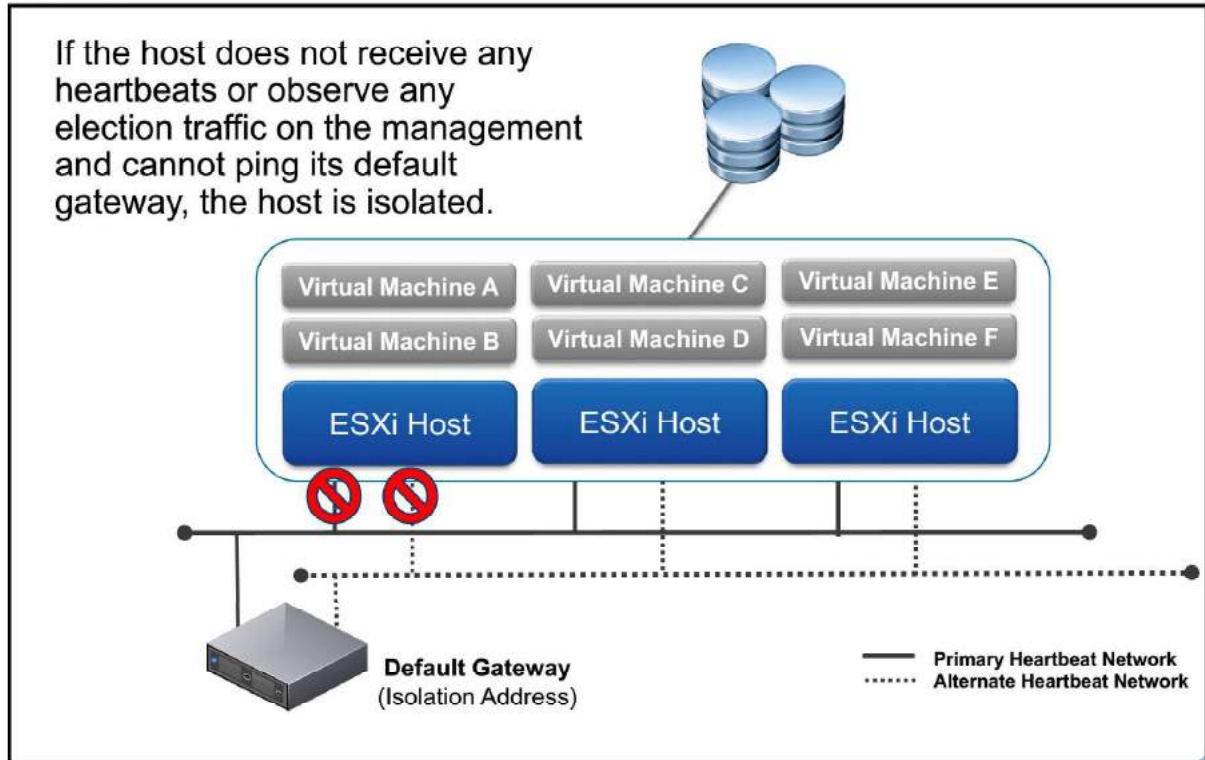
In this case, an election must take place to determine a new master host. The host with access to the greatest number of datastores is elected the master. If all slave hosts have equal datastore access, the election process selects a new master host using the highest lexically ordered managed object ID (MOID) assigned by vCenter Server when the host was added to the vCenter Server inventory.

If the master host fails, the slave participates in a new master election. When a new master is selected, it reads MAC and IP addresses of the hosts and virtual machines from a host list that is stored on a datastore. The host list is used to determine whether the master should accept a connection from a slave.

The new master reads a file on each datastore that contains the state of all virtual machines and determines which virtual machines are protected by vSphere HA and might have to be restarted. The new master first identifies which virtual machines are running on the slave hosts. Through the process of elimination it determine the virtual machines that require restarting. This same process is used to restart virtual machines after a total cluster failure.

Isolated Host

Slide 9-25



A host is declared isolated when the following two conditions occur:

- When the host is not receiving network heartbeats
- When the host cannot ping its isolation addresses

Several scenarios might result in host isolation. The slide illustrates one scenario. If a host loses connectivity to the primary heartbeat network and the alternate heartbeat network, the host no longer receives network heartbeats from the other hosts in the vSphere HA cluster. Furthermore, the slide depicts that this same host can no longer ping its isolation address.

If a host becomes isolated, the master vSphere HA agent must determine if that host is still alive, and merely isolated, by checking for datastore heartbeats. Datastore heartbeats are used by vSphere HA only when a host becomes isolated or partitioned.

Design Considerations

Slide 9-26

Host isolation events can be minimized through good design:

- Implement redundant heartbeat networks.
- Implement redundant isolation addresses.

If host isolation events do occur, good design enables vSphere HA to determine whether the isolated host is still alive.

Implement datastores so that they are separated from the management network by using one or both of the following approaches:

- Fibre Channel over fiber optic
- Physically separating your IP storage network from the management network

If a datastore is based on Fibre Channel, the datastore access is not disrupted by the network failure. When using datastores based on IP storage (for example, NFS, iSCSI, Fibre Channel over Ethernet), you must physically separate (or logically separate if physical separation is impossible) the IP storage network and the management network (the heartbeat network).

Virtual Machine Storage Failures

Slide 9-27


With an increasing number of virtual machines and datastores on each host, storage connectivity issues have high costs but are infrequent.

Connectivity problems due to:

- Network or switch failure
- Array misconfiguration
- Power outage

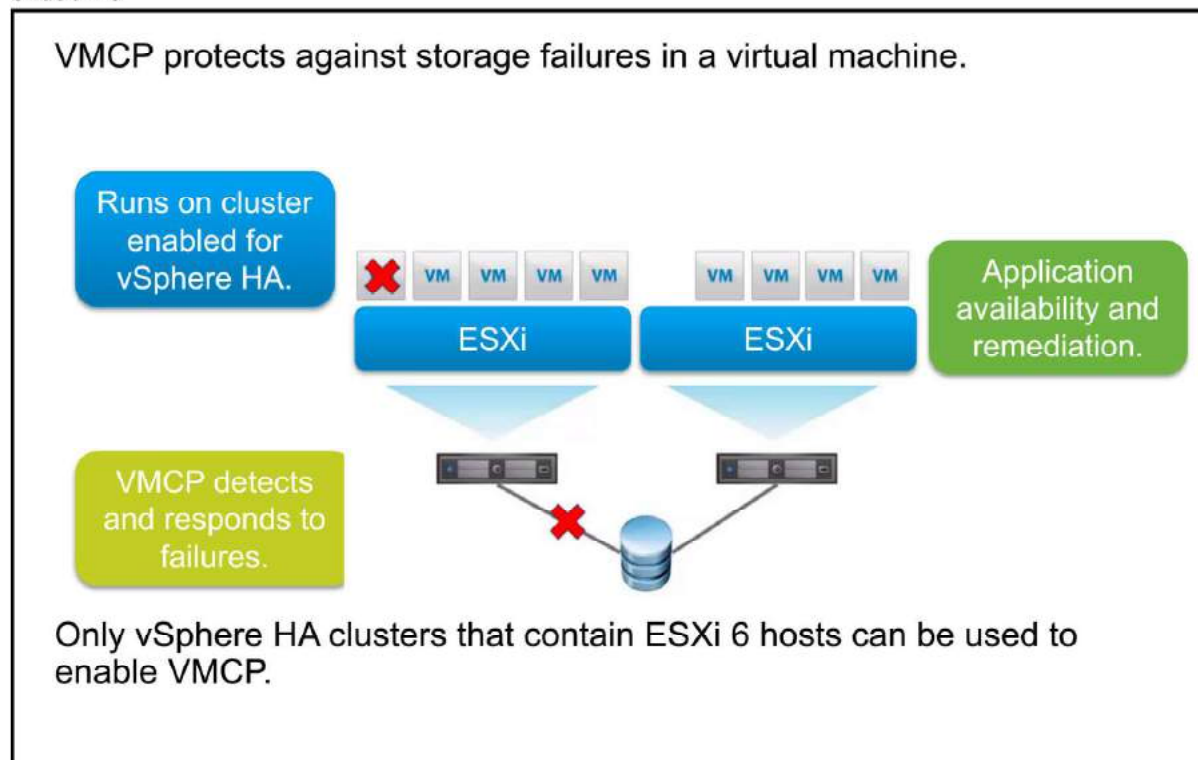
Virtual machine availability is affected:

- Virtual machines on affected hosts are difficult to manage.
- Applications with attached disks crash.



Virtual Machine Component Protection

Slide 9-28



If VMCP is enabled, vSphere HA can detect datastore accessibility failures and provide automated recovery for affected virtual machines.

VMCP provides protection against datastore accessibility failures that can affect a virtual machine running on a host in a vSphere HA cluster. When a datastore accessibility failure occurs, the affected host can no longer access the storage path for a specific datastore. You can determine the response that vSphere HA will make to such a failure, ranging from the creation of event alarms to virtual machine restarts on other hosts.

Only vSphere HA clusters that contain ESXi 6 hosts can be used to enable VMCP. Clusters that contain hosts from an earlier release cannot enable VMCP. Such hosts cannot be added to a cluster enabled for VMCP.

Review of Learner Objectives

Slide 9-29

You should be able to meet the following objectives:

- Describe the heartbeat mechanisms used by vSphere HA
- Identify and discuss other failure scenarios
- Recognize vSphere HA design considerations

Lesson 3: Configuring vSphere HA

Slide 9-30



Lesson 3: Configuring vSphere HA

Learner Objectives

Slide 9-31

By the end of this lesson, you should be able to meet the following objectives:

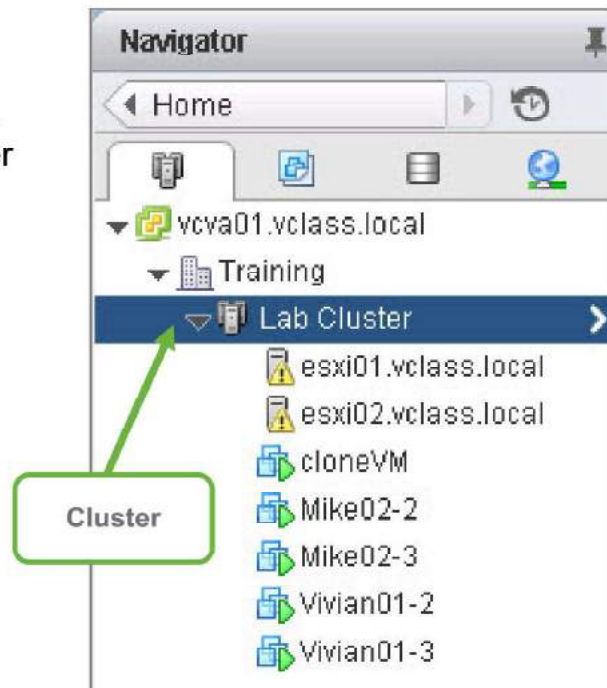
- Recognize the requirements for creating and using a vSphere HA cluster
- Configure a vSphere HA cluster

About Clusters

Slide 9-32

A cluster is used in vSphere to share physical resources between a group of ESXi hosts. vCenter Server manages cluster resources as a single pool of resources.

Features such as vSphere HA, vSphere DRS, and vSAN can be enabled in a cluster.



vSphere HA leverages multiple ESXi hosts that are configured as a cluster to provide rapid recovery from outages and high availability for applications running in virtual machines. The resources are managed by vCenter Server. Features such as vSphere HA, vSphere DRS, and vSAN can be enabled in a cluster.

vSphere HA Prerequisites

Slide 9-33

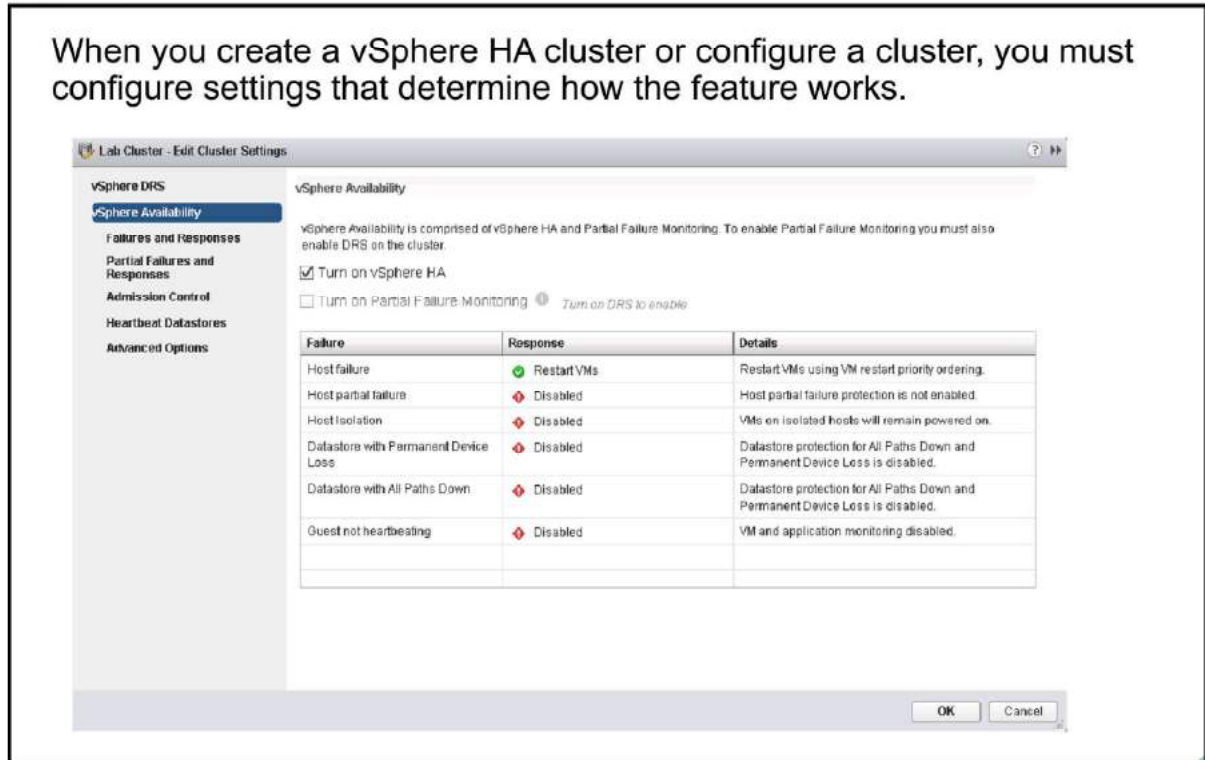
You must follow the guidelines to create a vSphere HA cluster:

- VMware vSphere® Essentials Plus Kit, VMware vSphere® Standard Edition™, VMware vSphere® Enterprise Edition™, and VMware vSphere® Enterprise Plus Edition™ licenses support vSphere HA.
- A cluster can contain between two and 64 ESXi hosts.
- All hosts must be configured with static IP addresses. If you are using DHCP, you must ensure that the address for each host persists across reboots.
- All hosts must have at least one management network in common.
- For virtual machine monitoring to work, VMware Tools must be installed in every virtual machine.
- Only vSphere HA clusters that contain ESXi 6.x hosts can be used to enable VMCP.

Configuring vSphere HA Settings

Slide 9-34

When you create a vSphere HA cluster or configure a cluster, you must configure settings that determine how the feature works.



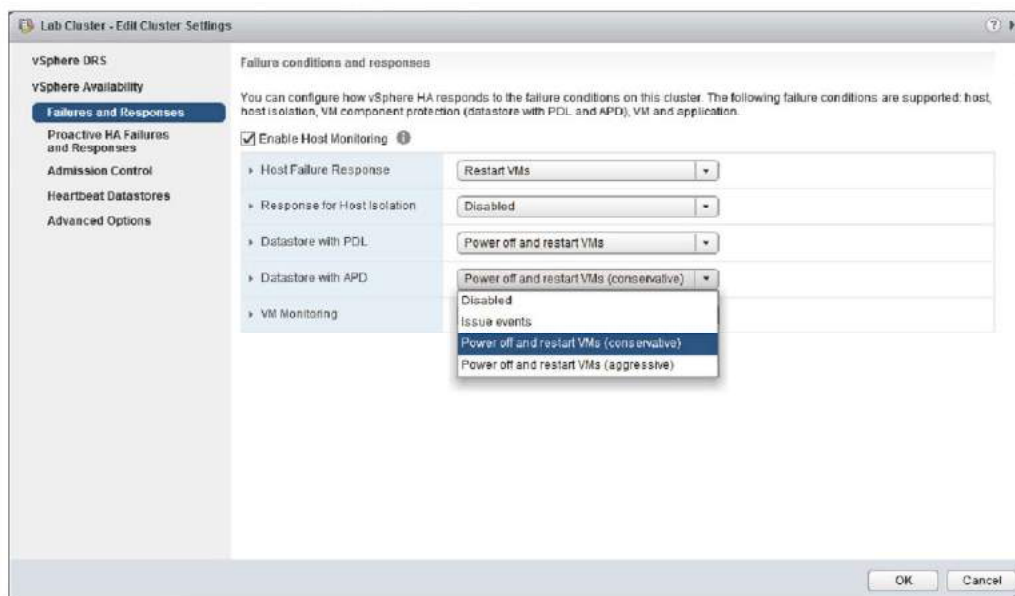
In vSphere Web Client or in vSphere Client, you can configure the following vSphere HA settings:

- Availability failure conditions and responses: Provide settings for host failure responses, host isolation, virtual machine monitoring, and virtual machine component protection.
- Admission Control: Enable or disable admission control for the vSphere HA cluster and choose a policy for how it is enforced.
- Heartbeat Datastores: Specify preferences for the datastores that vSphere HA uses for datastore heartbeating.
- Advanced Options: Customize vSphere HA behavior by setting advanced options.

vSphere HA Settings: Failure and Responses

Slide 9-35

You use the Failure and Responses pane to configure a cluster's response if a failure occurs.

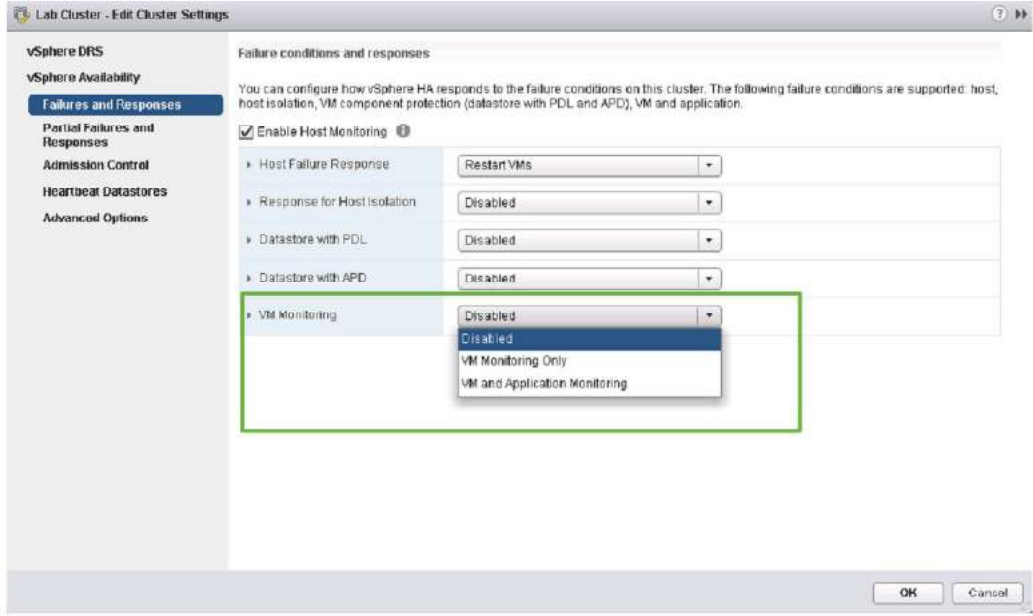


The Failure and Responses pane allows you to configure how your cluster should function when problems are encountered. You can specify the vSphere HA cluster's response for host failures and isolation. You can also configure VMCP actions when Permanent Device Loss and All Paths Down situations occur and enable VM monitoring

vSphere HA Settings: Virtual Machine Monitoring

Slide 9-36

You use **VM Monitoring** settings to control the monitoring of virtual machines.



The screenshot shows the 'Lab Cluster - Edit Cluster Settings' window. On the left, the 'vSphere Availability' section is expanded to 'Failures and Responses'. The 'VM Monitoring' setting is highlighted with a green box, and its dropdown menu is open, showing options: 'Disabled', 'VM Monitoring Only', and 'VM and Application Monitoring'. The 'VM Monitoring' option is currently selected.

Setting	Value
Enable Host Monitoring	<input checked="" type="checkbox"/>
Host Failure Response	Restart VMs
Response for Host Isolation	Disabled
Datastore with PDL	Disabled
Datastore with APD	Disabled
VM Monitoring	VM Monitoring Only

By default, **VM Monitoring** is set to **Disabled**.

Virtual machine monitoring restarts individual virtual machines if their VMware Tools heartbeats are not received or if the guest operating system has not issued an I/O for the last 2 minutes (by default). If neither criteria is met, it is most likely because the guest operating system has failed. In such a case, the virtual machine monitoring service determines that the virtual machine has failed and the virtual machine resets to restore service.

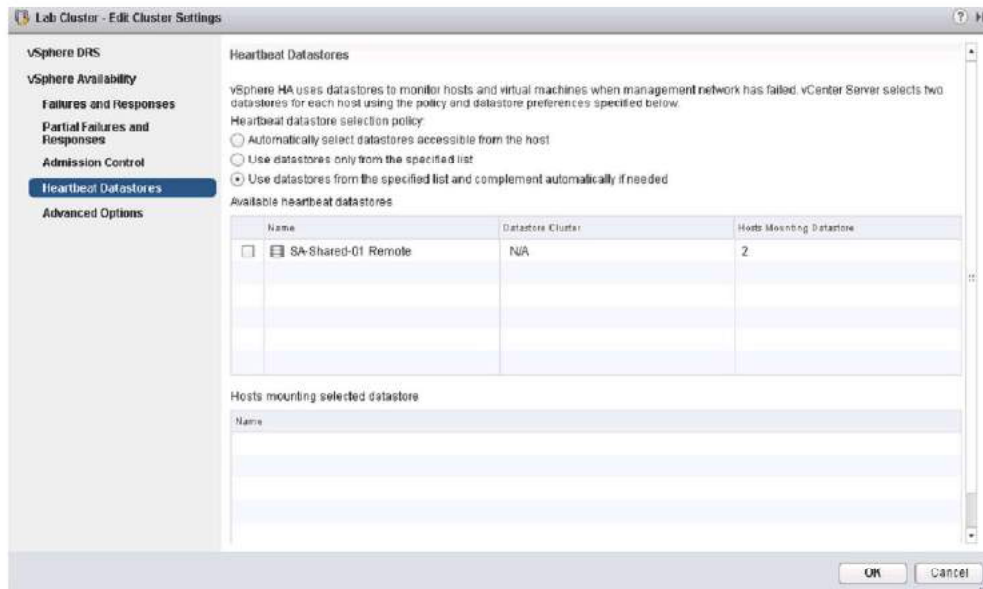
You can configure the level of monitoring sensitivity. Highly sensitive monitoring results in a more rapid conclusion that a failure has occurred. Although unlikely, highly sensitive monitoring might lead to falsely identifying failures when the virtual machine or application is still working but heartbeats have not been received because of factors like resource constraints. Low-sensitivity monitoring results in longer interruptions in service between actual failures and virtual machines being reset. Select an option that is an effective compromise for your needs.

Select **VM Monitoring Only** to reset individual virtual machines if their heartbeats are not received within a set time. You can also select **VM and Application Monitoring** to enable application monitoring.

vSphere HA Settings: Heartbeat Datastores

Slide 9-37

A heartbeat file is created on the selected datastores and is used if a management network fails.



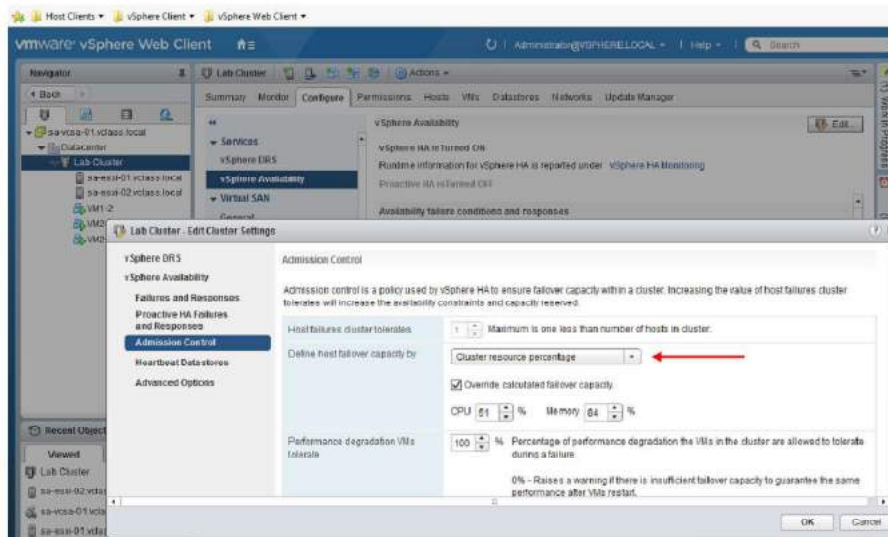
Using datastore heartbeating, the master host determines whether a host has failed, or a network isolation has occurred. If datastore heartbeating from the host stops, then the host is considered failed. In this case, the failed host's virtual machines are started on another host in the vSphere HA cluster.

Datastore heartbeating takes checking the health of a host to another level by checking more than the management network to determine a host's health. You may configure a list of datastores to monitor for a particular host. Alternatively, you may let vSphere HA make the decision. You may also use the combination of both methods.

vSphere HA Settings: Admission Control

Slide 9-38

vCenter Server uses admission control to ensure that sufficient resources are available in a cluster to provide failover protection and that virtual machine resource reservations are respected.



After you create a cluster, admission control allows you to specify whether virtual machines can be started if they violate availability constraints. The cluster reserves resources to allow failover for all running virtual machines on the specified number of hosts.

The admission control settings include:

- **Disabled:** This option disables admission control, allowing the virtual machines violating availability constraints to power on.
- **Slot Policy:** A slot is a logical representation of memory and CPU resources. With the slot policy option, vSphere HA calculates the slot size, determines how many slots each host in the cluster can hold, and consequently determines the current failover capacity of the cluster.
- **Cluster resource Percentage:** Specify a percentage of the cluster's CPU and Memory resources to be reserved as spare capacity to support failovers.
- **Dedicated failover hosts:** Select hosts to use for failover actions. If a default failover host does not have enough resources, failovers can still occur to other hosts in the cluster.

Example: Admission Control Using Cluster Resources Percentage

Slide 9-39

The Configured Failover Capacity for CPU and Memory are both set to 25 percent.


Total resource requirements: 7 GHz and 6 GB.

Current CPU failover capacity is 70 percent:

- ((24 GHz through 7 GHz)/24 GHz).

Current memory failover capacity is 71 percent:

- ((21 GB through 6 GB)/21 GB).




Example: Admission Control Using Slots

Slide 9-40

Slot is memory and CPU requirement required for any powered-on VM in the cluster.

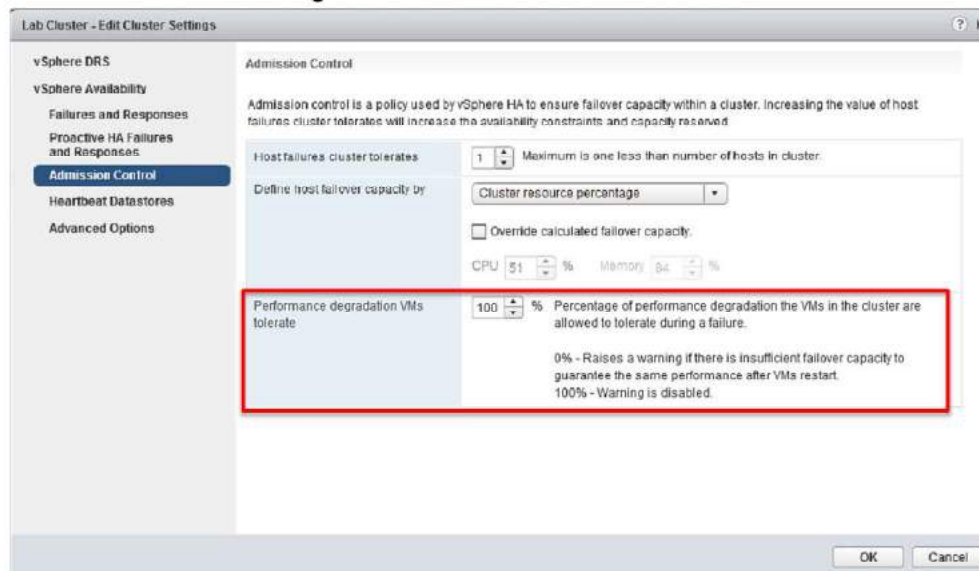
1. Calculate the slot size: 2 GHz CPU and 2 GB memory.
2. Calculate slot/host.
3. Current failover capacity is one.
 - If H1 fails, six slots remain in the cluster, which is sufficient for all five of the powered-on virtual machines.
 - If both H1 and H2 fail, only three slots remain, which is insufficient.
4. If current failover capacity < configured failover capacity, disallow further VM power on.



vSphere HA Settings: Performance Degradation VMs Tolerate

Slide 9-41

The Performance degradation VMs tolerate threshold specifies the percentage of performance degradation that the VMs in the cluster are allowed to tolerate during a failure.



Admission Control can also be configured to offer warnings when the actual usage exceeds the failover capacity percentage. The resource reduction calculation takes into account a virtual machine's reserved memory and memory overload to decide whether to permit it to power on, migrate, or have reservation changes.

Setting the **Performance degradation VMs tolerate** threshold allows you to specify when the occurrence of a configuration issue should arise. For example:

- The default value is 100 percent, which produces no warnings.
- If you reduce the threshold to 0 percent, a warning is generated as soon as cluster utilization exceeds the available capacity.
- If you reduce the threshold to 20 percent, the performance reduction that can be tolerated is calculated as: $\text{performance reduction} = \text{current utilization} \times 20 \text{ percent}$. When the current utilization minus the performance reduction exceeds the available capacity, a configuration notice is issued.

The **Performance degradation VMs tolerate** threshold is not available unless vSphere DRS is enabled.

vSphere HA Settings: Advanced Options

Slide 9-42

To customize vSphere HA behavior, you set advanced vSphere HA options.

To force a cluster not to use the default isolation address (default gateway):

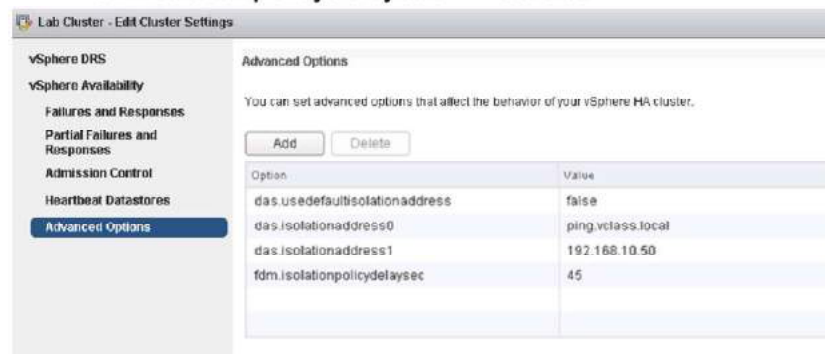
- `das.usedefaultisolationaddress = false`

To force a cluster to ping alternate isolation addresses:

- `das.isolationaddressX = ping reachable address`

To force a cluster to wait beyond the default 30-second isolation action window:

- `fdm.isolationpolicydelaysec = > 30 sec`



The screenshot shows the 'Lab Cluster - Edit Cluster Settings' window. On the left, a navigation pane lists 'vSphere DRS', 'vSphere Availability', 'Failures and Responses', 'Partial Failures and Responses', 'Admission Control', 'Heartbeat Datastores', and 'Advanced Options' (which is selected). The main area is titled 'Advanced Options' and contains the text: 'You can set advanced options that affect the behavior of your vSphere HA cluster.' Below this text are 'Add' and 'Delete' buttons. A table lists the configured options:

Option	Value
<code>das.usedefaultisolationaddress</code>	false
<code>das.isolationaddress0</code>	ping.vcless.local
<code>das.isolationaddress1</code>	192.168.10.50
<code>fdm.isolationpolicydelaysec</code>	45

You can set advanced options that affect the behavior of your vSphere HA cluster. For more details, see *vSphere Availability Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

vSphere HA Orchestrated Restart

Slide 9-43

Orchestrated restart in vSphere HA improves the current restart priority by creating restart tiers and VM-VM dependencies.

Conditions that must be met before a VM is considered ready:

- VM has resources secured.
- VM is powered on.
- VMware Tools heartbeat is detected.
- VMware Tools application heartbeat is detected.

You can optionally configure a delay when a certain condition is met.

Priority tier:

- VMs can be grouped into tiers indicating their startup priority. All VMs in the priority 1 tier receive their resources first and are powered on.
- Only when all the VMs in tier 1 have met their defined restart condition does vSphere HA continue to the VMs in the priority 2 tier, and so on.

After a host failure, virtual machines are assigned to other hosts with unreserved capacity, with the highest priority virtual machines placed first. The process continues to those virtual machines with lower priority until all have been placed or no more cluster capacity is available to meet the reservations or memory overhead of the virtual machines. A host then restarts the virtual machines assigned to it in priority order.

If there are insufficient resources, vSphere HA waits for more unreserved capacity to become available, for example, due to a host coming back online, and then retries the placement of these virtual machines. To reduce the chance of this situation occurring, configure vSphere HA admission control to reserve more resources for failures. Admission control allows you to control the amount of cluster capacity that is reserved by virtual machines, which is unavailable to meet the reservations and memory overhead of other virtual machines if a failure occurs.

Configuring vSphere HA Orchestrated Restart

Slide 9-44

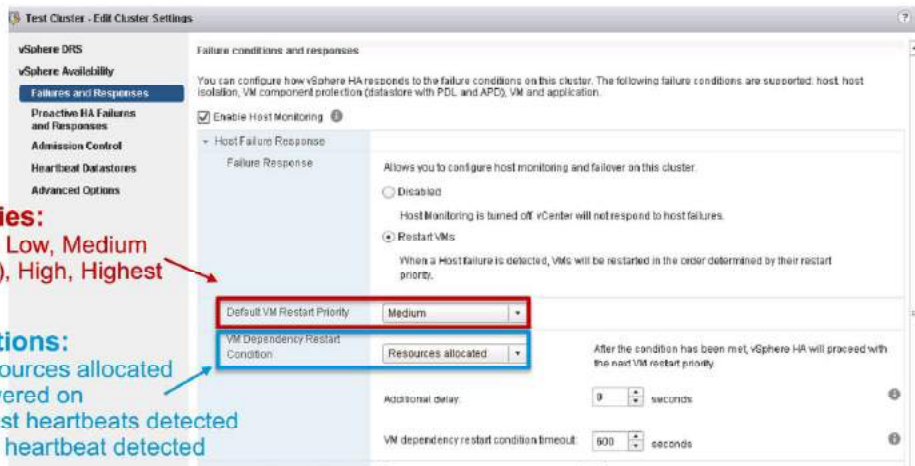
Regular VMs are put in the Medium restart priority by default, unless the restart priority is explicitly set using VM overrides.

Exceptions:

- Agent VMs always start first, and the restart priority is nonconfigurable.
- vSphere Fault Tolerance secondary VMs fail over before regular VMs. Primary VMs follow normal restart priority.

Priorities:
Lowest, Low, Medium (default), High, Highest

- Conditions:**
- Resources allocated
 - Powered on
 - Guest heartbeats detected
 - App heartbeat detected



VM Dependencies in Orchestrated Restart

Slide 9-45

VMs can depend only on other VMs of the same or higher priority. Only direct dependencies are supported. VM-to-VM dependency is a hard rule. Creating cyclical dependencies causes VM restart to fail.



In vSphere 6.5, vSphere HA restarts VMs only from a failed host. Configure affinity rules to keep VMs on the host if necessary.

B does not get restarted when C fails over.

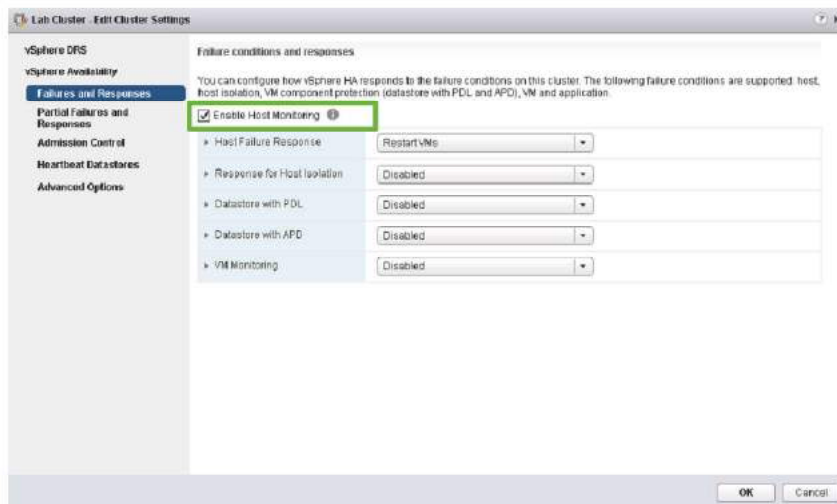
Restart in parallel.



Network Configuration and Maintenance

Slide 9-46

Disable host monitoring before modifying virtual networking components that involve the VMkernel ports configured for the management or vSAN traffic. This practice prevents unwanted attempts to fail over virtual machines.



The following network maintenance suggestions can help you avoid the false detection of host failure and network isolation due to dropped vSphere HA heartbeat:

- Changing your network hardware or networking settings can interrupt the heartbeats used by vSphere HA to detect host failures, and might result in unwanted attempts to fail over virtual machines. When making changes to the networks where your clustered ESXi hosts reside, suspend the Host Monitoring feature and place the host in maintenance mode.
- Host Monitoring is only required when modifying virtual networking components and properties that involve the VMkernel ports configured for the Management or vSAN traffic, which are used by the vSphere HA networking heartbeat service.
- After you change the networking configuration on ESXi hosts, for example, adding port groups, removing virtual switches, or suspending Host Monitoring, you must reconfigure vSphere HA on all hosts in the cluster. This reconfiguration causes the network information to be reinspected. Then, you must re-enable Host Monitoring.

Cluster Resource Reservation

Slide 9-47

The **Resource Reservation** tab reports total cluster CPU, memory, memory overhead, storage capacity, the capacity reserved by virtual machines, and how much capacity remains available.

The screenshot shows the vSphere interface for a 'Lab Cluster'. The 'Resource Reservation' tab is active, displaying a summary of cluster resources. A green box highlights the following data:

Cluster Total Capacity	12.00 GB
Total Reservation Capacity	5.33 GB
Used Reservation	120.00 MB
Available Reservation	5.21 GB

Below this summary is a table listing individual VM reservations:

Name	Reservation (MB)	Limit (MB)	Type
VM1-2	0	Unlimited	Fixed
VM2-1	0	Unlimited	Fixed
VM2-2	0	Unlimited	Fixed

vCenter Server uses vSphere HA admission control to ensure that sufficient resources are available in a cluster to provide failover protection and to ensure that virtual machine resource reservations are respected.

Monitoring Cluster Status

Slide 9-48

You can monitor the status of a vSphere HA cluster on the **Monitor** tab.

The screenshot displays the vSphere HA Monitor tab in a web interface. The interface includes a navigation bar with tabs for Summary, Monitor, Configure, Permissions, Hosts, VMs, Datastores, Networks, and Update Manager. Below this, there are sub-tabs for Issues, Performance, Tasks & Events, Profile Compliance, Resource Reservation, vSphere DRS, vSphere HA (selected), and Utilization. The main content area is divided into three panels:

- Summary:** A sidebar menu with options: Summary (selected), Heartbeat, Configuration Issues, and Datastores under APD or PDL.
- Hosts:** A table showing the status of hosts in the cluster.
- Virtual Machines:** A summary of VM protection status.

Hosts	
Master	sa-esxi-02.vclass.local
Hosts connected to master	1
Hosts not connected to master	0
vSphere HA agent not reachable	0
vSphere HA agent configuration error	0
Hosts failed	0
Network isolated	0
Network partitioned	0
vSphere HA agent initializing	0

Virtual Machines	
Protected	3
Unprotected	0

You can monitor the status of a vSphere HA cluster on the Summary page of the vSphere HA panel on the **Monitor** tab.

Configuration issues and other errors can occur for your cluster or its hosts that adversely affect the proper operation of vSphere HA. You can monitor these errors in the Configuration Issues page, which is accessible from the same area.

Lab 19: Using vSphere HA

Slide 9-49

Use vSphere HA functionality

1. Create a Cluster Enabled for vSphere HA
2. Add Your ESXi Host to a Cluster
3. Test vSphere HA Functionality
4. View the vSphere HA Cluster Resource Usage
5. Manage vSphere HA Slot Size
6. Configure a vSphere HA Cluster with Strict Admission Control
7. Prepare for the Next Lab

Review of Learner Objectives


Slide 9-50

You should be able to meet the following objectives:

- Recognize the requirements for creating and using a vSphere HA cluster
- Configure a vSphere HA cluster

Lesson 4: Introduction to vSphere Fault Tolerance

Slide 9-51



Learner Objectives

Slide 9-52

By the end of this lesson, you should be able to meet the following objectives:

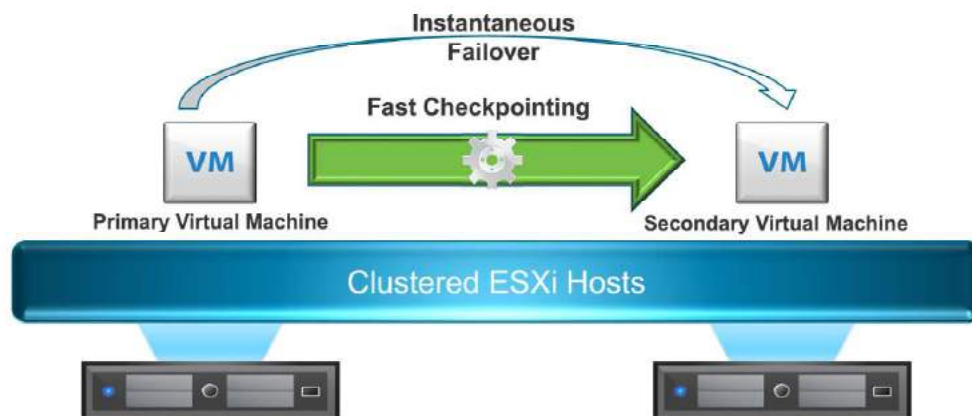
- Describe benefits of using vSphere Fault Tolerance
- Explain how vSphere Fault Tolerance works
- Describe main vSphere Fault Tolerance features
- Describe how vSphere Fault Tolerance works with vSphere HA and vSphere DRS
- Enable vSphere Fault Tolerance from vSphere Web Client

vSphere Fault Tolerance

Slide 9-53

vSphere Fault Tolerance provides instantaneous failover and continuous availability:

- Zero downtime
- Zero data loss
- No loss of TCP connections



vSphere Fault Tolerance Features

Slide 9-54

vSphere Fault Tolerance protects mission-critical, high-performance applications regardless of the operating system used.

vSphere Fault Tolerance:

- Supports virtual machines configured with up to 4 vCPUs and 64 GB memory
- Supports up to 4 vSphere Fault Tolerance virtual machines
- Supports hosts configured with up to 8 vCPUs in vSphere Fault Tolerance mode
- Supports vSphere vMotion migration for primary and secondary virtual machines
- Creates a secondary copy of all virtual machine files and disks
- Provides fast checkpoint copying to keep primary and secondary CPUs synchronized
- Supports multiple virtual machine disk formats: thin provision, thick provision lazy-zeroed, and thick provision eager-zeroed.
- Can be used with vSphere DRS only when EVC is enabled.
- Supports interoperability with vSAN

You can use vSphere Fault Tolerance for your virtual machines to ensure business continuity with higher levels of availability and data protection than is offered by vSphere HA.

vSphere Fault Tolerance is built on the ESXi host platform. vSphere Fault Tolerance provides continuous availability by having identical virtual machines run on separate hosts.

You can use vSphere Fault Tolerance for most mission-critical virtual machines. vSphere Fault Tolerance provides continuous availability for such a virtual machine by creating and maintaining another virtual machine that is identical and continuously available to replace it in the event of a failover situation.

The protected virtual machine is called the primary virtual machine. The duplicate virtual machine, the secondary virtual machine, is created and runs on another host. The secondary virtual machine's execution is identical to that of the primary virtual machine. The secondary virtual machine can take over at any point without interruption, thus providing fault tolerant protection.

The primary and secondary virtual machines continuously monitor the status of one another to ensure that fault tolerance is maintained. A transparent failover occurs if the host running the primary virtual machine fails, in which case the secondary virtual machine is immediately activated to replace the primary virtual machine. A new secondary virtual machine is started and fault tolerance redundancy is reestablished automatically. If the host running the secondary virtual machine fails, it is also immediately replaced. In either case, users experience no interruption in service and no loss of data.

vSphere Fault Tolerance with vSphere HA and vSphere DRS

Slide 9-55

vSphere Fault Tolerance works with vSphere HA and vSphere DRS.

vSphere HA:

- Is required for vSphere Fault Tolerance
- Restarts failed virtual machines
- Is vSphere Fault Tolerance aware

vSphere DRS:

- When a virtual machine is powered on, selects which hosts run the primary and secondary virtual machines.
- Does not automatically migrate fault-tolerant virtual machines.



You can use vSphere Fault Tolerance with vSphere DRS only when the Enhanced vMotion Compatibility (EVC) feature is enabled. This process allows fault tolerant virtual machines to benefit from better initial placement.


When a cluster has EVC enabled, vSphere DRS makes the initial placement recommendations for fault tolerant virtual machines and allows you to assign a vSphere DRS automation level to primary VMs. The secondary VM always assumes the same setting as its associated primary VM.

When vSphere Fault Tolerance is used for virtual machines in a cluster that has EVC disabled, the fault tolerant virtual machines are given vSphere DRS automation levels of disabled. In such a cluster, each primary VM is powered on only on its registered host and its secondary VM is automatically placed.

Redundant VMDKs

Slide 9-56

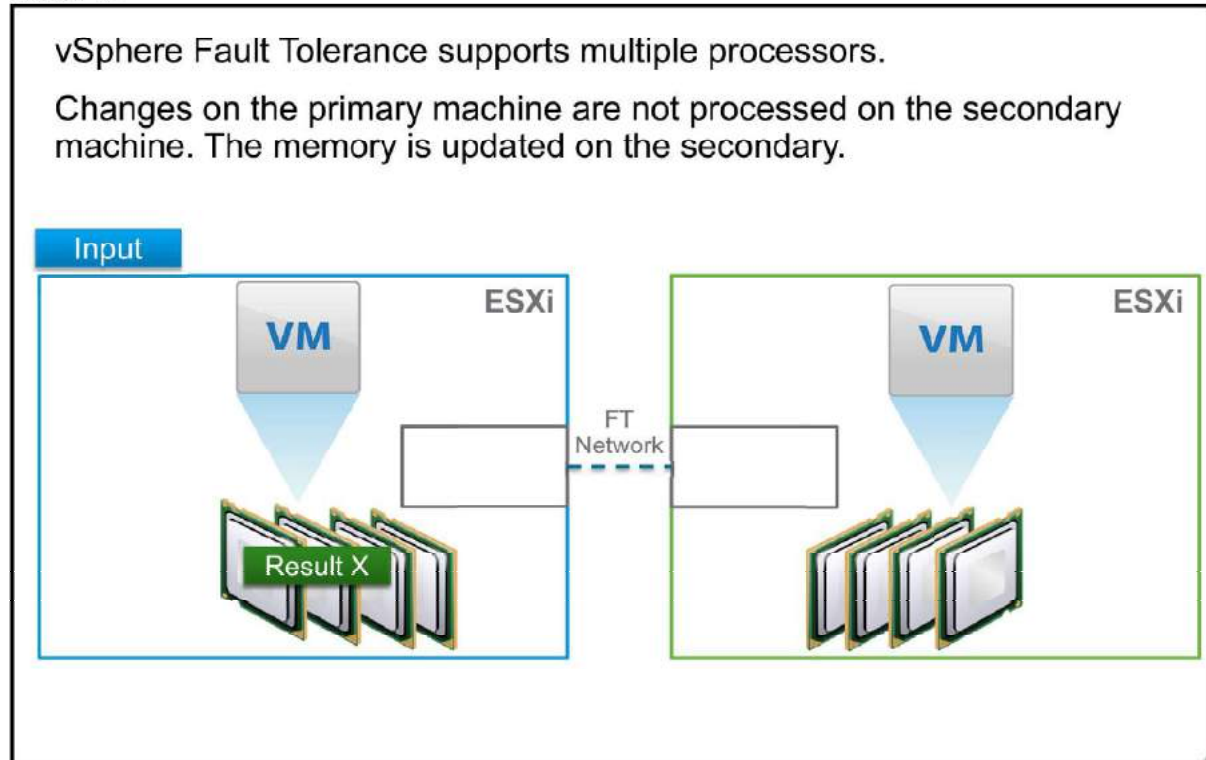
vSphere Fault Tolerance creates two complete virtual machines. Each virtual machine has its own `.vmx` configuration file and `.vmdk` files. Each of these virtual machines can be on a different datastore.



A fault-tolerant virtual machine and its secondary copy are not allowed to run on the same host. This restriction ensures that a host failure cannot result in the loss of both virtual machines.

vSphere Fault Tolerance Checkpoint

Slide 9-57



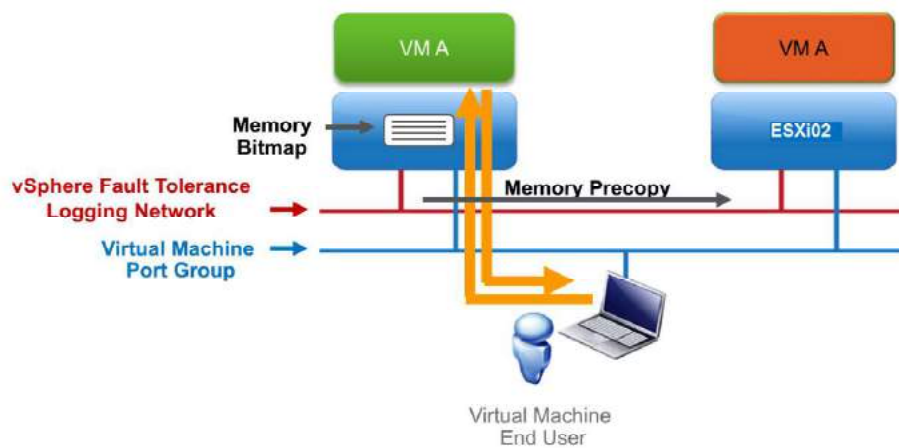
vSphere Fault Tolerance provides failover redundancy by creating two full virtual machine copies. The virtual machine files can be placed on the same datastore. However, VMware recommends that you place these files on separate datastores to provide recovery from datastore failures.

vSphere Fault Tolerance breaks the barriers of single processor and now it can accommodate symmetric multiprocessor (SMP) virtual machines with up to four vCPUs with the VMware vSphere® Enterprise Plus Edition™ license.

vSphere Fault Tolerance: Precopy

Slide 9-58

Using vSphere Fault Tolerance, a second VM is created on the secondary host. Then the memory of the source virtual machine is copied to the secondary host.



You can use vSphere Fault Tolerance for most mission-critical virtual machines. vSphere Fault Tolerance provides continuous availability for such a virtual machine by creating and maintaining another VM that is identical and continuously available to replace it in the event of a failover situation.

The protected virtual machine is called the primary VM. The duplicate virtual machine, the secondary VM, is created and runs on another host. The secondary VM's execution is identical to that of the primary VM and it can take over at any point without interruption, thereby providing fault tolerant protection.

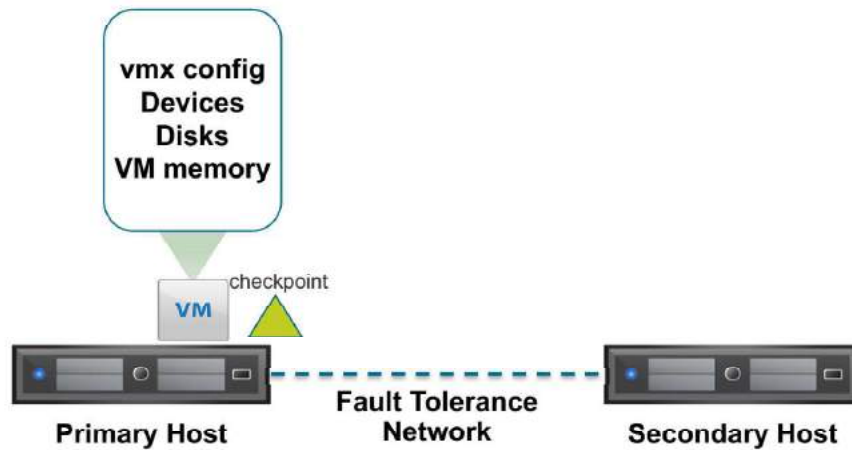
The primary and secondary VMs continuously monitor the status of one another to ensure that vSphere Fault Tolerance is maintained. A transparent failover occurs if the host running the primary VM fails, in which case the secondary VM is immediately activated to replace the primary VM. A new secondary VM is started and fault tolerance redundancy is re-established automatically. If the host running the secondary VM fails, it is also immediately replaced. In either case, users experience no interruption in service and no loss of data.

A fault tolerant virtual machine and its secondary copy are not allowed to run on the same host. This restriction ensures that a host failure cannot result in the loss of both VMs.

vSphere Fault Tolerance Fast Checkpointing

Slide 9-59

The SMP FT checkpoint interval is dynamic. It adapts to maximize the workload performance and can range from as small as a few milliseconds to as large as several hundred milliseconds.



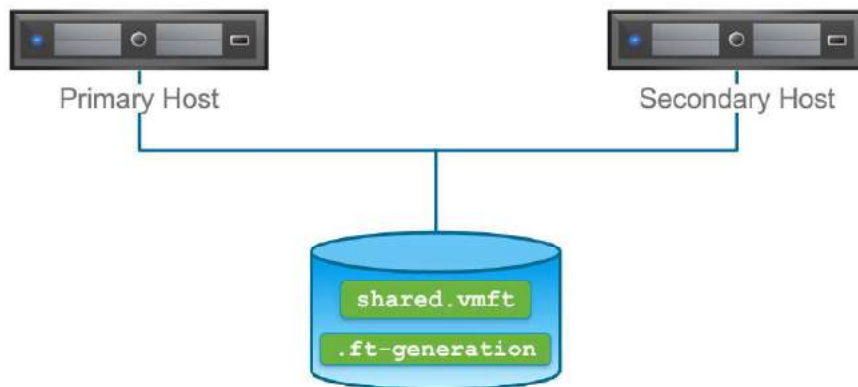
In vSphere Fault Tolerance, checkpoint data is the last changed pages of memory. The source virtual machine is paused to access this memory. This pause is typically under one second.

Shared Files

Slide 9-60

vSphere Fault Tolerance has shared files:

- `shared.vmft` prevents UUID change.
- `.ftgeneration` is for the split-brain condition.



To enable vSphere Fault Tolerance for SMP virtual machines, you must do very fast, continuous copying (checkpointing) of the primary host virtual machine. The primary host virtual machine is copied (checkpointed) periodically, and the copies are sent to a secondary host. If the primary host crashes, the virtual machine continues on the secondary host at the point of its last network send.

The goal is to take checkpoints of virtual machines at least every 10 milliseconds with small CPU overhead in the critical path. You can move noncritical path processing to another processor core. The primary virtual machine is continuously copied (checkpointed) and these copies (checkpoints) are sent to a secondary host.

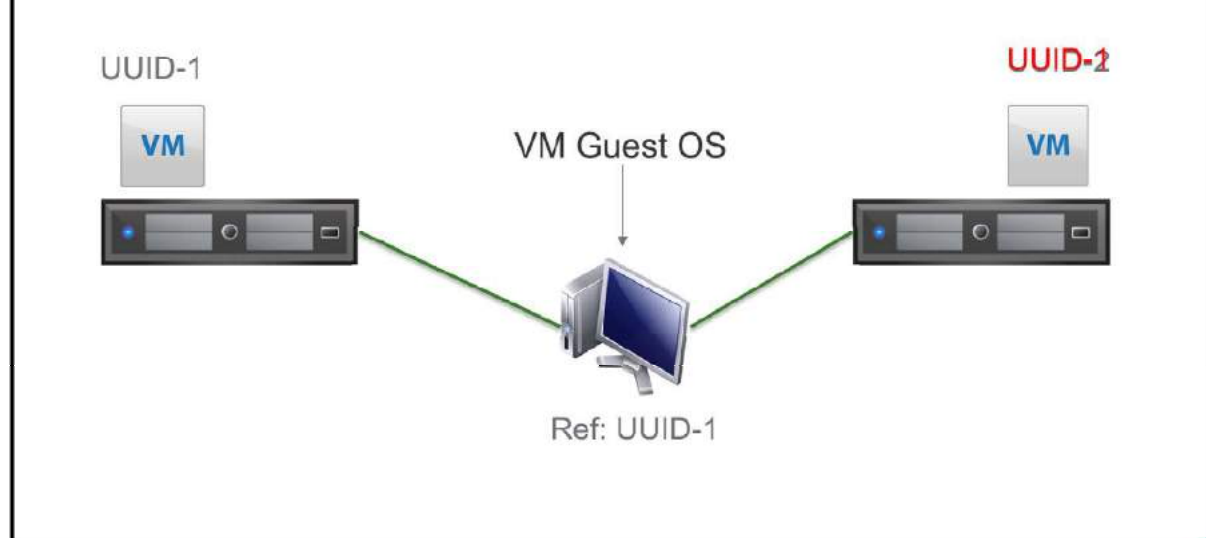
The initial complete copy (checkpoint) is created using a modified form of vSphere vMotion to the secondary host. The primary virtual machine holds each outgoing network packet until the following copy (checkpoint) has been sent to the secondary host.

If the primary host crashes, the virtual machine can be resumed on the backup from the last complete copy (checkpoint).

shared.vmtx File

Slide 9-61

The `shared.vmtx` file, which is found on a shared datastore, is the vSphere Fault Tolerance metadata file and contains the primary and secondary instance UUIDs and the primary and secondary vmx paths.

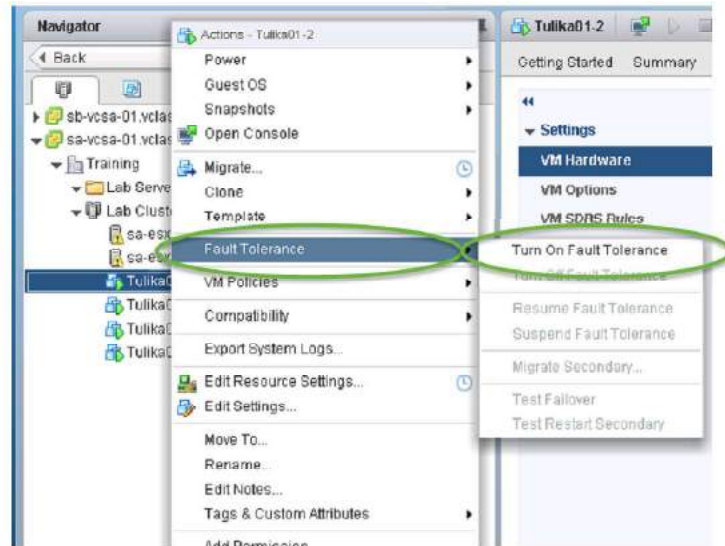


A virtual machine split-brain condition can occur when a host becomes isolated or partitioned from a master host, and the master host cannot communicate with it using heartbeat datastores. In this situation, the master host cannot determine that the host is alive and so declares the host as dead. The master host attempts to restart the virtual machines that are running on the isolated or partitioned host. This attempt succeeds if the virtual machines remain running on the isolated or partitioned host and that host lost access to the virtual machines' datastores when it became isolated or partitioned. A split-brain condition exists because there are two instances of the virtual machine. The `.ftgeneration` file ensures that only one instance of the virtual machine can read from or write to the virtual machine's virtual disks.

Enabling vSphere Fault Tolerance on a Virtual Machine

Slide 9-62

You can turn on vSphere Fault Tolerance for a virtual machine through vSphere Web Client.



After you have taken all of the required steps for enabling vSphere Fault Tolerance for your cluster, you can use the feature by turning it on for individual virtual machines.

Before vSphere Fault Tolerance can be turned on, validation checks are performed on a virtual machine.

After these checks are passed and you turn on vSphere Fault Tolerance for a virtual machine, new options are added to the **Fault Tolerance** section of its context menu. These options include turning off or disabling vSphere Fault Tolerance, migrating the secondary virtual machine, testing failover, and testing restart of the secondary virtual machine.

When vSphere Fault Tolerance is turned on, vCenter Server resets the virtual machine's memory limit and sets the memory reservation to the memory size of the virtual machine. While vSphere Fault Tolerance remains turned on, you cannot change the memory reservation, size, limit, number of virtual CPUs, or shares. You also cannot add or remove disks for the virtual machine. When vSphere Fault Tolerance is turned off, any parameters that were changed are not reverted to their original values.

Review of Learner Objectives

Slide 9-63

You should be able to meet the following objectives:

- Describe benefits of using vSphere Fault Tolerance
- Explain how vSphere Fault Tolerance works
- Describe main vSphere Fault Tolerance features
- Describe how vSphere Fault Tolerance works with vSphere HA and vSphere DRS
- Enable vSphere Fault Tolerance from vSphere Web Client

Lesson 5: vSphere Replication and vSphere Data Protection

Slide 9-64



Lesson 5: vSphere Replication and vSphere Data Protection

Learner Objectives

Slide 9-65

By the end of this lesson, you should be able to meet the following objectives:

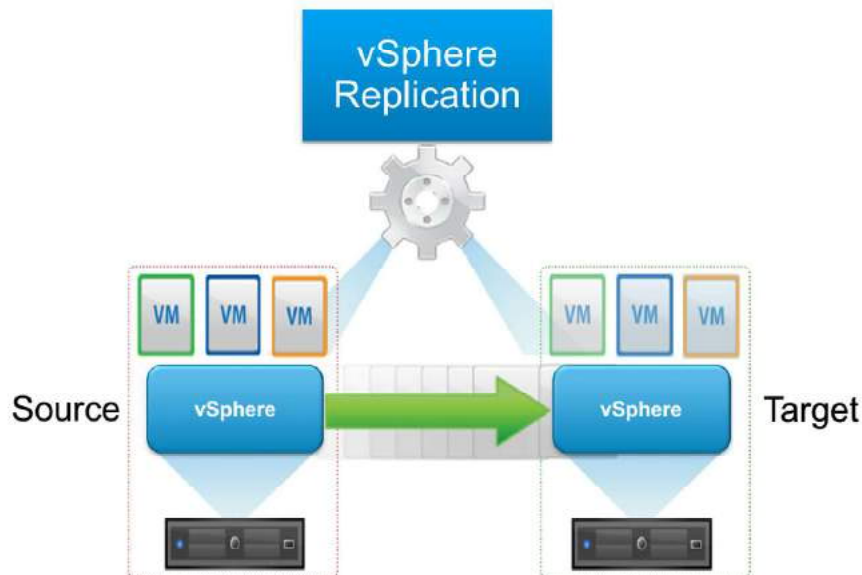
- Describe vSphere Replication
- Identify vSphere Data Protection requirements
- List vSphere Data Protection sizing guidelines
- Describe vSphere Data Protection installation and configuration
- Explain how to back up and restore data with vSphere Data Protection

About vSphere Replication

Slide 9-66

vSphere Replication is an extension to vCenter Server.

It provides hypervisor-based virtual machine replication and recovery.



vSphere Replication is an alternative to storage-based replication. vSphere Replication protects virtual machines from partial or complete site failures by replicating the virtual machines between the following sites:

From a source site to a target site:

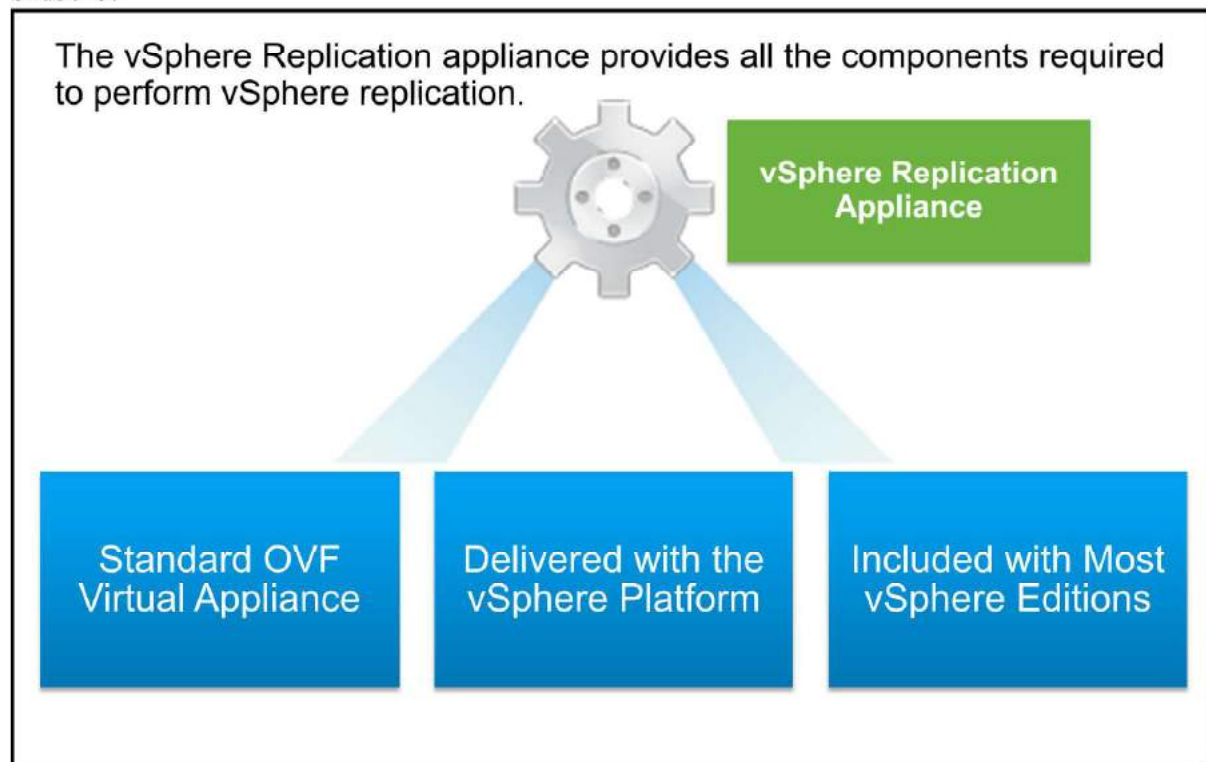
- Within a single site from one cluster to another
- From multiple source sites to a shared remote target site

vSphere Replication provides several benefits as compared to storage-based replication:

- Data protection at lower cost per virtual machine than storage array-based replication.
- A replication solution that allows flexibility in storage vendor selection at the source and target sites.
- Overall lower cost per replication than storage array-based replication

vSphere Replication Appliance

Slide 9-67



The contents of the vSphere Replication appliance include:

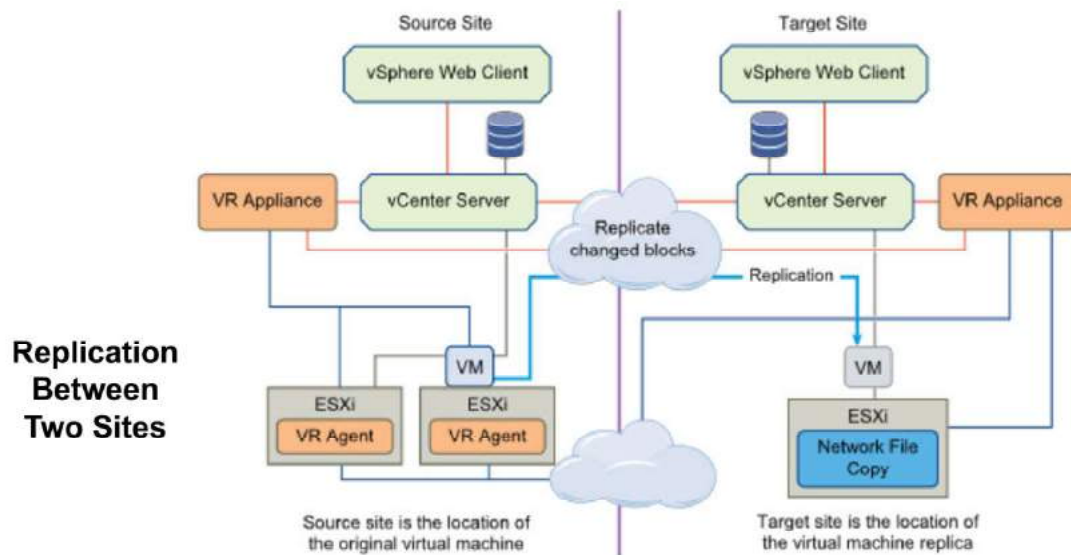
- A vSphere Replication server that provides the core of the vSphere Replication infrastructure.
- An embedded database that stores replication configuration and management information.
- A vSphere Replication management server:
 - Configures the vSphere Replication server
 - Enables, manages, and monitors replications
 - Authenticates users and checks their permissions to perform vSphere Replication operations
- A plug-in to vSphere Web Client that provides a user interface for vSphere Replication.

You can use vSphere Replication immediately after you deploy the appliance. The vSphere Replication appliance provides a virtual appliance management interface (VAMI) that is used to reconfigure the appliance after deployment. For example, you can use VAMI to change the appliance security settings, change the network settings, or configure an external database. You can deploy additional vSphere Replication servers by using a separate OVF package.

Replication Functions

Slide 9-68

vSphere Replication enables replication of a virtual machine from a source site to a target site, monitoring and managing the status of the replication, and recovering the virtual machine at the target site.

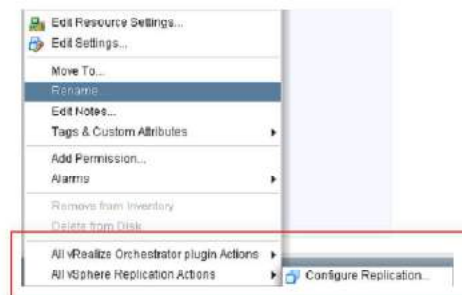


Configuring vSphere Replication for a Single Virtual Machine

Slide 9-69

To configure vSphere Replication for a virtual machine, do one of the following actions in vSphere Web Client:

- Select the virtual machine in the inventory and select **Actions > All vSphere Replication Actions > Configure**.
- Right-click the virtual machine in the inventory and select **All vSphere Replication Actions > Configure**.



vSphere Replication can protect individual virtual machines and their virtual disks by replicating them to another location.

When you configure replication, you set a recovery point objective (RPO) to determine the period of time between replications. For example, an RPO of 1 hour seeks to ensure that a virtual machine loses no more than 1 hour of data during the recovery. For smaller RPOs, less data is lost in a recovery, but more network bandwidth is consumed keeping the replica up to date.

Every time that a virtual machine reaches its RPO target, vSphere Replication records approximately 3800 bytes of data in the vCenter Server events database. If you set a low RPO period, this can quickly create a large volume of data in the database. To avoid creating large volumes of data in the vCenter Server events database, limit the number of days that vCenter Server retains event data. See [Configure Database Retention Policy](#) in the vCenter Server and Host Management Guide. Alternatively, set a higher RPO value.

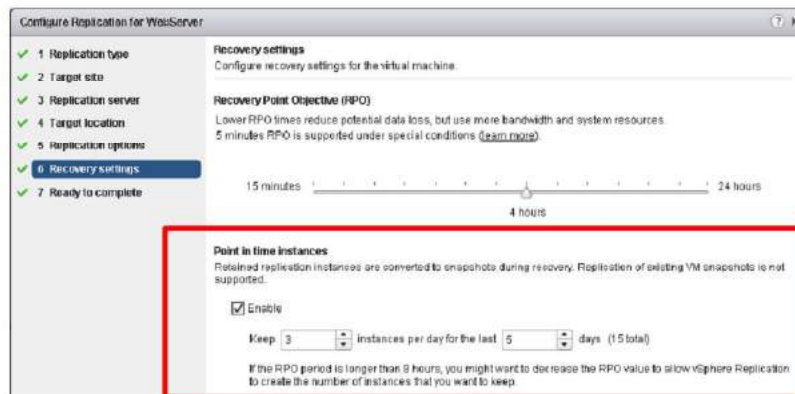
vSphere Replication guarantees crash consistency amongst all the disks that belong to a virtual machine. If you use VSS quiescing, you might obtain a higher level of consistency. The available quiescing types are determined by the virtual machine's operating system. See [Compatibility Matrixes for vSphere Replication 5.5 for Microsoft Volume Shadow Copy Service \(VSS\) quiescing support for Windows virtual machines](#).

You can use vSphere Replication with a Virtual SAN datastore on the source and target sites. See [Using vSphere Replication with Virtual SAN Storage](#) for the limitations when using vSphere Replication with Virtual SAN.

Configuring MPIT

Slide 9-70

Point-in-time instances are configured during vSphere Replication configuration using vSphere Web Client.



When you configure a virtual machine for replication, the vSphere Replication agent sends changed blocks in the virtual machine disks from the source site to the target site, where they are applied to the copy of the virtual machine. This process occurs independently of the storage layer. vSphere Replication performs an initial full synchronization of the source virtual machine and its replica copy. You can use replication seeds to reduce the amount of time and bandwidth required for the initial replication.

During replication configuration, you can set a recovery point objective (RPO) and enable retention of instances from multiple points in time (MPIT).

As administrator, you can monitor and manage the status of the replication. You can view information for incoming and outgoing replications, source and target site status, replication issues, warnings, and errors.

When you manually recover a virtual machine, vSphere Replication creates a copy of the virtual machine connected to the replica disk, but does not connect any of the virtual network cards to port groups. You can review the recovery and status of the replica virtual machine and attach it to the networks. You can recover virtual machines at different points in time, such as the last known consistent state. vSphere Replication presents the retained instances as ordinary virtual machine snapshots to which you can revert the virtual machine.

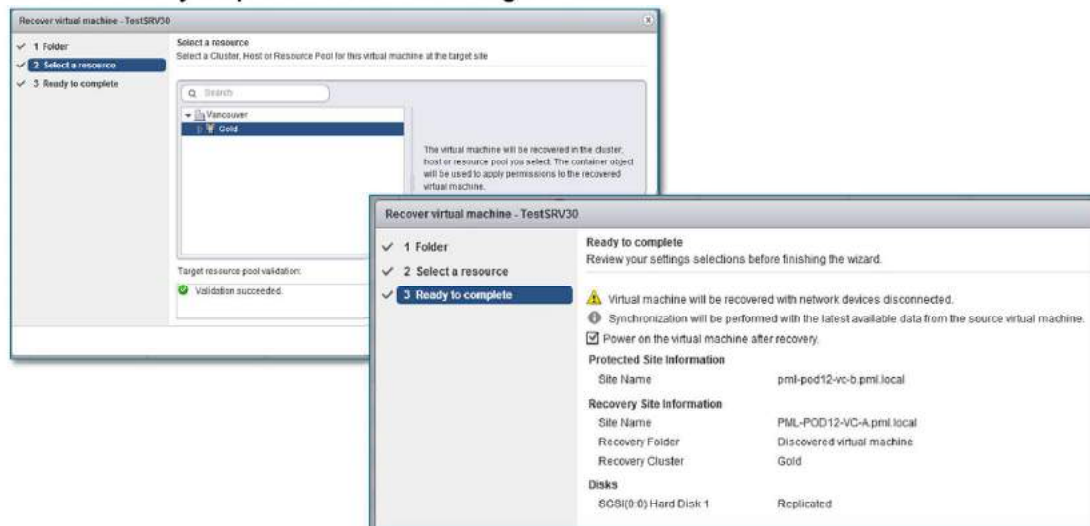
vSphere Replication stores replication configuration data in its embedded database. You can also configure vSphere Replication to use an external database.

You can replicate a virtual machine between two sites. vSphere Replication is installed on both source and target sites. Only one vSphere Replication appliance is deployed on each vCenter Server system. You can deploy additional vSphere Replication servers.

Recovering Virtual Machines

Slide 9-71

With vSphere Replication, you can recover virtual machines that were successfully replicated at the target site.



You can recover one virtual machine at a time on the **Incoming Replications** tab.

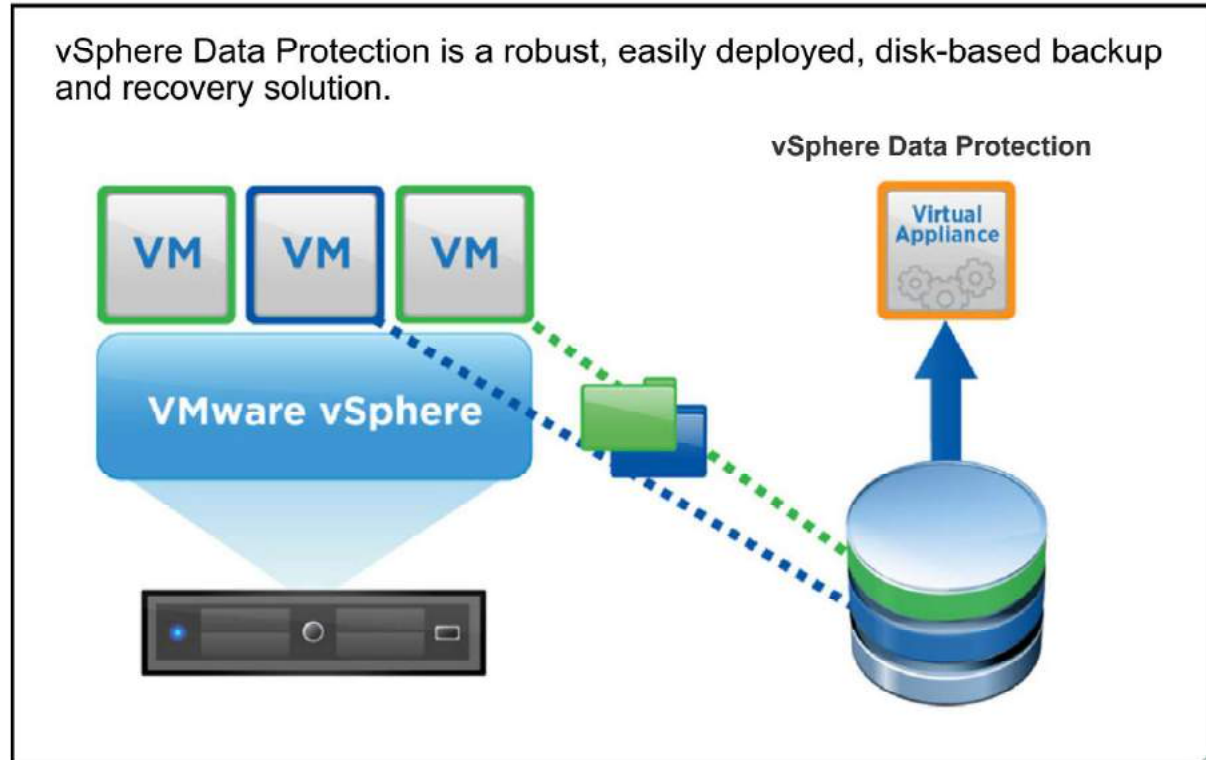
You use the Recover virtual machine wizard in vSphere Web Client at the target site to perform the recovery.

You are asked to select whether to recover the virtual machine with all the latest data, or to recover the virtual machine with the most recent data available on the target site. If you select **Recover with recent changes** to avoid data loss, vSphere Replication performs a full synchronization of the virtual machine from the source site to the target site before recovering the virtual machine. This option requires that the data of the source virtual machine be accessible. You can only select this option if the virtual machine is powered off. If you select **Recover with latest available data**, vSphere Replication recovers the virtual machine by using the data from the most recent replication on the target site, without performing synchronization. Selecting this option results in the loss of any data that has changed since the most recent replication. Select this option if the source virtual machine is inaccessible or if its disks are corrupted.

vSphere Replication validates the input that you provided and recovers the virtual machine. If successful, the virtual machine status changes to Recovered. The virtual machine appears in the inventory of the target site.

About vSphere Data Protection

Slide 9-72




vSphere Data Protection is a backup and recovery solution from VMware. vSphere Data Protection is fully integrated with vCenter Server and vSphere Web Client, providing disk-based backup of virtual machines and applications. vSphere Data Protection is based on the industry-leading EMC Avamar backup and recovery solution.

vSphere Data Protection Requirements and Architecture

Slide 9-73

vSphere Data Protection requires vCenter Server, either the Windows implementation or vCenter Server Appliance.

vSphere Data Protection Components



vSphere Data Protection has certain requirements.

vCenter Single Sign-On is also required. vSphere Data Protection supports backing up virtual machine backups on multiple versions of vSphere.

Web browsers must be enabled with Adobe Flash Player to access vSphere Web Client and vSphere Data Protection functionality. See vSphere documentation for a list of Web browsers that are currently supported with vSphere Web Client.

vSphere Data Protection is deployed as a prebuilt, Linux-based virtual appliance. A maximum of 20 vSphere Data Protection appliances can be deployed per vCenter Server instance. Each appliance is deployed by default with four virtual CPUs and 4 GB of memory. Storage capacity for deduplicated backup data is configured during deployment.

Optionally, as many as eight external proxies (virtual appliances) can be deployed per vSphere Data Protection virtual appliance. Proxies can be deployed to enable SCSI hot-add transport backups of virtual machines running on datastores not directly accessible by the vSphere Data Protection virtual appliance. Examples include vSphere hosts utilizing local direct-attached storage (DAS) and hosts deployed at remote locations. External proxies are required for the Linux logical volume manager (LVM) and Ext4 FLR. External proxies are deployed by using the vSphere Data Protection configure user interface.

vSphere Data Protection application agents are downloaded by using vSphere Web Client and are installed in the guest operating system (OS) of the virtual machines running Exchange Server, SQL Server, and SharePoint.

vSphere Data Protection supports as much as 8 TB of deduplicated backup data capacity per appliance. Assuming average virtual machine sizes, average data change rates, and a 30-day retention policy, approximately 150 to 200 virtual machines can be protected with a vSphere Data Protection appliance. Every environment is different, so actual results vary.

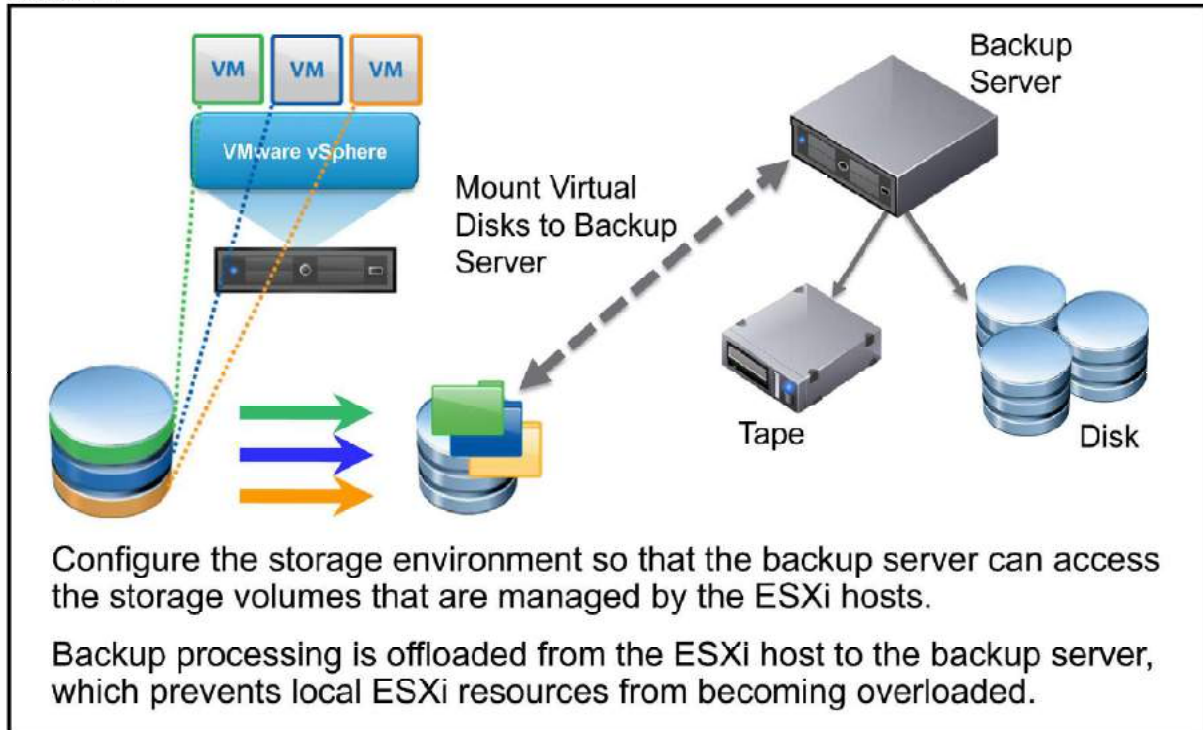
vSphere Data Protection virtual appliances can be deployed to vSAN, VMFS, vSphere Virtual Volumes, and NFS datastores. The virtual machine disk (VMDK) files for a vSphere Data Protection virtual appliance can be stored together on the same datastore or distributed across multiple vSphere datastores. You can detach VMDK files that are included in an existing vSphere Data Protection virtual appliance backup data partition and attach these files to a newly deployed appliance.

More storage capacity can be added after the appliance has been deployed, up to a maximum of 8 TB. For example, a vSphere Data Protection appliance originally deployed with 2 TB of backup data storage capacity can be expanded by 6 TB, for a total of 8 TB of capacity.

When determining storage capacity requirements, several factors, including the number of protected virtual machines, amount and formats of data being backed up, retention periods, data change rates, and others, should be considered.

Offloaded Backup Processing

Slide 9-74



Perhaps one of the biggest bottlenecks can be the backup server that is handling all the backup coordination tasks. One of these backup tasks is copying data from point A to point B. Other backup tasks do a lot of CPU processing. For example, tasks are performed to determine what data to back up and what not to back up. Other tasks are performed to deduplicate data and compress data that is written to the target.

A server with insufficient CPU resources can greatly reduce backup performance. It is important not to skimp on resources for your backup server. A physical server or virtual machine with an ample amount of memory and CPU capacity is necessary for the best backup performance possible.

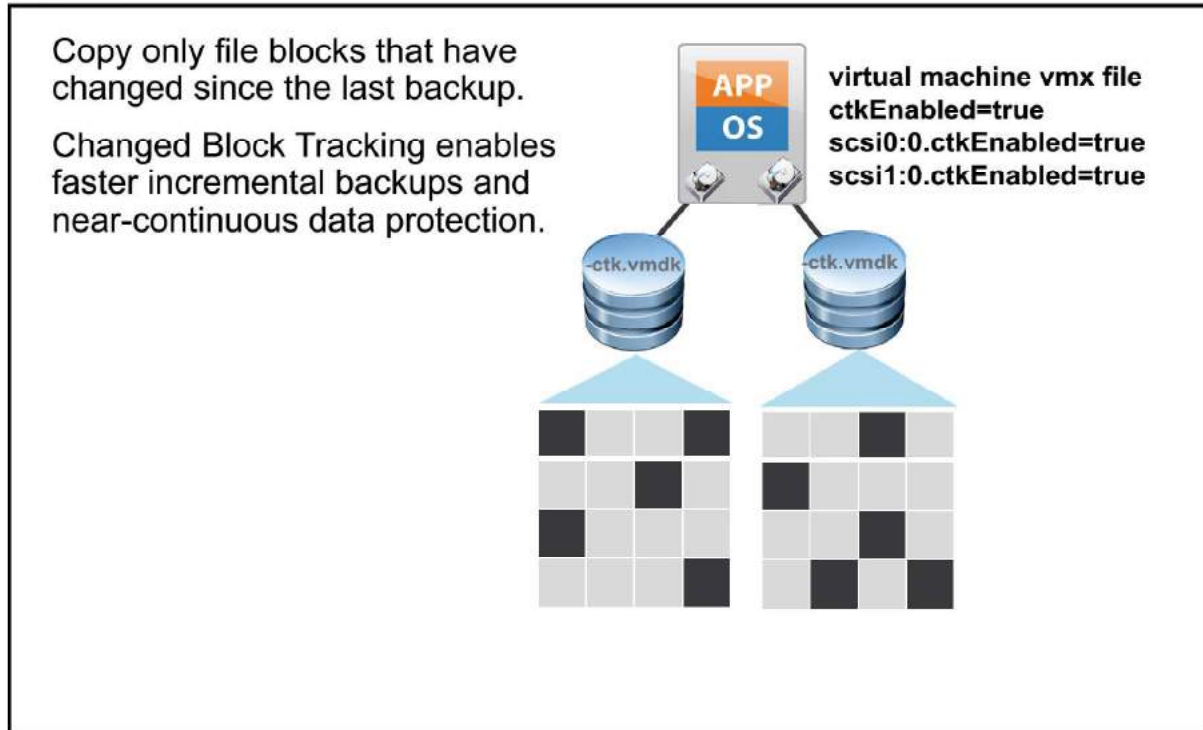
The motivation to use LAN-free backups is to reduce the stress on the physical resources of the ESXi host when virtual machines are backed up. LAN-free backups reduce the stress by offloading backup processing from the ESXi host to a backup proxy server.

You can configure your environment for LAN-free backups to the backup server, also called the backup proxy server. For LAN-free backups, the backup server must be able to access the storage managed by the ESXi hosts on which the virtual machines to back up are running.

If you are using NAS or direct-attached storage, ensure that the backup proxy server is accessing the volumes with a network-based transport. If you will be running a direct SAN backup, zone the SAN and configure the disk subsystem host mappings. The host mappings must be configured so that all ESXi hosts and the backup proxy server access the same disk volumes.

Changed Block Tracking

Slide 9-75



Changed Block Tracking (CBT) is a VMkernel feature that keeps track of the storage blocks of virtual machines as they change over time. The VMkernel keeps track of block changes on virtual machines, which enhances the backup process for applications that have been developed to take advantage of VMware vSphere® vStorage APIs.

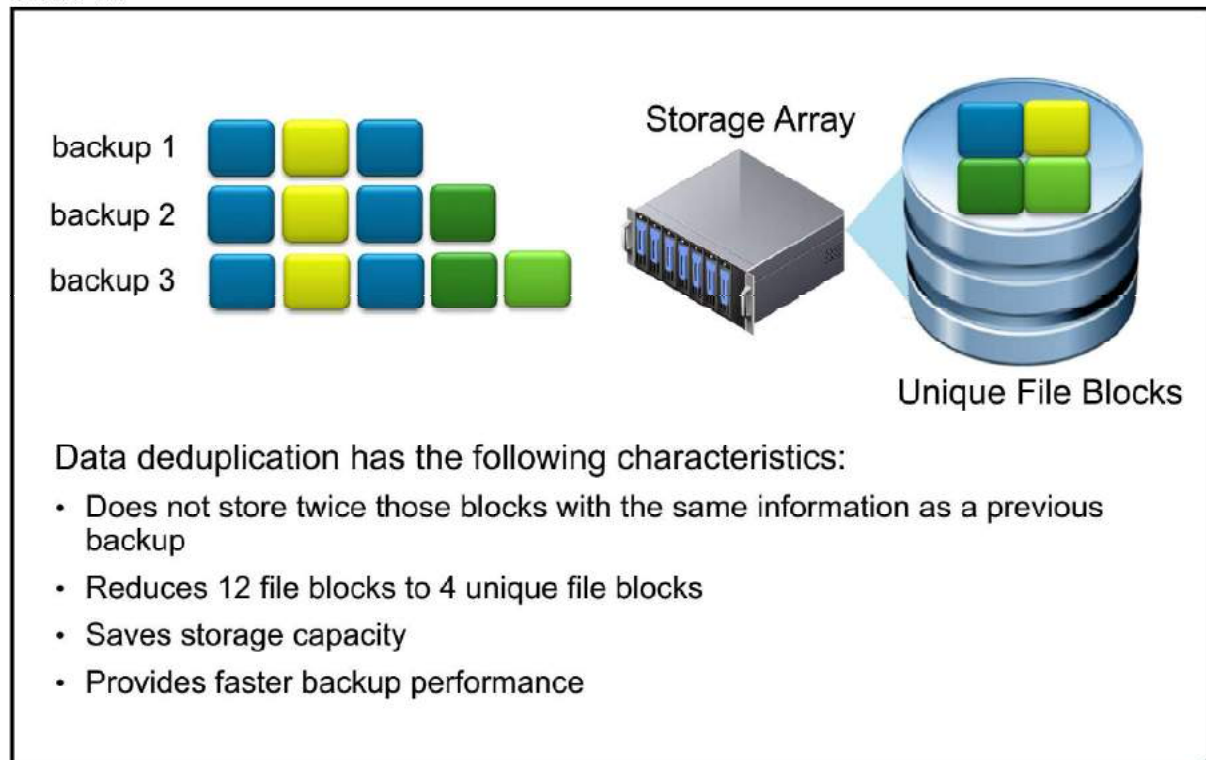
To increase the efficiency of image-level backups, vSphere Data Protection utilizes the CBT feature, which greatly reduces the backup time of a given virtual machine image and provides the ability to process a large number of virtual machines within a particular backup window.

By leveraging CBT during restores, vSphere Data Protection offers fast and efficient recoveries of virtual machines to their original location. During a restore process, vSphere Data Protection uses CBT to determine which blocks have changed since the last backup. The use of CBT reduces data transfer within the vSphere environment during a recovery operation and more importantly reduces the recovery time.

Additionally, vSphere Data Protection evaluates the workload between both restore methods (full image restore or a recovery leveraging CBT) and performs the method resulting in the fastest restore time. This is useful in scenarios where the change rate since the last backup in a virtual machine being restored is very high and the overhead of a CBT analysis operation would be more costly than a direct full-image recovery. vSphere Data Protection determines which method results in the fastest image recovery times for virtual machines in the environment.

Data Deduplication

Slide 9-76



Deduplication Store Benefits: Enterprise data is highly redundant, with identical files or data stored within and across systems (for example, OS files or documents sent to multiple recipients). Edited files also have tremendous redundancy with previous versions. Traditional backup methods magnify this by storing all of the redundant data over and over again. vSphere Data Protection uses patented deduplication technology to eliminate redundancy at both the file and the subfile data segment level.

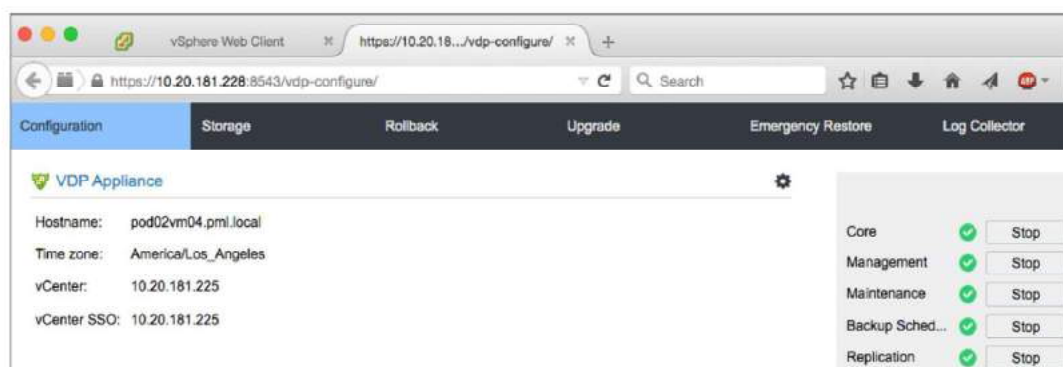
During a backup, vSphere Data Protection creates a quiesced snapshot of the virtual machine. Deduplication is automatically performed with every backup operation.

vSphere Data Protection Deployment and Configuration

Slide 9-77

vSphere Data Protection is deployed using vSphere Web Client or vSphere Client from a prepackaged Open Virtualization Archive (OVA) file.

vSphere Data Protection: Configuring the UI to Run in Maintenance Mode



After the appliance is deployed and powered on, a Web browser is used to access the vSphere Data Protection configure utility to perform the initial configuration. The first time a user connects to the vSphere Data Protection configure UI, it runs in Install Mode. With the Install Mode wizard, items such as IP address, host name, DNS, time zone, vCenter Server connection information, and storage are configured.

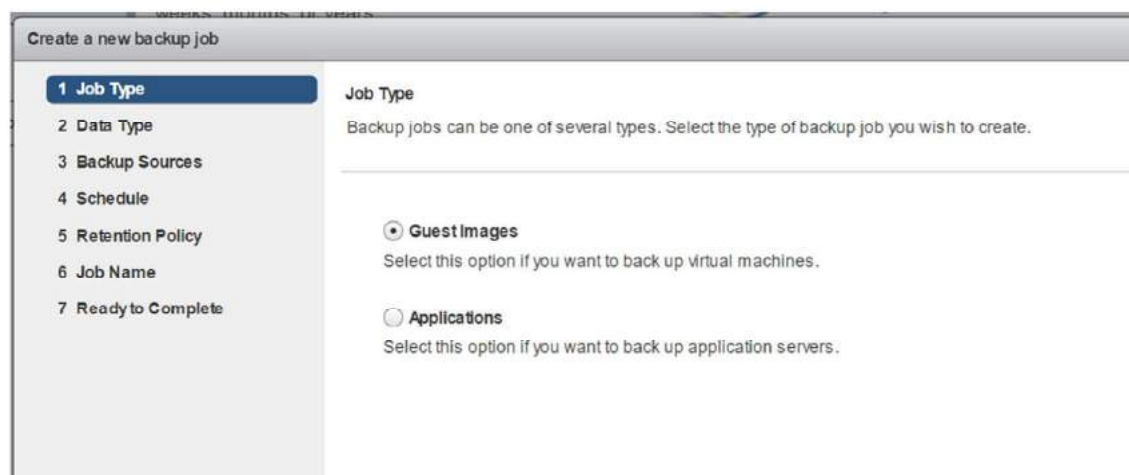
A storage performance test can also be run at this time, which is highly recommended to validate that the storage on which vSphere Data Protection is running meets or exceeds recommended performance levels. Upon successful completion of these tasks, the appliance must be rebooted. The reboot takes several minutes as the appliance automatically finalizes its initial configuration.

After initial configuration, the vSphere Data Protection configure utility runs in maintenance mode. This mode is used to perform functions, such as starting and stopping services in the appliance, deploying proxies, collecting logs, performing emergency restores, upgrading the vSphere Data Protection appliance, and rolling back the appliance to a previous valid configuration state.

Virtual Machine Backup

Slide 9-78

You create and edit a backup job on the **Backup** tab of the vSphere Data Protection UI in vSphere Web Client or vSphere Client.



Individual virtual machines or specific VMDK files can be selected for backup. Containers of virtual machines such as data centers, clusters, and resource pools can also be selected for backup. When a virtual machine is added to the protected container, it is backed up. Likewise, when a virtual machine is removed from the container, it is no longer included in the backup job. Restore points are preserved until expired by the retention policy.

Backup jobs can be scheduled daily, weekly, or monthly. Each job starts at its scheduled time and runs once on the day it is scheduled.

The retention policy can be defined in a few ways, for example, retention for 30 days or until a specific date. A custom retention policy also can be defined.

After a backup job has been created, it can be edited or deleted. You can also clone a backup job. Cloning can be useful, for example, if the backup administrator wants to easily duplicate an existing custom retention policy for a new set of virtual machines.

The initial backup of a virtual machine can take some time because all data blocks that make up that virtual machine must be backed up. Subsequent backups typically take much less time because vSphere Data Protection utilizes CBT in vSphere.

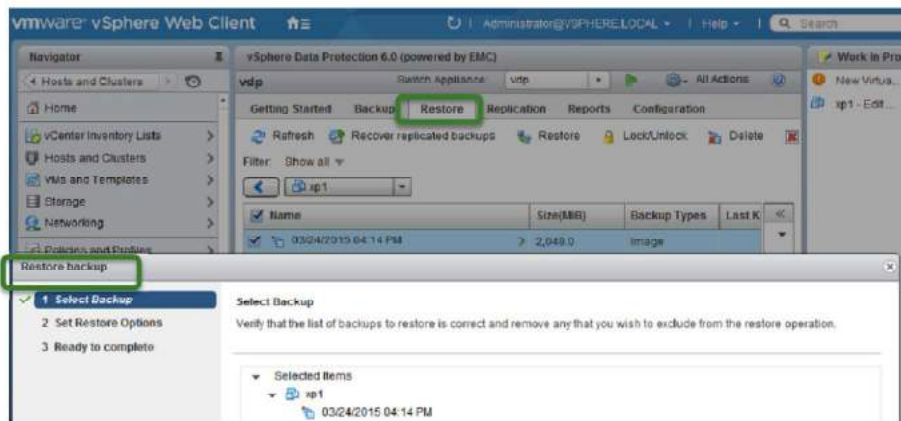
Performing Restores with vSphere Data Protection

Slide 9-79

You can restore the backups to the original or an alternate location.

You can restore an entire virtual machine from the **Restore** tab in the vSphere Data Protection user interface in vSphere Web Client.

- The administrator can browse the list of protected virtual machines and select one or more restore points.



- Individual VMDKs can also be restored.

After you back up virtual machines, you can restore virtual machines from replicated backup data locally or at another target location. For example, a vSphere Data Protection virtual appliance protects virtual machines in a primary data center. Backup data is replicated by vSphere Data Protection from the primary data center to a vSphere Data Protection virtual appliance at a disaster recovery data center. Disaster strikes the primary data center and virtual machines, including the vSphere Data Protection virtual appliance, are lost. When the primary data center is back online, a new vSphere Data Protection virtual appliance is deployed and connected to vSphere Data Protection at the disaster recovery site. The new vSphere Data Protection virtual appliance can retrieve backup data from the disaster recovery site and perform restores at the primary data center.

vSphere Data Protection offers fast and efficient recovery by leveraging CBT. When restoring an entire virtual machine to its original location, the workloads of both a full image restore and a restore leveraging CBT are evaluated. vSphere Data Protection intelligently determines which method results in the faster virtual machine recovery time.

Restore operations are performed on the **Restore** tab. This tab displays a list of virtual machines that have been backed up by the vSphere Data Protection appliance. By navigating through the list of backups, you can select and restore specific backups.

Before you select a backup to restore, note the expiration date of the backup. Over time, the information that is displayed on the **Restore** tab might become out of date. You can use the **Refresh** button to see the most up-to-date information on backups that are available for restore.

Review of Learner Objectives

Slide 9-80

You should be able to meet the following objectives:

- Describe vSphere Replication
- Identify vSphere Data Protection requirements
- List vSphere Data Protection sizing guidelines
- Describe vSphere Data Protection installation and configuration
- Explain how to back up and restore data with vSphere Data Protection

Key Points

Slide 9-81


- vSphere HA restarts virtual machines on the remaining hosts in the cluster.
- Hosts in vSphere HA clusters have a master-slave relationship.
- You implement redundant heartbeat networks either with NIC teaming or by creating additional heartbeat networks.
- vSphere Fault Tolerance provides zero downtime for applications that must be available at all times.
- vSphere Replication can be used to protect virtual machines as part of a disaster recovery strategy.
- vSphere Replication is the only hypervisor-based replication solution that operates at the individual VMDK level, enabling replication between datastores hosted on any storage.
- vSphere Data Protection is a backup and recovery solution from VMware.

Questions?

MODULE 10

vSphere DRS

Slide 10-1



You Are Here

Slide 10-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
- 10. vSphere DRS**
11. vSphere Update Manager

Importance

Slide 10-3

As you scale your vSphere environment, you must be aware of the vSphere features and functions that will help you manage the hosts in your environment.

Learner Objectives

Slide 10-4

By the end of this module, you should be able to meet the following objectives:

- Describe the functions of a vSphere DRS cluster
- Create a vSphere DRS cluster
- View information about a vSphere DRS cluster
- Remove a host from a vSphere DRS cluster

vSphere DRS Cluster Prerequisites

Slide 10-5

vSphere DRS works best when the virtual machines meet vSphere vMotion migration requirements.

To use vSphere DRS for load balancing, the hosts in the cluster must be part of a vSphere vMotion migration network. If not, vSphere DRS can still make initial placement recommendations.

To use shared storage, configure all hosts in the cluster: Datastores must be accessible by source and destination hosts.



An ESXi host that is added to a vSphere DRS cluster must meet certain prerequisites to use cluster features:

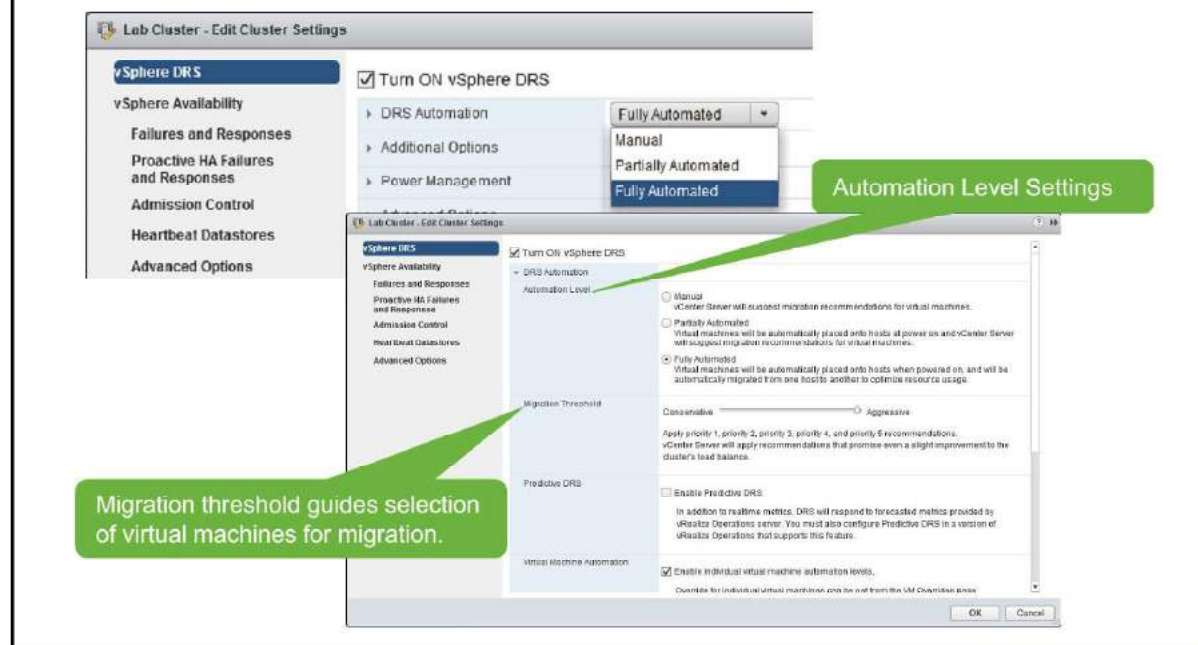
- vSphere DRS works best if the virtual machines meet vSphere vMotion requirements.
- To use vSphere DRS for load balancing, the hosts in your cluster must be part of a vSphere vMotion network.
- Configure all managed hosts to use shared storage: VMFS, vSAN, vSphere Virtual Volumes, or NFS datastores.
- Place the disks of all virtual machines on shared storage that is accessible by source and destination hosts.

vSphere DRS clusters can be created, or vSphere DRS can be enabled for existing vSphere HA or vSAN clusters.

vSphere DRS Cluster Settings: Automation Level

Slide 10-6

Configure the automation level for the initial placement of virtual machines and dynamic balancing while virtual machines are running.



To create a vSphere DRS cluster, right-click your data center in the inventory and select **New Cluster**. The New Cluster dialog box appears. Enter a descriptive name for your cluster and select the **Turn On VMware DRS** check box.

On the vSphere DRS window that appears (shown on the slide), you define the automation level. The automation level determines whether vSphere DRS makes migration recommendations or automatically places virtual machines on hosts. vSphere DRS makes placement decisions when a virtual machine powers on and when virtual machines must be rebalanced across hosts in the cluster. The following automation levels are available:

- **Manual:** When you power on a virtual machine, vSphere DRS displays a list of recommended hosts on which to place the virtual machine. When the cluster becomes imbalanced, vSphere DRS displays recommendations for virtual machine migration.
- **Partially automated:** When you power on a virtual machine, vSphere DRS places it on the best-suited host. When the cluster becomes imbalanced, vSphere DRS displays recommendations for manual virtual machine migration.
- **Fully automated:** When you power on a virtual machine, vSphere DRS places it on the best-suited host. When the cluster becomes imbalanced, vSphere DRS migrates virtual machines from overutilized hosts to underutilized hosts to ensure a balanced use of cluster resources.

The migration threshold determines how aggressively vSphere DRS selects to migrate virtual machines:

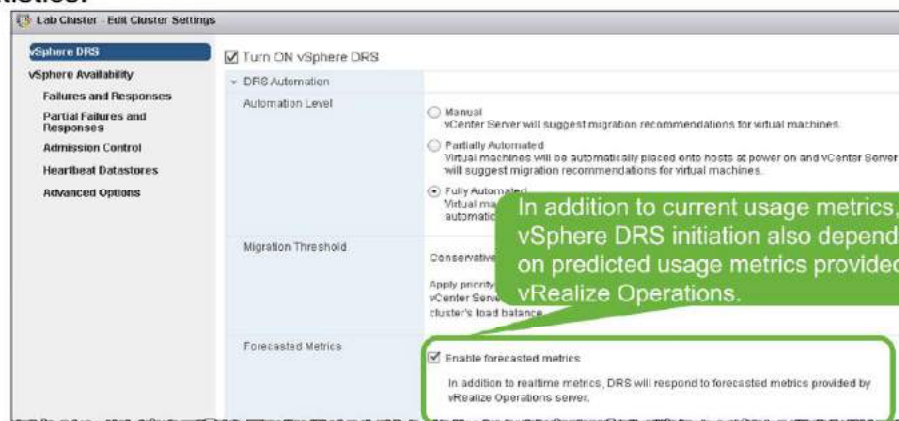
- Level 1 (Conservative): Applies only priority 1 recommendations. vCenter Server applies only recommendations that must be taken to satisfy cluster constraints, such as affinity rules and host maintenance.
- Level 2: Apply priority 1 and priority 2 recommendations. vCenter Server applies recommendations that promise a significant improvement to the cluster's load balance.
- Level 3 (default): Apply priority 1, priority 2, and priority 3 recommendations. vCenter Server applies recommendations that promise at least good improvement to the cluster's load balance.
- Level 4: Apply priority 1, priority 2, priority 3, and priority 4 recommendations. vCenter Server applies recommendations that promise even a moderate improvement to the cluster's load balance.
- Level 5 (Aggressive): Apply all recommendations. vCenter Server applies recommendations that promise even a slight improvement to the cluster's load balance.

vSphere DRS Cluster Settings: Forecasted Metrics

Slide 10-7

When forecasted metrics is enabled:

- The vSphere DRS data collector retrieves statistics from the following sources:
 - Resource usage statistic from hosts
 - Predictive statistic from VMware vRealize® Operations™ server (stored in a separate cache)
- The predicted usage statistics always take precedence over current usage statistics.



When the demands for resources on a virtual machine changes, vSphere DRS itself does not have the ability to predict these resource demands. However, by utilizing statistics from vRealize Operations, vSphere DRS can analyze expected workload on virtual machines in the near future. vSphere DRS can then use the predicted workload values with the current usage values to make placement decisions in advance of the resources being requested.

If the current usage is higher than the predicted usage, the current usage statistics take precedence and these values are used to generate recommendations.

If the predicted usage is higher than the current usage, the predicted usage value is used as current demand of the VM. Thus, the resources are available in advance of the VM needing them, avoiding any performance impact that might occur.

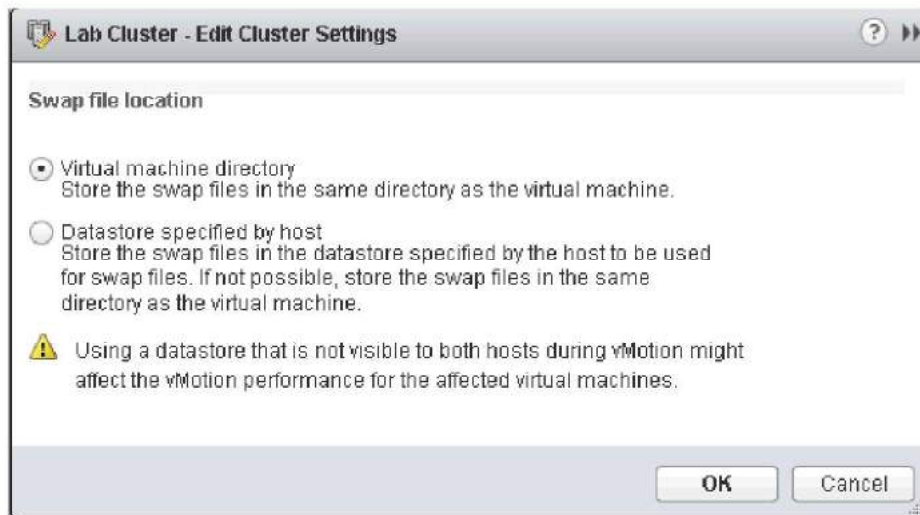
By default, Proactive DRS uses values for one hour in advance to ensure timely migration.

Other Cluster Settings: Swap File Location

Slide 10-8

By default, swap files for a virtual machine are located on a datastore in the folder containing other virtual machine files.

You may place virtual machine swap files on an alternative datastore.



By default, the swap file is created in the same location as the virtual machine's configuration file, which could either be on a VMFS datastore, a vSAN datastore, or a vSphere Virtual Volumes datastore. On a vSAN datastore or a vSphere Virtual Volumes datastore, the swap file is created as a separate vSAN or vSphere Virtual Volumes object.

A swap file is created by the ESXi host when a virtual machine is powered on. If this file cannot be created, the virtual machine cannot power on. Instead of accepting the default, you can also use the following options:

- Use per-virtual machine configuration options to change the datastore to another shared storage location.
- Use host-local swap, which allows you to specify a datastore stored locally on the host. This allows you to swap at a per-host level, saving space on the SAN. However, it can lead to a slight degradation in performance for vSphere vMotion because pages swapped to a local swap file on the source host must be transferred across the network to the destination host. Currently, vSAN and vSphere Virtual Volumes datastores cannot be specified for host-local swap.

vSphere DRS Cluster Settings: Virtual Machine Affinity

Slide 10-9

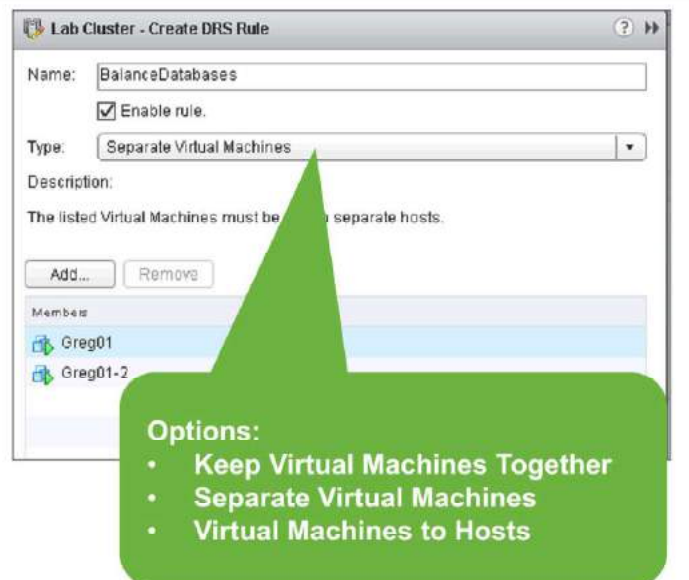
vSphere DRS affinity rules specify that selected virtual machines be placed either on the same host (affinity) or on separate hosts (anti-affinity).

Affinity rules:

- Use for multivirtual machine systems where virtual machines communicate heavily with one another.

Anti-affinity rules:

- Use for multivirtual machine systems where load balance or high availability is desired.



After a vSphere DRS cluster is created, you can edit its properties to create rules that specify affinity. The following types of rules exist:

- Affinity rules: vSphere DRS should try to keep certain virtual machines together on the same host (for example, for performance reasons).
- Anti-affinity rules: vSphere DRS should try to make sure that certain virtual machines are not together (for example, for availability reasons).

The slide shows an anti-affinity rule that requires two virtual machines (Greg01 and Greg01-2) to be placed on different hosts, possibly due to availability or performance reasons. Conversely, affinity rules can be used to keep certain virtual machines on the same host because of increased locality or performance benefits, for example, if virtual machines are communicating heavily with one another.

If two rules conflict, you are prevented from enabling both.

When you add or edit a rule, and the cluster is immediately in violation of that rule, the system continues to operate and tries to correct the violation.

For vSphere DRS clusters that have a default automation level of manual or partially automated, migration recommendations are based on both rule fulfillment and load balancing.

vSphere DRS Cluster Settings: DRS Groups

Slide 10-10

VM or host groups are used in defining VM-Host affinity rules.

Types of VM or host groups:

- VM group:
 - A virtual machine can belong to multiple virtual machine DRS groups.
- Host group:
 - A host can belong to multiple host DRS groups.

A virtual machine can belong to multiple VM or host groups.

A host can belong to multiple VM or host groups.

The screenshot shows the vSphere DRS configuration interface for VM-Host Groups. It features a table with columns for Name and Type. The table lists three groups: 'Database server' (VM Group), 'Application server' (VM Group), and 'Group B' (Host Group). Below the table are buttons for 'Add...', 'Edit...', and 'Delete'. A second section titled 'VM-Host Group Members' contains 'Add' and 'Remove' buttons. Underneath, a sub-section for 'Database server Group Members' lists two hosts: 'Tullika01-2' and 'Tullika01-3'.

Name	Type
Database server	VM Group
Application server	VM Group
Group B	Host Group

VM-Host Group Members

Database server Group Members

- Tullika01-2
- Tullika01-3

The third vSphere DRS affinity rule option is to set a Virtual Machines to Hosts affinity rule. This type of rule specifies whether virtual machines can or cannot be run on a host. For ease of administration, virtual machines can be placed in VM or host groups. You can create one or more VM or host groups in a vSphere DRS cluster, each consisting of one or more virtual machines. On the slide, Database Server is a virtual machine DRS group. A virtual machine can belong to more than one virtual machine DRS group.

Likewise, a host DRS group consists of one or more ESXi hosts. A host can belong to more than one host DRS group. On the slide, Group B is a host DRS group.

The main use for DRS groups is to assist in defining DRS rules known as the Virtual Machines to Hosts affinity rules.

vSphere DRS Cluster Settings: VM-Host Affinity Rules

Slide 10-11

A VM-Host affinity rule:

- Specifies an affinity (or anti-affinity) relationship between a virtual machine DRS group and a host DRS group
- Is either a required rule or a preferential rule

Other options:
Must run on hosts in group,
Must not run on hosts in group,
Should not run on hosts in group.

Lab Cluster - Create DRS Rule

Name:

Enable rule.

Type:

Description:

Virtual machines that are members of the Cluster DRS VM Group Group A should run on host group Group B.

VM Group:

Host Group:

A Virtual Machines to Hosts affinity or anti-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 an affinity rule for virtual machines, which specifies affinity (or anti-affinity) between individual virtual machines, a Virtual Machines to Hosts affinity rule specifies an affinity relationship between a group of virtual machines and a group of hosts. Rules are either required or preferential.

A Virtual Machines to Hosts affinity rule includes the following components:

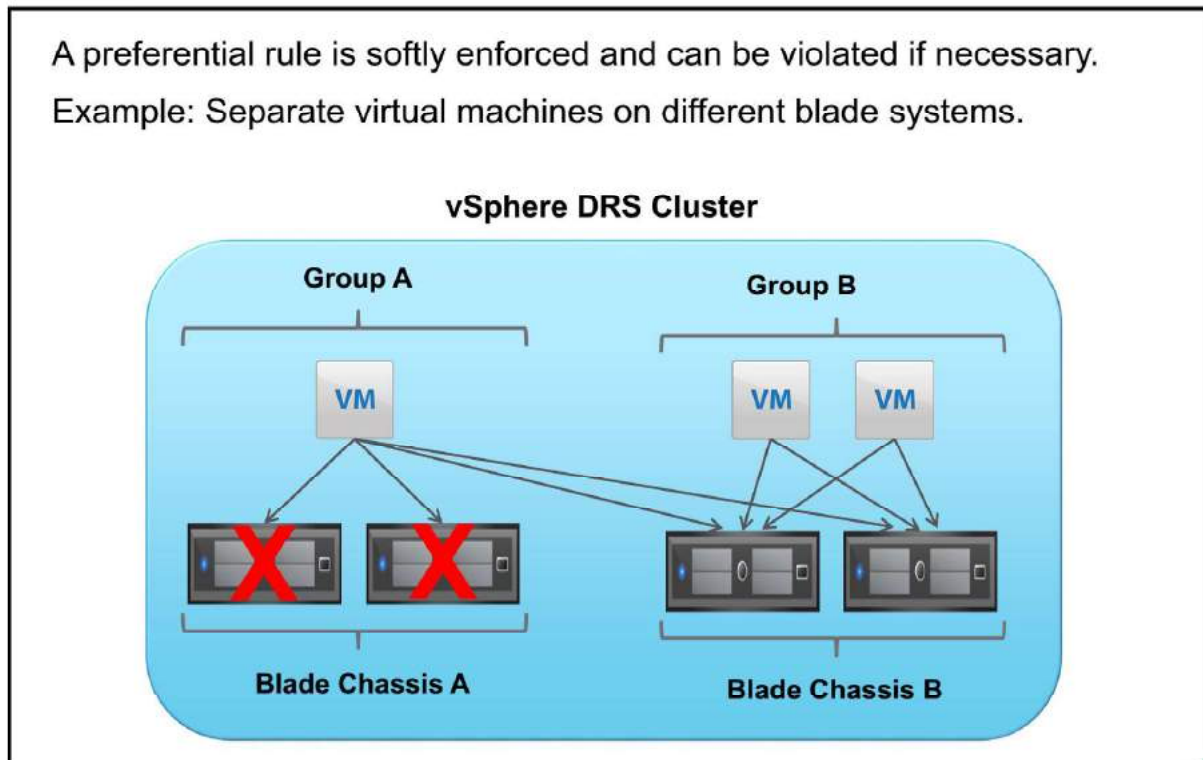
- A virtual machine DRS group
- A host DRS group
- A designation of whether the rule is a requirement (must) or a preference (should) and whether it is affinity (run on) or anti-affinity (not run on).

Because Virtual Machines to Hosts affinity rules are cluster-based, the virtual machines and hosts that are included in a rule must all reside in the same cluster. If a virtual machine is removed from the cluster, it loses its membership from all virtual machine groups, including the DRS group affiliation, even if it is later returned to the cluster.

VM-Host Affinity Rule: Preferential

Slide 10-12

A preferential rule is softly enforced and can be violated if necessary.
Example: Separate virtual machines on different blade systems.



A preferential rule is softly enforced. Preferential rules can be violated to allow the proper functioning of vSphere DRS, vSphere HA, and VMware vSphere® Distributed Power Management™.

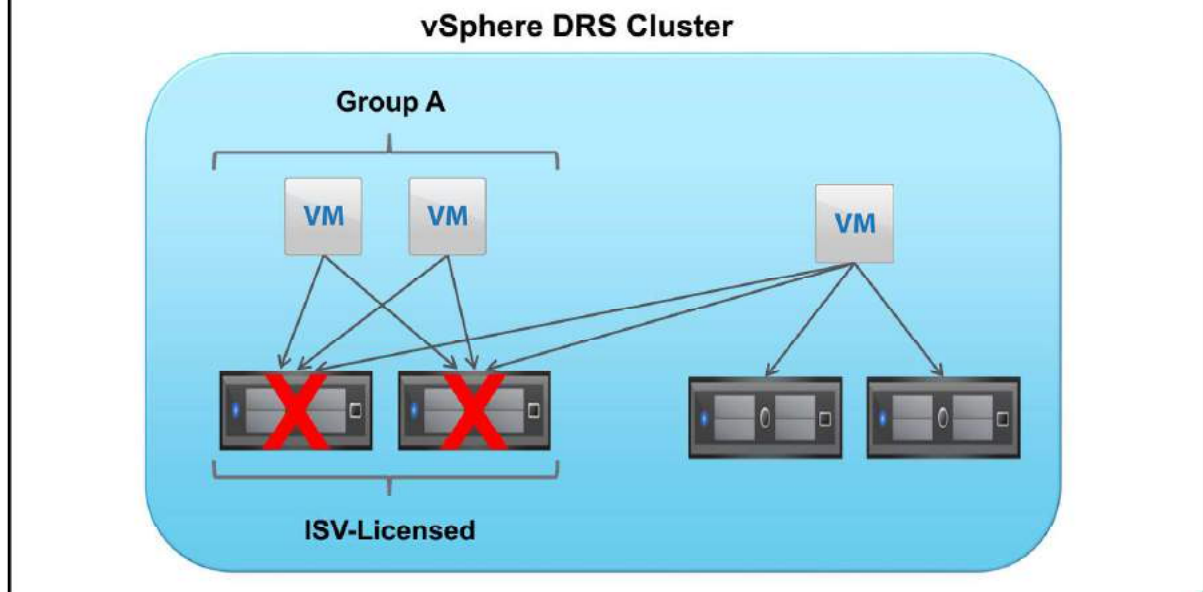
A preferential rule might be used for separating virtual machines onto different blade systems for better performance. On the slide, Group A and Group B are virtual machine DRS groups. Blade Chassis A and Blade Chassis B are host DRS groups. The goal is to force the virtual machines in Group A to run on the hosts in Blade Chassis A and to force the virtual machines in Group B to run on the hosts in Blade Chassis B. If the hosts fail, vSphere HA restarts the virtual machines on the other hosts in the cluster. If the hosts are put into maintenance mode or become overutilized, vSphere DRS moves the virtual machines on the other hosts in the cluster.

VM-Host Affinity Rule: Required

Slide 10-13

A required rule is strictly enforced and can never be violated.

Example: Enforce host-based ISV licensing.



A Virtual Machines to Hosts affinity rule that is required, instead of preferential, can be used when the software running in your virtual machines has licensing restrictions. You can enforce this rule when the software running in your virtual machines has licensing restrictions. You can place such virtual machines in a DRS group. Then create a rule that requires the virtual machines to run on a host DRS group, which contains hosts with the required licenses.

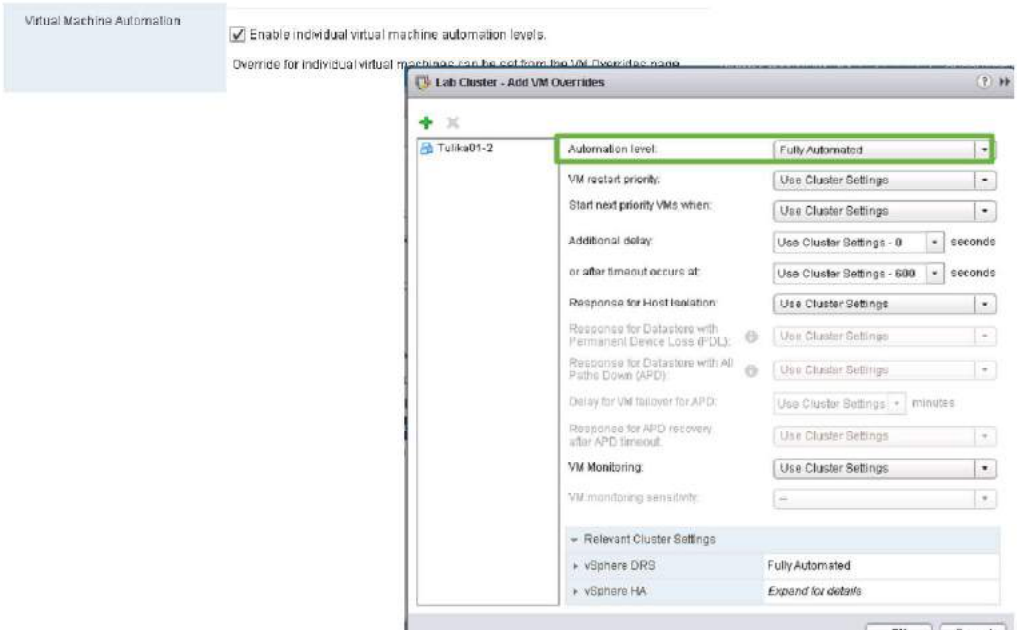
When you create a Virtual Machines to Hosts affinity rule that is based on the licensing or hardware requirements of the software running in your virtual machines, you are responsible for ensuring that the groups are properly set up. The rule does not monitor the software running in the virtual machines nor does it know which third-party licenses are in place on which ESXi hosts.

On the slide, Group A is a virtual machine DRS group. You can force Group A to run on hosts in the group called ISV-Licensed to ensure that the virtual machines in Group A run on hosts that have the required licenses. But if the hosts in the group ISV-Licensed fail, vSphere HA cannot restart the virtual machines in Group A on hosts that are not in the group. If the hosts in the group ISV-Licensed are put into maintenance mode or become overutilized, vSphere DRS cannot move the virtual machines in Group A to hosts that are not in the group.

vSphere DRS Cluster Settings: Automation at the Virtual Machine Level

Slide 10-14

You can customize the automation level for individual virtual machines in a cluster to override the automation level set on the entire cluster.



Virtual Machine Automation

Enable individual virtual machine automation levels.

Override for individual virtual machines can be set from the VM Description page.

Lab Cluster - Add VM Overrides

Tulika01-2

Automation level: Fully Automated

VM restart priority: Use Cluster Settings

Start next priority VMs when: Use Cluster Settings

Additional delay: Use Cluster Settings - 0 seconds

or after timeout occurs at: Use Cluster Settings - 600 seconds

Response for Host Isolation: Use Cluster Settings

Response for Datastore with Permanent Device Loss (PDL): Use Cluster Settings

Response for Datastore with All Paths Down (APD): Use Cluster Settings

Delay for VM failover for APD: Use Cluster Settings - minutes

Response for APD recovery after APD timeout: Use Cluster Settings

VM Monitoring: Use Cluster Settings

VM monitoring sensitivity: -

Relevant Cluster Settings

vSphere DRS	Fully Automated
vSphere HA	Expand for details

OK Cancel

Setting the automation level for individual virtual machines allows you to fine-tune automation to suit your needs. For example, you might have a virtual machine that is especially critical to your business and you would like more control over its placement. Set its automation level to manual. If a virtual machine's automation level is set to disabled, vCenter Server does not migrate that virtual machine or provide migration recommendations for it.

As a best practice, enable automation. Select the automation level based on your environment and level of comfort. For example, if you are new to vSphere DRS clusters, you might select **Partially Automated** because you want control over the placement of virtual machines. When you are comfortable with what vSphere DRS does and how it works, you might set the automation level to **Fully Automated**. Set the automation level of **Manual** on virtual machines over which you want to exercise more control, such as your business-critical virtual machines.

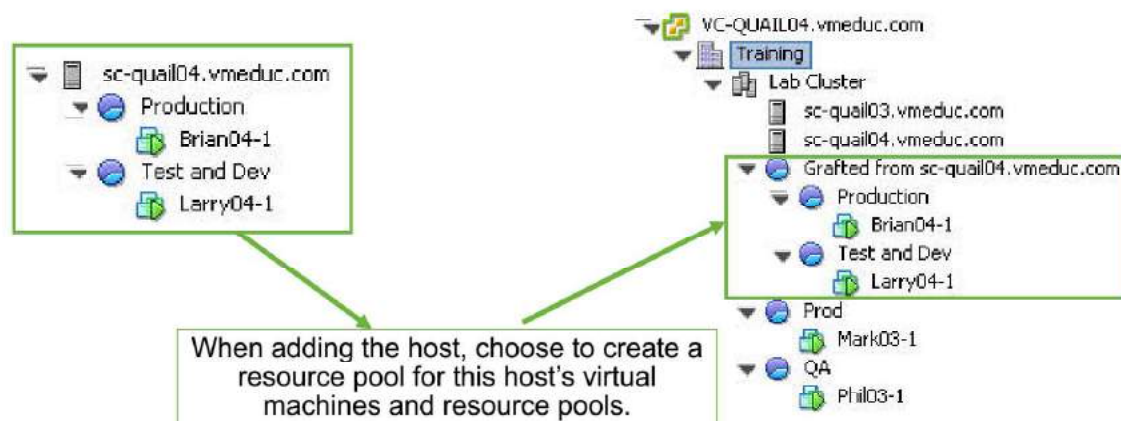
Adding a Host to a Cluster

Slide 10-15

When adding a host or moving a host into a vSphere DRS cluster, you can keep the resource pool hierarchy of the existing host:

- If vSphere DRS is not enabled, host resources pools are lost.

For example, add sc-quail04 to Lab Cluster.



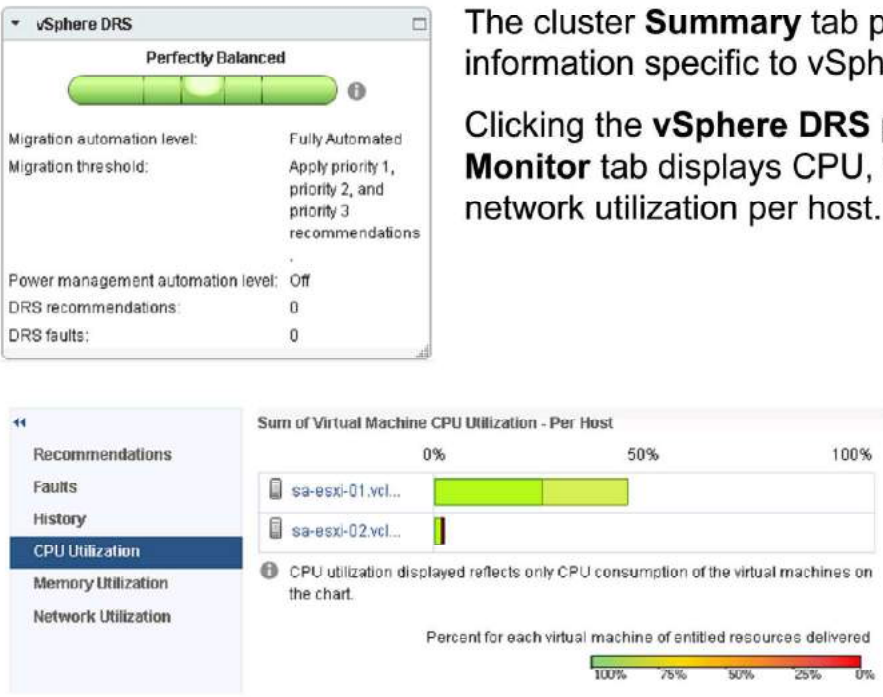
To add a host to a vSphere DRS cluster, drag an ESXi host onto the cluster object in the inventory. Use the Add Host wizard to complete the process.

When adding a host with resource pools to a vSphere DRS cluster, you must decide on the placement of the resource pools. By default, the resource pool hierarchy is discarded and the host is added at the same level as the virtual machines. However, you can choose to retain the resource pool hierarchy of the existing host by ensuring that vSphere DRS is enabled. You can also choose to graft the host's resource pools onto the cluster's resource pool hierarchy. Grafting refers to adding the branches of the host's tree to the branches of the cluster's tree, as fruit tree branches are grafted onto rootstock.

By default, the resource pool created to represent the host's resources is named as Grafted from *host_name*, but you can provide a different name to it.

Viewing vSphere DRS Cluster Information

Slide 10-16



The screenshot shows two parts of the vSphere interface. The top part is the 'vSphere DRS' Summary tab, which displays a 'Perfectly Balanced' status with a green progress bar. Below this, it lists configuration details: Migration automation level (Fully Automated), Migration threshold (Apply priority 1, priority 2, and priority 3 recommendations), Power management automation level (Off), DRS recommendations (0), and DRS faults (0). The bottom part shows the 'CPU Utilization' view, which includes a sidebar with navigation options (Recommendations, Faults, History, CPU Utilization, Memory Utilization, Network Utilization) and a chart titled 'Sum of Virtual Machine CPU Utilization - Per Host'. The chart shows two hosts: 'sa-esxi-01.vcl...' with a green bar at approximately 50% utilization and 'sa-esxi-02.vcl...' with a very thin green bar. A legend below the chart indicates that green represents 100% of entitled resources delivered, with a color gradient from green to red representing 100%, 75%, 50%, 25%, and 0%.

The cluster **Summary** tab provides information specific to vSphere DRS.

Clicking the **vSphere DRS** page on the **Monitor** tab displays CPU, memory, and network utilization per host.

The vSphere DRS pane in the cluster's **Summary** tab appears only when vSphere DRS is enabled. This section provides the following information related to vSphere DRS:

- The selected automation levels
- The number of vSphere DRS recommendations and faults
- The configured migration threshold

A priority level for each migration recommendation is computed using the load imbalance metric of the cluster. This metric is displayed as Current host load standard deviation in the cluster's **Summary** tab in vSphere Web Client. A higher load imbalance leads to higher-priority migration recommendations.

Under the **Monitor** tab, click the **vSphere DRS** page. Click the **CPU Utilization**, **Network Utilization** or the **Memory Utilization** link to view a chart of all the hosts in the cluster, and how their CPU and memory resources are allocated to each virtual machine.

For CPU use, the virtual machine information is represented by a colored box. If you point to the colored box, the virtual machine's CPU use information appears. If the virtual machine is receiving the resources that it is entitled to, the box is green. Green means that 100 percent of the virtual machine's entitled resources have been delivered to it. If the box is not green (for example, entitled

resources are 80 percent or less) for an extended time, you might want to investigate what is causing this shortfall (for example, unapplied recommendations).

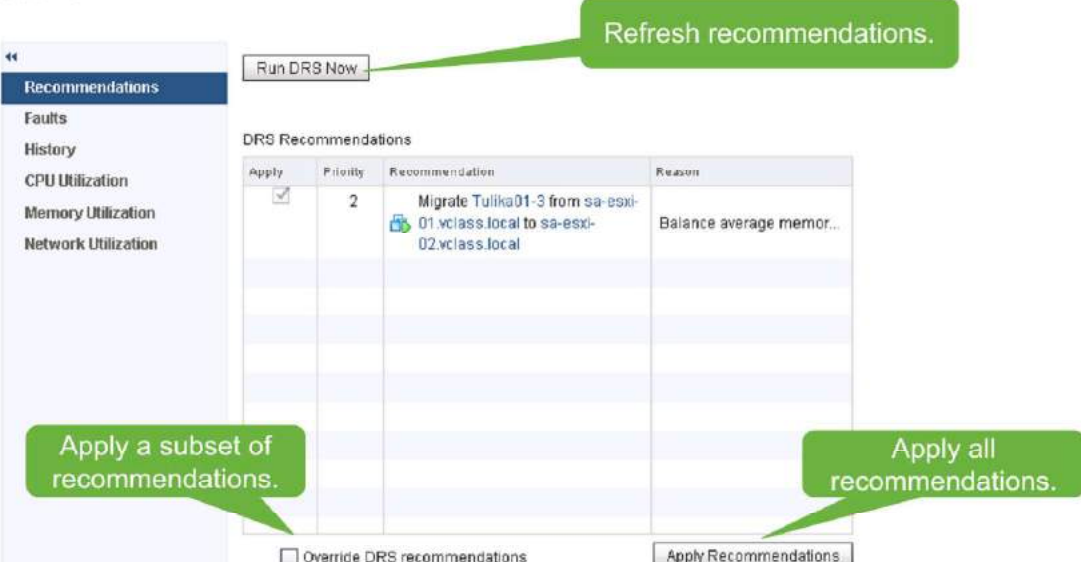
For memory usage, the virtual machine boxes are not color-coded, because the relationship between consumed memory and entitlement is often not easily categorized.

For network usage, the displayed network data reflects all traffic across physical network interfaces on the host.

Viewing vSphere DRS Recommendations

Slide 10-17

The **vSphere DRS** tab displays information about the vSphere DRS recommendations made for the cluster, the faults that occurred in applying such recommendations, and the history of vSphere DRS actions.



The screenshot shows the vSphere DRS Recommendations interface. On the left is a navigation pane with options: Recommendations (selected), Faults, History, CPU Utilization, Memory Utilization, and Network Utilization. At the top right is a 'Run DRS Now' button with a callout 'Refresh recommendations.'. Below is a table titled 'DRS Recommendations' with columns: Apply, Priority, Recommendation, and Reason. The first row has a checked 'Apply' box, a priority of '2', a recommendation to 'Migrate Tulika01-3 from sa-esxi-01.vclass.local to sa-esxi-02.vclass.local', and a reason 'Balance average memor...'. At the bottom left is an 'Override DRS recommendations' checkbox with a callout 'Apply a subset of recommendations.'. At the bottom right is an 'Apply Recommendations' button with a callout 'Apply all recommendations.'.

Apply	Priority	Recommendation	Reason
<input checked="" type="checkbox"/>	2	Migrate Tulika01-3 from sa-esxi-01.vclass.local to sa-esxi-02.vclass.local	Balance average memor...
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			

When you select the cluster object, click the **Monitor** tab and the **vSphere DRS** tab. The first three views are: Recommendations, Faults, and History.

In the Recommendations view, you can see the current set of recommendations generated for optimizing resource use in the cluster through either migrations or power management. Only manual recommendations awaiting user confirmation appear on this list.

To refresh the recommendations, click **Run DRS Now**.

To apply all recommendations, click **Apply Recommendations**.

To apply a subset of the recommendations, select the **Override DRS recommendations** check box. Select the check box next to each desired recommendation and click **Apply Recommendations**.

Monitoring Cluster Status

Slide 10-18

Select the cluster object in the inventory to view the state of the cluster and all its hosts and virtual machines.

You can view the cluster's **Tasks** and **Events** tabs for more information.

The screenshot displays the vSphere Web Client interface. The top section shows a table of events with columns for Description, Type, Date/Time, Target, and Task. Below the table, there is a detailed view of a selected event.

Description	Type	Date/Time	Target	Task
Alarm 'Virtual machine CPU usage' on Server 01 changed fr...	Information	3/18/2015 1:25...	Server 01	
DRS powered On Mail Server 01 on esxi01.vclass.local in T...	Information	3/18/2015 1:24...	Mail Server 01	
DRS powered On Domain Controller 01 on esxi01.vclass.lo...	Information	3/18/2015 1:24...	Domain Controll...	
DRS powered On Web Server 01 on esxi01.vclass.local In ...	Information	3/18/2015 1:24...	Web Server 01	
Message on Mail Server 01 on esxi01.vclass.local in Traini...	Information	3/18/2015 1:24...	Mail Server 01	
Message on Domain Controller 01 on esxi01.vclass.local L...	Information	3/18/2015 1:24...	Domain Controll...	
Message on Web Server 01 on esxi01.vclass.local in Traini...	Information	3/18/2015 1:24...	Web Server 01	
Web Server 01 on host esxi01.vclass.local in Training is sta...	Information	3/18/2015 1:24...	Web Server 01	

Date Time: 3/18/2015 1:24:55 PM
User: VSPHERE_LOCAL\Administrator
Description: 3/18/2015 1:24:55 PM DRS powered On Mail Server 01 on esxi01.vclass.local in Training
Event Type Description: A virtual machine was powered on by the user and DRS choose a host for the virtual machine based on the current cluster load distribution combined with the virtual machine's resource requirements

vSphere Web Client indicates whether a vSphere DRS cluster is valid, overcommitted (yellow), or invalid (red). vSphere DRS clusters become overcommitted or invalid for several reasons:

- A cluster might become overcommitted if a host fails.
- A cluster becomes invalid if vCenter Server is unavailable and you power on virtual machines using vSphere Web Client.
- A cluster becomes invalid if the user reduces the reservation on a parent resource pool while a virtual machine is in the process of failing over.
- If changes are made to hosts or virtual machines by using vSphere Web Client while vCenter Server is unavailable, those changes take effect. When vCenter Server becomes available again, you might find that clusters have turned red or yellow because cluster requirements are no longer met.

For more information about cluster states, see *vSphere Resource Management Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Maintenance Mode and Standby Mode

Slide 10-19

To service a host in a cluster, for example, to install more memory, or remove a host from a cluster, you must place the host in maintenance mode:

- Virtual machines on the host should be migrated to another host or shut down.
- You cannot power on virtual machines or migrate virtual machines to a host entering maintenance mode.
- While in maintenance mode, the host does not allow you to deploy or power on a virtual machine.

When a host is placed in standby mode, it is powered off:

- This mode is used by VMware vSphere® Distributed Power Management™ to optimize power usage.

You place a host in maintenance mode when you need to service it, for example, to install more memory. A host enters or leaves maintenance mode only as the result of a user request.

Virtual machines that are running on a host entering the maintenance mode must be shut down or migrated to another host, either manually or automatically by vSphere DRS. The host continues to run the Enter Maintenance Mode task until all virtual machines are powered down or moved away. You cannot power on virtual machines or migrate virtual machines to a host entering maintenance mode.

When no more running virtual machines are on the host, the host's icon indicates that it has entered maintenance mode. The host's **Summary** tab indicates the new state. While in maintenance mode, the host prevents you from deploying or powering on a virtual machine.

When a host machine is placed in standby mode, it is powered off.

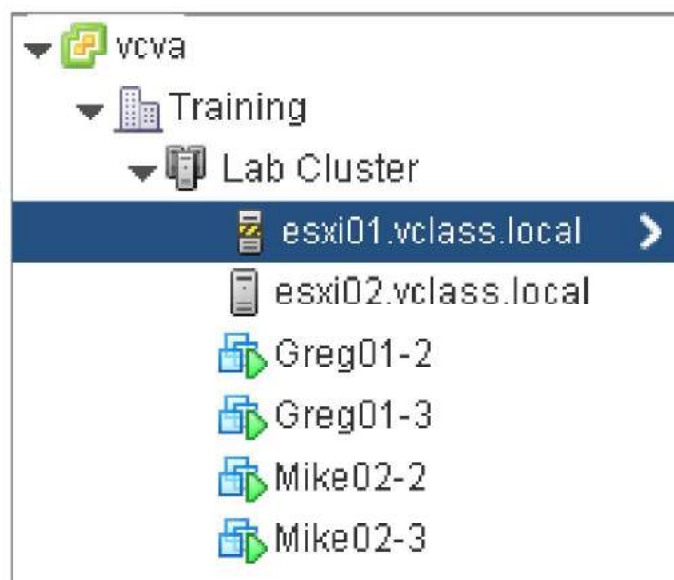
Normally, host machines are placed in the standby mode by vSphere DPM to optimize power usage. When a host is placed in the standby mode, it is powered off. You can also place a host in standby mode manually. However, when vSphere DRS runs next time, it might undo your change or recommend that you undo the changes. If you want a host to remain powered off, place it in the maintenance mode and turn it off.

Removing a Host from the vSphere DRS Cluster

Slide 10-20

Before removing a host from a vSphere DRS cluster, consider the following issues:

- The resource pool hierarchy remains with the cluster.
- Because a host must be in maintenance mode, all virtual machines running on that host are powered off.
- The resources available for the cluster decrease.



To remove a host from a cluster, right-click the host in the inventory and select **Enter Maintenance Mode**. After the host is in maintenance mode, drag it to a different inventory location, for example, the data center or another cluster.

Before you remove a host from a vSphere DRS cluster, consider the following issues:

- When you remove a host from a cluster, the host retains only the root resource pool, even if you used a vSphere DRS cluster and grafted the host resource pool when you added the host to the cluster. The hierarchy remains with the cluster. You can create a host-specific resource pool hierarchy.
- When a host is put into maintenance mode, all its running virtual machines must be shut down, suspended, or migrated to other hosts by using vSphere vMotion. Virtual machines with disks on local storage must be powered off, suspended, or migrated to another host and datastore using shared-nothing vSphere vMotion migration. When you remove the host from the cluster, the virtual machines that are currently associated with the host are also removed from the cluster.
- If you remove a host from a cluster, the resources available for the cluster decrease. If the cluster still has enough resources to satisfy the reservations of all virtual machines and resource pools in the cluster, the cluster adjusts resource allocation to reflect the reduced amount of resources. If the cluster lacks the resources to satisfy the reservations of all resource pools but has enough resources to satisfy the reservations for all virtual machines, an alarm is issued. vSphere DRS continues to run.

Disabling vSphere DRS and Restoring a Resource Pool Tree

Slide 10-21

You may turn off vSphere DRS for a cluster.

If you disable vSphere DRS, the resource pools are removed from the cluster.

When vSphere DRS is disabled, the cluster's resource pool hierarchy and affinity rules are not reestablished when DRS is turned back on.

To avoid losing the resource pools, save a snapshot of the resource pool tree on your local machine.

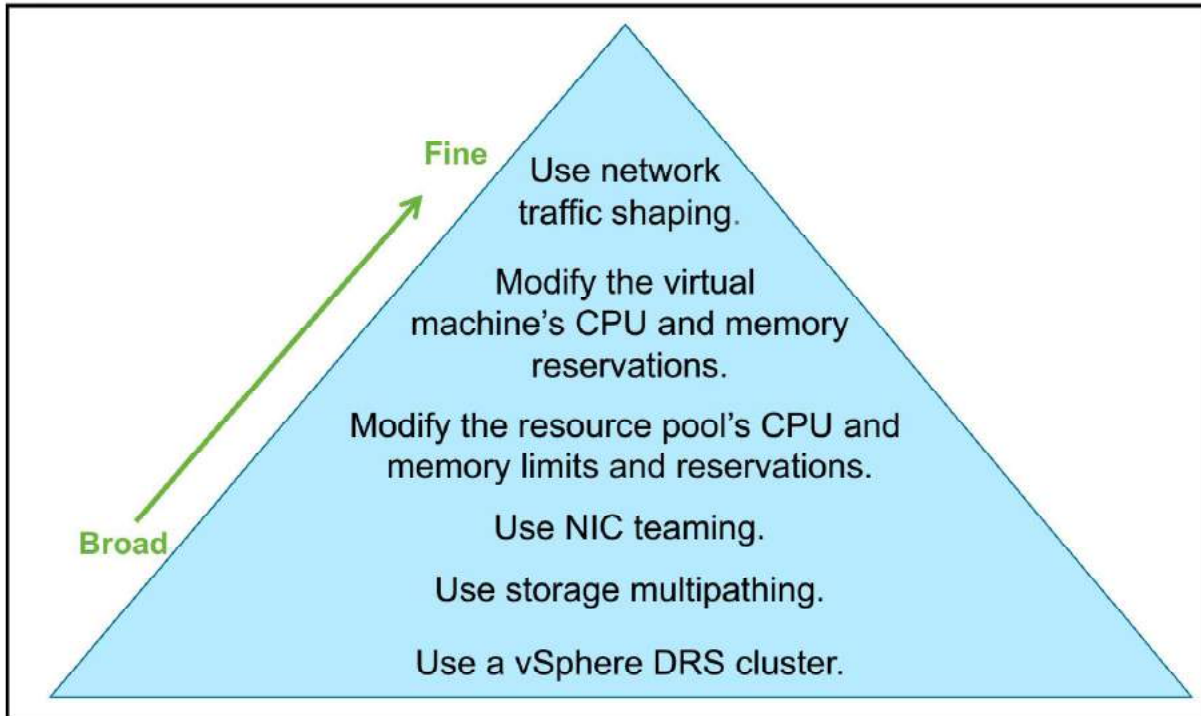
You can restore a previously saved resource pool tree snapshot when you enable vSphere DRS.

You can use the snapshot to restore the resource pool when you enable vSphere DRS. To restore a resource pool tree, you need to meet the following prerequisites:

- vSphere DRS must be turned on.
- You can restore a snapshot only on the same cluster that it was taken.
- No other resource pools are present in the cluster.

Improving Virtual Machine Performance Methods

Slide 10-22



Review the methods used to improve a virtual machine's performance. The methods are listed from specific to broad. Specific methods affect a particular virtual machine. Broad methods affect several entities. The following methods can be used to improve the performance of a virtual machine:

- If a virtual machine is network-constrained, use network traffic shaping to give a virtual machine more network bandwidth during its peak hours.
- If a virtual machine is constrained by memory, add memory shares or increase the virtual machine's memory reservation.
- If a virtual machine is constrained by CPU, add CPU shares or increase the virtual machine's CPU reservation.
- If a virtual machine is constrained by CPU or memory, increase the limits or reservations of the resource pool that the virtual machine belongs to.
- If the network load is imbalanced, you can use the network interface card (NIC) teaming to balance the network load across multiple physical network adapters.
- If the disk I/O load is imbalanced, you can use storage multipathing to balance the disk I/O load across multiple paths to a datastore.
- If the virtual machine load is imbalanced, you can place hosts in a vSphere DRS cluster to balance the virtual machine load across all hosts in the cluster.

Using vSphere HA with vSphere DRS

Slide 10-23

Reasons why vSphere HA might not be able to fail over virtual machines:

- vSphere HA admission control is disabled, and resources are insufficient in the remaining hosts to power on all of the failed VMs.
- Required VM-Host affinity rules prevent vSphere HA from failing over.
- Sufficient aggregated resources exist, but they are fragmented across hosts. In such cases, vSphere HA uses vSphere DRS to try to adjust the cluster by migrating virtual machines to defragment the resources.

When vSphere HA performs failover and restarts virtual machines on different hosts, its first priority is immediate availability of all virtual machines. After the virtual machines are restarted, those hosts in which they were powered on are usually heavily loaded, and other hosts are comparatively lightly loaded.

vSphere HA is closely integrated with vSphere DRS. When a failover occurs, vSphere HA first checks whether resources are available on that host for the failover. If resources are not available, vSphere HA asks vSphere DRS to accommodate for these where possible. Inefficient use of resources might lead to defragmented resources throughout the cluster. vSphere HA requests, but cannot guarantee defragmentation of resources to accommodate for this virtual machine's resource requirements.

In addition, vSphere DRS flattens shares and limits on virtual machines before failover. This flattening process ensures that virtual machines get their entitled resources if they would have been failed over to the correct resource pool.

Lab 20: Implementing a vSphere DRS Cluster

Slide 10-24

Implement a vSphere DRS cluster

1. Create a Load Imbalance
2. Create a vSphere DRS Cluster
3. Verify Proper vSphere DRS Cluster Functionality
4. Create, Test, and Disable a VM-VM Affinity Rule
5. Create, Test, and Disable an Anti-Affinity Rule
6. Create, Test, and Disable a VM-Host Affinity Rule

Review of Learner Objectives

Slide 10-25

You should be able to meet the following objectives:

- Describe the functions of a vSphere DRS cluster
- Create a vSphere DRS cluster
- View information about a vSphere DRS cluster
- Remove a host from a vSphere DRS cluster

Key Points

Slide 10-26


- vSphere DRS clusters provide automated resource management for multiple ESXi hosts.
- vSphere DRS works best if the virtual machines meet vSphere vMotion migration requirements.

Questions?

MODULE 11

vSphere Update Manager

Slide 11-1



You Are Here

Slide 11-2

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Creating Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere HA, vSphere Fault Tolerance, and Protecting Data
10. vSphere DRS
- 11. vSphere Update Manager**

Importance

Slide 11-3

Over time, your vSphere environment might undergo changes in its hardware or software configuration, or in the form of software updates or patches.

From a manageability and scalability perspective, you should implement changes to your vSphere environment in an orderly, controlled, and systematic fashion.

Learner Objectives

Slide 11-4

By the end of this module, you should be able to meet the following objectives:

- Describe vSphere Update Manager functionality
- List the steps to install vSphere Update Manager
- Use vSphere Update Manager to create and attach a baseline

About vSphere Update Manager

Slide 11-5

vSphere Update Manager enables centralized, automated patch and version management for ESXi hosts, virtual machine hardware, VMware Tools, and virtual appliances.

vSphere Update Manager reduces security risks:

- Reduces the number of vulnerabilities.
- Eliminates many security breaches that exploit older vulnerabilities.

vSphere Update Manager reduces the diversity of systems in an environment:

- Makes management easier.
- Reduces security risks.

vSphere Update Manager keeps machines running more smoothly:

- Patches include bug fixes.
- Makes troubleshooting easier.

You can use vSphere Update Manager with either vCenter Server that runs on Windows or with vCenter Server Appliance.

vSphere Update Manager enables centralized, automated patch and version management for vSphere and supports ESXi hosts, virtual machine hardware, VMware Tools, and virtual appliances. Updates that you specify can be applied to ESXi hosts, virtual machine hardware, and virtual appliances that you scan. With vSphere Update Manager, you can perform the following tasks:

- Scan for compliance and apply updates to virtual machine hardware, appliances, and hosts
- Directly upgrade hosts, virtual machine hardware, VMware Tools, and virtual appliances
- Apply third-party software on hosts

Keeping the patch versions up to date for virtual machine hardware and ESXi hosts helps reduce the number of vulnerabilities in an environment and the range of problems requiring solutions. All systems require ongoing patching and reconfiguration or other solutions. Reducing the diversity of systems in an environment and keeping them in compliance are security best practices. Additionally, because patches include bug fixes, vSphere Update Manager keeps environments operating properly and without service interruption or errors.

vCenter Server Appliance delivers vSphere Update Manager as an optional service. vSphere Update Manager is bundled in vCenter Server Appliance.

VMware Tools Integration with vSphere Update Manager

Slide 11-6

vSphere Update Manager relies on the host and its shared product locker designated datastore. The product locker contains three ISOs distributed with vSphere 6.5 and any downloaded ISOs from My VMware for older guest operating system versions:

- An ESXi host acquires VMware Tools through a Tools Live VMware Infrastructure Bundle (VIB).
- Using vSphere Update Manager to install or repair VMware Tools on operating systems before compatible Windows and Linux versions requires the `productLocker` directory and the appropriate VMware Tools files downloaded from My VMware.


VMware maintains its infrastructure through the use of an automated tool called vSphere Update Manager. vSphere Update Manager can be used to apply VMware patches to hosts and guest updates to VMware Tools in an automated manner. To apply patches, vSphere Update Manager uses a user-created baseline to apply to the chosen infrastructure element. These baseline updates can be applied to a single object or many objects simultaneously.

VMware Tools Product Locker

Slide 11-7

The VMware Tools product locker is distributed with the ESXi host and is copied from the host file system to a shared datastore:

- The product locker is a directory that contains two subdirectories: `/vmtools` and `/floppies`.
- The intent is to create a known location in which to store the distributed VMware Tools ISOs and thereby reduce the hosts size.
- The product locker is accessible by vSphere Update Manager.



In vSphere 6.5, VMware Tools includes a product locker for off-host storing of virtual machine VMware Tools ISO files. This new feature provides the following benefits:

- The product locker is a known location for storing ISO files.
- The product locker reduces the size of the ESXi hosts.
- The product locker is compatible with vSphere Update Manager.
- The product locker provides the ability to shut down the virtual machine.

To use the product locker

1. Copy the `/productLocker` directory from an ESXi host to a shared data store.
2. Change the ESXi host `UserVars.ProductLockerLocation` variable to point to the newly created `/productLocker` directory.

You can use the vSphere Client UI to configure the variable by clicking the **Manage** tab and navigating to **Settings > Advanced System Settings**.

3. Filter the settings listing with `UserVars` or `productLocker`.

You must restart the host for the changes to take effect.

vSphere Update Manager Capabilities

Slide 11-8

Automated patch downloading:

- Begins with information-only downloading.
- Is scheduled at regular configurable intervals.

Creation of baselines and baseline groups

Scanning:

- Inventory systems are scanned for baseline compliance.

Remediation:

- Inventory systems that are not compliant can be automatically patched.

Reduces the number of reboots required after VMware Tools updates

vSphere Update Manager uses a set of operations to ensure effective patch and upgrade management.

This process begins by downloading information about a set of security patches. One or more of these patches are aggregated to form a baseline. Multiple baselines can be added to a baseline group. You can use baseline groups to combine different types of baselines and scan and remediate an inventory object against all of them as a whole. If a baseline group contains both upgrade and patch baselines, the upgrade runs first.

A collection of virtual appliances and ESXi hosts can be scanned for compliance with a baseline or a baseline group and remediated (updated or upgraded). These processes can be started manually or through scheduled tasks.

For detailed information regarding backward compatibility, see *vSphere Update Manager 6.5 Release Notes* at <https://www.vmware.com/support/vsphere6/doc/vsphere-update-manager-65-release-notes.html>.

vSphere Update Manager Components

Slide 11-9

vSphere Update Manager includes several components and requires network connectivity with vCenter Server:

- vSphere Update Manager server component:
 - Windows vCenter Server:
 - Install on the same computer as Windows vCenter Server or on a different computer.
 - vCenter Server Appliance:
 - vSphere Update Manager 6.5 is integrated with vCenter Server Appliance and is delivered as an optional service.
 - Starting in vSphere 6.5, you can no longer connect the vSphere Update Manager instance that is installed on a Windows Server machine with vCenter Server Appliance.
- Client component:
 - The Update Manager Web Client appears as an **Update Manager** tab under the **Configure** tab in vSphere Web Client.
- Database:
 - Used to store and organize server data.
 - Oracle or Microsoft SQL Server database.

vSphere Update Manager includes the following components:

- vSphere Update Manager server component:
 - You can install the vSphere Update Manager server component either on the same Windows server where vCenter Server is installed or on a separate machine. To install vSphere Update Manager, you must have Windows administrator credentials for the computer on which you install vSphere Update Manager.
 - You can deploy vSphere Update Manager in a secured network without Internet access. You can use the vSphere Update Manager download service to download update metadata and update binaries.
- vSphere Update Manager client component:
 - The vSphere Update Manager client component is a plug-in that runs on vSphere Web Client. This plug-in is automatically enabled after installation of the vSphere Update Manager server component on Windows, and after deployment of vCenter Server Appliance.
 - The vSphere Update Manager Web Client plug-in appears as an **Update Manager** tab under the **Configure** tab in vSphere Web Client.

- Database component:
 - The vSphere Update Manager server and vSphere Update Manager download service require a database to store and organize server data.
 - vSphere Update Manager supports Oracle and Microsoft SQL Server databases.

Configuring vSphere Update Manager Settings

Slide 11-10

You can modify the vSphere Update Manager configuration only if you have the correct privileges:

- Network Connectivity Settings
- Download Settings
- Proxy Settings
- Checking for Updates (Download Schedule) Settings
- Notification Check Schedule Settings
- Virtual Machine Settings
- Host and Cluster Settings

You can modify the vSphere Update Manager settings only if you have the privileges to configure the vSphere Update Manager settings and service. These permissions must be assigned on the vCenter Server system with which vSphere Update Manager is registered.

Connect the vSphere Update Manager Client to a vCenter Server system with which vSphere Update Manager is registered. On the Home page, click the **Update Manager** icon to access administrative settings:

- **Network Connectivity:** The network ports are configured during installation. You can modify the IP address or host name for the patch store in the vSphere Update Manager network connectivity settings.
- **Download Source:** If your deployment system is connected to the Internet, you can directly download ESXi patches and extensions, as well as virtual appliance upgrades.
- **Proxy Settings:** You can configure vSphere Update Manager to download updates from the Internet by using a proxy server.
- **Checking for Updates (Download Schedule):** vSphere Update Manager checks for virtual appliance upgrades, host patches, and extensions at regular intervals. Generally, the default schedule settings are sufficient, but you can change the schedule if your environment requires more or less frequent checks.

- **Notification Check Schedule:** By default vSphere Update Manager checks for notifications about patch recalls, patch fixes, and alerts at certain time intervals. You can modify this schedule. By default the task to check for notifications and to send notifications alerts is enabled and is called the VMware vSphere Update Manager Check Notification task. By modifying this task, you can configure the time and frequency at which vSphere Update Manager checks for patch recalls or for the release of patch fixes, and sends notifications to the email addresses you specify.
- **Virtual Machine Settings:** By default, vSphere Update Manager is configured to take snapshots of virtual machines before applying updates. If the remediation fails, you can use the snapshot to return the virtual machine to the state before the remediation. vSphere Update Manager does not take snapshots of fault tolerant virtual machines and virtual machines that are running virtual machine hardware version 3. If you decide to take snapshots of such virtual machines, the remediation might fail. You can choose to keep snapshots indefinitely or for a fixed period of time.
- **Host and Cluster Settings:** You determine how you want vSphere Update Manager to behave with hosts and clusters.

For more information, see *Installing and Administering VMware vSphere Update Manager* at https://www.vmware.com/support/pubs/vum_pubs.html. You can also see *vSphere Upgrade Guide* at <https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>.

Baseline and Baseline Groups

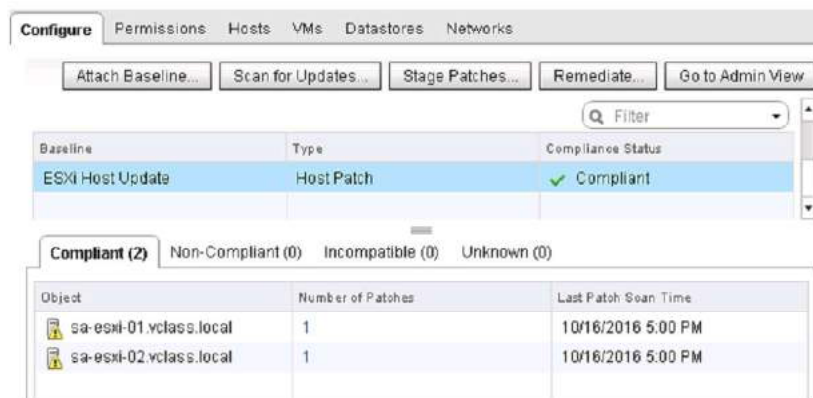
Slide 11-11

A baseline includes one or more patches, extensions, or upgrades:

- vSphere Update Manager includes two default dynamic patch baselines and three upgrade baselines.

A baseline group includes multiple baselines:

- Can contain one upgrade baseline and one or more patch and extension baselines.



Baselines can be upgrade, extension, or patch baselines. Baselines contain a collection of one or more patches, extensions, or upgrades.

Baseline groups are assembled from existing baselines, and might contain one upgrade baseline and one or more patch and extension baselines, or might contain a combination of multiple patch and extension baselines. When you scan hosts, virtual machines, and virtual appliances, you evaluate them against baselines and baseline groups to determine their level of compliance.

To create, edit, or delete baselines and baseline groups, you must have the Manage Baseline privilege. To attach baselines and baseline groups, you must have the Attach Baseline privilege. Privileges must be assigned on the vCenter Server system with which vSphere Update Manager is registered.

vSphere Update Manager includes the following default dynamic patch and upgrade baselines:

- **Critical Host Patches (Predefined):** Checks ESXi hosts for compliance with all critical patches.
- **Non-Critical Host Patches (Predefined):** Checks ESXi hosts for compliance with all optional patches.

- **VMware Tools Upgrade to Match Host (Predefined):** Checks virtual machines for compliance with the latest VMware Tools version on the host. vSphere Update Manager supports upgrading of VMware Tools for virtual machines on hosts that are running ESXi 5.0 and later.
- **VM Hardware Upgrade to Match Host (Predefined):** Checks the virtual hardware of a virtual machine for compliance with the latest version supported by the host. vSphere Update Manager supports upgrading to virtual hardware version vmx-11 on hosts that are running ESXi 6.
- **VA Upgrade to Latest (Predefined):** Checks virtual appliance compliance with the latest released virtual appliance version.

In vSphere Web Client, default baselines are displayed on the **Baselines and Groups** tab of the vSphere Update Manager server Administration view.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain and you have a vSphere Update Manager instance for each vCenter Server system in the group, the baselines and baseline groups you create and manage are applicable only to inventory objects managed by the vCenter Server system with which the selected vSphere Update Manager instance is registered. You can use a vSphere Update Manager instance only with a vCenter Server system on which the instance is registered.

Creating and Editing Patch or Extension Baselines

Slide 11-12

You can create custom patch, extension, and upgrade baselines to meet the needs of your specific deployment by using the New Baseline wizard:

- Create a fixed patch baseline:
 - Fixed baselines consist of a set of patches that do not change as patch availability changes.
- Create a dynamic patch baseline:
 - Dynamic baselines consist of a set of patches that meet certain criteria.
- Create a host extension baseline:
 - Extension baselines contain additional software for ESXi hosts. This additional software might be VMware software or third-party software.
- Filter patches or extensions in the New Baseline wizard:
 - When you create a patch or extension baseline, you can filter the patches and extensions available in the vSphere Update Manager repository to find specific patches and extensions to exclude or include in the baseline.

You create and manage baselines in the **Update Manager** tab. You can remediate hosts against baselines that contain patches or extensions. Depending on the patch criteria you select, patch baselines can be either dynamic or fixed.

Dynamic patch baselines contain a set of patches, which updates automatically according to patch availability and the criteria that you specify. Fixed baselines contain only patches that you select, regardless of new patch downloads.

Extension baselines contain additional software modules for ESXi hosts. This additional software might be VMware software or third-party software. You can install additional modules by using extension baselines, and update the installed modules by using patch baselines.

If your vCenter Server system is part of a connected group in vCenter Linked Mode, and you have more than one vSphere Update Manager instance, patch and extension baselines that you create are not applicable to all inventory objects managed by other vCenter Server systems in the group. Baselines are specific for the vSphere Update Manager instance you select.

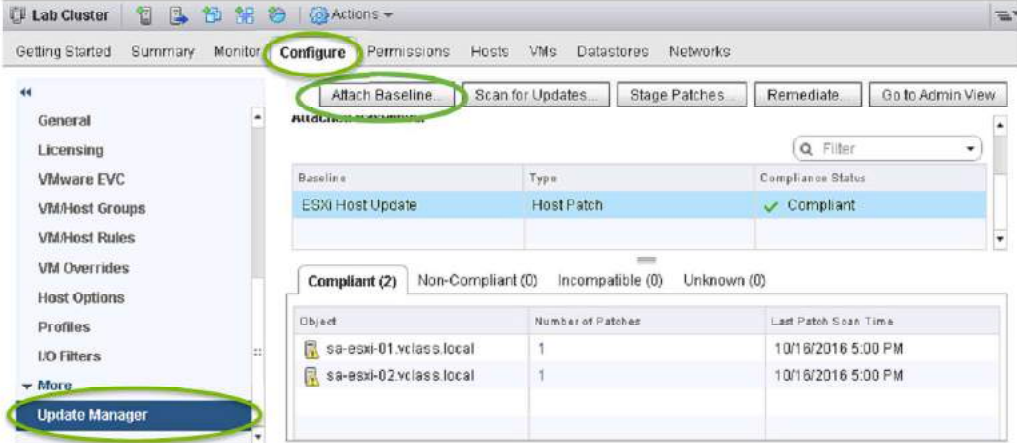
Attaching a Baseline

Slide 11-13

To view compliance information and scan objects in the inventory against baselines or baseline groups, you must first attach baselines or baseline groups to these objects.

You can attach baselines or baseline groups to objects from the **Update Manager** tab in vSphere Web Client:

- Click **Configure > Update Manager > Attach Baseline**.



The screenshot shows the vSphere Web Client interface for a 'Lab Cluster'. The 'Configure' tab is selected, and the 'Update Manager' sub-tab is active in the left navigation pane. The main area displays the 'Attach Baseline' button, which is circled in green. Below it, a table shows the attached baseline 'ESXi Host Update' with a 'Compliant' status. A summary bar indicates 'Compliant (2)', 'Non-Compliant (0)', 'Incompatible (0)', and 'Unknown (0)'. A table below lists the objects and their patch scan results.

Baseline	Type	Compliance Status
ESXi Host Update	Host Patch	✓ Compliant

Object	Number of Patches	Last Patch Scan Time
sa-esxi-01.vclass.local	1	10/16/2016 5:00 PM
sa-esxi-02.vclass.local	1	10/16/2016 5:00 PM

To view compliance information and remediate objects in the inventory against specific baselines and baseline groups, you must first attach baselines and baseline groups to these objects.

To use baselines and baseline groups, you must attach baselines and baseline groups to selected inventory objects, such as virtual machines, virtual appliances, hosts, or container objects.

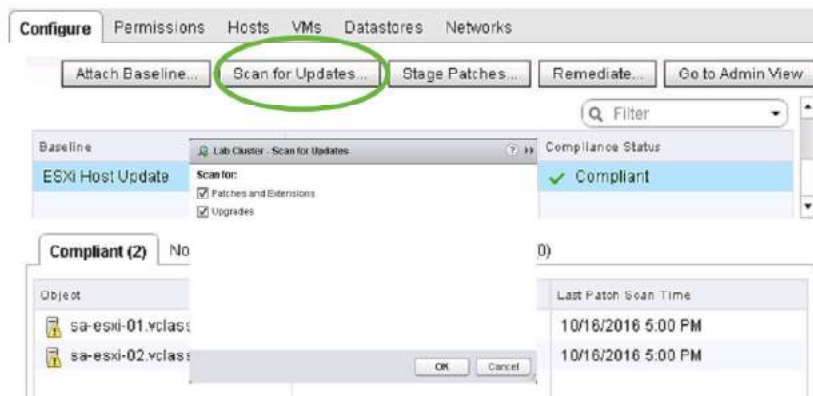
Although you can attach baselines and baseline groups to individual objects, a more efficient method is to attach them to container objects, such as folders, vApps, clusters, and data centers. Individual vSphere objects inherit baselines attached to the parent container object. Removing an object from a container removes the inherited baselines from the object.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, you can attach baselines and baseline groups to objects managed by the vCenter Server system with which vSphere Update Manager is registered. Baselines and baseline groups that you attach are specific to the vSphere Update Manager instance that is registered with the vCenter Server system.

Scanning for Updates

Slide 11-14

Scanning evaluates the inventory object against the baseline or baseline group.



Scanning is the process in which attributes of a set of hosts, virtual machines, or virtual appliances are evaluated against patches, extensions, and upgrades in the attached baselines and baseline groups. You can configure vSphere Update Manager to scan virtual machines, virtual appliances, and ESXi hosts against baselines and baseline groups by scheduling or manually initiating scans to generate compliance information.

If the object that you select is a container object, all child objects are also scanned. The larger the virtual infrastructure and the higher up in the object hierarchy that you begin the scan, the longer the scan takes.

After you have an inventory object attached to a baseline, perform a scan by right-clicking the object and selecting **Scan for updates**.

Viewing Compliance for vSphere Objects

Slide 11-15

You can review compliance information for the virtual machines, virtual appliances, and hosts against baselines and baseline groups that you attach.

Baseline	Type	Compliance Status
ESXi Host Update	Host Patch	✓ Compliant

Compliant (2) Non-Compliant (0) Incompatible (0) Unknown (0)

Object	Number of Patches	Last Patch Scan Time
sa-esxi-01.vclass.local	1	10/16/2016 5:00 PM
sa-esxi-02.vclass.local	1	10/16/2016 5:00 PM

When you select a container object, you view the overall compliance status of the attached baselines, as well as all the individual compliance statuses. If you select an individual baseline attached to the container object, you see the compliance status of the baseline.

If you select an individual virtual machine, appliance, or host, you see the overall compliance status of the selected object against all attached baselines and the number of updates. If you further select an individual baseline attached to this object, you see the number of updates grouped by the compliance status for that baseline.

Compliant state indicates that a vSphere object is compliant with all baselines in an attached baseline group or with all patches, extensions, and upgrades in an attached baseline. Compliant state requires no further action. If a baseline contains patches or upgrades that are not relevant to the target object, the individual updates, and baselines or baseline groups that contain them, are treated as not applicable, and represented as compliant.

Non-compliant state indicates that one or more baselines in a baseline group, or one or more patches, extensions, or upgrades in a baseline are applicable to the target object, but are not installed (missing) on the target. You must remediate the target object to make it compliant.

Remediating Objects

Slide 11-16

You can remediate virtual machines, templates, virtual appliances, and hosts:

- You can perform the remediation immediately or schedule it for a later date.
- Host remediation runs in different ways, depending on the types of baselines that you attach and whether the host is in a cluster or not.
- For ESXi hosts in a cluster, the remediation process is sequential by default.
- Remediation of hosts in a cluster requires that you temporarily disable cluster features such as vSphere DPM and vSphere HA admission control.

You can remediate virtual machines, virtual appliances, and hosts using either user-initiated remediation or scheduled remediation at a time that is convenient for you.

You can remediate virtual machines and appliances together.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, you can remediate only the inventory objects managed by the vCenter Server system with which vSphere Update Manager is registered.

Host remediation runs in different ways depending on the types of baselines you attach and whether the host is in a cluster or not.

For ESXi hosts in a cluster, the remediation process is sequential by default. With vSphere Update Manager 6 and later, you can select to run host remediation in parallel.

When you remediate a cluster of hosts sequentially and one of the hosts fails to enter maintenance mode, vSphere Update Manager reports an error, and the process stops and fails. The hosts in the cluster that are remediated stay at the updated level. The ones that are not remediated after the failed host remediation are not updated. If a host in a cluster enabled for vSphere DRS runs a virtual machine on which vSphere Update Manager or vCenter Server are installed, vSphere DRS first attempts to migrate the virtual machine running vCenter Server or vSphere Update Manager to

another host, so that the remediation succeeds. In case the virtual machine cannot be migrated to another host, the remediation fails for the host, but the process does not stop. vSphere Update Manager proceeds to remediate the next host in the cluster.

The host upgrade remediation of ESXi hosts in a cluster proceeds only if all hosts in the cluster can be upgraded.

Remediation of hosts in a cluster requires that you temporarily disable cluster features, such as vSphere DPM and vSphere HA admission control. You should also turn off vSphere Fault Tolerance if it is enabled on any of the virtual machines on a host, and disconnect the removable devices connected to the virtual machines on a host, so that they can be migrated with vSphere vMotion. Before you start a remediation process, you can generate a report that shows which cluster, host, or virtual machine has the cluster features enabled.

When you remediate a cluster of hosts in parallel, vSphere Update Manager remediates multiple hosts concurrently.

During parallel remediation, if vSphere Update Manager encounters an error when remediating a host, it ignores the host and the remediation process continues for the other hosts in the cluster. vSphere Update Manager continuously evaluates the maximum number of hosts it can remediate concurrently without disrupting vSphere DRS settings.

You can limit the number of concurrently remediated hosts to a specific number.

vSphere Update Manager remediates hosts that are part of a vSAN cluster sequentially even if you select the option to remediate them in parallel. By design, only one host from a vSAN cluster can be in a maintenance mode at any time.

Patch Recall Notification

Slide 11-17

At regular intervals, vSphere Update Manager contacts VMware to download notifications about patch recalls, new fixes, and alerts:

- **Notification Check Schedule** is selected by default.

On receiving patch recall notifications, vSphere Update Manager takes the following actions:

- Generates a notification in the notification tab
- No longer applies the recalled patch to any host:
 - Patch is flagged as recalled in the database.
- Deletes the patch binaries from its patch repository

vSphere Update Manager does not uninstall recalled patches from ESXi hosts. It waits for a newer patch and applies that patch to make a host compliant.

At regular intervals, vSphere Update Manager contacts VMware to download information (notifications) about patch recalls, new fixes, and alerts. You can change the schedule by modifying the **Notification Check Schedule** setting in the vSphere Update Manager **Configuration** tab.

When patches with problems or potential problems are released, these patches are recalled in the metadata, and vSphere Update Manager marks them as recalled. If you try to install a recalled patch, vSphere Update Manager notifies you that the patch is recalled and does not install it on the host. If you have already installed such a patch, vSphere Update Manager notifies you that the recalled patch is installed on certain hosts. vSphere Update Manager also deletes all the recalled patches from the vSphere Update Manager patch repository.

When a new patch is released, vSphere Update Manager downloads it and prompts you to install it to fix the problems that the recalled patch might cause. If you try to install the recalled patch, vSphere Update Manager alerts you that the patch is recalled and that you must install a fix.

Lab 21: Using vSphere Update Manager

Slide 11-18

Install, configure, and use vSphere Update Manager

1. Modify the Cluster Settings
2. Configure vSphere Update Manager
3. Create a Patch Baseline
4. Attach a Baseline and Scan for Updates
5. Stage the Patches onto the ESXi Hosts
6. Remediate the ESXi Hosts

Review of Learner Objectives

Slide 11-19

You should be able to meet the following objectives:

- Describe vSphere Update Manager functionality
- List the steps to install vSphere Update Manager
- Use vSphere Update Manager to create and attach a baseline

Key Points

Slide 11-20

- vSphere Update Manager reduces security vulnerabilities by keeping systems up to date and by reducing the diversity of systems in an environment.

Questions?