

VMware vSphere: What's New

Lecture Manual
ESXi 7 and vCenter Server 7



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

<https://t.me/learningnets>

VMware vSphere: What's New [V6.7 to V7]

Lecture Manual

ESXi 7 and vCenter Server 7

Part Number EDU-EN-VSWN7-LECT (04/2020)

Copyright © 2020 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. VMware, VMware Remote Console™, VMware Site Recovery Manager™, VMware Site Recovery™, VMware Skyline Advisor™, VMware Skyline Collector™, VMware Skyline™, VMware Tanzu™, VMware Tanzu™ Kubernetes Grid™, VMware Tanzu™ Kubernetes Grid™ service, VMware Tanzu™ Kubernetes Grid™ service for vSphere®, VMware Tools™, VMware Validated Design™, VMware Validated Design™ for DevOps-Ready IT, VMware Validated Design™ for Horizon®, VMware Validated Design™ for Intelligent Operations™, VMware Validated Design™ for IT Automating IT, VMware Validated Design™ for Management and Workload Consolidation™, VMware Validated Design™ for Micro-Segmentation, VMware Validated Design™ for Remote Office Branch Office, VMware Validated Design™ for SAP, VMware vCenter Server®, VMware vCenter Server® High Availability, VMware vCenter® Lifecycle Manager™, VMware vCenter® Log Insight™, VMware vCenter® Server Appliance™, VMware vCenter® Single Sign-On, VMware Verify™, VMware View®, VMware vRealize®, VMware vRealize® Operations Manager™, VMware vRealize® Operations Manager™ for Horizon®, VMware vRealize® Operations™, VMware vRealize® Operations™ Advanced, VMware vRealize® Operations™ Enterprise, VMware vRealize® Operations™ Standard, VMware vRealize® Suite Lifecycle Manager™, VMware vSAN™, VMware vSphere®, VMware vSphere® Auto Deploy™, VMware vSphere® Client™, VMware vSphere® DirectPath I/O™, VMware vSphere® Distributed Resource Scheduler™, VMware vSphere® High Availability, VMware vSphere® Trust Authority™, VMware vSphere® Update Manager™, VMware vSphere® Virtual Volumes™, VMware vSphere® VMFS, VMware vSphere® vMotion®, and VMware vSphere® with Kubernetes are registered trademarks or trademarks 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.

www.vmware.com/education

<https://t.me/learningnets>

Contents

Module 1	Course Introduction	1
1-2	Course Introduction.....	1
1-3	Importance	1
1-4	Learner Objectives (1).....	2
1-5	Learner Objectives (2)	2
1-6	vSphere 7 Overview	3
1-7	Course Outline	4
1-8	Typographical Conventions	5
1-9	References (1).....	6
1-10	References (2).....	7
1-11	VMware Online Resources.....	8
1-12	VMware Education Overview	9
1-13	VMware Certification Overview	10
1-14	VMware Badge Overview	11
Module 2	Installation and Upgrade	13
2-2	Importance	13
2-3	Module Lessons.....	14
2-4	Lesson 1: Upgrading and Migrating vCenter Server.....	15
2-5	Learner Objectives	15
2-6	Deployment Models.....	16
2-7	vCenter Server 7.0 Deployment Enhancements (1)	17
2-8	vCenter Server 7.0 Deployment Enhancements (2)	18

2-9	Enhanced Linked Mode with an Embedded vCenter Server	19
2-10	Benefits of vCenter Server 7.0 Deployment Model.....	20
2-11	vSphere Version Upgrade Compatibility	22
2-12	vCenter Server Upgrade Process	23
2-13	Upgrade Paths to vCenter Server 7.0	24
2-14	Migration Paths to vCenter Server 7.0	25
2-15	Upgrade or Migration to vCenter Server 7.0 with Shared Platform Services Controller	26
2-16	Enhancements to vCenter Server Appliance and CLI Installers	28
2-17	Resource Requirements for Deploying vCenter Server Appliance (1)	29
2-18	Resource Requirements for Deploying vCenter Server Appliance (2)	30
2-19	vCenter Server Upgrade and Migration Prerequisites	31
2-20	Upgrade Sequence for vSphere Components	32
2-21	Direct Upgrade or Migration to vCenter Server Appliance 7.0.....	33
2-22	Converge Upgrade or Migration to vCenter Server Appliance 7.0	34
2-23	Demonstration: Upgrading vCenter Server Appliance 6.7 to 7.0	35
2-24	Review of Learner Objectives.....	36
2-25	Lesson 2: vCenter Server Virtual Appliance Enhancements	37
2-26	Learner Objectives	37
2-27	vCenter Server Appliance Architecture	38
2-28	vCenter Server 7.0 Scalability	39
2-29	vCenter Server Appliance Management Interface	40
2-30	File-Based Backups of vCenter Server Appliance	41
2-31	Backup and Restore Enhancements	42
2-32	Scheduling File-Based Backups	43
2-33	Restoring File-Based Backups of vCenter Server Appliance	44
2-34	Review of Learner Objectives.....	46
2-35	Lesson 3: Upgrading ESXi	47
2-36	Learner Objectives	47
2-37	ESXi Host Upgrade Process	48
2-38	Hardware and System Requirements for ESXi Upgrade or Installation	49
2-39	Storage Requirements for ESXi 7.0 Installation or Upgrade	50
2-40	Direct Upgrade to ESXi 7.0	51

2-41	vSphere 7 Core Licensing	52
2-42	Review of Learner Objectives	53
2-43	Key Points	54
Module 3	Management Enhancements	55
3-2	Importance	55
3-3	Module Lessons	56
3-4	Lesson 1: vSphere Client and vCenter Server Appliance Enhancements	57
3-5	Learner Objectives	57
3-6	vCenter Server Client	58
3-7	About vCenter Server Profiles	59
3-8	vCenter Server Profiles Management Benefits	60
3-9	vCenter Server Profiles Management APIs	61
3-10	vCenter Server Profiles Operations	62
3-11	About the DDNS Client	63
3-12	DDNS Client Use Cases	64
3-13	About Multihoming	65
3-14	About PNID	66
3-15	PNID Change Use Cases	67
3-16	Changing the PNID	68
3-17	PNID Change Considerations (1)	70
3-18	PNID Change Considerations (2)	71
3-19	Prerequisites for Changing the PNID	72
3-20	MAC Address Allocation in vSphere	73
3-21	VMware OUI MAC Allocation	74
3-22	Prefix-Based MAC Address Allocation	75
3-23	Range-Based MAC Address Allocation	76
3-24	Selecting Prefix-Based MAC Address Allocation	77
3-25	vSphere Client Developer Center	78
3-26	Code Capture Enhancements (1)	79
3-27	Code Capture Enhancements (2)	80
3-28	vSphere Client Networking Enhancements	81
3-29	Overview of NSX-T Data Center with vCenter Server	82

3-30	NSX-T Data Center Architecture	83
3-31	Identifying NSX Port Groups	84
3-32	Lab 1: Reconfiguring the Primary Network Identifier	86
3-33	Review of Learner Objectives	87
3-34	Lesson 2: Deploying and Managing Templates in the Content Library	88
3-35	Learner Objectives	88
3-36	VM Template In-Place Update Use Cases	89
3-37	VM Template In-Place Update	90
3-38	Content Library Interface	91
3-39	Registered VM Template (1)	92
3-40	Registered VM Template (2)	93
3-41	Template Versioning	94
3-42	Discarding Versions	95
3-43	Registered VM Template with Updates	96
3-44	Unregistered VM Template	97
3-45	Lab 2: Deploying and Managing Templates in the Content Library	98
3-46	Review of Learner Objectives	99
3-47	Lesson 3: vCenter Server Identity Federation	100
3-48	Learner Objectives	100
3-49	ADFS as an Identity Provider	101
3-50	About vCenter Server IDP Federation for ADFS	102
3-51	vCenter Server IDP Federation Use Cases	103
3-52	vCenter Server IDP Federation Components	104
3-53	ADFS Application Groups	105
3-54	Prerequisites for vCenter Server IDP Federation	106
3-55	Configuring vCenter Server IDP Federation in vSphere (1)	107
3-56	Configuring vCenter Server IDP Federation in vSphere (2)	108
3-57	Configuring vCenter Server IDP Federation in vSphere (3)	109
3-58	Configuring vCenter Server IDP Federation in vSphere (4)	110
3-59	Enabling ADFS on an Existing Enhanced Linked Mode Configuration	111
3-60	Login Workflow of the vCenter Server IDP Federation for ADFS	112
3-61	Lab 3: Configuring Identity Federation to Use Windows ADFS	113
3-62	Review of Learner Objectives	114

3-63	Lesson 4: vCenter Certificate Enhancements	115
3-64	Learner Objectives	115
3-65	vSphere Security Certificates	116
3-66	Use Cases for Certificate Management	117
3-67	Managing Certificates (1)	118
3-68	Managing Certificates (2)	119
3-69	Managing Certificates (3)	120
3-70	Viewing Certificate Details	121
3-71	Renewing Certificates	122
3-72	Importing and Replacing (1)	123
3-73	Importing and Replacing (2)	124
3-74	Importing and Replacing (3)	125
3-75	Generating a Certificate Signing Request	126
3-76	Lab 4: Managing vCenter Server Certificates	127
3-77	Review of Learner Objectives	128
3-78	Lesson 5: VMware Skyline Overview	129
3-79	Learner Objectives	129
3-80	About VMware Skyline	130
3-81	VMware Skyline Family	131
3-82	Skyline Health for vSphere (1)	133
3-83	Skyline Health for vSphere (2)	134
3-84	Skyline Health for vSAN	135
3-85	Skyline Components: Collector and Advisor	136
3-86	VMware Skyline Networking Design	137
3-87	How VMware Skyline Works	138
3-88	Skyline Advisor	139
3-89	Skyline Advisor: Proactive Findings	140
3-90	Log Assist	141
3-91	Skyline Features Comparison by Support Level	142
3-92	Skyline Privacy and Security Policy	143
3-93	Review of Learner Objectives	144
3-94	Key Points	145

Module 4	ESXi Enhancements	147
4-2	Importance	147
4-3	Module Lessons	148
4-4	Lesson 1: ESXi Security	149
4-5	Learner Objectives	149
4-6	CPU Vulnerability Background	150
4-7	CPU Vulnerability Mitigation (1)	151
4-8	CPU Vulnerability Mitigation (2)	152
4-9	vSphere vMotion Encryption	153
4-10	Internet Protocol Security Policy Updates	154
4-11	Review of Learner Objectives	155
4-12	Lesson 2: Hardware Support in ESXi	156
4-13	Learner Objectives	156
4-14	Support for AMD EPYC Processors	157
4-15	Support for Multiple NVIDIA vGPUs	158
4-16	New Device Support	159
4-17	Enhanced Support for vSphere DirectPath I/O Network Devices	160
4-18	Driver Updates (1)	162
4-19	Driver Updates (2)	163
4-20	Receive Side Scaling	164
4-21	Multi-RSS Support for nvmxnet3 Driver	165
4-22	Data Center Bridging	166
4-23	ESXi Native Data Center Bridging	167
4-24	Review of Learner Objectives	168
4-25	Lesson 3: Other ESXi Enhancements	169
4-26	Learner Objectives	169
4-27	ESXi System Storage	170
4-28	Recommendations and Requirements for ESXi System Storage	171
4-29	Precision Time Protocol Support	172
4-30	Configuring Precision Time Protocol	173
4-31	Review of Learner Objectives	174
4-32	Key Points	175

Module 5	Virtual Machine Enhancements	177
5-2	Importance	177
5-3	Learner Objectives	177
5-4	VM Compatibility Version	178
5-5	About Virtual Watchdog Timer	179
5-6	Virtual Watchdog Timer Use Cases	180
5-7	Virtual Watchdog Timer Requirements	181
5-8	Adding Virtual Watchdog Timer to a VM	182
5-9	Virtual Watchdog Timer Status	183
5-10	About Precision Clock	184
5-11	Precision Clock Use Cases	185
5-12	Precision Clock Requirements	186
5-13	Adding and Configuring Precision Clock	187
5-14	About SGX	188
5-15	SGX Terminology	189
5-16	Virtual SGX Use Cases	190
5-17	Intel SGX Architecture	191
5-18	Virtual SGX Requirements	192
5-19	Virtual SGX Restrictions	193
5-20	Enabling and Disabling Virtual SGX for a VM	194
5-21	Guest Customization Improvements	195
5-22	About Instant Clone Guest Customization	196
5-23	Instant Clone Guest Customization Use Cases	197
5-24	Instant Clone Guest Customization Requirements	198
5-25	Latency Sensitivity Advantages	199
5-26	Selective Latency Sensitivity	200
5-27	Selective Latency Sensitivity Use Cases	201
5-28	Selective Latency Sensitivity Requirements	202
5-29	Selective Latency Sensitivity Restrictions	203
5-30	Migration of VMs with Remote Devices	204
5-31	Virtual Machine Maximums for Large Workloads	205
5-32	Lab 5: Using vSphere vMotion to Migrate a VM with Attached Devices	206
5-33	Review of Learner Objectives	207

5-34	Key Points	208
Module 6	Storage Enhancements	209
6-2	Importance	209
6-3	Module Lessons	210
6-4	Lesson 1: Enhancements to NVMe Storage Systems	211
6-5	Learner Objectives	211
6-6	Enhanced Support for NVMe	212
6-7	Supported NVMe Topologies	213
6-8	FC-NVMe and NVMe over RDMA (RoCE v2)	214
6-9	Support for FC-NVMe	215
6-10	Requirements for FC-NVMe	216
6-11	Support for NVMe over RDMA (RoCE v2)	217
6-12	Requirements for NVMe over RDMA (RoCE v2)	218
6-13	Configuring ESXi for NVMe over RDMA (RoCE v2) (1)	219
6-14	Configuring ESXi for NVMe over RDMA (RoCE v2) (2)	220
6-15	Configuring ESXi for NVMe over RDMA (RoCE v2) (3)	221
6-16	Configuring ESXi for NVMe over RDMA (RoCE v2) (4)	222
6-17	Pluggable Storage Architecture (1)	223
6-18	Pluggable Storage Architecture (2)	225
6-19	High-Performance Plug-In for NVMe	226
6-20	Support for High-Performance Plug-In Multipathing	227
6-21	High-Performance Plug-In Best Practices	228
6-22	Review of Learner Objectives	229
6-23	Lesson 2: Support for iSER	230
6-24	Learner Objectives	230
6-25	Support for iSER-Enabled vSphere Virtual Volumes Arrays	231
6-26	Configuring iSER	232
6-27	Requirements for iSER	233
6-28	Review of Learner Objectives	234
6-29	Lesson 3: vSAN Enhancements	235
6-30	Learner Objectives	235
6-31	Enterprise Storage in Native vSphere Architecture	236

6-32	Unified HTML5 User Interface for vSphere and vSAN	238
6-33	vSAN 6.7 Update 1: New Features and Enhancements	239
6-34	Simplified Upgrade Through vSphere Update Manager	240
6-35	Improved Resource Decommissioning Safeguards	242
6-36	Improved Insight with Cluster Capacity Monitoring.....	244
6-37	Enhancements to vRealize Operations.....	245
6-38	Improved Space Efficiency Using Storage Reclamation.....	246
6-39	Improved Isolation of vSAN Traffic Types	248
6-40	Improved Flexibility for Stretched Cluster Topologies.....	250
6-41	Improved Guidance of Health Checks and Recommendations	252
6-42	vSAN 6.7 Update 3: New Features and Enhancements (1)	254
6-43	vSAN 6.7 Update 3: New Features and Enhancements (2)	256
6-44	Scalable Storage Provisioning for DevOps	257
6-45	Creating Hybrid Cloud Solutions	258
6-46	Data Management for Stateful Cloud Native Applications	259
6-47	Predictive Analysis for Maintenance Operations	261
6-48	Improved Performance Metrics and Tools	263
6-49	vSAN 7: New Features and Enhancements (1)	265
6-50	vSAN 7: New Features and Enhancements (2)	266
6-51	vSAN File Services Architecture	267
6-52	Review of Learner Objectives.....	269
6-53	Key Points	270
Module 7	Security Enhancements	271
7-2	Importance	271
7-3	Module Lessons.....	272
7-4	Lesson 1: Virtual Machine Encryption Enhancements	273
7-5	Learner Objectives	273
7-6	Improvements to VM Encryption Functionality	274
7-7	Encryption Use Cases	275
7-8	VM Encryption Operations During VM Cloning	276
7-9	Encrypting or Decrypting the Destination VM	277
7-10	System Requirements for Encryption, Decryption, and Recryption	278

7-11	Cross vCenter Cloning of Encrypted VMs	279
7-12	Cross vCenter Migration of Encrypted VMs	280
7-13	Architecture for Cross vCenter Server Migrations and Clones (1)	281
7-14	Architecture for Cross vCenter Server Migrations and Clones (2)	282
7-15	Architecture for Cross vCenter Server Migrations and Clones (3)	283
7-16	Shallow Recrypts of VMs with Snapshots	284
7-17	Encrypted VMs with Instant Clone Technology	285
7-18	Pushing Keys to Hosts	286
7-19	VM Encryption Events and Alarms	287
7-20	Warning: Encrypted Core Dump Exists (1)	288
7-21	Warning: Encrypted Core Dump Exists (2)	289
7-22	Critical Event: Insufficient Disk Space	290
7-23	Critical Event: Key Generation Failed	291
7-24	Critical Event: KMS Unreachable	292
7-25	Lab 6: Cloning and Decrypting VMs	293
7-26	Review of Learner Objectives	294
7-27	Lesson 2: vSphere Trust Authority	295
7-28	Learner Objectives	295
7-29	About vSphere Trust Authority	296
7-30	Use Case for vSphere Trust Authority	297
7-31	vSphere Trust Authority	298
7-32	vSphere Trust Authority Components.....	299
7-33	About vSphere Trust Authority Administrator.....	300
7-34	ESXi Components	301
7-35	Attestation Service.....	302
7-36	Key Provider Service	303
7-37	Trusted Infrastructure Agent	304
7-38	Ports Used by vSphere Trust Authority	305
7-39	vCenter Server Topologies (1)	306
7-40	vCenter Server Topologies (2)	307
7-41	vSphere Trust Authority Workflow.....	308
7-42	Enabling vSphere Trust Authority	309
7-43	VM Encryption with vSphere Trust Authority	310

7-44	Identifying VM Encryption Key Providers	311
7-45	Migrating a VM Encrypted with a Trusted Key Provider (1)	312
7-46	Migrating a VM Encrypted with a Trusted Key Provider (2)	313
7-47	Labs	314
7-48	Lab 7: Assigning a vSphere Trust Authority Administrator	315
7-49	Lab 8: Enabling and Configuring vSphere Trust Authority	316
7-50	Lab 9: Encrypting a VM with a Trusted Key Provider	317
7-51	Review of Learner Objectives	318
7-52	Key Points	319
Module 8	Cluster Enhancements	321
8-2	Importance	321
8-3	Module Lessons	322
8-4	Lesson 1: Cluster Quickstart Workflow	323
8-5	Learner Objectives	323
8-6	Cluster Quickstart Workflow	324
8-7	Creating a Cluster	325
8-8	Using Cluster Quickstart Workflow to Add Hosts (1)	326
8-9	Using Cluster Quickstart Workflow to Add Hosts (2)	327
8-10	Using Cluster Quickstart Workflow to Configure a Cluster	328
8-11	Configuring a Cluster: Distributed Switches	329
8-12	Configuring a Cluster: vMotion Traffic	330
8-13	Configuring a Cluster: Storage Traffic	331
8-14	Configuring a Cluster: Advanced Options	332
8-15	Configuring a Cluster: Claim Disks	333
8-16	Configuring a Cluster: Proxy Settings	334
8-17	Additional Cluster Configuration	335
8-18	Using Quickstart to Expand Clusters	336
8-19	Review of Learner Objectives	337
8-20	Lesson 2: vSphere DRS Enhancements	338
8-21	Learner Objectives	338
8-22	DRS in vSphere 6.x	339
8-23	vSphere 6.x DRS Summary	340

8-24	DRS in vSphere 7	341
8-25	VM DRS Score.....	342
8-26	Cluster DRS Score	344
8-27	DRS Settings and Score List.....	345
8-28	VM DRS Score List.....	346
8-29	Cluster DRS Score in Performance Charts	347
8-30	Cluster DRS Score Chart.....	348
8-31	Agent VMs.....	349
8-32	Agent VM Management in vSphere 6.x	350
8-33	Agent VM Management in vSphere 7	351
8-34	Resource Pool Shares	352
8-35	Diluted Shares	353
8-36	About Scalable Shares	354
8-37	Enabling Scalable Shares.....	355
8-38	vSphere DirectPath I/O.....	356
8-39	Dynamic DirectPath I/O	357
8-40	Adding a Hardware Label to PCI Devices.....	358
8-41	Adding Dynamic vSphere DirectPath I/O Devices to a VM.....	359
8-42	Review of Learner Objectives.....	360
8-43	Key Points	361

Module 9	vSphere Lifecycle Management	363

9-2	Importance	363
9-3	Module Lessons.....	364
9-4	Lesson 1: vCenter Server Update Planner	365
9-5	Learner Objectives	365
9-6	Overview of vCenter Server Update Planner	366
9-7	Update Planner Requirements	367
9-8	Generating Update Planner Reports	368
9-9	Exporting Update Planner Reports.....	369
9-10	Upgrading vCenter Server	370
9-11	Review of Learner Objectives.....	371
9-12	Lesson 2: Overview of vSphere Lifecycle Manager	372

9-13	Learner Objectives	372
9-14	Introduction to vSphere Lifecycle Manager	373
9-15	vSphere Update Manager and vSphere Lifecycle Manager: Functionalities	374
9-16	Elements of ESXi Images	375
9-17	About ESXi Base Image	376
9-18	About Components	377
9-19	About Vendor Add-Ons	378
9-20	About Firmware and Driver Add-Ons	379
9-21	Review of Learner Objectives	380
9-22	Lesson 3: Working with vSphere Lifecycle Manager	381
9-23	Learner Objectives	381
9-24	vSphere Lifecycle Manager Home View	382
9-25	Image Depot	383
9-26	Importing Updates	385
9-27	Patch Settings	386
9-28	Using Images to Perform ESXi Host Life Cycle Operations	387
9-29	Creating an ESXi Image for a New Cluster	388
9-30	Customizing Cluster Images	390
9-31	Checking Image Compliance	391
9-32	Running a Remediation Precheck	393
9-33	Hardware Compatibility	394
9-34	Standalone VIBs	395
9-35	Remediating a Cluster Against a Single Image	396
9-36	Reviewing Remediation Impact	398
9-37	Recommended Images	399
9-38	Viewing Recommended Images	400
9-39	Selecting a Recommended Image	402
9-40	Esxcli Support for vSphere Lifecycle Manager Images	403
9-41	Lab 10: Updating ESXi Hosts with vSphere Lifecycle Manager	404
9-42	Review of Learner Objectives	405
9-43	Key Points	406

Module 10	Introduction to vSphere with Kubernetes	407
10-2	Importance	407
10-3	Module Lessons	408
10-4	Lesson 1: Introduction to Containers	409
10-5	Learner Objectives	409
10-6	Virtual Machines and Containers (1)	410
10-7	Virtual Machines and Containers (2)	411
10-8	About Container Hosts	412
10-9	Container Workflow Example	413
10-10	About Container Engines	414
10-11	About Dockerfile	416
10-12	Container Images	417
10-13	Image Registry	418
10-14	Review of Learner Objectives	419
10-15	Lesson 2: Introduction to Kubernetes	420
10-16	Learner Objectives	420
10-17	Problems Solved by Kubernetes	421
10-18	Kubernetes Architecture	422
10-19	About Manifests	424
10-20	About Pods	425
10-21	About Deployments	426
10-22	About ReplicaSets	427
10-23	About Services	428
10-24	About Network Policies	429
10-25	Kubernetes Workflow Example	430
10-26	Review of Learner Objectives	432
10-27	Lesson 3: Introduction to vSphere with Kubernetes	433
10-28	Learner Objectives	433
10-29	vSphere with Kubernetes and VMware Cloud Foundation 4	434
10-30	vSphere with Kubernetes and VMware Tanzu (1)	436
10-31	vSphere with Kubernetes and VMware Tanzu (2)	437
10-32	About vSphere with Kubernetes	439

10-33	About Supervisor Clusters (1)	440
10-34	About Supervisor Clusters (2)	441
10-35	About Control Plane VMs (1)	442
10-36	About Control Plane VMs (2)	443
10-37	Control Plane Networking	444
10-38	About Spherelet	445
10-39	About CRX	446
10-40	About vSphere Pods	447
10-41	vSphere Pods Network and Storage	448
10-42	Comparing vSphere Pods with Kubernetes	449
10-43	About Namespaces	450
10-44	Container Registry	452
10-45	Tanzu Kubernetes Grid Service for vSphere	453
10-46	vSphere with Kubernetes Objects	454
10-47	Review of Learner Objectives	455
10-48	Key Points	456

Module 1

Course Introduction

1-2 Course Introduction

1-3 Importance

By becoming familiar with the new features in vSphere 7, you can help support your software-defined data center. Using the new and enhanced features, you can improve the availability, scalability, security, and automated management of your environment. With native Kubernetes support, you can create and manage your Kubernetes cluster.

1-4 Learner Objectives (1)

After completing this course, you should be able to meet the following objectives:


- Recognize the Importance of key features and enhancements in vSphere 7
- Migrate a vCenter Server system for Windows to vCenter Server Appliance 7.0
- Upgrade a vCenter Server instance to vSphere 7
- Update templates in the content library
- Use identity federation to configure vCenter Server to use external identity sources
- Use the vSphere Client to manage certificates
- Use vSphere Trust Authority to secure the infrastructure for encrypted VMs

1-5 Learner Objectives (2)

- Use vSphere Lifecycle Manager to manage vSphere deployments and updates
- Create a basic vSphere with Kubernetes configuration
- Describe the cloning and migration operations on encrypted VMs
- Describe the purpose of scalable shares and the improvements to vSphere DRS
- Describe how to change the IP address or FQDN of vCenter Server
- Identify new network hardware device support
- Recognize the enhancements available with VM compatibility version 17

1-6 vSphere 7 Overview

vSphere 7 introduces many new features and enhancements.



1-7 Course Outline

1. Course Introduction
2. Installation and Upgrade
3. Management Enhancements
4. ESXi Enhancements
5. VM Enhancements
6. Storage Enhancements
7. Security Enhancements
8. Cluster Enhancements
9. vSphere Lifecycle Management
10. Introduction to vSphere with Kubernetes

1-8 Typographical Conventions

The following typographical conventions are used in this course.

Conventions	Usage and Examples
Monospace	Identifies command names, command options, parameters, code fragments, error messages, filenames, folder names, directory names, and path names: • Run the <code>esxtop</code> command. • ... found in the <code>var/log/messages</code> file.
Monospace Bold	Identifies user inputs: • Enter <code>ipconfig/release</code>.
Boldface	Identifies user interface controls: • Click the Configuration tab.
<i>Italic</i>	Identifies book titles: • <i>vSphere Virtual Machine Administration</i>
< >	Indicates placeholder variables: • <ESXi_host_name> • ... the Settings/<Your_Name>.txt file

1-9 References (1)

Title	Location
<i>vSphere Installation and Setup</i>	https://docs.vmware.com/en/VMware-vSphere/index.html
<i>vCenter Server and Host Management</i>	https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vcenterhost.doc/GUID-3B5AF2B1-C534-4426-B97A-D14019A8010F.html https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vcenterhost.doc/GUID-3B5AF2B1-C534-4426-B97A-D14019A8010F.html
<i>vSphere Virtual Machine Administration</i>	https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html
<i>vSphere Storage</i>	https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html
<i>vSphere Networking</i>	https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.networking.doc/GUID-35B40B0B-0C13-43B2-BC85-18C9C91BE2D4.html

1-10 References (2)

Title	Location
<i>vSphere with Kubernetes Configuration and Management</i>	https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-kubernetes/GUID-152BE7D2-E227-4DAA-B527-557B564D9718.html
<i>vSphere Security</i>	https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-52188148-C579-4F6A-8335-CFBCE0DD2167.html
<i>VMware Compatibility Guide</i>	https://www.vmware.com/resources/compatibility/search.php
<i>VMware Product Interoperability Matrices</i>	https://www.vmware.com/resources/compatibility/sim/interop_matrix.php

1-11 VMware Online Resources

Documentation for vSphere 7: <https://docs.vmware.com/>

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

- Start a discussion.
- Access the knowledge base.
- Access documentation, technical papers, and compatibility guides.
- Access communities.
- Access user groups.

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

VMware Hands-on Labs: <http://hol.vmware.com>

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

- Access course catalog and worldwide course schedule.


1-12 VMware Education Overview

Your instructor will introduce other Education Services offerings available to you:

- VMware Learning Paths:
 - Help you find the course that you need based on the product, your role, and your level of experience
 - Can be accessed at <https://vmware.com/education>
- VMware Learning Zone, which is the official source of digital training, includes the following options:
 - On Demand Courses: Self-paced learning that combines lecture modules with hands-on practice labs
 - VMware Lab Connect: Self-paced, technical lab environment that lets you practice skills learned during instructor-led training
 - Certification Exam Prep: Comprehensive video-based reviews of exam topics and objectives to help you take your certification exam
- For more information, see <https://vmwarelearningzone.vmware.com>.

1-13 VMware Certification Overview

VMware certifications validate your expertise and recognize your technical knowledge and skills with VMware technology.



VMware certification sets the standards for IT professionals who work with VMware technology. Certifications are grouped into technology tracks. Each track offers one or more levels of certification (up to five levels).

For the complete list of certifications and details about how to attain these certifications, see <https://vmware.com/certification>.

1-14 VMware Badge Overview

VMware badges are digital emblems of skills and achievements.



Digital badges have the following features:

- Easy to share in social media (LinkedIn, Twitter, Facebook, blogs, and so on)
- Tethered to VMware to validate and verify achievement
- Contain metadata with skill tags and accomplishments
- Based on Mozilla's Open Badges standard

For the complete list of digital badges, see <http://www.pearsonvue.com/vmware/badging>.

Module 2

Installation and Upgrade

2-2 Importance

vSphere 7 reduces the complexity of deployments by simplifying the topology options for vCenter Server. When planning a vSphere deployment, you must know how to use the new installation and configuration options. These options can simplify and improve your deployment.

2-3 Module Lessons

1. Upgrading and Migrating vCenter Server
2. vCenter Server Appliance Enhancements
3. Upgrading ESXi

2-4 **Lesson 1: Upgrading and Migrating vCenter Server**


2-5 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize vCenter Server 6.x deployment models
- Describe vCenter Server 7.0 deployment enhancements
- Identify installer enhancements
- Identify vCenter Server upgrade steps
- Describe the direct upgrade and migration processes

2-6 Deployment Models

Since their introduction with vSphere 6.0, Platform Services Controller and vCenter Server share an architecture. The architecture has two deployment models.



vCenter Server with embedded Platform Services Controller: The Platform Services Controller services are installed in the same appliance as the vCenter Server services.


vCenter Server with external Platform Services Controller: The Platform Services Controller services are installed as a separate appliance to the vCenter Server services.

vCenter Server deployment models define how the vCenter Server services are made available to users:

- You can install vCenter Server with an embedded or external Platform Services Controller instance on a Microsoft Windows VM or physical server.
- You can deploy vCenter Server Appliance with an embedded or external Platform Services Controller instance.

2-7 vCenter Server 7.0 Deployment Enhancements (1)


vCenter Server 7.0 supports only vCenter Server Appliance with embedded Platform Services Controller. It does not support the deployment of an external Platform Services Controller node.



vCenter Server 7.0 supports only an embedded architecture where Platform Services Controller and vCenter Server systems are deployed to the same appliance.

2-8 vCenter Server 7.0 Deployment Enhancements (2)

With the release of vSphere 7, you cannot install vCenter Server to a Windows OS.




Users who want to move from Windows vCenter Server (6.5 or 6.7) can migrate to vCenter Server 7.0.

Older versions of vCenter Server must be upgraded to either 6.5 or 6.7 before you can migrate to vCenter Server 7.0.

2-9 Enhanced Linked Mode with an Embedded vCenter Server

As with previous versions of vSphere, vCenter Server can be deployed in Enhanced Linked Mode:

- Enhanced Linked Mode connects multiple vCenter Server systems.
- Enhanced Linked Mode provides a single view for all linked vCenter Server systems.




Before vSphere 6.5 U2, you could not use vCenter Server in Enhanced Linked Mode when using an embedded Platform Services Controller node.

Because an embedded configuration is the only deployment model, Enhanced Linked Mode is supported. It also works with an external Platform Services Controller node.

2-10 Benefits of vCenter Server 7.0 Deployment Model

Simplifying the deployment model has the following benefits:

- No extra VM for Platform Services Controller or its life cycle overhead
- No extra Windows license for the Platform Services Controller node
- No load balancer for high availability



With fewer deployment options, the complexity of vCenter Server management is reduced, saving cost and time.

You require a Windows license for the Microsoft Windows version of vCenter Server, and, if using a SQL database, you require an extra license.

If vCenter Server is configured with an external Platform Services Controller node, you need another Windows license.

Having a distributed configuration also makes management more complicated because you must manage, patch, and back up additional VMs.

2-11 vSphere Version Upgrade Compatibility

The VMware Product Interoperability Matrix provides the supported upgrade paths for vSphere and other VMware products.


vSphere	6.0 U1	6.0 U2	6.0 U3	6.5	6.5 U1	6.5 U2	6.7 (All)	7.0
6.0	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
6.0 U1		Yes	Yes	Yes	Yes	Yes	Yes	No
6.0 U2			Yes	Yes	Yes	Yes	Yes	No
6.0 U3				No	Yes	Yes	Yes	No
6.5					Yes	Yes	Yes	Yes
6.5 U1						Yes	Yes	Yes
6.5 U2							Yes	Yes
6.7 (All)								Yes

See the VMware Product Interoperability Matrices at
https://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php#upgrade.

https://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php#upgradeThe matrix covers both upgrade of vCenter Server Appliance and the migration a Windows vCenter Server system.

2-12 vCenter Server Upgrade Process

You can upgrade or migrate vCenter Server 6.5 or 6.7 to 7.0. Verify you meet the upgrade requirements before commencing the upgrade process.




You follow several high-level steps when upgrading or migrating vCenter Server:

1. Select your upgrade goal. Are you upgrading vCenter Server Appliance or migrating vCenter Server for Windows to vCenter Server Appliance?
2. Verify that your system meets the hardware and software requirements.
3. Prepare your environment for the upgrade or migration.
4. Upgrade or migrate vCenter Server for Windows or vCenter Server Appliance.
5. Complete the required post-upgrade or post-migration tasks.

2-13 Upgrade Paths to vCenter Server 7.0

The installer upgrades vCenter Server 6.5 or 6.7, with an embedded or with an external Platform Services Controller node, to vCenter Server 7.0.




The installer upgrades vCenter Server 6.5 and 6.7 with an embedded Platform Services Controller node to vCenter Server 7.0.

The installer upgrades vCenter Server 6.5 and 6.7 with an external Platform Services Controller node to vCenter Server 7.0.

2-14 Migration Paths to vCenter Server 7.0

You can migrate vCenter Server 6.5 or 6.7, which runs on Windows, to vCenter Server Appliance 7.0.

You can migrate vCenter Server 6.5 or 6.7 with an external Platform Services Controller node to vCenter Server Appliance 7.0.




If VMware Update Manager is registered with vCenter Server, its configuration, including baselines and metadata, are migrated to the vCenter Server Appliance 7.0.

2-15 Upgrade or Migration to vCenter Server 7.0 with Shared Platform Services Controller

If multiple vCenter Server instances are connected to a shared Platform Services Controller node, the upgrade or migration includes the following steps:

- With the first upgrade or migration, a convergence is performed.
- All other vCenter Server instances remain connected to the shared Platform Services Controller node.
- When the last vCenter Server instance is upgraded or migrated, Platform Services Controller can be decommissioned.



If you upgrade a vSphere deployment where multiple vCenter Server instances are registered to a shared Platform Services Controller node, a phased convergence can be performed.

When the first upgrade or migration is performed, vCenter Server is upgraded, and Platform Services Controller is converged into an embedded system.

Replication agreements are created to ensure the vCenter Server instances remain in Enhanced Linked Mode.


The vCenter Server instances that are yet to be upgraded remain in a distributed configuration and registered with the shared Platform Services Controller node.

All vCenter Server instances can be upgraded in a staged fashion.

When the last vCenter Server instance is upgraded, Platform Services Controller can be decommissioned.

2-16 Enhancements to vCenter Server Appliance and CLI Installers

The external Platform Services Controller deployment option is not available in the vCenter Server Appliance 7.0 UI and is not a valid option in the `.json` files that are used by the CLI installer.



The installer is simplified: Only one deployment type is available. You must still specify whether you want to create single sign-on (SSO) domain or join an existing one.

Unlike in previous vCenter Server versions, you can apply Enhanced Linked Mode to an embedded vCenter Server instance.

2-17 Resource Requirements for Deploying vCenter Server Appliance (1)

Resource requirements depend on the deployment size of vCenter Server Appliance 7.0.

1 Introduction

2 End user license agreement

3 vCenter Server deployment target

4 Set up vCenter Server VM

5 Select deployment size

6 Select datastore

7 Configure network settings

8 Ready to complete stage 1

Select deployment size

Select the deployment size for this vCenter Server.

For more information on deployment sizes, refer to the vSphere 7.0 documentation.

Deployment size Tiny

Storage size Default ⓘ

Resources required for different deployment sizes

Deployment Size	vCPUs	Memory (GB)	Storage (GB)	Hosts (up to)	VMs (up to)
Tiny	2	12	415	10	100
Small	4	19	480	100	1000
Medium	8	28	700	400	4000
Large	16	37	1065	1000	10000
X-Large	24	56	1805	2000	35000

vSphere 7 has increased memory requirements compared with version 6.x. This increase is handled automatically during the upgrade.

You can increase the default storage size for the deployment if you have larger storage requirements.

You can modify compute and storage resources that are allocated to the vCenter Server Appliance after deployment.

2-18 Resource Requirements for Deploying vCenter Server Appliance (2)

To upgrade or migrate from supported vCenter Server instances to vCenter Server Appliance 7.0, you require the following user permissions:

- Source and destination vCenter Single Sign-On: User name and password
- Source and destination root password for ESXi (If the source or destination vCenter Server is hosted on the ESXi server)
- Source Windows server administrator password (for migration)



2-19 vCenter Server Upgrade and Migration Prerequisites

A direct upgrade and a converge upgrade or migration from vCenter Server systems earlier than 6.5 to vCenter Server 7.0 are not supported.

You must upgrade vCenter Server and ESXi servers to vSphere version 6.5 or 6.7 before performing a direct or converge update to vCenter Server 7.0.

You must back up the source vCenter Server and all the required data from the database.

To perform a converge or migrate update to vCenter Server 7.0, you must be running vCenter Server 6.5 or later. If you are running an earlier version, you must upgrade vCenter Server to 6.5 or 6.7.

Check the vSphere compatibility matrix for the migration paths that are supported.

Before starting the upgrade process, back up the vCenter Server guest and the vCenter Server database.

2-20 Upgrade Sequence for vSphere Components

In a large vSphere environment, you might need to upgrade multiple components.


Understanding the required sequence of tasks is key to a successful vSphere upgrade.

When you upgrade vSphere, you must perform all procedures in the specified sequence to avoid possible data loss and to minimize downtime. You can perform the upgrade process for each component in only one direction. For example, after you upgrade to vCenter Server 7.0, you cannot revert to vCenter Server 6.5 or 6.7. With backups and some planning, however, you can restore your original vCenter Server environment, if needed.

Upgrading vSphere includes the following tasks:

- Reviewing the vSphere 7 documentation and release notes
- Backing up your configuration and VMs
- Reviewing the VMware Product Interoperability Matrix and ensuring that all solutions are compatible
- Upgrading vCenter Server
- Upgrading ESXi hosts
- Updating VMware Tools, if needed


In vSphere 7, you create VMs using hardware version 17 by default. However, you can select older versions. And you can upgrade the hardware on existing VMs, if needed. Customers might want to perform this operation to access a feature that is only available on the newer ESXi compatibility version.



2-21 Direct Upgrade or Migration to vCenter Server Appliance 7.0

To upgrade or migrate from vCenter Server 6.5 or 6.7 directly to vSphere 7:


1. Upgrade from vCenter Server 6.5 or 6.7 to vCenter Server Appliance 7.0.
2. Decommission the source vCenter Server system.
3. Upgrade ESXi 6.5 or 6.7 to ESXi 7.0.



2-22 Converge Upgrade or Migration to vCenter Server Appliance 7.0

To converge upgrade or migrate from vCenter Server version 6.5 or 6.7 directly to vSphere 7.0:

1. Converge upgrade or migrate from vCenter Server 6.5 or 6.7 to vCenter Server Appliance 7.0.
2. Decommission the source vCenter Server system and Platform Services Controller.
3. Upgrade ESXi 6.5 or 6.7 to ESXi 7.0.



2-23 Demonstration: Upgrading vCenter Server Appliance 6.7 to 7.0

Your instructor will run a demonstration.

The lab documentation and lab vApp contain all the required resources if you want to perform the steps on your own.

2-24 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize vCenter Server 6.x deployment models
- Describe vCenter Server 7.0 deployment enhancements
- Identify installer enhancements
- Identify vCenter Server upgrade steps
- Describe the direct upgrade and migration processes

2-25 **Lesson 2: vCenter Server Virtual Appliance Enhancements**

2-26 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Identify enhancements to vCenter Server Appliance
- Recognize vCenter Server and ESXi support and trusted maximums
- Explain the vCenter Server Architecture

2-27 vCenter Server Appliance Architecture

vCenter Server Appliance is a preconfigured, Linux-based VM that is optimized for running the vCenter Server application. vCenter Server Appliance architecture has the following characteristics:

- Runs on Photon OS
- Configured with a PostgreSQL embedded database
- Managed using the vCenter Server Appliance Management Interface
- Can manage ESXi 6.5, 6.7, and 7.0 hosts

vCenter Server 7.0 runs on Photon OS 3.0.

PostgreSQL version 11.0 is the configured database in the vCenter Server Appliance 7.0.

vCenter Server 7.0 manages previous versions of the ESXi host so that only hosts running version 6.5 or later are supported.

The vSphere Web Client is retired. The vSphere Client is the only supported client. It is HTML5-based and does not require Flash.

2-28 vCenter Server 7.0 Scalability

vCenter Server can scale to meet large workloads.

Metric	vCenter Server 6.7	vCenter Server 7.0
Hosts per vCenter Server	2,000	2,500
Hosts per cluster	64	64
Powered on VMs per vCenter Server	25,000	40,000
Registered VMs per vCenter Server	35,000	45,000
VM per cluster	8,000	8,000
Linked vCenter Server instances	15	15

In Enhanced Linked Mode, vCenter Server can support 15,000 hosts and 150,000 powered-on VMs.


Although cluster size has increased to 96 hosts, the maximum number of hosts per vSAN cluster is still 64 hosts.

2-29 vCenter Server Appliance Management Interface

A virtual appliance management interface (VAMI) provides administrative control of vCenter Server Appliance.

Using VAMI, you can configure the following settings:

- SSH, console, and shell access
- Networking
- Firewall
- NTP source and time zone
- Appliance patching
- Syslog
- Appliance backup



A virtual appliance management interface (VAMI) provides administrative functionality such as configuration and monitoring of the Photon OS and vCenter Server services.

Under Networking, you can change the FQDN and IP addresses of vCenter Server.

2-30 File-Based Backups of vCenter Server Appliance

You use the VAMI to perform a file-based backup. You can select what gets backed up, for example:

- Configuration
- Inventory
- Historical data

The backup is not stored on vCenter Server Appliance but can be streamed to a remote repository using the following protocols:

- FTP/FTPS
- HTTP/HTTPS
- SCP
- NFS
- SMB

Configurations of all vCenter Server and Platform Services Controller services are backed up.

You can back up vCenter Server at the file level by using NFS and SMB, as of vSphere 6.7 U2.

You can perform a file-based restore only for a vCenter Server Appliance instance that you previously backed up by using the VAMI. You can perform the restore operation by using the GUI installer of vCenter Server.

2-31 Backup and Restore Enhancements

vCenter Server Appliance 7.0 supports new protocols for vCenter Server Appliance backup and restore:

- NFSv3 and later
- Samba (SMBv2 and later)

Create Backup Schedule

Backup location ⓘ

Backup server credentials

Schedule ⓘ : P.M.

Encrypt backup (optional)

Number of backups to retain ☒ Retain all backups ☐ Retain last backups

Data ☒ Inventory and configuration ☐

☒ Stats, Events, and Tasks

Total size (compressed)

Tooltip:

1. Supported protocols for backup are FTPS, HTTPS, SFTP, FTP, NFS, SMB and HTTP.
2. Use the following format to provide the address and location:
protocol://server-address<port-number>/folder/subfolder

2-32 Scheduling File-Based Backups

When scheduling file-based backups, follow these practices:

- Schedule only one periodic backup at a time.
- Encrypt vCenter Server backups.

To store the backup, you must have an FTP, FTPS, HTTP, HTTPS, NFS, SCP, or SMB server running with sufficient disk space.

You provide the backup location, including the protocol for connecting to your backup server, and the port, the server address, and the backup folder for storing the backup files.

Use one of the following protocols: FTP, FTPS, HTTP, HTTPS, NFS, SCP, or SMB.


For FTP, FTPS, HTTP, or HTTPS, the path is relative to the home directory configured for the service.

2-33 Restoring File-Based Backups of vCenter Server Appliance

You restore vCenter Server by using the vCenter Server Appliance installer.

The restore process works in the following ways:

- The vCenter Server and Platform Services Controller configuration is restored.
- The restore operation is performed for a newly deployed appliance.
- The version that is restored must match the appliance version.



When you back up a vCenter Server Appliance instance using the VAMI, you can restore that instance by using the GUI installer of vCenter Server Appliance.

The restore process involves deploying a new vCenter Server Appliance instance and copying the data from the file-based backup to the new appliance.

If you back up a vCenter Server Appliance high availability cluster, the backup operation backs up only the primary vCenter Server instance.

Before restoring a vCenter Server Appliance high availability cluster, you must power off the active, passive, and witness nodes. The restore operation restores vCenter Server in non vCenter Server high availability mode. You must reconstruct the cluster after the restore operation completes successfully.

You cannot restore a backup from a different version of vCenter Server. For example, you cannot restore a vCenter Server 6.7 U3 backup to vCenter Server Appliance 7.0.

2-34 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Identify enhancements to vCenter Server Appliance
- Recognize vCenter Server and ESXi support and trusted maximums
- Explain the vCenter Server Architecture

2-35 **Lesson 3: Upgrading ESXi**

2-36 Learner Objectives

After completing this lesson, you should be able to meet the following objective:


- Recognize how to upgrade ESXi 6.5 or 6.7 to ESXi 7.0

2-37 ESXi Host Upgrade Process

VMware provides several ways to upgrade ESXi 6.5.x and 6.7.x hosts to ESXi 7.0.

High-level steps for upgrading ESXi are as follows:

1. Verify that the host meets the upgrade requirements.
2. Prepare your environment before upgrading.
3. Upgrade ESXi using one of the supported methods.
4. After upgrading, verify the ESXi configuration and reconnect to vCenter Server.



2-38 Hardware and System Requirements for ESXi Upgrade or Installation

To install or upgrade ESXi, you must meet hardware and system resources requirements:

- You must have a supported server platform.
- The host must have at least two CPU cores.
- The NX/XD feature must be enabled for the CPU in the BIOS.
- You require a minimum of 4 GB of physical RAM.
- Hardware virtualization (Intel VT-x or AMD RVl) must be enabled on x64 CPUs for 64-bit VMs.
- Ethernet controllers must be 1 Gigabit or faster.
- You require a disk of at least 4 GB for USB or SD devices, and 32 GB for other device types such as HDD, SSD, or NVMe.

vSphere 7 supports starting ESXi hosts from the Unified Extensible Firmware Interface (UEFI). With UEFI, you can start systems from hard drives, CD-ROM drives, or USB media.

vSphere Auto Deploy supports network starting and provisioning of ESXi hosts with UEFI.

ESXi can start from a disk larger than 2 TB if the system firmware and the firmware on any add-in card that you use support it. See the vendor documentation.

You can validate hardware and platform support with the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/>.

2-39 Storage Requirements for ESXi 7.0 Installation or Upgrade

For an ESXi 7.0 installation, the following requirements apply:


- A 4 GB USB or SD and an additional 32 GB local disk
- A local disk with a minimum of 32 GB for boot partition and ESX-OSData
- For a local disk larger than 128 GB, a VMFS partition is created.

For an ESXi 7.0 installation, you should meet the following requirements:

A 4 GB USB or SD and an additional 32 GB local disk: The ESXi boot partitions reside on the USB or SD.

A local disk with a minimum of 32 GB: The disk contains the boot partitions and ESX-OSData volume.

A local disk larger than 128 GB: The disk contains the boot partitions, ESX-OSData volume, and VMFS datastore.



ESXi 7

2-40 Direct Upgrade to ESXi 7.0

For a direct upgrade to ESXi 7.0, several methods or tools are supported:

- The interactive GUI installer from ISO, CD, DVD, or USB drive
- Scripted upgrade
- The `esxcli` command-line interface
- vSphere Auto Deploy
- vSphere Update Manager

2-41 vSphere 7 Core Licensing

In releases before vSphere 7, an ESXi host uses socket-based licensing. One socket requires one license to be applied regardless of the number of cores.

vSphere 7 implements a core-based model where one vSphere license is required for up to 32 cores (1 CPU).

Administration

Access Control

Roles

Global Permissions

Licensing

Licenses

Solutions

Client Plugins

vCenter Server Extensions

Deployment

System Configuration

Customer Experience Improvement Progr...

Support

Upload File to Service Request

Licenses

Licenses

Products

Assets

GO TO MY VMWARE

Product/License	Usage	Capacity
VMware vSphere 7 Enterprise Plus	0 CPUs (up to 32 cores)	33 CPUs (up to 32 cores)

EXPORT

1 item

Summary

Features

Licenses

Hosts

2-42 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objective:

- Recognize how to upgrade ESXi 6.5 or 6.7 to ESXi 7.0

2-43 Key Points

- vCenter Server 7.0 supports only an embedded architecture.
- vCenter Server 7.0 is available only as a Linux virtual appliance.
- vCenter Server 7.0 can manage ESXi 6.5, 6.7, and 7.0 hosts.
- You can use several methods to upgrade ESXi 6.5.x and 6.7.x hosts to ESXi 7.0.

Questions?

Module 3

Management Enhancements

3-2 Importance

You use vCenter Server as a central point of management in your vSphere environment. To work with vSphere components effectively, you must understand how to use the vSphere Client. You can manage vCenter Server Appliance by using the virtual appliance management interface (VAMI). You can also proactively monitor your vSphere environment using VMware Skyline.

3-3 Module Lessons

1. vSphere Client and vCenter Server Appliance Enhancements
2. Deploying and Managing Templates in the Content Library
3. vCenter Server Identity Federation
4. vCenter Certificate Enhancements
5. VMware Skyline Overview

3-4 **Lesson 1: vSphere Client and vCenter Server Appliance Enhancements**

3-5 Learner Objectives


After completing this lesson, you should be able to meet the following objectives:

- Recognize how to create and apply vCenter Server profiles
- Identify the scaling and performance improvements for vCenter Server Appliance
- Describe the support provided for dynamic DNS
- Recognize how to change the IP address or FQDN of vCenter Server
- Identify enhancements to the vCenter Server file-based backup
- Identify enhancements to tools provided by the vSphere Client Developer Center
- Describe prefix-based MAC address allocation
- Identify NSX-T Data Center logical switches and their properties in the vSphere Client
- Use virtual distributed port groups to connect to NSX-T Data Center logical switches

3-6 vCenter Server Client

The vSphere Client (HTML5) is the only client that is required to manage vCenter Server. You do not need a flash-enabled browser.

The vSphere Client extensibility platform caters to various VMware and third-party plug-ins.




Using the HTML5-based vSphere Client does not depend on the installation of Adobe Flex. No dedicated or separate installation is necessary.

You can access the vSphere Client from vCenter Server Appliance at https://vCenter_FQDN/UI.

https://vCenter_FQDN/UIThe Web Client is no longer available. Full functionality is available in the vSphere Client (HTML5).

3-7 About vCenter Server Profiles

In vSphere 7, the vCenter Server Profiles feature helps you maintain configuration consistency across multiple vCenter Server instances.




3-8 vCenter Server Profiles Management Benefits

With vCenter Server profiles, you can export an existing vCenter Server configuration and import the exported configuration data to another vCenter Server instance.

The vCenter Server Profiles feature offers many benefits:

- Simplified setup and change management for vCenter Server instances
- Easy detection of noncompliance with standard configurations
- Automated remediation
- Export and import vCenter Server configuration in a user-readable format



With vCenter Server profile, you can import and export existing vCenter Server configurations. You can use the exported data for configuring another vCenter Server setup.

You can use the same configuration to maintain versions of configurations.

vCenter Server Profiles APIs can be used by many scripting tools, including Ansible, Chef, and Puppet, to automate the process.

3-9 vCenter Server Profiles Management APIs

vCenter Server profiles use a modular action-based web framework. You can access vCenter Server profiles through a set of APIs.

With vCenter Server profiles, you can access the following API interfaces:

- List
- Export
- Validate
- Import

List API: Returns a list of components or services in vCenter Server

Export API: Exports the desired configuration profiles

Validate API: Validates the desired profile configuration that can be imported to vCenter Server

Import API: Imports the desired profile configuration JSON string into vCenter Server

Every plug-in owner implements a basic set of mandatory APIs and creates a Java JAR file.

3-10 vCenter Server Profiles Operations

A vCenter Server profile is a JSON file that consists of the vCenter Server configuration.

vCenter Server profiles can perform the following tasks:

- Exporting service configuration
- Validating service configuration
- Importing service configuration


For upgrade, backups, and restores, the vCenter Server profiles retain changes made to the profile.

To access vCenter Server APIs, see https://www.vmware.com/support/pubs/sdk_pubs.html.

3-11 About the DDNS Client

vCenter Server Appliance 7.0 introduces the Dynamic Domain Name Service (DDNS) client. The DDNS client performs the following tasks:

- Updating the DNS server with changes to IP addresses
- Sending keep alive messages to the DNS server every day at 12:00 AM UTC



The Dynamic Domain Name Service (DDNS) client updates the DNS server only about changes to the vCenter Server IPs. No other infrastructure IP changes are captured. For example, changes to VMs or ESXi host IPs are not captured.

3-12 DDNS Client Use Cases

The DDNS client helps in these situations:

- Changing the IP address of vCenter Server Appliance and updating the DNS server with the changed IP address
- Sending keep-alive messages to the DNS server so that entries in the DNS server do not become stale and get removed by DNS aging and scavenging processes

DNS aging and scavenging processes automatically clean up and remove stale resource records in the DNS server.

3-13 About Multihoming

With vCenter Server Appliance 7.0 multihoming, you can configure multiple NICs to manage network traffic.

Network Settings		EDIT
Hostname	sb-vcsa-01.vclass.local	
DNS Servers	172.20.110.10	
- NIC 0		
Status	Up	
MAC Address	00:0C:29:8A:9C:F5	
IPv4 Address	172.20.110.84 / 24 (Static)	
IPv4 Default Gateway	172.20.110.10	
- NIC 2		
Status	Down	
MAC Address	00:0C:29:8A:9C:09	
IPv4 Address	172.20.110.95 / 24 (Static)	
- NIC 1 (Reserved for VCHA)		
Status	Down	
MAC Address	00:0C:29:8A:9C:FF	
IPv4 Address	/ (Unconfigured)	
Proxy Settings		
HTTPS	Disabled	EDIT
FTP	Disabled	
HTTP	Disabled	

A maximum of four NICs are supported for multihoming. All four multihoming-supported NIC configurations are preserved during the upgrade, backup, and restore processes.

vCenter might need to communicate with entities on two different subnets.

3-14 About PNID

The primary network identifier (PNID) is the unique system name of an individual vCenter Server instance.

vCenter Server uses the FQDN or IP address as the system name during vCenter Server installation.

If you leave the optional FQDN field empty on the installation wizard, the consequences are as follows:

- The reverse lookup tries to identify the host name associated with the IP address.
- If a hostname is not found, the IP address is used as the system name.

The system name is encoded in the SSL certificate of the system. The vCenter Server and Platform Services Controller components communicate with each other by using the system name.

To verify the current primary network identifier (PNID), run the following command in vCenter Server Appliance:

```
#/usr/lib/vmware-vmafd/bin/vmafd-cli get-pnid --server-name localhost
```


3-15 PNID Change Use Cases

You might need to change the PNID in the following situations:

- The network configuration changes in your vSphere infrastructure.
- The domain name changes as the result of an organization change such as a merger, split, or branch out.
- You are configuring vCenter Server before deployment to a branch office.

In previous versions, the PNID was static. If the host name of the system hosting vCenter Server changed, you could not easily change the PNID.

The vCenter Server PNID might change for several reasons. One common reason is when the domain where vCenter Server is hosted changes. This change might be because of a company buyout or merger.

3-16 Changing the PNID

You can use the VAMI to change the host name (system name):

- Changing the host name results in the change of the PNID of vCenter Server.
- HOSTNAME information of the OS is updated.

In vSphere 7, minimum user effort is required for actions such as regenerating custom certificates, rejoining to Active Directory (AD), recreating replication partners, and redeploying vCenter Server High Availability.

Use FQDN as the vCenter Server PNID. However, vCenter Server can also be deployed with the IP address as its PNID.

In releases earlier than vSphere 7, when a user changes the PNID, the only solution is to reinstall vCenter Server.

The advantages of changing the PNID are as follows:

- You do not need to deploy additional Platform Services Controller or vCenter Server instances.
- The minimum downtime is less than 15 minutes.
- SSO data is not lost.
- It is easy to recover from failures.

3-17 PNID Change Considerations (1)

Before changing the PNID of vCenter Server:

- Remove vCenter HA Configuration
- Unregister solutions (Site Recovery Manager, vRealize Operations Manager, Log Insight, and so on)

3-18 PNID Change Considerations (2)

After changing the PNID:

- Regenerate and replace vCenter Server certificates.
- Configure vCenter Server High Availability (if necessary).
- Register any solutions to vCenter Server.
- Add vCenter Server with an updated host name to Active Directory (AD) if it was configured with AD before the PNID change.


After the PNID change, the vCenter Server certificates are reset to VMware CA generated certificates.

3-19 Prerequisites for Changing the PNID

Before changing the PNID, take a backup or snapshot of vCenter Server Appliance.

You require the following user permissions:

- vCenter Single Sign-On: User name and password
- VAMI access



To log into the virtual appliance management interface (VAMI), get the root password. The VAMI listens on port 5480. If you access vCenter Server Appliance from an external network, verify that the network traffic is not blocked.

3-20 MAC Address Allocation in vSphere

vCenter Server provides schemes of MAC address generation:

- VMware Organizationally Unique Identifier (OUI) allocation
- Range-based allocation
- Prefix-based allocation

In vSphere 7, prefix-based MAC address allocation is selected as the default.

You can select the scheme that best suits your OUI requirements for locally or universally administered addresses, avoiding MAC address duplication.

vCenter Server provides the following schemes of MAC address generation:


- VMware Organizationally Unique Identifier (OUI) allocation is the default scheme in versions earlier than vSphere 7. The scheme uses the MAC address prefix 00:50:56 and supplements the balance of the MAC address as a calculation for the remaining six digits.
- Prefix-based allocation is an option for setting a custom, locally administered MAC address OUI.
- Range-based allocation is an option for programming a locally administered MAC address range into vCenter Server using the `vpxd.cfg` file. vCenter Server checks for MAC address conflicts, for example, an overlap with another vSphere installation in the network.

vCenter Server relies on the prefixes to avoid MAC address duplication issues.

3-21 VMware OUI MAC Allocation

The VMware OUI allocation scheme uses the MAC address prefix 00:50:56:

- The two-digit hexadecimal numbers represent the vCenter Server ID + 80.
- The final two-digit hexadecimal numbers are randomly generated.




The allocation works with up to 64 vCenter Server instances, and each vCenter Server can assign up to 64,000 unique MAC addresses.

3-22 Prefix-Based MAC Address Allocation

You can use prefix-based allocation to specify an OUI other than the default, or to introduce Locally Administered MAC Addresses (LAA) for a larger address space.

You might use prefix-based MAC address allocation for the following reasons:


- To overcome the MAC allocation limitations of VMware OUI
- To customize MAC addresses



3-23 Range-Based MAC Address Allocation

You can use range-based allocation to include or exclude ranges of LAA:

- You specify one or more ranges using a starting and ending MAC address.
- You can specify multiple ranges.
- vCenter Server checks for MAC address conflicts within its ranges.
- You must provide different instances of vCenter Server with ranges that do not overlap.



3-24 Selecting Prefix-Based MAC Address Allocation

You can select prefix-based MAC address allocation to provide unique addresses in large-scale deployments.

To configure prefix-based MAC address allocation, use the vSphere Client to connect the vCenter Server instance, select the **Configure** tab, expand **Settings**, and select **Advanced Settings**.

Key	Value
config.vpxd.macAllocScheme.prefixScheme.prefix	005026
config.vpxd.macAllocScheme.prefixScheme.prefixLength23	


Prefix-based MAC address allocation overcomes the limits of the default VMware allocation to provide unique addresses in large-scale deployments.

For more information about prefix-based MAC address allocation and management, see *vSphere Networking* at <https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.networking.doc/GUID-35B40B0B-0C13-43B2-BC85-18C9C91BE2D4.html>.

3-25 vSphere Client Developer Center

The vSphere Client Developer Center provides tools for managing API structure and capturing user actions to translate them into executable code:


- Using the API Explorer, you can browse and invoke vSphere REST APIs that are supported by the system. And you can access information and context related to the API calls.
- Code Capture records user actions and translates them into executable code, which can be downloaded and copied into a script.



3-26 Code Capture Enhancements (1)

In vSphere 7, the Code Capture feature on the vSphere Client Developer Center page is enhanced to provide translation of new languages (in addition to PowerCLI):

- Python
- vRealize Operations JavaScript



Code Capture is an existing function of the vSphere Client Developer Center. It was introduced in vSphere 6.7 U2.

With the Code Capture feature, you can see the code that is run when the user performs a task in the vSphere Client. Users can use the code in their own scripts to automate some vSphere tasks.

You do not require any special privileges to use Code Capture. However, you must have the correct privileges to perform the tasks to be recorded.

3-27 Code Capture Enhancements (2)

In vSphere 7, the Code Capture feature can record virtual API calls that are performed in the vSphere Client for the following tasks:

- Content library operations
- vCenter Server High Availability operations
- Creation of users and roles

Code Capture does not automatically log information to a log file when a recording starts or stops.

3-28 vSphere Client Networking Enhancements

In vSphere 7, you can use NSX-T Data Center logical switches in the vSphere Client.

With NSX-T Data Center logical switches, also known as NSX port groups, you can perform the following tasks:

- Bind NSX port groups to VMkernel adapters.
- Assign VM network adapters to NSX port groups.
- Migrate VMs to and from NSX port groups.

3-29 Overview of NSX-T Data Center with vCenter Server

In vSphere versions earlier than 7, NSX-T Data Center is decoupled from vCenter Server to avoid scalability limitations. For this reason, in the vSphere Client, NSX-T Data Center networks have the following characteristics:

- Opaque, read-only networks
- Difficult to differentiate from other networks


An opaque network is created and managed by a separate entity outside of vSphere. For example, logical networks that are created and managed by NSX-T Data Center appear in vCenter Server as opaque networks of the type `nsx.LogicalSwitch`.

You can select an opaque network to use with a VM network adapter. To manage an opaque network, use the management tools associated with the opaque network, such as NSX Manager or the NSX API management tools.

In vSphere 7, NSX-T Data Center logical switches behave as transparent virtual distributed port groups.

3-30 NSX-T Data Center Architecture

A single NSX-T Data Center logical switch can span multiple distributed switches.




Multiple NSX-T Data Center logical switches can coexist in the same vSphere Distributed Switch instance.


Also, one NSX-T Data Center logical switch can be present in multiple vSphere Distributed Switch instances at the same time.

3-31 Identifying NSX Port Groups

In the vSphere Client, you can easily distinguish NSX port groups from traditional distributed port groups.



Summary for a Distributed Port Group



Summary for an NSX Port Group

vCenter Server requires that networking interfaces use unique names. For example, two or more distributed port groups with the same name cannot exist in the same system.

However, NSX allows multiple logical switches to use the same name. The same NSX port group can span multiple distributed switches.

To avoid confusion between seemingly identical NSX port groups in the vSphere Client, additional properties are provided for NSX port groups to differentiate them from virtual distributed port groups:

- Segment ID: Equivalent to a logical switch, the segment ID provides network connections for virtual machines. It encapsulates the VM traffic and carries it over the physical IP network. A segment ID can span multiple virtual distributed switches.
- Transport zone: A transport zone can span one or more clusters. Transport zones dictate which hosts and, therefore, which VMs can participate in the use of a specific network. If two VMs are attached to switches that are in different transport zones, the VMs cannot communicate with each other.

The icon that represents NSX port groups is different from the icon for virtual distributed port groups.

For more information about transport zones, see *NSX-T Data Center Installation Guide* at <https://docs.vmware.com/en/VMware-NSX-T-Data-Center/2.3/com.vmware.nsx.install.doc/GUID-3E0C4CEC-D593-4395-84C4-150CD6285963.html>.

When an NSX-T Data Center logical switch is on more than one VDS in the same environment, the new NSX port group properties in the vSphere Client are used to identify them.

3-32 Lab 1: Reconfiguring the Primary Network Identifier

Reconfigure the primary network identifier (PNID):

1. Change the FQDN and IP (PNID) Address of vCenter Server Appliance

3-33 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize how to create and apply vCenter Server profiles
- Identify the scaling and performance improvements for vCenter Server Appliance
- Describe the support provided for dynamic DNS
- Recognize how to change the IP address or FQDN of vCenter Server
- Identify enhancements to the vCenter Server file-based backup
- Identify enhancements to tools provided by the vSphere Client Developer Center
- Describe prefix-based MAC address allocation
- Identify NSX-T Data Center logical switches and their properties in the vSphere Client
- Use virtual distributed port groups to connect to NSX-T Data Center logical switches

3-34 **Lesson 2: Deploying and Managing Templates in the Content Library**

3-35 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe the VM template in-place update
- Describe version control and management of VM templates in vSphere content libraries
- Identify use cases for the VM template in-place update

3-36 VM Template In-Place Update Use Cases

Before vSphere 7, you manually maintain VM templates, for example, when applying guest operating system updates and performing security patching or other template maintenance tasks:

1. Convert the VM template in the content library to a VM.
2. Take a snapshot in case a rollback is required.
3. Perform updates to the guest operating system or other VM maintenance operations.
4. Convert the VM to a VM template.
5. Copy the updated VM template to the content library.
6. Remove the old VM template from the content library and the worker VM where changes were made.

The new VM template in-place update feature simplifies this procedure, reducing the number of steps and introducing version control with traceability.

3-37 VM Template In-Place Update

The VM template in-place update improves VM template management through the content library.

When a VM template resides in the content library, you can work with it in the following ways:

- Deploy VMs while simultaneously patching the VM template.
- Get a history of changes made to the VM template.
- Access two copies of the VM template, the previous and current versions.
- Roll back the VM template.
- Publish templates and synchronize them between publisher and subscriber libraries.

You use linked clone technology in VM deployments to clone a VM template on checkout. This VM clone is used for changes and updates, whereas the original VM template remains available in the content library to deploy VMs. When the VM template changes to the cloned VM are complete, the VM clone is remerged with the VM template, the VM is destroyed, and the VMTX template is updated.

Patching: Customers can power on the checked-out VM and patch it.

VM template history: On the **Versioning** tab, you can view the history of changes made to a VM. The tab shows the metadata for each change made to a VMTX template. It captures time and date, user account information, and any notes that you might enter during the check-in. This metadata remains in the vCenter Server database for the lifetime of the VMTX template.


VM template rollbacks: Rollbacks reverse changes made to a VM through a snapshot process. A button on the UI facilitates rollbacks.

Most VM operations that the customer can perform are available with the VM template, with the following exceptions:

- Cloning the checked-out VM as a new VM does not copy the item ID to the new VM.
- Migrating the checked-out VM to a different vCenter Server instance is not allowed.
- Converting the checked-out VM to a regular VM or template is not allowed.

3-38 Content Library Interface

The content library interface provides several panes for managing templates.



3-39 Registered VM Template (1)

The **Summary** tab provides information about the template, including where the template is managed.

The screenshot displays the vSphere Web Client interface for a registered VM template. The top navigation bar shows the template name 'sa-production-template-01' and an 'ACTIONS' dropdown menu. Below this, a tabbed interface includes 'Summary' (selected), 'Monitor', 'Configure', 'Permissions', 'Datastores', 'Versioning', and 'Updates'. The 'Summary' tab contains a list of properties: Guest OS (VMware Photon OS (64-bit)), Compatibility (ESXi 7.0 and later (VM version 17)), VMware Tools (Not running, not installed), DNS Name, IP Addresses, Host (sa-esxi-01.vclass.local), and Managed By (sa-content-library-01). The 'Managed By' field is highlighted with a red rectangle. Below the properties is a 'Versioning' section with a sub-header 'sa-production-template-01' and a timestamp '11/21/2019, 12:00:06 PM' by 'VSPHERE.LOCAL\Administrator'. A message box states 'No notes available.' and a button labeled 'CHECK OUT VM FROM THIS TEMPLATE' is visible.

sa-production-template-01 | ACTIONS ▾

Summary | Monitor | Configure | Permissions | Datastores | Versioning | Updates

Guest OS: VMware Photon OS (64-bit)
Compatibility: ESXi 7.0 and later (VM version 17)
VMware Tools: Not running, not installed
[More info](#)

DNS Name:
IP Addresses:
Host: sa-esxi-01.vclass.local
Managed By: sa-content-library-01

Versioning

sa-production-template-01
11/21/2019, 12:00:06 PM
VSPHERE.LOCAL\Administrator

No notes available.

CHECK OUT VM FROM THIS TEMPLATE ⓘ

After you perform a checkout and check-in on the VM template, the version information is populated on the VM template Summary page.

3-40 Registered VM Template (2)

On the **Versioning** tab, you can find check-in and checkout information.

The screenshot displays the VMware vSphere Web Client interface for a registered VM template named "sa-production-template-01 (3)". The "Versioning" tab is selected and highlighted with a red box. The interface shows a list of versions on the left and a details panel on the right.

sa-production-template-01 (3) | ACTIONS ▾

Summary Monitor Configure Permissions Datastores **Versioning** Updates

sa-production-template-01 (3)
11/21/2019, 12:14:19 PM
VSPHERE.LOCAL\Administrator

sa-production-template-01 (2)
11/21/2019, 12:00:06 PM
VSPHERE.LOCAL\Administrator

Template updated with latest patches.

CHECK OUT VM FROM THIS TEMPLATE

No notes available. ...

3-41 Template Versioning

You can use the checkout and check-in functionalities to manage multiple versions of templates and to track changes.

The screenshot displays the 'sa-production-template-01 (3)' versioning page in vSphere. The interface includes a top navigation bar with tabs: Summary, Monitor, Configure, Permissions, Datastores, Versioning (selected), and Updates. A vertical timeline on the left shows three versions of the template, each with a document icon and a red dot indicating its status. The first version, 'sa-production-template-01_u3', is checked out and has a 'CHECK IN VM TO TEMPLATE' button. The second version, 'sa-production-template-01 (3)', is the current template and has a 'CHECK OUT VM FROM THIS TEMPLATE' button. The third version, 'sa-production-template-01 (2)', has no notes available. The right side of the interface contains three panels corresponding to these versions, each with a message and a button or status indicator.

Version	Timestamp	User	Status	Action	Notes
sa-production-template-01_u3	11/25/2019, 8:45:20 AM	VSPHERE.LOCAL\Administrator	Checked Out	CHECK IN VM TO TEMPLATE	Navigate to the checked out VM to make updates.
sa-production-template-01 (3)	11/21/2019, 12:14:19 PM	VSPHERE.LOCAL\Administrator	Current	CHECK OUT VM FROM THIS TEMPLATE	Template updated with latest patches.
sa-production-template-01 (2)	11/21/2019, 12:00:06 PM	VSPHERE.LOCAL\Administrator	Archived	-	No notes available.

You can check out the VM template to make hardware or software updates.

You can have multiple versions of the template.

You can add notes to each version of the template to track the changes.

The version summary shows whether the VM template is checked out.

3-42 Discarding Versions

After you check out a template version, you can check it in or discard it.

The screenshot displays the VMware vSphere Web Client interface for the 'sa-production-template-01 (3)' template. The 'Versioning' tab is active, showing a list of versions on the left and a main content area on the right. The versions listed are:

- sa-production-template-01_u3**: 11/25/2019, 8:45:20 AM, VSPHERE.LOCAL\Administrator
- sa-production-template-01 (3)**: 11/21/2019, 12:14:19 PM, VSPHERE.LOCAL\Administrator
- sa-production-template-01 (2)**: 11/21/2019, 12:00:06 PM, VSPHERE.LOCAL\Administrator

The main content area shows the following actions and messages:

- Navigation instruction:** "Navigate to the checked out VM to make updates." with a "CHECK IN VM TO TEMPLATE" button.
- Template update message:** "Template updated with latest patches." with a "CHECK OUT VM FROM THIS TEMPLATE" button.
- No notes available:** "No notes available."

Red boxes highlight the "CHECK IN VM TO TEMPLATE" button and the "Discard Checked Out VM" button.

3-43 Registered VM Template with Updates

You can update or change a VM template.

The screenshot displays the VMware vSphere Web Client interface for a VM template named "sa-production-template-01 (4)". The "Versioning" tab is selected, showing a list of versions and their details.

Version	Created By	Created At	Notes
sa-production-template-01 (4)	VSPHERE.LOCAL\Administrator	11/25/2019, 9:35:45 AM	Updated VMware Tools to 10.3.5 -TC
sa-production-template-01 (3)	VSPHERE.LOCAL\Administrator	11/21/2019, 12:14:19 PM	Template updated with latest patches.
		11/21/2019, 12:00:06 PM	

The "sa-production-template-01 (4)" version is highlighted with a red box. The "Updated VMware Tools to 10.3.5 -TC" note is also highlighted. A "CHECK OUT VM FROM THIS TEMPLATE" button is visible next to the first version. A context menu is open for the second version, showing "Revert to This Version" and "Delete Version" options, which are also highlighted with a red box.


The versioning information highlights the following details:

- Who made the updates
- When the updates were made
- Notes or information added during the updates

The quality of the notes on the **Versioning** tab can vary, depending on who enters them. Whatever the quality, the notes are captured and retained until the VM template is deleted.

3-44 Unregistered VM Template

The **Versioning** tab does not provide version information about VM templates that are not registered with the content library.



To add a VM template into content library management, follow the directions provided in the message on the **Versioning** tab.

After the VM template is under content library management, it becomes a VMTX template.

3-45 Lab 2: Deploying and Managing Templates in the Content Library

Manage template versioning in the vCenter Server content library:

1. Check Out a VM Template in the vCenter Server Content Library
2. Make Changes to a VM Template
3. Check In a VM Template to the vCenter Server Content Library
4. Revert to a Previous Version of a VM Template

3-46 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe the VM template in-place update
- Describe version control and management of VM templates in vSphere content libraries
- Identify use cases for the VM template in-place update

3-47 **Lesson 3: vCenter Server Identity Federation**

3-48 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe identity federation
- Recognize use cases for identity federation
- Describe the architecture of identity federation
- Configure identity federation

3-49 ADFS as an Identity Provider

Active Directory Federation Services (ADFS) is a software component to provide users with single sign-on access to systems and applications that are located across organizational boundaries:

- ADFS uses claims-based access control authorization to maintain application security and implement federated identity.
- SAML tokens, which are a collection of XML-formatted security information describing users, are created and consumed during a federated access request.
- The identity provider (IDP) is the federation service that creates security tokens for users.

3-50 About vCenter Server IDP Federation for ADFS

vCenter Server 7.0 supports identity provider (IDP) federation for Active Directory Federation Services (ADFS).

In a vCenter Server identity provider federation configuration, the external IDP interacts with the identity source on behalf of vCenter Server.

In this federation authentication setup, vCenter Server does not perform the following functions:

- Handling the user's credentials
- Creating the initial identity token
- Managing or updating user and group information in the identity source

3-51 vCenter Server IDP Federation Use Cases

You can use an external IDP to replace vCenter Server as the IDP and enable single sign-on (SSO) with federated infrastructure and applications.

You can use the vCenter Server IDP federation in the following scenarios:

- When using SSO with the existing ADFS-integrated infrastructure and applications
- For improving data center security, because vCenter Server never handles the user's credentials
- When using the authentication mechanisms that are supported by the external IDP

3-52 vCenter Server IDP Federation Components

The vCenter Server 7.0 IDP federation configuration includes the following vSphere and third-party components:

- vCenter Server
- An IDP service that is configured on vCenter Server
- An ADFS server and associated Microsoft AD domain
- An ADFS application group
- AD groups and users that map to vCenter Server groups and users

3-53 ADFS Application Groups

In ADFS, OAuth services are called application groups:

- Each client request must be associated with an application group.
- When creating application groups, you provide details about the servicing application.

3-54 Prerequisites for vCenter Server IDP Federation

vCenter Server supports ADFS as an external IDP.

ADFS requirements are as follows:

- ADFS for Windows Server 2016 R2 or greater must be deployed.
- An application group for vCenter Server must be created in ADFS as part of the configuration process.
- The ADFS server certificate should be signed by a trusted certificate authority.

vCenter Server supports only one external identity provider (IDP) (one ADFS source) and the localos or vsphere.local identity source.

You cannot use multiple external IDPs.

The vCenter Server IDP federation uses OpenID Connect (OIDC) for user login to vCenter Server.

3-55 Configuring vCenter Server IDP Federation in vSphere (1)

The vCenter Server IDP federation is not enabled by default. After installing or upgrading to vSphere 7, you must manually configure it.


The high-level steps to configure IDP federation are as follows:

1. Obtain the redirect URIs from vCenter Server.
2. Create an ADFS application group and configure it for vCenter Server.
3. Modify the IDP configuration on the vCenter Server.
4. Configure group memberships or privileges through global or object permissions in vCenter Server for ADFS authorization.
5. Verify login of an AD user to vCenter Server.

For more information about configuring vCenter Server IDP federation for vCenter Server, see VMware knowledge base article at <https://kb.vmware.com/s/article/78029>.


3-56 Configuring vCenter Server IDP Federation in vSphere (2)

In the vSphere Client, you obtain the redirect URIs from vCenter Server.



3-57 Configuring vCenter Server IDP Federation in vSphere (3)

In the vSphere Client, you can change the vCenter Server identity provider federation.



The screenshot shows the vSphere Client interface. The left sidebar contains the navigation menu with categories like Administration, Access Control, Licensing, Solutions, Deployment, Support, Certificates, and Single Sign On. The main content area is titled 'Configuration' and has tabs for 'Identity Provider', 'Local Accounts', and 'Login Message'. The 'Identity Provider' tab is active. In the top right corner of this tab, there is a button labeled 'CHANGE IDENTITY PROVIDER' which is highlighted with a red rectangular box. Below this, there is a section for 'Identity Sources' with a table listing the configured identity providers.

Name	Server URL	Type	Domain	Alias
--	--	System Domain	vsphere.local	--
--	--	Local OS (Default)	localos	--

2 items

3-58 Configuring vCenter Server IDP Federation in vSphere (4)

In the vSphere Client, you provide application group information.

The screenshot shows the vSphere Client interface with the 'Configure Main Identity Provider' dialog box open. The dialog box has a sidebar with four steps: 1. Identity Provider, 2. ADFS server (selected), 3. Users and Groups, and 4. Review. The main area of the dialog box is titled 'ADFS server' and contains three input fields, each with a red border and a red 'X' icon to its left:

- Client Identifier *
- Shared secret *
- OpenID Address *

At the bottom right of the dialog box, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.


3-59 Enabling ADFS on an Existing Enhanced Linked Mode Configuration

To enable ADFS on an existing Enhanced Linked Mode configuration:

1. Deploy a minimum of two vCenter Server nodes in an Enhanced Linked Mode configuration.
2. Configure ADFS on one of the linked vCenter Server nodes.
The ADFS configuration is replicated to the other vCenter Server node.
3. Add all redirect URIs for both vCenter Server nodes to the configured OAuth Application Group in ADFS.

3-60 Login Workflow of the vCenter Server IDP Federation for ADFS

vCenter Server, ADFS, and AD interact during the login workflow for vCenter Server IDP federation for ADFS.



vCenter Server, ADFS, and AD interact during the login workflow as follows:

1. The user logs in to vCenter Server on the usual vCenter Server landing page.
2. vCenter Server redirects the authentication request to ADFS.
3. If needed, ADFS prompts the user to log in with Active Directory credentials.
4. ADFS authenticates the user with Active Directory.
5. ADFS issues a security token with group information from Active Directory.
6. vCenter Server uses the token to log in the user.

If the external IDP is unreachable, the login process returns to the vCenter Server landing page, showing an appropriate information message. Administrative users can still log in using their local accounts (localos or vsphere.local identity sources).

3-61 Lab 3: Configuring Identity Federation to Use Windows ADFS

Configure the vCenter Server identity provider federation with Microsoft ADFS:

1. Configure vCenter Server Identity Provider Federation
2. Log In to vCenter Server Using an AD Account

3-62 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe identity federation
- Recognize use cases for identity federation
- Describe the architecture of identity federation
- Configure identity federation

3-63 **Lesson 4: vCenter Certificate Enhancements**

3-64 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize use cases for certificate management through the vSphere Client
- Describe how to manage certificates with the vSphere Client

3-65 vSphere Security Certificates

vSphere components use SSL to communicate securely with each other and with ESXi.

SSL communications ensure data confidentiality and integrity. Data is protected and cannot be modified in transit without detection.

In vSphere 6.x and later, VMware CA provisions each ESXi host and each vCenter Server service with a certificate that is signed by VMware CA by default.

You can replace the existing certificates with new VMware CA signed certificates, make VMware CA a subordinate certificate authority, or replace all certificates with custom certificates.


3-66 Use Cases for Certificate Management

You can use certificate management through the vSphere Client to perform several tasks:

- Renew a self-signed vCenter Server certificate.
- Replace a self-signed vCenter Server certificate.
- Generate a certificate signing request (CSR) for use with a third-party CA.
- Replace a VMware CA issued self-signed certificate with a CA-signed certificate.
- Import a third-party CA-signed certificate from a third-party-generated CSR.

3-67 Managing Certificates (1)

To manage certificates, select **Menu > Administration**.




3-68 Managing Certificates (2)

Under Administration in the navigation pane, select **Certificate Management**.

The Certificate Management page provides the following options:

- View details of and manage the vCenter Server machine SSL certificate.
- View details of and add trusted root certificates.




In the vSphere 7 UI, solution user certificates are hidden. All services can use solution user certificates as of the vSphere 7 release.

3-69 Managing Certificates (3)

The options in the **ACTIONS** menu for the machine SSL certificate are as follows:

- Renew the certificate using VMware CA for the self-signed machine certificate.
- Replace self-signed certificates, or install third-party CA-signed certificates.
- Generate a CSR for use with a third-party CA.



Custom certificates can be overwritten using the **Import and Replace Certificate** option.

3-70 Viewing Certificate Details

When you select **View Details** on the vCenter Server machine SSL certificate, you can view the properties of the certificate, including that time frame that the certificate is valid for.

[< BACK TO CERTIFICATE MANAGEMENT](#)

MACHINE_CERT

sa-vcsa-01.vclass.local

Certificate Information

Common name	sa-vcsa-01.vclass.local
Issued by	sa-vcsa-01
Status	Valid
Valid from	10/1/19, 10:29 AM
Valid until	9/30/21, 10:29 AM
Signature Algorithm	SHA256withRSA
Thumbprint	CB62CF3F2262FE06F95838CC5F6C9111B0D3F2FB
Organization	VMware
Organizational Unit	VMware Engineering
Locality	Palo Alto
State/Province	California
Country	US

Issuer Information


Issuer Name	sa-vcsa-01
Organization	sa-vcsa-01.vclass.local
Organizational Unit	VMware Engineering
State/Province	California
Country	US
Serial Number	f27cade5f43cff65
Version	3

If VMware issues the certificate, it is probably a VMware CA issued self-signed certificate. In this situation, VMware CA is your certificate authority.

3-71 Renewing Certificates

To renew a vCenter Server VMware CA issued self-signed certificate, select **RENEW**. VMware CA renews the self-signed certificate.

This action causes vCenter Server services to restart, and vCenter Server goes offline for some time.




vCenter Server services go offline during restart. vCenter Server dependent processes are unavailable during this time. A scheduled maintenance window is suggested for certificate replacement operations.

3-72 Importing and Replacing (1)

To import or replace a vCenter Server VMware CA signed machine certificate, select **Import or Replace Certificate**. VMware CA replaces the self-signed certificate.

This action causes vCenter Server services to restart, and vCenter Server goes offline for some time.



The screenshot shows a 'Replace Certificate' dialog box. On the left, a sidebar lists two steps: '1 Choose appropriate option' (highlighted) and '2 Replace with VMCA'. The main area is titled 'Choose type of certificate to replace' and contains a yellow warning banner stating: 'vCenter server services will be automatically restarted after successful replacement of the machine SSL certificate.' Below the banner, it says 'Please choose the most appropriate option for you' and lists three radio button options: 'Replace with VMCA certificate' (selected), 'Replace with certificate generated from vCenter server', and 'Replace with external CA certificate(requires private key)'.

vCenter Server services go offline during restart. vCenter Server dependent processes are unavailable during this time. A scheduled maintenance window is suggested for certificate replacement operations.

3-73 Importing and Replacing (2)

To import or replace a vCenter Server self-signed machine certificate whose CSR was created by VMware CA, select **Import or Replace Certificate**. A third-party CA-signed certificate and key is used to update vCenter Server.

Replace Certificate

1 Choose appropriate option

2 Replace with signed certificate

Choose type of certificate to replace

vCenter server services will be automatically restarted after successful replacement of the machine SSL certificate.

Please choose the most appropriate option for you

☐ Replace with VMCA certificate

☒ Replace with certificate generated from vCenter server

☐ Replace with external CA certificate(requires private key)

vCenter Server retains the private key from when the CSR is generated.

3-74 Importing and Replacing (3)

To import or replace the current vCenter Server certificate with a third-party-issued certificate and key combination, select **Import or Replace Certificate**. You import a third-party CA-signed certificate whose CSR is generated externally.

Replace Certificate

1 Choose appropriate option

2 Replace with signed certificate

Choose type of certificate to replace

⚠ vCenter server services will be automatically restarted after successful replacement of the machine SSL certificate.

Please choose the most appropriate option for you

☐ Replace with VMCA certificate

☐ Replace with certificate generated from vCenter server

☒ Replace with external CA certificate(requires private key)

3-75 Generating a Certificate Signing Request

To generate a CSR, you must enter information, such as common name, organization, organizational unit, and country.

The screenshot shows a web-based wizard titled "Generate CSR". On the left, a sidebar contains two steps: "1 Enter Info" (highlighted) and "2 Generate CSR". The main area is titled "Enter Info" and contains several input fields. The "Common name" field is filled with "sa-vcsa-01.vclass.local". The "Organization" field is filled with "VMware Inc.". The "Organizational Unit" field is filled with "VMware". The "Country" field is a dropdown menu set to "United States". The "State/Province" field is filled with "CA". The "Locality" field is filled with "Palo Alto". The "Email Address" field is filled with "noreply@vmware.com". The "Host" field is filled with "sa-vcsa-01.vclass.local". The "Subject Alternative Name (Optional)" field is a text area with placeholder text "Enter optional IP addresses or FQDN separated by a comma". The "Key Size" field is a dropdown menu set to "2048". At the bottom right, there are two buttons: "CANCEL" and "NEXT".

Field	Value
Common name	sa-vcsa-01.vclass.local
Organization	VMware Inc.
Organizational Unit	VMware
Country	United States
State/Province	CA
Locality	Palo Alto
Email Address	noreply@vmware.com
Host	sa-vcsa-01.vclass.local
Subject Alternative Name (Optional)	
Key Size	2048

Some wizard information is required to complete this procedure. You usually take this step in preparation for implementing a vCenter Server certificate that is not generated by VMware CA.

Generating a CSR does not take vCenter Server services offline. You can perform this action anytime. You do not require a maintenance window.

3-76 Lab 4: Managing vCenter Server Certificates

Use the vSphere Client to replace and renew VMware CA signed machine SSL certificates and to generate a certificate signing request:

1. Renew a VMware CA Signed Machine SSL Certificate
2. Replace the VMware CA Signed Machine SSL Certificate with a Custom Certificate
3. Generate a CSR for a Machine SSL Certificate

3-77 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize use cases for certificate management through the vSphere Client
- Describe how to manage certificates with the vSphere Client

3-78 **Lesson 5: VMware Skyline Overview**

3-79 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe the benefits and capabilities of VMware Skyline
- Explain how VMware Skyline works
- Identify the types of health information provided by Skyline Health
- Recognize uses for Skyline Advisor

3-80 About VMware Skyline

VMware Skyline is a support technology that provides predictive analysis and proactive recommendations to help avoid problems. VMware Skyline provides the following benefits:

- Issue avoidance:
 - Proactively identifies potential issues based on environment-specific configuration, details, and usage.
 - Resolves issues before they occur, improving environment reliability and stability.
- Shortens time to resolution:
 - Environment-specific, data-driven analytics accelerate problem resolution.
- Personalized recommendations:
 - Resolution is specific to your environment.
- No additional cost:
 - You receive additional value with your current support subscription (Basic, Production, or Premier support).

VMware Skyline shortens the time it takes to resolve a problem so that you can get back to business quickly. The VMware Technical Support engineers can use VMware Skyline to view your environment's configuration and the specific, data-driven analytics to help speed up problem resolution.

3-81 VMware Skyline Family

The VMware Skyline family includes Skyline Health and Skyline Advisor.

Skyline Health

All VMware Customers

Key capabilities:

- vSphere and vSAN findings
 - Available in the vSphere Client
 - Supports vSphere 6.7 and later
-

Skyline Advisor

Production and Premier Support Customers

Key capabilities:

- Supports vSphere, vSAN, NSX for vSphere, vRealize Operations Manager, and VMware Horizon
 - Supports vSphere 5.5 and later
 - Identification of VMware Validated Design, VxRail, and VMware Cloud Foundation deployments
 - Automates log transfers with Log Assist
 - Uses cloud-based ID and access
-

Premier Support Customers

Key capabilities:

- Advanced findings and reporting
 - Tailored remediation plans
-

With Basic Support, you can access Skyline findings and recommendations for vSphere and vSAN by using Skyline Health in the vSphere Client (version 6.7 and later).

With Production or Premier Support, you must use Skyline Advisor and the full functionality of Skyline (including Log Assist).

With Premier Support, you receive additional Skyline features that are not available with Production Support, for example:

- An advanced set of findings and recommendations
- Scheduled and custom operational summary reports that provide an overview of the proactive findings and recommendations
- All additional benefits of Premier Support, including a designated support team, direct access to senior-level technical support engineers, assistance with multivendor troubleshooting and onsite support services, such as Mission Critical Support (MCS), Healthcare Critical Support (HCS), and Carrier Grade Support (CGS)

Skyline supports vSphere, NSX for vSphere, vSAN, VMware Horizon, and vRealize Operations Manager. A Skyline management pack for vRealize Operations Manager is also available. If you install this management pack, you can see Skyline proactive findings and recommendations within the vRealize Operations Manager client.

The identification and tagging of VxRail and VMware Validated Design deployments help you and VMware Technical Support to better understand and support multiproduct solutions.

Skyline identifies all ESXi 5.5 objects within a vCenter Server instance and provides additional information in VMware knowledge base article 51491 at <https://kb.vmware.com/s/article/51491>. This article details the end of general support for vSphere 5.5.

For versions of vSphere, vSAN, NSX for vSphere, VMware Horizon, and vRealize Operations Manager that are supported by Skyline, see the Skyline Collector Release Notes at <https://docs.vmware.com>.

3-82 Skyline Health for vSphere (1)

Skyline Health provides findings for your vSphere environment:

- Findings are based on telemetry data collected from your vSphere environment.
- Data is analyzed and then provided in the vSphere Client.
- Discovered findings can be related to stability and incorrect configurations in vSphere.


In vSphere 6.7 environments, you can use vSphere health and vSAN health services to monitor your vSphere 6.7 environments for configuration consistency, best practice adherence, vulnerability discovery, and proactive analytics. vSphere health and vSAN health services are part of VMware Skyline Health.

3-83 Skyline Health for vSphere (2)

You access Skyline Health with the vSphere Client.

For Skyline Health to function correctly:

- You must join the Customer Experience Improvement Program (CEIP).
- Your vCenter Server instance must have Internet connectivity.



To view vSphere health information, click the **Monitor** tab for a vCenter Server Appliance instance and select **Skyline Health**.

To view vSAN health information, click the **Monitor** tab for a vSAN cluster or vSAN enabled ESXi host and select **Skyline Health**.

3-84 Skyline Health for vSAN

For vSAN, Skyline Health findings have the following characteristics:


- Findings are based on data from thousands of vSAN deployments.
- Findings draw on VMware knowledge base articles and best practices.
- Discovered findings can be related to misconfigurations and the general health and well-being of hosts in a vSAN cluster.

vSAN Health includes in-product health checks that do not require you to join CEIP. Nor do they require internet connectivity.

3-85 Skyline Components: Collector and Advisor

Skyline Advisor and Skyline Collector work together to provide proactive support:

- Skyline Collector is a standalone virtual appliance that collects diagnostic data about your vSphere, NSX for vSphere, vSAN, VMware Horizon, and vRealize Operations Manager environments.
- Skyline Advisor is a web-based portal where you can view proactive findings and recommendations on demand.




You can access Skyline Advisor directly using this link: <https://skyline.vmware.com/advisor>.

You can also access Skyline Advisor from the list of services in the console of your Cloud Services organization.

3-86 VMware Skyline Networking Design

This sample networking design has the following features:

- Skyline Collector instance has routed access to the management network.
- An upstream router provides external connectivity.
- An HTTP proxy server can be used for outbound connectivity.



3-87 How VMware Skyline Works

1. Install and configure:
 - You install the Skyline Collector appliance into your environment.
 - You configure Skyline Collector with one or more endpoints (for vCenter Server, NSX Manager, Horizon View, vRealize Operations).
2. Collect:
 - Skyline Collector starts collecting product usage data and capturing patterns, events, trends, configuration, design-compliance, and cross-product information.
3. Analyze:
 - VMware Skyline analyzes the environment details and product usage and compares this data against VMware best practices and the VMware knowledge base.
4. Act:
 - VMware Skyline provides proactive findings and recommendations.
 - You access this information through the Skyline Advisor, a web-based portal.


Skyline Collector is a lightweight appliance that takes 15 to 30 minutes to install. The installation involves adding Skyline Collector to a new or existing Cloud Services organization, associating your support entitlement with the organization, downloading the Skyline Collector appliance from My VMware, and using the vSphere Client to deploy the appliance from an OVF template.

For complete details on how to install and configure Skyline Collector, see VMware Skyline Collector Installation and Configuration Guide at <https://docs.vmware.com/en/VMware-Skyline/index.html>.

Skyline Collector gathers and aggregates product usage information such as configuration, feature, and performance data while listening for changes and events within the customer's environment. This information is encrypted and then sent to VMware for analysis.

3-88 Skyline Advisor

Skyline Advisor is a self-service web application with which you can view proactive findings and recommendations on-demand, within a web browser.




The following information is available within Skyline Advisor.

- Account Details, including total number of vCenter Servers, ESXi hosts and virtual machines being analyzed for potential issues.
- Deployed Skyline Collectors.
- Inventory details for VMware vSphere, NSX-V, Horizon 7, and vRealize Operations Manager.
- Download scheduled and custom Operational Summary Reports (OSR) (Premier Support customers only).
- Streamlined support log bundle upload capabilities with Skyline Log Assist.

3-89 Skyline Advisor: Proactive Findings

Proactive findings are at the heart of VMware Skyline.



On the **Findings & Recommendations** tab, you can find all potential issues in your environment. These issues should be remediated to keep your environment secure, stable, and productive. Findings are ranked by severity, and you can filter by several different options.

3-90 Log Assist

When using VMware Skyline Log Assist, you do not need to upload your log files manually. Support log bundles can be automatically uploaded to VMware Technical Support.

Skyline Advisor

Skyline Demo Premier

Home Account Details Inventory Proactive Findings OSR Library Log Assist

Selected Inventory

- ✓ vcenter01.core.hypervizor...
- esx01-s1.core.hyperv...

Log Assist requires additional privileges on the Skyline service account. [Click here](#) for more details.

Step 1 of 2 - Select Inventory

Inventory

Select inventory to initiate log transfer from below topology view

Last Updated: 04/12/2019 06:32:59 EDT

Search by Object Name and hit enter

Inventory	Version	Build	Add-Ons	Collector Name
> ☐ manager	6.3.2	5672532	-	Skyline_Demo
✓ ☐ vcenter01.core.hypervizor.com	6.0.0	3634793	-	Skyline_Demo
✓ ☐ Cairo Datacenter	-	-	-	Skyline_Demo
✓ ☐ CAI Cluster	-	-	-	Skyline_Demo
☒ esx01-s1.core.hypervizor.com	6.0.0	3620759	vSAN Enabled NSX Enabled	Skyline

NEXT

FEEDBACK

Log Assist helps to streamline the time to resolution from when you open a support request.

VMware Technical Support engineers can request the log file bundles associated with your support request. When you approve their request, the log files are automatically uploaded to the Technical Support team, and they can begin troubleshooting right away.

You can also use Log Assist to initiate the automatic upload yourself. Log Assist eliminates the time-consuming process of manually uploading log files using SSH File Transfer Protocol (SFTP) or My VMware.

3-91 Skyline Features Comparison by Support Level

Features	Basic Support	Production Support	Premier Support
Skyline Health	X	X	X
Basic findings and recommendations	X	X	X
Moderated Skyline community	X	X	X
Skyline Advisor		X	X
Log Assist		X	X
Proactive support for all supported products		X	X
Export of findings, recommendations, and affected objects to a comma-separated file		X	X
Advanced findings and recommendations			X
Custom operational summary reports			X
Guided remediation from VMware Global Support Services			X

3-92 Skyline Privacy and Security Policy

VMware is committed to keeping your data private and secure:

- The Skyline Collector appliance is separate from your environment:
 - Your production infrastructure is not exposed to the Internet.
- Collected data is forwarded to VMware using an encrypted channel:
 - The data is only temporarily stored in the Skyline Collector appliance.
- Data is warehoused in a secure back-end system in the United States and is managed by VMware:
 - Access to your data is restricted to permitted users only.
- Information is limited to data such as host names:
 - Data that is known to identify an individual person is not collected.

Skyline collects data in a private and secure manner. Your production infrastructure is not exposed to the Internet. All collected data is tagged with legally identifiable customer information. No customer data is permanently stored in the Skyline Collector appliance or shared with third parties.

3-93 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe the benefits and capabilities of VMware Skyline
- Explain how VMware Skyline works
- Identify the types of health information provided by Skyline Health
- Recognize uses for Skyline Advisor

3-94 Key Points

- In vSphere 7, you can change the host name and networking information of vCenter Server.
- Version control is available for templates stored in the content library.
- vCenter Server identity federation extends authorization to Microsoft ADFS.
- In vSphere 7, certificate management simplifies the process of certificate replacement.
- In vSphere 7, NSX-T Data Center port groups behave like traditional distributed port groups. After NSX-T Data Center port groups are configured, vSphere administrators can assign VMkernel and VM network adapters to them.
- VMware Skyline is a proactive support technology that provides predictive analysis and recommendations to help you avoid problems.

Questions?

Module 4

ESXi Enhancements

4-2 Importance

ESXi 7.0 includes enhanced and new features related to availability, security, and performance. Using these enhancements and new features, you can plan and design your vSphere-based software-defined data center.

4-3 Module Lessons

1. ESXi Security
2. Hardware Support in ESXi
3. Other ESXi Enhancements

4-4 **Lesson 1: ESXi Security**


4-5 Learner Objectives

After completing this lesson, you should be able to meet the following objective:

- Recognize the enhancements made to ESXi security

4-6 CPU Vulnerability Background

ESXi hosts are exposed to CPU vulnerabilities that are known to affect some Intel processors when vCPUs from different VMs are scheduled on the same core.



The default ESXi host security boundary prevents information sharing between hosts.

For maximum CPU performance:


- Enable hyperthreading in the system BIOS.
- Use hyperthreading if your system supports this technology.

With the host security boundary, malicious guests can obtain the credentials, crypto keys, and secrets that are used by the hypervisor or the VMs.

To learn more about these vulnerabilities, see VMware knowledge base article 55636 at <https://kb.vmware.com/s/article/55636>.

4-7 CPU Vulnerability Mitigation (1)

In vSphere 6.7, a patch offers the Side-Channel-Aware (SCA) CPU Scheduler to prevent information sharing between in-guest processes.




SCA version 1 schedules processes on only one thread for each CPU core, offering the highest security.

To configure the process security boundary:

- Set the advanced option `hyperthreadingMitigation` to `true`.
- Set the advanced option `hyperthreadingMitigationIntraVM` to `true`.

4-8 CPU Vulnerability Mitigation (2)

In vSphere 6.7 U2, SCA version 2 prevents information sharing between the VMs on a host or between a VM and the hypervisor.



This option provides a balance of performance and security for environments where the VM is considered the information security boundary.


Configure the VM security boundary:

- Set the advanced option `hyperthreadingMitigation` to `true`.
- Set the advanced option `hyperthreadingMitigationIntraVM` to `false`.

To learn more about performance in relation to the Side-Channel-Aware (SCA) CPU Scheduler, see "Performance of vSphere 6.7 Scheduling Options" at <https://www.vmware.com/techpapers/2018/scheduler-options-vsphere67u2-perf.html>.

4-9 vSphere vMotion Encryption

ESXi encrypts vSphere vMotion traffic inside the hypervisor using vmkcrypto, a VMkernel module.



You can also use Internet Protocol Security (IPsec) to secure vSphere vMotion traffic:

- IPsec uses third-party software to perform software-based hashing, encryption, and decryption.
- In earlier versions of vSphere, this functionality is limited to IPv6 traffic.

4-10 Internet Protocol Security Policy Updates

In vSphere 7, IPsec uses the vmkcrypto module for cryptographic operations over IPv4 and IPv6 traffic.

Benefits of using vmkcrypto:

- It uses the AES-NI encryption standard for increased speed and efficiency.
- It is fully FIPS 140-2 certified.

4-11 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objective:

- Recognize the enhancements made to ESXi security

4-12 **Lesson 2: Hardware Support in ESXi**

4-13 Learner Objectives

After completing this lesson, you should be able to meet the following objective:

- Identify hardware support updates in ESXi 7.0

4-14 Support for AMD EPYC Processors

vSphere 6.7 U3 and later versions are compatible with the second generation of AMD EPYC processors.

AMD and VMware deliver support for new security and other features of these processors in vSphere.


The jointly developed security capability makes it more difficult for bad actors at all levels to steal data.

4-15 Support for Multiple NVIDIA vGPUs

In vSphere 6.7 U3, you can deploy up to four NVIDIA virtual GPUs (vGPUs) in a single virtual machine.

Benefits:

- Scale graphics and compute application performance.
- Speed up machine learning workflows.



This feature provides NVIDIA GRID with the capability that is available in vSphere DirectPath I/O and vSphere Bitfusion FlexDirect.

4-16 New Device Support

vSphere 7 introduces support for the following devices:

- USB 3.1 devices with 10 Gbps speed using the xHCI controller
- Marvell's 64 G FC HBA: Tested in 32 G mode only
- On-chip NIC in Coffee Lake processor
- Broadcom 100 G NIC with bnxtnet driver
- Solarflare NICs with sfvmk driver
- Marvell's RDMA NIC with qedentv and qedrntv driver
- Broadcom's Tri-Mode (SATA/SAS/NVMe) controller with lsi_mr3 and lsi_msgpt35 driver

4-17 Enhanced Support for vSphere DirectPath I/O Network Devices

With vSphere DirectPath I/O, a guest operating system on a VM can directly access physical PCI and PCIe devices connected to a host.

You can enable vSphere DirectPath I/O passthrough for a PCI network device on a host:

- In earlier versions of vSphere, you must reboot the host to make the PCI network device available for use.
- In vSphere 7, you are not required to reboot the host to enable vSphere DirectPath I/O passthrough.

The screenshot shows the vSphere Client interface with the 'Configure' tab selected. The left sidebar shows the 'Hardware' section expanded, with 'PCI Devices' selected. The main pane displays a table titled 'DirectPath I/O PCI Devices Available to VMs'.

ID	Status	Vendor Name	Device Name
0000:04:00.0	Not Configurable	Pericom Semico...	<class> PCI bridge
0000:05:01.0	Not Configurable	Pericom Semico...	<class> PCI bridge
0000:06:00.0	Available	Realtek Semicon...	RTL8111/8168/8...
0000:05:02.0	Not Configurable	Pericom Semico...	<class> PCI bridge
0000:07:00.0	Available	Realtek Semicon...	RTL8111/8168/8...

Below the table, the details for the selected device '0000:06:00.0' are shown. It states 'This device is available for VMs to use'.

Name	RTL8111/8168/8411 PCI Express Gigabit Ethernet Controller	Vendor Name	Realtek Semiconductor Co., Ltd.
Device ID	8168	Vendor ID	10EC
Subdevice ID	123	Subvendor ID	10EC
Class ID	200		

The 'Bus Location' is also displayed at the bottom.

To configure passthrough, follow these steps:

1. In the vSphere Client, select an ESXi host.
2. On the **Configure** tab, expand **Hardware** and click **PCI Devices**.

3. Select **All PCI devices**.
4. Select a device and click **CONFIGURE PASSTHROUGH**.
5. Click **OK**.

The PCI network device is available for use.

4-18 Driver Updates (1)

In vSphere 6.7 U2, several drivers are updated:

- Solarflare network adapters (sfvmk)
- Broadcom 100 G network adapters (bnxtnet)
- Mellanox 40 G and 100 G
- RoCE Mellanox 40 G and 100 G
- Broadcom lpfc, brcmfcoe, and various lsi
- Cavium qlnativefc
- Cisco nfnic

4-19 Driver Updates (2)

In vSphere 7, several drivers are updated:

- VMware nvme
- Microchip smartpqi
- Marvell qlnativefc
- Broadcom lpfc/brcmfcoe
- Broadcom lsi_msgpt2
- Broadcom lsi_msgpt35
- Broadcom lsi_msgpt3
- Broadcom lsi_mr3
- Intel i40en
- Intel ixgben
- Cisco nenic
- Broadcom bnxtnet


With these native network drivers in vSphere 7, you do not need the legacy VMKLinux driver stack. If servers still use IO devices that depend on VMKLinux APIs, you must upgrade those devices before installing ESXi.

4-20 Receive Side Scaling

Receive side scaling (RSS) is a network driver technology that supports the efficient distribution of network receive-processing across multiple CPUs in multiprocessor systems.

RSS distributes packets among different queues based on flows.

As a result, a single virtual NIC with multiple flows can use parallelism in the physical NIC and in the ESXi stack by using multiple queues.



The main advantages of using RSS are as follows:

- Throughput is higher because more CPUs can be used independently.
- Load balancing is offloaded to the NIC's hardware instead of doing software-based load-balancing.

Not all ESXi drivers support this feature.

4-21 Multi-RSS Support for nvmxnet3 Driver

nvmxnet3 is the native driver for VMXNET3 devices.

With vSphere 6.7 U1, the nvmxnet3 driver can use multi-RSS to improve the network stack performance.

Virtual machines running on ESXi hosts with a VMXNET3 device can benefit from multi-RSS by using multiple queues.

VM Hardware



> CPU	2 CPU(s)
> Memory	 12 GB, 3.84 GB memory active
> Hard disk 1	12 GB
Total hard disks	16 hard disks
▼ Network adapter 1	
Adapter Type	VMXNET 3
MAC Address	00:50:56:9d:02:8f
DirectPath I/O	Inactive
Network	VM Network (connected)
CD/DVD drive 1	Disconnected 
> Video card	4 MB

4-22 Data Center Bridging

Data center bridging (DCB) is a collection of standards that supports a lossless transport over Ethernet and a converged network where all data center applications can run over the same physical infrastructure.

The potential benefits of a converged data center network include:

- Simpler management with only one fabric to deploy, maintain, and upgrade
- Reliable Ethernet transport services for workloads that do not tolerate loss
- Lower costs because you can use common adapters, cables, and switches

4-23 ESXi Native Data Center Bridging

In vSphere versions earlier than 7, DCB is confined to Fibre Channel over Ethernet (FCoE) adapters.

In vSphere 7, native DCB support is enabled in ESXi.

ESXi can take advantage of this feature with any DCB-capable network adapter.

To find out whether your software or hardware DCB adapter is supported, see the VMware Compatibility Guide.

Two new VMK APIs handle the hardware DCB support:

- `vmk_UplinkDCBGetCfg`: Gets the current DCB configuration settings from the uplink
- `vmk_UplinkDCBSetCfg`: Sets the uplink DCB configuration

The DCB configuration is maintained at run time.

4-24 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objective:

- Identify hardware support updates in ESXi 7.0

4-25 **Lesson 3: Other ESXi Enhancements**

4-26 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:


- Recognize layout changes in the ESXi 7.0 system storage partitions
- Describe the enhancements made to ESXi timekeeping options in vSphere 7

4-27 ESXi System Storage

ESXi system storage represents storage space that is used by the system for operations related to software images, system configuration, cache, logs, and core dumps.

The new unified ESX-OSData partition in ESXi 7.0 provides more flexibility for future demands.

After you upgrade to vSphere 7, you cannot roll back to vSphere 6.x because all partitions are rebuilt during the upgrade process.



Boot-bank 1

4-28 Recommendations and Requirements for ESXi System Storage

When installing ESXi 7.0, follow these recommendations:

- For system storage, use a minimum disk space of 32 GB.
- For the ESXi boot device, use high-endurance devices.

When upgrading to ESXi 7.0, you must meet the following requirements:


- Upgrade paths to ESXi 7.0 require at least ESXi 6.5 or later.
- The boot media size must be at least 4 GB.

High-endurance devices include HDD, SSD, NVMe, and iSCSI or Fibre Channel LUNs.

4-29 Precision Time Protocol Support

To ensure precise timekeeping and synchronization between the ESXi host and the other components of the vSphere network, you can use these options:

- Configure the time settings on the ESXi host manually.
- Synchronize the time and date of the host to an NTP server.
- Synchronize the time and date of the host to a PTP server.



The screenshot shows the vSphere ESXi host configuration interface for 'sa-esxi-01.vclass.local'. The 'Configure' tab is active, and the 'Time Configuration' section is selected in the left-hand navigation pane. The main content area displays three configuration sections:

- Manual Time Configuration**: A table with one row showing 'Status' as 'Inactive (Either NTP or PTP client is enabled. To set time manually, disable the active time synchronization client.)'. An 'EDIT...' button is present.
- Network Time Protocol**: A table with three rows: 'NTP Client' (Enabled), 'NTP Service Status' (Running), and 'NTP Servers' (172.20.10.10). An 'EDIT...' button is present.
- Precision Time Protocol**: A table with three rows: 'PTP Client' (Disabled), 'PTP Service Status' (Stopped), and a collapsed 'Network Interface' row. An 'EDIT...' button is present.

Manual Time Configuration	
Status	Inactive (Either NTP or PTP client is enabled. To set time manually, disable the active time synchronization client.)

Network Time Protocol	
NTP Client	Enabled
NTP Service Status	Running
NTP Servers	172.20.10.10

Precision Time Protocol	
PTP Client	Disabled
PTP Service Status	Stopped
> Network Interface	--

4-30 Configuring Precision Time Protocol

To synchronize the time and date of the host to a PTP server:

1. Verify that the NTP client is disabled.
2. In the vSphere Client, navigate to a host.
3. On the **Configure** tab, select **System** > **Time Configuration**.
4. Click the **Edit** button next to Precision Time Protocol.
5. In the Edit Precision Time Protocol dialog box, select the **Enable** check box to synchronize the time and date of the host with a PTP server.
6. Select a network interface from the drop-down menu and click **OK**.

The screenshot shows the vSphere Client interface for host 'sa-esxi-01.vclass.local'. The 'Configure' tab is active, and the 'Time Configuration' section is selected. The 'Edit Precision Time Protocol' dialog box is open, showing the following settings:

- Enable:** ☒ (with a help icon)
- Networking Settings:**
 - Network interface:** vmk0 (selected from a dropdown)
 - IPv4:** 172.20.10.54
 - Subnet mask:** 255.255.255.0

The 'OK' button is highlighted. In the background, the 'Precision Time Protocol' configuration page is visible, showing the following status:

Precision Time Protocol	
PTP Client	Disabled
PTP Service Status	Stopped
> Network Interface	--

4-31 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize layout changes in the ESXi 7.0 system storage partitions
- Describe the enhancements made to ESXi timekeeping options in vSphere 7

4-32 Key Points

- VMware addresses existing hardware vulnerabilities with ESXi enhancements. You must constantly update your hardware firmware to keep it secure.
- ESXi supports additional hardware devices and drivers to improve the availability and performance of diverse vSphere environments.
- The new layout of ESXi system storage improves the flexibility for core operations such as storing system configuration, log files, and core dumps.

Questions?

Module 5

Virtual Machine Enhancements

5-2 Importance

To create and modify virtual machines in vSphere, you must understand the VM settings and features that are available. vSphere 7 introduces a new ESXi compatibility version that adds new virtual devices and features that you can use to enhance the functionality of VMs.

5-3 Learner Objectives

After completing this module, you should be able to meet the following objectives:

- Identify the characteristics of VM compatibility level for ESXi 7.0 and later
- Describe the support provided for new virtual devices in vSphere 7
- Recognize improvements to guest customization
- Migrate VMs with attached devices

5-4 VM Compatibility Version

vSphere 7 introduces the VM compatibility level for ESXi 7.0 and later, adding new functionalities:

- New EVC modes for additional processors
- Virtual SGX for runtime protection of the VM
- A virtual watchdog timer (WDT) to assist clustering software that implements heartbeats
- Selective latency sensitivity
- A precision clock to improve timekeeping in the guest OS

The VM compatibility level for ESXi 7.0 and later also adds support for new devices:

- Local NVMe devices
- USB 3.1 devices
- PCIe storage device hot-plug support

vSphere 7 introduces hardware version 17, which provides native instruction support for the Intel Cascade Lake and AMD Rome processors.

5-5 About Virtual Watchdog Timer

The virtual WDT detects and recovers from OS problems.

If a failure occurs, the WDT attempts to reset or power off the VM.

This feature is based on Microsoft specifications: Watchdog Resource Table (WDRT) and Watchdog Action Table (WDAT).

The WDT is a virtual device that detects and recovers from OS problems. The device's timeout can be initialized with VM power-on, or by the guest OS. The guest OS or an application in the guest OS can reset the device's timeout based on a predefined logic. If the virtual WDT is not reset before its timeout, it resets or powers off the guest OS.

How the guest OS or the application in the guest OS interacts with the virtual WDT differs depending on the guest OS and the application that interacts with the virtual device. Configuring the guest OS or the application that interacts with the virtual WDT is beyond the scope of this module.

5-6 Virtual Watchdog Timer Use Cases

The virtual WDT is useful in the following scenarios:

- With high availability solutions such as Red Hat High Availability and the MS SQL failover cluster
- On VMware Cloud and in hosted environments for implementing custom failover logic to reset or power off VMs

Many high availability solutions require watchdog devices. These high availability solutions can use the virtual WDT. The timer's use in high availability solutions might differ from one vendor to another.

The Red Hat High Availability Add-On is a clustered system that provides reliability, scalability, and availability to critical production services. RHEL HA uses watchdog devices in their clustering solution. It supports watchdog timer devices (hardware devices) and virtualization-emulated watchdogs (virtual devices).

RHEL HA does not support software-emulated watchdogs.

With the introduction of the virtual WDT, RHEL HA clusters can be created on the vSphere infrastructure.

Using the virtual WDT, users can implement their own failover logic to reset or restart a VM.

5-7 Virtual Watchdog Timer Requirements

You must meet the following requirements for the virtual WDT:

- VM must include hardware version 17 or above.
- VM must be powered off before adding the virtual WDT.
- Guest OS support must be provided for the watchdog timer:
 - Windows Server 2003 (WDRT)
 - Windows Server 2008 onward (WDAT)
 - Linux (kernel version 4.9 or above)

Other guest operating systems, such as FreeBSD, MAC OS X, are not supported.

For Linux operating systems with kernel 4.9+, verify that the module `wdat_wdt.ko` is installed.

Check your Linux distribution document for further details.

No additional drivers are required in Windows and the Linux OS.

Linux support for Windows Server 2003 (WDRT) is no longer available.

5-8 Adding Virtual Watchdog Timer to a VM

You can use the vSphere Client to add the virtual WDT to a VM.

The screenshot shows the 'Virtual Hardware' tab in the vSphere Client. The main table lists various hardware components for a VM. A red box highlights the 'ADD NEW DEVICE' button in the top right corner. A dropdown menu is open, showing categories like 'Disks, Drives and Storage', 'Controllers', 'Other Devices', and 'Network'. The 'Watchdog Timer' option under 'Other Devices' is highlighted with a red box. The 'New Watchdog timer' row in the main table is also highlighted in blue.

Virtual Hardware		VM Options
> Memory *	1 GB	
> New Hard disk *	1 GB	
> New SCSI controller *	VMware Paravirtual	
> New Network *	VM Network	
> New CD/DVD Drive *	Client Device	
> Video card *	Specify custom settings	
> Security devices	Not Configured	
VMCI device		
New SATA Controller	New SATA Controller	
✓ New Watchdog timer	☐ Start with BIOS/EFI boot ⓘ	
> Other	Additional Hardware	

Compatibility: ESXi


ADD NEW DEVICE

- Disks, Drives and Storage
 - Hard Disk
 - Existing Hard Disk
 - RDM Disk
 - Host USB Device
 - NVDIMM
 - CD/DVD Drive
- Controllers
 - NVMe Controller
 - SATA Controller
 - SCSI Controller
 - USB Controller
- Other Devices
 - PCI Device
 - Watchdog Timer**
 - Precision Clock
 - Serial Port
- Network
 - Network Adapter

Use the Edit Settings page to add the virtual WDT to a VM.

5-9 Virtual Watchdog Timer Status

You can view the status of the virtual WDT in the vSphere Client.



The screenshot shows the 'VM Hardware' settings in the vSphere Client. The 'Watchdog timer' is highlighted with a red box, indicating its status is 'Running'. Other hardware components listed include CPU, Memory, Hard disk, Network adapter, CD/DVD drive, Video card, and VMCI device.

VM Hardware	
> CPU	1 CPU(s)
> Memory	1 GB, 0 GB memory active
> Hard disk 1	1 GB
> Network adapter 1	VM Network (connected)
CD/DVD drive 1	Disconnected
> Video card	8 MB
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface
Watchdog timer	Running
> Other	Additional Hardware
Compatibility	ESXi 7.0 and later (VM version 17)

[Edit Settings...](#)

To check the virtual WDT state, take these steps:

1. In the navigation pane, select a VM.
2. In the **Summary** tab, verify the watchdog timer state.

5-10 About Precision Clock

vSphere 7 introduces a virtual device called Precision Clock:

- The guest OS uses the Precision Clock time as reference time.
- Precision Clock is not directly involved in guest OS time synchronization.
- Precision Clock presents the ESXi host's system time to the guest OS.
- To provide better time synchronization, Precision Clock bypasses the virtual networking stack and the guest networking stack.
- Precision Clock helps the guest OS to achieve clock accuracy in a 1 millisecond range.

VMware Tools provides a plug-in called Tools TimeSync. This plug-in synchronizes VM time with ESXi time. This tool is not recommended over guest native agents such as ntpd, chrony, and W32Time.

5-11 Precision Clock Use Cases

Precise timekeeping is a requirement for many applications, for example, financial services applications.

Precise time stamps are required on events that track financial transactions.

Precision Clock helps ensure that the guest OS is accurate and that the applications running in the guest have accurate timestamps.

In financial services applications, the time stamps on transactions and records must be precise.

The financial sector must adhere to many regulations to time-stamp records and transactions. These time stamps must be accurate within a 1 second range.


5-12 Precision Clock Requirements

Precision Clock requires ESXi compatibility level for ESXi 7.0. Other requirements include:

- Precision Clock is available only for Linux guests.
- On Linux VMs, only the chrony time sync tool is supported.
- The Linux guest must have the vmwptp driver installed.
- The ESXi host must be synchronized with either the Network Time Protocol (NTP) or the Precision Time Protocol (PTP).

5-13 Adding and Configuring Precision Clock

You add and configure Precision Clock using the vSphere Client.



Precision Clock has the following optional device settings:

- Host system time (Auto): Precision Clock automatically detects the ESXi host's time-sync protocol and presents the time to the guest OS.
- Host system time (NTP): Precision Clock requires that the ESXi host's time-sync protocol is the Network Time Protocol (NTP).
- Host system time (PTP): Precision Clock requires that the ESXi host's time-sync protocol is the Precision Time Protocol (PTP).

For more information about how to set up NTP on the ESXi host, see VMware knowledge base article 57147 at <https://kb.vmware.com/s/article/57147>.

5-14 About SGX

Intel Software Guard Extensions (SGX) technology meets the needs of the trusted computing industry:

- With Intel SGX, software programs can create private memory regions called enclaves.
- Data in enclaves can be accessed only by the intended program.
- The enclave region is isolated from other programs, operating systems, hypervisors, and so on.
- Enclaves are used by the software program to secure data.
- Virtual SGX exposes Intel's SGX to VMs running in a vSphere environment.

Trusted computing refers to various technologies and proposals for resolving computer security problems. The problems are resolved through hardware enhancements and associated software modifications. Several major hardware manufacturers and software vendors, collectively known as the Trusted Computing Group (TCG), are cooperating in this venture and have come up with specific plans.

For more information about the TCG, see <https://trustedcomputinggroup.org>.

Data inside an application can be compromised by rogue software or a bug in the operating system or in the hypervisor.

Programmers can use Intel SGX to create private memory regions called enclaves. Data in enclaves is protected from unauthorized access and modification by rogue software running at higher privilege levels. Intel SGX provides this protection without disrupting the ability of legitimate software to schedule and manage the use of platform resources.

Virtual SGX is implemented as part of the vSphere core virtualization stack. With virtual SGX, applications running on VMs can create their own enclaves.

5-15 SGX Terminology

Enclave Page Cache (EPC) is a static portion in RAM that is used for storing running enclaves.

The launch control configuration determines which enclaves can be started.

Enclave Page Cache (EPC) is a static portion of physical RAM. This portion is allocated at boot time by the BIOS.

EPC stores running enclaves. EPC is a limited resource, typically not more than 93 MB.

The launch control configuration is an architectural component from which the platform can control the enclaves to be started.

5-16 Virtual SGX Use Cases


You can use virtual SGX in several scenarios:

- For applications such as secure remote computing, secure web browsing, and digital rights management
- For VMware virtual appliance components
- For applications that must conceal proprietary algorithms and encryption keys
- Where you do not want cloud service providers to inspect the client's code and data

Many third-party applications can use virtual SGX.

5-17 Intel SGX Architecture

SGX reduces the attack surface of a malicious application.



Applications frequently work with private information such as passwords, encryption keys, bank account numbers, and so on. This private data should be accessed only by the designated recipient. The job of the operating system is to protect such private data from other applications and users.

Operating systems and applications often employ safeguards to protect this private data. Despite these protections, most computer systems are still vulnerable.

Malware with administrative privileges can access all operating system resources, including all applications running on that system. Malware can target an application's protected data to extract encryption keys and secret data directly from memory.

With Intel SGX, applications can create enclaves, which protect the confidentiality and integrity of applications from privileged malware.

Privileged malware is malware that gets privileges by manipulating the underlying system.

5-18 Virtual SGX Requirements

To use virtual SGX, you must meet the following requirements:

- The ESXi host must have Intel Coffee Lake CPUs or later.
- SGX must be enabled in the BIOS of the ESXi host.
- Hyperthreading must be disabled on certain Coffee Lake SKUs.
- Virtual hardware version 17 or later is required.
- VMs must be powered off.
- Supported guest operating systems are:
 - Windows 10
 - Windows Server 2016 or above
 - Linux

All versions of Linux with the required kernel and modules are supported.

A special kernel module published by Intel must be loaded. This kernel module is available for only some Linux distributions.

For more information about the kernel module, see Intel's Open Source Technology Center at <https://01.org/>.

5-19 Virtual SGX Restrictions

Virtual SGX has the following restrictions:

- For VMs configured with virtual SGX, the following operations are not supported:
 - Live migration with vSphere
 - Suspend and resume
 - Snapshots
 - Fault tolerance
- For virtual SGX enabled VMs, the following VM configuration attributes are fixed:
 - Guest firmware type: EFI.
 - Standby response: Put the guest OS in standby mode and leave the VM powered on.

Cold migrations can be performed.

5-20 Enabling and Disabling Virtual SGX for a VM

In the vSphere Client, you can enable or disable virtual SGX on the **Virtual Hardware** tab under Security devices.

New NVMe Controller	
▼ Security devices	SGX (32 MB)
▼ SGX (Unlocked)	☒ Enable (Maximum: 92 MB)
Enclave page cache size (MB)	32 MB ▼
Restrictions	 Some operations and features are restricted 
Launch control configuration	Unlocked ▼

You can use the **Edit Settings** menu to enable and disable virtual SGX for a VM. The VM must be powered off to enable or disable the feature.

5-21 Guest Customization Improvements

vSphere 7 includes new REST APIs for guest customization:

- Create, Update, Set, Delete spec
- List, Get, Apply spec
- Import/Export spec to/from JSON

You can now create pre- and post-customization scripts in your preferred scripting language (sh, bash, Python, and so on).

With the new REST APIs, you can interact with guest customization outside of the vSphere Client.

5-22 About Instant Clone Guest Customization

With instant clone guest customization, you can modify the network settings of a running VM:

- Customization is done without powering off or restarting the VM.
- This customization feature is only available for Linux VMs.
- Instant clone guest customization is an API-only feature.

You can customize networking settings such as IP address, DNS server, and gateway. You can change these settings without powering off or restarting the VM.

5-23 Instant Clone Guest Customization Use Cases

You can use the instant clone guest customization in various ways:

- Customize network settings of an instant clone.
- Customize network settings of a normal running VM.

5-24 Instant Clone Guest Customization Requirements

To use instant clone guest customization, you must meet several requirements:

- The customization engine must be installed on the guest OS.
- The feature is supported only for Linux VMs.
- Version restrictions:
 - vCenter version 7.0 or later.
 - ESXi version 6.0 or later.
 - VMware Tools version 10.2 or later.

5-25 Latency Sensitivity Advantages

Latency sensitivity offers the following advantages:

- The latency sensitivity feature is available from vSphere 5.5.
- This feature reduces latency of a given VM.
- The reduction in latency improves the performance of the VM.
- You can use different latency sensitivity settings:
 - High latency sensitivity lowers the latency for a VM.
 - Normal latency sensitivity is the default latency for a VM.
 - Low and medium values are experimental and are not be used in production.

Latency sensitivity improves the performance of a latency-sensitive application on a VM.

With virtualization, VMs do not have direct access to physical resources.

But with latency sensitivity, VMs can access the preferred physical resources, bypassing the virtualization layer. This bypass eliminates extra processing.

When you create a VM, the default latency sensitivity on that VM is normal.

5-26 Selective Latency Sensitivity

Selective latency sensitivity allows mixed latency sensitivity settings across virtual CPUs within the same VM.

vCPUs that are configured with high-latency sensitivity have the same priority as existing VMs with high latency sensitivity settings.

vCPUs that are not configured with a high-latency sensitivity setting have the default scheduling priority.

5-27 Selective Latency Sensitivity Use Cases

Selective latency sensitivity is useful for the deployment of networking services in the telecommunication industry:

- Networking service applications have different performance requirements.
- CPUs and vCPUs are exclusively dedicated to networking service applications.
- Noncritical tasks use other CPUs and vCPUs.

In telecommunications, multiple applications often coexist in a single physical server. These applications perform operations such as data plane packet forwarding, management and control plane request processing, health checks, and so on.

Such applications often have different performance requirements. To maximize throughput for the data plane traffic and minimize jitters caused by other applications, service providers often dedicate the physical core exclusively to the data plane applications.

All noncritical tasks are processed on the remaining shared physical cores.

With selective latency sensitivity, you can fine-tune the latency sensitivity of applications.

5-28 Selective Latency Sensitivity Requirements

To use selective latency sensitivity, several requirements must be met:

- Virtual ESXi compatibility level for ESXi 7.0 and later is required.
- VM-level latency sensitivity must be set to high so that you can set per vCPU latency sensitivity.
- Full memory reservation for the VM is required.
- VMs with insufficient CPU reservation are rejected.

5-29 Selective Latency Sensitivity Restrictions

Several restrictions apply to the selective latency sensitivity feature:

- This feature is only available through the vSphere 7 SDK and API.
- VMs with high-latency sensitivity cannot be migrated using vSphere vMotion, either by DRS or by the administrator.
- When you set VM latency sensitivity to high, all the vCPUs for that VM have a per vCPU latency sensitivity of high.
- The guest CPU ID might not be consistent with the vCPU numbering scheme used by ESXi.

When using selective latency sensitivity, you must ensure that the mapping between the guest CPU IDs and guest APIC IDs is correct.

This mapping allows the user to properly select the correct vCPUs to set the latency sensitivity.

5-30 Migration of VMs with Remote Devices

In vSphere 7, you can migrate a VM with a device attached through a remote console.

Requirements for this feature are:

- vCenter Server 7.0
- ESXi 7.0
- VMware Remote Console with reconnect support

Previously, you could not use vSphere vMotion to migrate a VM with a floppy or CD-ROM device attached through a remote console (VMware Remote Console, Workstation, Workstation Player, and Fusion). These remote devices might be physical devices or on disk images on the client machine running a remote console.

This limitation causes issues when you try to evacuate all VMs from a host to reduce the time for performing host upgrades.

Earlier versions of VMware Remote Console, vSphere, and ESXi do not support migration of VMs with remote floppies or CD-ROMs.

5-31 Virtual Machine Maximums for Large Workloads

With the introduction of VM compatibility level for ESXi 7.0 in vSphere 7, the CPU and RAM limits per VM are increased.

A VM running a compatibility version of ESXi 7.0 supports:

- 768 vCPUs
- 24 TB of RAM

When using vSphere vMotion to perform a migration of a large VM, you should have 40 Gb links on the vSphere vMotion network.

5-32 Lab 5: Using vSphere vMotion to Migrate a VM with Attached Devices

Use the vSphere Client to migrate a virtual machine with vSphere vMotion while a device is attached through VMware Remote Console:

1. Use VMware Remote Console to Attach an ISO Image to a VM
2. Migrate the VM with vSphere vMotion

5-33 Review of Learner Objectives

After completing this module, you should be able to meet the following objectives:

- Identify the characteristics of VM compatibility level for ESXi 7.0 and later
- Describe the support provided for new virtual devices in vSphere 7
- Recognize improvements to guest customization
- Migrate VMs with attached devices

5-34 Key Points

- vSphere 7 introduces VM compatibility ESXi 7.0 , which increases maximums for CPU and memory.
- The Precision Clock device for VMs provides better time synchronization.
- Virtual SGX can be used to reduce the attack vector on VMs.
- With guest customization, you can use scripting from your preferred programming languages.
- Migrations of VMs with remote devices can be performed in vSphere 7.

Questions?

Module 6

Storage Enhancements

6-2 Importance

vSphere supports various storage options and devices to give you the flexibility to set up the storage solution that integrates with your environment.

Using vSphere storage and its enhanced features can help meet the demands of your workloads and improve the overall management experience of your virtual data center.

6-3 Module Lessons

1. Enhancements to NVMe Storage Systems
2. Support for iSER
3. vSAN Enhancements

6-4 **Lesson 1: Enhancements to NVMe Storage Systems**

6-5 Learner Objectives

After completing this module, you should be able to meet the following objectives:

- Describe the enhancements to NVMe support
- Configure ESXi for NVMe over RDMA (RoCE v2)
- Describe the enhancements to VMware Pluggable Storage Architecture
- Describe the purpose of the High-Performance Plug-in

6-6 Enhanced Support for NVMe

Non-Volatile Memory Express (NVMe) is a standardized protocol for high-performance multiqueue communication with NVM devices. ESXi supports the NVMe protocol to connect to local and networked storage devices.

vSphere 7 introduces the following features to enhance NVMe support

- Pluggable Storage Architecture (PSA) stack using a High-Performance Plug-in (HPP) with multipathing
- Hot-add and hot-remove of local NVM devices, if the server system supports it
- Support for devices that use Remote Direct Memory Access (RDMA) or Fibre Channel

ESXi uses Pluggable Storage Architecture (PSA) as a VMkernel layer for managing storage multipathing.

The PSA is an open and modular framework. It coordinates software modules that are responsible for multipathing operations.

Modules include generic multipathing modules (MPPs) that VMware provides, Native Multipathing Plug-in (NMP) and High-Performance Plug-in (HPP), and third-party MPPs.

6-7 Supported NVMe Topologies

In vSphere 7, the NVMe storage can be directly attached to a host using a Peripheral Component Interconnect Express (PCIe) interface or indirectly through different fabric transports.


NVMe Transport	ESXi 7.0 Support
NVMe over PCIe	Local storage
NVMe over RDMA	Shared NVMe over Fabrics (NVMe-oF) storage with RDMA over Converged Ethernet (RoCE) v2 technology
NVMe over Fibre Channel (FC-NVMe)	Shared NVMe-oF storage with Fibre Channel

ESXi 7.0 supports the NVMe protocol to connect to local and networked storage devices. For local devices, you can use a PCIe interface. For devices on a shared storage array, use NVMe over Fabrics (NVMe-oF)

For shared NVMe storage, you can use NVMe over Fibre Channel (FC-NVMe) or NVMe over RDMA with the RoCE v2 technology.

6-8 FC-NVMe and NVMe over RDMA (RoCE v2)

vSphere 7 supports remote access to shared NVMe storage, such as NVMe over Fibre Channel (FC-NVMe) and NVMe over RDMA with RoCE v2 technology.



For more details about VMware NVMe concepts, see *vSphere Storage* at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html>.

6-9 Support for FC-NVMe

FC-NVMe is a technology that maps NVMe onto the Fibre Channel protocol.

With FC-NVMe, data and commands can be transferred between a host computer and a target storage device.


To access the FC-NVMe storage, you must install a Fibre Channel storage adapter that supports NVMe on your ESXi host.

You do not need to configure the controller. After you install the required hardware NVMe adapter, it automatically connects to all targets and controllers that are reachable.

Your ESXi host detects the device and displays it in the vSphere Client as a standard Fibre Channel adapter (vmhba) with the storage protocol indicated as NVMe.

You do not need to configure the controller. After you install the required hardware NVMe adapter, it automatically connects to all targets and controllers that it can reach.

You can later reconfigure the adapter and disconnect its controllers or connect other controllers that were not available during the host boot.



6-10 Requirements for FC-NVMe

When you use FC-NVMe, the following components are required:

- Fibre Channel storage array that supports NVMe.
- NVMe controller on your storage array.
- ESXi 7.0 or later.
- Hardware NVMe adapter on your ESXi host: This adapter is typically a Fibre Channel host bus adapter (HBA) that supports NVMe.

Configuration maximums for NVMe-oF:

- 32 namespaces
- 128 paths
- 4 paths per namespace per host

For more information about NVMe concepts, see *vSphere Storage* at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html>.

6-11 Support for NVMe over RDMA (RoCE v2)

The NVMe over RDMA (RoCE v2) technology uses RDMA for transport between two systems on the network.

Data exchange occurs in the main memory, bypassing the operating system and processor of both systems.


ESXi 7.0 supports RoCE v2 technology, which enables RDMA over an Ethernet network.

To access storage, the ESXi host uses an RDMA network adapter installed on your host and a software NVMe over RDMA storage adapter.

NVMe over RDMA with RoCE v2 technology provides ultra-efficient storage to optimize performance-sensitive workloads over your existing Ethernet fabric.

You must configure both adapters to use them for storage discovery. The adapter configuration process involves setting up VMkernel binding for an RDMA network adapter and enabling a software NVMe over RDMA adapter.

For more information about configuring adapters for NVMe over RDMA (RoCE v2) storage, see *vSphere Storage* at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html>.



6-12 Requirements for NVMe over RDMA (RoCE v2)

NVMe over RDMA (RoCE v2) requires the following components:

- NVMe storage array with NVMe over RDMA (RoCE v2) transport support
- NVMe controller on your storage array
- ESXi 7.0 or later
- Ethernet switches supporting a lossless network
- Network adapter on your ESXi host that supports RoCE v2
- Software NVMe over RDMA adapter enabled on your ESXi hosts


6-13 Configuring ESXi for NVMe over RDMA (RoCE v2) (1)

To configure your ESXi host to use NVMe over RDMA (RoCE v2):

1. On your ESXi host, install an adapter that supports RDMA (RoCE v2).
2. In the vSphere Client, verify that the RDMA adapter is discovered by your host.
3. Configure VMkernel binding for the RDMA adapter.
4. Verify the VMkernel binding configuration for the RDMA adapter.
5. Enable the software NVMe over RDMA adapter.
6. Discover NVMe controllers on your array (automatic or manual).
7. Rescan storage to discover NVMe namespaces.
8. Create your VMFS datastores.

6-14 Configuring ESXi for NVMe over RDMA (RoCE v2) (2)

Identify the physical network adapter that corresponds to the RDMA adapter.



The screenshot shows the vSphere Client configuration page for the host 'sa-esxi-01.vclass.local'. The 'Configure' tab is selected, and the left sidebar shows the 'RDMA adapters' section under 'Networking'. The main content area displays a table of RDMA adapters:

Name	Driver	State	Paired Uplink
vmrdma0	nmlx5_rdma	Active	vmnic1

Below the table, the 'RDMA Device: vmrdma0' section is expanded, showing the 'Vmkernel adapters binding' table:

VMkernel Adapter	TCP/IP Stack	IP Address
vmk1	Default	192.168.30.204

6-15 Configuring ESXi for NVMe over RDMA (RoCE v2) (3)

Configure the VMkernel binding for the RDMA adapter.

The screenshot shows the vSphere Client interface for host 'sa-esxi-01.vclass.local'. The 'Configure' tab is selected, and the left-hand navigation pane shows 'RDMA adapters' under the 'Networking' section. The main content area displays the configuration for the 'RDMA Device: vmrdma0'. Under the 'Vmkernel adapters binding' sub-tab, a table shows the binding for 'vmk1'.

VMkernel Adapter	TCP/IP Stack	IP Address
vmk1	Default	192.168.30.204

6-16 Configuring ESXi for NVMe over RDMA (RoCE v2) (4)

Select **Add software NVMe over RDMA adapter** and select an appropriate RDMA adapter (vmrdma) from the drop-down menu.

Add Software Adapter | hwe-drv-10.eng.vmware.com

☒ Add software iSCSI adapter
A new software iSCSI adapter will be added to the list. After it has been added, select the adapter and use the Adapter Details section to complete the configuration.

☒ Add software NVMe over RDMA adapter
Discover software NVMe adapters associated with the following RDMA devices.

RDMA device: vmrdma0/MT27630 Family [ConnectX-4 LX]

☐ Add Software FCoE Adapter
Discover software FCoE adapters associated with the following physical network adapters.

Physical Network Adapter:

VLAN ID: Range: 0 - 4094

Priority Class: Range: 0 - 7

Controller MAC Address:


CANCEL **OK**

The software NVMe over RDMA adapter appears on the list as a vmhba storage adapter. You can remove the adapter if you must free the underlying RDMA network adapter for other purposes.

For more information about configuring adapters for NVMe over RDMA (RoCE v2) storage, see *vSphere Storage* at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html>.

6-17 Pluggable Storage Architecture (1)

Pluggable Storage Architecture (PSA) is an open and modular framework that coordinates various software modules that are responsible for multipathing operations.



Acronym	Definition
MPP	Multipathing Plug-in: A multipathing module developed and provided by a third party
NMP	Native Multipathing Plug-in: Generic VMware multipathing module that is used with SCSI storage devices
HPP	High-Performance Plug-in provided by VMware: Used with ultra-fast local and networked flash devices

VMware provides generic native multipathing modules, called VMware NMP and VMware HPP.


In addition, the PSA offers a collection of VMkernel APIs that third-party developers can use.

Software developers can create their own load-balancing and failover modules for a particular storage array. These third-party multipathing modules (MPPs) can be installed on the ESXi host and run in addition to the VMware native modules, or as their replacement.

Multiple third-party MPPs can run in parallel with the VMware NMP or HPP.
When installed, the third-party MPPs can replace the behavior of the native modules.

6-18 Pluggable Storage Architecture (2)

The PSA software modules use different mechanisms to handle the path selection and failover.



Acronym	Definition
SATP	Storage Array Type Plug-in: Handles path failover for a given SCSI storage array
PSP	Path Selection Plug-in: Handles path selection for a SCSI storage device
PSS	Path Selection Scheme: Handles multipathing for NVMe storage devices

6-19 High-Performance Plug-In for NVMe

In vSphere 7, the HPP performs the following functions:

- Acts as the default plug-in that claims NVMe-oF targets
- Supports only active-active and implicit Asymmetric Logical Unit Access (ALUA) targets
- Improves the performance of ultra-fast flash devices installed on ESXi
- Can replace the Native Multipathing Plug-in for local NVMe devices
- Supports multipathing for NVMe-oF

In vSphere 7, HPP integrates path selection and failover logic within a single module, reducing cross-module calls.

To support multipathing, HPP uses a Path Selection Scheme (PSS) when selecting physical paths for I/O requests. You can use vSphere ESXi Shell or the vSphere CLI to configure the HPP and PSS.

For local NVMe devices, NMP remains the default plug-in, but you can replace it with HPP.

For more information about configuring and using HPP and path selection schemes, see *vSphere Storage* at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html>.

6-20 Support for High-Performance Plug-In Multipathing

The HPP improves the performance of NVMe devices on your ESXi host.

HPP Support	vSphere 7
Storage devices	Local NVMe PCIe Shared NVMe-oF (active-active and implicit ALUA targets only)
Multipathing	Yes
Second-level plug-ins	No Path Selection Scheme (PSS)
SCSI-3 persistent reservations	No
4Kn devices with software emulation	No
vSAN	No

6-21 High-Performance Plug-In Best Practices

To achieve the fastest throughput from a high-speed storage device, follow these recommendations:

- Use the HPP for NVMe local or networked devices.
- Do not activate the HPP for HDDs or slower flash devices.
- When using FC-NVMe, follow general recommendations for Fibre Channel storage.
- When using NVMe-oF, do not mix transport types to access the same namespace.
- Before registering NVMe-oF namespaces, ensure that active paths are presented to the host.
- Configure your VMs to use VMware Paravirtual controllers.
- Configure your VMs to use high-latency sensitivity.
- If a single VM drives a significant share of the device's I/O workload, consider spreading the I/O across multiple virtual disks. Attach the disks to separate virtual controllers in the VM.

Do not activate the HPP for HDDs or slower flash devices. The HPP does not provide any performance benefits with devices incapable of at least 200,000 IOPS.

If you do not attach the disks to separate virtual controllers in the VM, I/O throughput might be limited because of saturation of the CPU core that is responsible for processing I/O operations on a virtual storage controller.

For information about device identifiers for NVMe devices that support only NGUID ID format, see *vSphere Storage* at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html>.

6-22 Review of Learner Objectives

After completing this module, you should be able to meet the following objectives:

- Describe the enhancements to NVMe support
- Configure ESXi for NVMe over RDMA (RoCE v2)
- Describe the enhancements to VMware Pluggable Storage Architecture
- Describe the purpose of the High-Performance Plug-in

6-23 **Lesson 2: Support for iSER**

6-24 Learner Objectives

After completing this module, you should be able to meet the following objectives:

- Describe support for iSER in vSphere 7
- Configure iSER on ESXi hosts

6-25 Support for iSER-Enabled vSphere Virtual Volumes Arrays

iSCSI Extension for RDMA (iSER) provides high-performance iSCSI connectivity for standard RDMA NICs.

In addition to supporting traditional iSCSI, ESXi supports the iSER protocol. When the iSER protocol is enabled, the iSCSI framework on the ESXi host can use the RDMA transport instead of TCP/IP.

vSphere 7 also adds iSER support for vSphere Virtual Volumes.

In vSphere 7, support for iSER requires the following:

- RoCE-enabled devices
- A target array that supports iSER on RoCE

6-26 Configuring iSER

Configure iSER on your ESXi host so that the iSCSI framework on the host can use Remote Direct Memory Access (RDMA) transport instead of TCP/IP.

Configuration Step	Description
Enable the VMware iSER adapter.	Use the <code>esxcli</code> command to enable the VMware iSER adapter.
Modify general properties for iSCSI or iSER adapters.	If needed, change the default name and alias assigned to your adapter.
Configure port binding for iSCSI or iSER.	You must create network connections to bind the iSER engine and the RDMA-capable network adapter.
Configure dynamic or static discovery for iSCSI and iSER.	With the dynamic discovery, each time the initiator contacts a specified iSER storage system, it sends the <code>SendTargets</code> request to the system. With static discovery, manually enter information for the targets.
Set up CHAP for iSCSI or iSER.	This step is necessary only if your environment uses the Challenge Handshake Authentication Protocol (CHAP).
Enable jumbo frames for networking.	If your environment supports jumbo frames, enable them for the adapter.

When installed on the host, the RDMA-capable adapter appears in vCenter Server as a network adapter (vmnic).

To make the adapter functional, you must enable the VMware iSER component and connect the iSER adapter to the RDMA-capable vmnic. You can then configure properties, such as targets and CHAP, for the iSER adapter.

NOTE: iSER does not support NIC teaming. When configuring port binding, use only one RDMA adapter per virtual switch.

6-27 Requirements for iSER

To enable the VMware iSER adapter, use the `esxcli` command and ensure that you meet the following requirements:

- Verify that your iSCSI storage supports the iSER protocol.
- Install the RDMA-capable adapter on your ESXi host.
- Use the RDMA-capable switch.
- Enable flow control on the ESXi host. To enable flow control for the host, use the `esxcli system module parameters` command.
- Configure RDMA switch ports to create lossless connections between the iSER initiator and target.

To enable flow control, see VMware knowledge base article 1013413 at <http://kb.vmware.com/kb/1013413>.

6-28 Review of Learner Objectives

After completing this module, you should be able to meet the following objectives:

- Describe support for iSER in vSphere 7
- Configure iSER on ESXi hosts

6-29 **Lesson 3: vSAN Enhancements**

6-30 Learner Objectives


After completing this module, you should be able to meet the following objective:

- Identify vSAN enhancements for vSphere 7

6-31 Enterprise Storage in Native vSphere Architecture

Enterprise storage has several characteristics:

- Integrated into the hypervisor
- Pools HDD/SDD into a single cluster-wide shared datastore
- Easily scalable
- Managed through storage policies in vCenter Server
- Runs enterprise and next-generation cloud-native applications



vSAN is VMware's software-defined storage solution.

It abstracts and aggregates locally attached disks in a vSphere cluster to create a storage solution that can be provisioned and managed from vCenter Server and the vSphere Client.

vSAN is hypervisor-converged, which means that storage and compute for VMs are delivered from the same x86 server platform running the hypervisor.

vSAN integrates with the entire VMware stack, including features such as vSphere vMotion, vSphere HA, and vSphere DRS.


You can control VM storage provisioning and the daily management of storage SLAs through VM-level policies that can be set and modified as needed.

vSAN delivers enterprise-class features, scale, and performance, making it the ideal storage platform for VMs.

6-32 Unified HTML5 User Interface for vSphere and vSAN

vSphere and vSAN use an HTML5 UI for vCenter Server.

This UI simplifies workflows and provides new integration capabilities



vSAN 6.7 introduces support for the new HTML5-based vSphere Client, which is based on the Clarity framework used for other VMware solutions.

VMware products that already use the Clarity user interface include vRealize Operations and vRealize Log Insight.

This improved UI integrates simpler workflows to configure vSAN, manage disk groups, and monitor vSAN health, performance, and capacity.

For more information about the new UI, see "vSAN 6.7 in the HTML5 vSphere Client" at <https://blogs.vmware.com/virtualblocks/2018/04/18/vsan-67-html5-vsphere-client/>.

6-33 vSAN 6.7 Update 1: New Features and Enhancements

vSAN 6.7 U1 introduces several new features and enhancements:

- Additions and enhancements to the UI in the vSphere Client
- Ability to upgrade storage IO controller firmware using vSphere Update Manager
- Enhancements to maintenance mode for ESXi hosts within vSAN clusters
- Additional capacity reporting capabilities
- Full integration with vRealize Operations 7.0
- Support for guest OS TRIM/UNMAP
- Health check enhancements
- Ability to specify multiple MTU sizes with witness traffic separation
- Integration with cluster Quickstart

The vSphere Client provides an updated advanced settings dialog box where you can perform the following tasks:

- Adjust the component repair delay timer
- Enable and disable thin swap files
- Enable and disable site read locality

To find the advanced settings in the vSphere Client, select your vSAN cluster and select **Configure > vSAN > Services > Advanced Options**.


For a complete list of new features and enhancements, see the VMware vSAN 6.7 Update 1 Release Notes at <https://docs.vmware.com/en/VMware-vSphere/index.html>.

6-34 Simplified Upgrade Through vSphere Update Manager

ESXi, driver, and firmware updates are centralized through vSphere Update Manager.

vSphere Update Manager supports:

- Custom ISOs for OEM-specific builds
- vCenter Server without Internet connectivity
- Per-host updating with additional safeguards



vSAN 6.7 U1 centralizes all updating of vSAN and associated hardware through vSphere Update Manager.

All ESXi, driver, and firmware updates are available in a single location.

Also, vSphere Update Manager can use custom ISOs from vendors to manage vSAN clusters running specific OEM builds of ESXi hosts.

Some environments prevent vCenter Server from accessing the Internet. Guided workflows help vCenter Server run without Internet connectivity.

The upgrade process immediately checks for issues after an update.

If an issue is identified, the upgrade process stops, and the host in question remains in maintenance mode.

When the ESXi cluster is enabled with vSAN, vSphere Update Manager generates and manages a system baseline.

This baseline is read-only and cannot be removed.


It provides controller driver recommendations and controller firmware recommendations.

6-35 Improved Resource Decommissioning Safeguards

Enter Maintenance Mode (EMM) has a precheck simulation. It determines the success or failure of an action with no unnecessary data movement.

EMM provides new warnings for the following issues:

- Hosts already in maintenance mode
- Ongoing resyncs
- Multiple decommission actions



Because each vSAN host in a cluster contributes to the cluster storage capacity, entering a host into maintenance mode takes on an additional set of tasks when compared to hosts in a traditional architecture.

vSAN 6.7 U1 improves the safeguards for maintenance and decommissioning activities on vSAN hosts.


vSAN 6.7 U1 improves the safeguards for maintenance and decommissioning activities on vSAN hosts.

- vSAN performs a full precheck simulation of data movement to determine whether the Enter Maintenance Mode (EMM) action will succeed or fail. This simulation prevents unnecessary data movement.
- New warnings for EMM activities ensure that no other hosts are in maintenance mode and that resync activity is not ongoing.

6-36 Improved Insight with Cluster Capacity Monitoring

vSAN 6.7 U1 introduces a historical capacity dashboard that reports on capacity use over time. It includes the following features:

- Historical capacity reports
- Usable capacity estimator



vSAN 6.7 U1 introduces a historical capacity dashboard that reports on total, used, and free capacity use over time.

You can access an ongoing history of deduplication ratios to understand how changes in an environment might affect deduplication ratios.

Historical reports are limited to a 90-day time frame.


This release also includes a usable capacity estimator:

- Shows free datastore space available based on a selected storage policy
- Estimates what-if scenarios when space-efficient RAID-5 or 6 is used

6-37 Enhancements to vRealize Operations

Native vRealize Operations dashboards that are built into vCenter Server can display intelligence for vSAN stretched clusters.

The deployment through the vRealize Operations Manager plug-in in the vSphere Client supports using virtual distributed switches.




The deployment process supports virtual distributed switches and is fully compatible with vRealize Operations 7.0.

6-38 Improved Space Efficiency Using Storage Reclamation

The vSAN and TRIM/UNMAP integration results in storage efficiency:

- Reclamation of guest OS capacity
- Reclamation of storage for vSAN datastore
- Support for guest OS-initiated TRIM/UNMAP commands



After data is deleted in a guest OS, the OS can reclaim the space that is no longer in use.

Using commands known as TRIM/UNMAP for the respective ATA and SCSI protocols, guest operating systems are more efficient with storage space use.

Guest OS initiated TRIM/UNMAP commands are supported in the following systems:

- Windows Server 2012 or Windows 8 and newer
- Linux supporting ext4, xfs, btrfs, and so on

In vSphere 6.7 U1 and later versions, vSAN has full awareness of TRIM/UNMAP commands that are sent from guest operating systems in the cluster and can reclaim previously allocated storage.

To unmap or reclaim released guest OS storage using vSAN, you must enable this feature through the API.

You can enable or disable the feature through the Ruby vSphere Console (RVC) command `vsan.unmap_support --enable` or `--disable`.

VMs require a power cycle after you enable the feature.

TRIM/UNMAP is applicable to only thin-provisioned VMDKs.

6-39 Improved Isolation of vSAN Traffic Types


You can set up witness traffic separation for stretched clusters so that witness traffic is fully separated from vSAN data traffic.

This separation prevents data traffic from being routed through the VMKNIC for witness host traffic when a link fails.

You can configure isolation of traffic by specifying an alternate VMKNIC:

- You enable the feature by using the ESXCLI command.
- A UI option is not available.

Isolation was previously supported only in two-node configurations.



You can set up vSAN stretched clusters with a witness host.

The witness host is a stripped-down vSphere host running vSAN. It cannot host VMs but is used exclusively to store cross-site witness components.

In vSAN 6.7 and previous versions, when the link between two sites fails, the data traffic can be routed through the VMKNIC for witness host traffic.

vSAN 6.7 introduces witness traffic separation on stretched cluster configurations.


A VMkernel interface can be designated to carry traffic destined for the witness rather than the vSAN tagged VMkernel interface.

This improvement helps you design and maintain stretched clusters that are more efficient and predictable in various conditions.

6-40 Improved Flexibility for Stretched Cluster Topologies

vSAN supports mixed MTU sizes for stretched clusters with witness traffic separation:

- You can use different MTU sizes for the inter-site link (ISL) that is used by vSAN data traffic and linked to the witness host.
- You can use jumbo frames between data sites.
- The witness site link maintains its simplicity.
- You configure MTU sizes when defining VMkernel traffic types.



vSAN 6.7 U1 improves flexibility on stretched cluster environments in the following ways:

- Support for different MTU sizes that are configured for witness traffic
- Support for the inter-site link (ISL) that is used for vSAN data traffic


With these improvements, you can perform the following tasks:

- Configure vSAN to use larger frame sizes on the inter-site link (ISL).
- Keep a standard MTU size in the traffic between witness uplinks and the witness site.
- Reduce potential network issues.

6-41 Improved Guidance of Health Checks and Recommendations

vSAN 6.7 UI enhances health and diagnostic services:

- The storage controller firmware health check includes multiple approved firmware levels.
- You can run a new unicast network performance test.
- You can use new UI controls:
 - Silence desired health checks
 - Purge inaccessible swap objects



The vSAN Health Check ensures that your environment meets hardware and software configuration requirements.

vSAN 6.7 U1 adds the following enhancements:

- vSAN handles multiple approved firmware levels for storage controllers at the same time.

You can recognize the latest approved version even when running an earlier approved version.

- A new unicast network performance health check and test ensures that proper continuity is achieved between hosts in a vSAN cluster.

The unicast network performance test uses iperf to test continuity across ESXi hosts in a vSAN cluster and reports effective bandwidth between nodes.

- Health checks can be silenced granularly in the user interface.

Silence alerts appear for different health checks.

When silenced, a Restore alert replaces the Silence alert.

- vSAN purges inaccessible swap objects that are no longer needed.

When you select **Purge Inaccessible Swap**, you can find out when the health check for vSAN object health fails.

Purging inaccessible swap objects helps to unblock operations previously blocked by an inaccessible object.

Sometimes, swap objects are orphaned and left over from VMs that are no longer in the cluster or are already using another swap object.

These orphaned swap objects tend to be problematic during Enter Maintenance Mode events.

Cleaning up these orphaned swap objects was previously only accessible through the RVC using `vsan.purge_inaccessible_vswp_objects`.

Purging inaccessible swap objects is a safe operation. The VM regenerates a new swap during the next power-on of a VM.

- The All hosts have matching subnets health check is removed.

This health check verifies that all vSAN hosts use matching subnets.

Since vSAN 6.0, the support for a layer 3 network, such as routed connections, and the popularity of topologies such as stretched clusters, eliminate the need for this health check.

6-42 vSAN 6.7 Update 3: New Features and Enhancements (1)

vSAN 6.7 Update 3 introduces several new features and enhancements:

- Performance enhancements
- Enhanced resync performance and monitoring
- Ability to directly upload and download VMDKs to and from a vSAN datastore
- Ability to automate proactive rebalancing operations
- Enhanced capacity monitoring

vSAN performance enhancements in vSAN 6.7 U3 are as follows:

- Improved performance and availability of SLAs on all-flash configurations with deduplication enabled. Rebuild times on disk and node failures decrease considerably, which provides better availability SLAs.
- Enhancements in the vSAN cache destaging algorithm and task parallelization across the I/O stack. As a result, latency sensitive applications perform better in terms of predictable I/O latencies and increased sequential I/O throughput.

Enhanced resync monitoring offers the following benefits:

- When a hardware device, host, or network fails, or if a host is placed into maintenance mode, vSAN initiates resynchronization in the vSAN cluster.
- The Resyncing Objects dashboard in vSAN 6.7 U3 introduces new logic to improve the accuracy of resync completion times and granular visibility into different types of resyncing activities, such as rebalancing and policy compliance.

vSAN 6.7 U3 includes parallel resynchronization:

- Optimized resynchronization behavior runs additional data streams per resyncing component when resources are available.
- This new behavior adaptively runs in the background and provides greater I/O management and performance for workload demands.

vSAN 6.7 U3 includes the ability to upload and download VMDKs to and from the vSAN datastore:

- This new capability is a simple way to protect and recover VM data during capacity-strained scenarios.

vSAN 6.7 U3 includes proactive rebalancing enhancements:

- vSAN can automate all rebalancing activities with a new cluster-wide configuration and threshold settings.
- Before this release, you manually initiated proactive rebalancing after you were alerted by vSAN health checks.

Enhanced capacity monitoring is another enhancement:

- vSAN 6.7 U3 introduces a redesigned capacity-monitoring dashboard for improved visibility of overall use and detailed breakdown and simplified capacity alerting.
- Capacity-related health checks are also more visible and consistent.
- Details about capacity use are now available per site, fault domain, and at the host and disk group level.

6-43 vSAN 6.7 Update 3: New Features and Enhancements (2)

- Support for Windows Server Failover Clustering (WSFC) on native vSAN VMDKs
- Ability to dynamically resize vSAN iSCSI LUNs
- Support for stateful cloud native applications with new cloud native storage platform
- Detailed precheck reporting to maintenance mode for ESXi hosts in vSAN clusters
- Addition of `vsantop` command-line utility

vSAN 6.7 U3 introduces native support for SCSI-3 persistent reservation (PR):

- Windows Server Failover Clusters can be deployed directly on VMDKs as first-class workload:
- With this new capability, you can migrate legacy deployments on physical raw device mappings (RDMs) or external storage protocols to vSAN.

With this release, the vSAN iSCSI service dynamically resizes iSCSI LUNs without disruption.

For a complete list of new features and enhancements, see *VMware vSAN 6.7 Update 3 Release Notes* at <https://docs.vmware.com/en/VMware-vSphere/index.html>.

6-44 Scalable Storage Provisioning for DevOps

With granular visibility into container volumes in vSAN 6.7 U3, you can control and monitor health and compliance information for each container volume.

**1** Sys Admin creates SPBM-based Storage Classes

**2** Developer claims vSAN datastore for persistent volumes with desired Storage Class

Direct access to volume properties


Single-UI view for multiple Orchestrators

Per-volume datastore placement information

Unified file and block volume listing

Rich filtering capabilities

SPBM policy compliance status




Volume name	Type	Label	Datastore	Compliance status	Volume ID
pmongo	File	function: no-intelligence	vsanDatastore	compliant	h3726e00-180f-1f6f-8fa1-00505683cc9
pmongo-2	File	function: no-intelligence	vsanDatastore	non-compliant	h3726e00-180f-1f6f-8fa1-00505683cc8
pmongo-3	File	function: no-intelligence	vsanDatastore	out dated	h3726e00-180f-1f6f-8fa1-00505683cc7
pmongo-4	File	function: no-intelligence	vsanDatastore	compliant	h3726e00-180f-1f6f-8fa1-00505683cc6
pmongo-5	File	Test: no-intelligence	vsanDatastore	compliant	h3726e00-180f-1f6f-8fa1-00505683cc5
pmongo-6	File	Test: no-intelligence	test-vpx-1519173591-54321	compliant	h3726e00-180f-1f6f-8fa1-00505683cc4
pmongo-7	File	environment: sandbox	SharedVnfs-1	compliant	h3726e00-180f-1f6f-8fa1-00505683cc3
pmongo-8	File	environment: sandbox	SharedVnfs-1	compliant	h3726e00-180f-1f6f-8fa1-00505683cc2
pmongo-9	Block	environment: sandbox	SharedVnfs-1	compliant	h3726e00-180f-1f6f-8fa1-00505683cc1
pmongo-10	Block	environment: sandbox	vsanDatastore	compliant	h3726e00-180f-1f6f-8fa1-00505683cc0

With granular visibility into containers, you can troubleshoot and remediate activities more rapidly. This feature helps DevOps teams to implement their container-based applications.

6-45 Creating Hybrid Cloud Solutions

With vSAN, you can use the same tools to manage your public and private cloud environments.



Docker containers, released in 2013, changed the value proposition for containers.

With minimal effort, developers can deploy containers on their laptops and then migrate them to a production environment on premises or to a public cloud provider.

However, stateful technologies need data to easily migrate between environments, and administrators want familiar tools and processes to manage a multicloud environment.

With vSAN, you can use the same tools to manage your public and private cloud environments.

vSAN has native services with three of the four largest public cloud providers: Amazon, Microsoft, and IBM.


Because these cloud environments use the same tools and processes as private clouds, data migration is as simple as using Storage vMotion, and Kubernetes can use your public cloud environment as your container environment scales.

6-46 Data Management for Stateful Cloud Native Applications

With vSAN 6.7 U3, you can provision and manage persistent volumes for container orchestrators such as Kubernetes.

Requirements are easy to assign:

- The vSAN administrator creates the Storage Policy Based Management (SPBM) policy.
- The developer consumes with storage class.



Containerized applications that require persistent storage can be deployed by the Kubernetes administrator or developers without having to understand the underlying storage.

vSAN administrators associate storage policies to Kubernetes storage classes.

As a result, vSAN presents this storage as first-class disks (FCD) in the vCenter Server UI for easy management and troubleshooting.

The example illustrates the following points:

- The container storage interface (CSI) replaces the vSphere cloud provider.
- The vSphere cloud provider provides dynamic volume provisioning (through a StorageClass API).
- The CSI is an out-of-tree driver, for example, a plug-in, which can be installed on any version of Kubernetes after 1.12. With this feature, decoupling from a Kubernetes release cadence is rapid.

You require virtual hardware version 14 (compatible with ESXi 6.7 U2 and later) or later.


Volume resizing and snapshots are not supported for this edition.

6-47 Predictive Analysis for Maintenance Operations

You can perform a detailed precheck analysis when you enter a host into maintenance mode (EMM).

This precheck has the following features:

- Different report views:
 - Object Compliance and Accessibility
 - Cluster Capacity
 - Predicted Health
- Dedicated view under Monitor
- Can be tested against alternate migration options



The Precheck Remediation is performed on the host or the cluster. The check displays a table that lists possible problems that might prevent a successful remediation and recommendations for how to fix the issues.

When you perform a maintenance task in vSAN 6.7 U3, the precheck analyzes the potential behavior of the following aspects of the cluster:

- Object compliance
- Object accessibility
- Cluster capacity
- Health of the cluster

The resulting report gives you an informed view before any action is taken.

To view detailed information about vSAN Health, follow these steps:

1. Select a vSAN cluster.
2. Click the **Monitor** tab.
3. Click **vSAN**.
4. Click **Health**.

If you start a remediation task in a vSAN cluster that failed the vSAN health precheck, the hosts enter maintenance mode, get upgraded, but might fail to exit maintenance mode, and the remediation eventually fails.

6-48 Improved Performance Metrics and Tools


vSAN 6.7 U3 provides new performance metrics and troubleshooting utility.

You can access dedicated vSAN CPU metrics to understand overheads:

- Available through API
- Exposed in vRealize Operations

You can run a new `vsantop` command for host-based vSAN monitoring and troubleshooting:

- Similar to `esxtop`
- Frequent collection interval



vSAN 6.7 U3 introduces a vSAN CPU metric through the performance service and provides a new command-line utility (`vsantop`) for real-time performance statistics for vSAN.

The vSAN CPU metrics help you understand the number of resources consumed for the processing of I/O operations and other data services.

CPU usage for various vSAN modules such as LSOM, DOM, and CMMDS is collected.

Each module is summed to provide an aggregate usage of the CPU by vSAN.

Module statistics are persisted in the statistics database and can be accessed in the following ways:

- vSAN support insight
- Log bundles
- API for solutions such as vRealize Operations

To view the details of vSAN CPU metrics, follow these steps:

1. Select a vSAN cluster.
2. Click the **Monitor** tab.
3. Click **vSAN**.
4. Click **Support**.
5. Click **Performance for Support**.

vSAN displays performance charts with the % Used and % Ready metrics for clients running on the cluster.

The command-line tool `vsantop` that runs on ESXi hosts shows you the real-time vSAN performance metrics.

You can use this tool to monitor vSAN performance.

`vsantop` collects and persists statistical data in a RAM disk.

Based on the configured interval rate, the metrics are retrieved and displayed on the secure shell console.

This interval is configurable. You can reduce or increase it depending on the amount of detail that you require.

For more information about how to use the `vsantop` command-line tool, see *vSAN Monitoring and Troubleshooting* at <https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vsan-monitoring.doc/GUID-B5123101-975F-494F-95DB-D15C9A91B8F1.html>.

6-49 vSAN 7: New Features and Enhancements (1)

vSAN 7 introduces several new features and enhancements:

- Native support for planned and unplanned NVMe hot plug
- Ability to override default gateway settings on the vSAN enabled VMkernel port
- Increased maximum drive capacity
- Addition of vSAN file services with NFS 4.1 support for cloud native storage

NVMe hot plug is natively supported in vSphere 7 and vSAN 7.0.

vSAN 7.0 supports overriding the default gateway on the vSAN traffic VMkernel and the witness traffic VMkernel.

vSAN 7.0 extends support to 32 TB for physical capacity drives and extends to 1 PB for logical capacity when deduplication is enabled. In this way, large capacity drives achieve better deduplication ratios.

For a complete list of new features and enhancements, see *VMware vSAN 6.7 Update 3 Release Notes* at <https://docs.vmware.com/en/VMware-vSphere/index.html>.

6-50 vSAN 7: New Features and Enhancements (2)


- Integration with vSphere Lifecycle Manager
- Integration with Skyline
- File-based persistent volumes (cloud native storage)
- Integration with vSphere with Kubernetes (cloud native storage)

6-51 vSAN File Services Architecture

The vSAN file service provisions NFS 4.1 file shares from your vSAN data center.

NFS 4.1 file servers are automatically created as agent VMs on select ESXi hosts in your vSAN cluster.

The NFS 4.1 file servers are federated, forming a single vSAN distributed file system (VDFS) with a unified namespace.



With vSAN file services, you can provision a file share from you vSAN data center.

You can access the file share by using SMB or NFS.

The basic design is to run SMB or NFS file servers in guest VMs on top of a file system service that is available to all VMs.

These file servers are federated to form a unified namespace that consists of multiple separate independent file servers called shards.

The file system service is a vSAN distributed file system. The shard manager manages the SMB or NFS shards and brings them together into a unified namespace.

The shard manager must manage all the shards and shares in it. The shard manager must also tie natively into the existing vSphere management UI that is used to control vSAN.

6-52 Review of Learner Objectives

After completing this module, you should be able to meet the following objective:

- Identify vSAN enhancements for vSphere 7

6-53 Key Points

- ESXi supports the NVMe protocol to connect to local and networked storage devices that use RDMA or Fibre Channel.
- Virtual volumes can use iSER on RoCE enabled devices if the target array supports it.
- The system storage layout in ESXi 7.0 addresses the sizing issues associated with fixed partition sizes.
- vSAN adds support to features like container persistent storage and file services to become nearly functionally equivalent to the traditional storage

Questions?

Module 7

Security Enhancements

7-2 Importance

vSphere components are secured by many built-in features such as certificates, authorization, and limited access. To maintain protection of your VM files, disks, and in-flight data, you must recognize the ways that vSphere 7 builds on the existing encryption features that vSphere 6.5 introduced.

7-3 Module Lessons

1. Virtual Machine Encryption Enhancements
2. vSphere Trust Authority

7-4 **Lesson 1: Virtual Machine Encryption Enhancements**

7-5 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe enhancements to cloning, migration, and encryption operations
- Identify the enhancements to the VM encryption architecture
- Describe how to enable encryption mode on ESXi hosts in vSphere 7
- Recognize new VM encryption events and alarms

7-6 Improvements to VM Encryption Functionality

vSphere 7 includes several enhancements to the VM encryption feature that was introduced in vSphere 6.5:

- Some VM operations now support encrypted VMs.
- With vSphere HA enabled and at least one encrypted VM, encryption mode is enabled on all hosts in the cluster.
- New APIs improve the management of the VM encryption environment.
- New events and alarms monitor the VM encryption environment.

These enhancements improve the general functionality of VM encryption. They are based on competitive analysis and feedback from VMware customers and partners.

7-7 Encryption Use Cases

By enabling encryption on VMs, you can protect your confidential data:

- VMDKs are encrypted and can be accessed only with a digital key.
- vSphere access control restricts access to the digital key for decryption.
- In-flight data is encrypted.

To mitigate the risk of an administrator accessing data without authorization, you can encrypt VMs and provide access to encrypt and decrypt to a subset of administrators. To provide this access, you use the built-in cryptographic privileges in vSphere.

A key management server (KMS) generates the digital key for encryption and decryption.

7-8 VM Encryption Operations During VM Cloning

In vSphere 7, cryptographic operations are supported during the cloning of a VM or the creation of a VM from a template.

When cloning an unencrypted VM, you can encrypt the destination VM by changing the storage policy of the destination VM to the VM encryption policy (or custom encryption policy).

When cloning an encrypted VM, you can decrypt the destination unencrypted VM by changing the storage policy of the destination VM to a nonencryption policy.

When cloning an encrypted VM, you rekey the destination encrypted VM by using PowerCLI to rekey VMs. You can perform either a shallow or a deep rekey.

A shallow rekey (or recrypt) changes the key encryption key (KEK) associated with the VM. The data encryption key (DEK) is wrapped with the new KEK.

Deep rekey (or recrypt) changes the DEK and recrypts all data using the new DEK. This function requires temporary disk space and is slower than a shallow rekey because all data must be rewritten.

7-9 Encrypting or Decrypting the Destination VM

To encrypt or decrypt the destination VM, change the VM's storage policy.

sa-sourceVM_01 - Clone Existing Virtual Machine

✓ 1 Select a name and folder

✓ 2 Select a compute resource

3 Select storage

4 Select clone options

5 Ready to complete

Select storage

Select the storage for the configuration and disk files

Select virtual disk format: Same format as source

VM Storage Policy:

Name	Capacity	Provisioned	Free	Type	Cluster
Storage Compatibility: Compatible					
ISCSI-01	58.75 GB	512.96 GB	18.14 GB	VMFS 6	
ISCSI-02	58.75 GB	61.86 GB	48.42 GB	VMFS 6	
ISCSI-03	58.75 GB	291.18 GB	34.94 GB	VMFS 6	

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

Decrypting a VM can be a security risk. Therefore, as a best practice, restrict most users from the privilege to decrypt VMs.

7-10 System Requirements for Encryption, Decryption, and Recryption

For encrypting, decrypting, and recrypting VMs during cloning, the following requirements must be met:

- A key management server (KMS) server must be configured in vCenter Server.
- If cloning occurs across ESXi hosts, both source and destination ESXi hosts must be vSphere 7.

7-11 Cross vCenter Cloning of Encrypted VMs

In vSphere 7, you can perform cryptographic operations across vCenter Server instances.

When cloning an unencrypted VM across vCenter Server instances, you can encrypt the destination VM.

When cloning an encrypted VM across vCenter Server instances, you can decrypt or reencrypt the destination VM.

VMs can be powered on or powered off when cloning across vCenter Server instances.

7-12 Cross vCenter Migration of Encrypted VMs

For migrations, certain rules apply:

- You can migrate an encrypted VM across vCenter Server instances.
- You cannot decrypt the VM during the migration.
- You cannot reencrypt the VM during the migration.

VMs can be powered on or powered off when migrating across vCenter Server instances.

During the migration operation, you cannot change keys and, therefore, reencrypting the VM during the migration is not supported.


You also cannot change the storage policy, which means that decrypting the VM during the migration is not supported.

7-13 Architecture for Cross vCenter Server Migrations and Clones (1)

Source and destination vCenter Server instances must be connected to the same KMS cluster.


Encrypted VMs, with keys in the shared KMS cluster, can be migrated or cloned between vCenter Server instances.

Both source and destination ESXi hosts must be vSphere 7.



7-14 Architecture for Cross vCenter Server Migrations and Clones (2)

Encrypted VMs, with keys in a dedicated KMS cluster, cannot be migrated or cloned to another vCenter Server instance.




VM keys must reside in the shared key management server (KMS) cluster for VM migrations and cloning to occur between vCenter Server instances.

7-15 Architecture for Cross vCenter Server Migrations and Clones (3)

If the destination host is not in encryption mode, verify that the appropriate KMS is also available to the vCenter Server instance that manages the destination host.

You can set the default KMS cluster to the shared KMS cluster or a different KMS cluster.



7-16 Shallow Recrypts of VMs with Snapshots

In vSphere 7, you can perform a shallow rekey of an encrypted VM with snapshots while the VM is powered on or off.

You can encrypt only VMs with a single snapshot branch (multiple branches are not supported).

In versions earlier than vSphere 7, you can perform a shallow rekey (recrypt) of an encrypted VM without snapshots while the VM is powered on or off.

A deep recryption replaces all keys in a VM with new keys, from top to bottom. This operation requires rewriting all encrypted data in the VM.

A shallow recryption replaces only the top-level managed keys in a VM, namely, the VMDiskKey. Only a small amount of encrypted data must be rewritten.

As in vSphere 6.7, you use PowerCLI to perform the shallow rekey.

If the recryption operation fails before all snapshots are recrypted in the chain, some snapshots are encrypted with the new key and other snapshots use the old key. To fix this problem, perform the shallow recryption again.

7-17 Encrypted VMs with Instant Clone Technology

In vSphere 7, Instant Clone Technology supports cloning an encrypted VM (the parent VM).

You can create an instant clone of a VM by using PowerCLI or API calls.

By default, the encrypted instant clone VM inherits the same keys as the parent VM.

This cloning process has limitations:

- During the instant clone operation, the encrypted VM cannot be reencrypted or decrypted.
- After you create an instant clone of an encrypted VM, you cannot reencrypt (shallow or deep) or decrypt the source VM or the destination VM.

You cannot change the storage policy or change encryption keys during an instant clone operation.

7-18 Pushing Keys to Hosts

In versions before vSphere 7, a cryptographic operation on an ESXi host works like this:

- The host must be enabled for encryption.
- When one host in a cluster enters crypto-safe mode, all the other hosts in the cluster must enter crypto-safe mode.
- vCenter Server retrieves the necessary keys for all hosts in the cluster.

In vSphere 7, a cryptographic operation works like this:

- If vSphere HA is enabled on a cluster, with at least one encrypted VM on a host in the cluster, vCenter Server pushes the necessary VM keys to all hosts in the cluster.
- If vSphere HA is not enabled, vCenter Server pushes VM keys only to the ESXi hosts that VM and its disks are registered on.

vSphere DRS does not affect the cryptographic operation. Only vSphere HA is checked.

7-19 VM Encryption Events and Alarms

vSphere 7 logs warnings and errors that provide information about the running state of vCenter Server and the VM encryption environment:

- Warning event:
 - Encrypted core dump exists.
- Critical events:
 - Insufficient disk space to perform the cryptographic operation
 - Key creation failed on the KMS cluster.
 - KMS server is unreachable.

These events help you find possible issues.

7-20 Warning: Encrypted Core Dump Exists (1)

The encrypted core-dump event is generated whenever an encrypted hostd, vpxa, or application core dump is detected.

The encrypted core-dump event is useful in the following situations:

- You need a core dump to troubleshoot a problem: The event provides the name and location of the core-dump file.
- You perform periodic maintenance and want to free up disk space: The event identifies a core dump, and you can delete the core dump if it is not needed.


Log entries are written to the vCenter logs.

For an encrypted hostd core dump, a log entry is written to the hostd logs on the ESXi host. `/var/log/hostd.log`.

As with vSphere 6.x, you can decrypt or reencrypt an encrypted core dump on your ESXi host by using the `crypto-util` CLI.

7-21 Warning: Encrypted Core Dump Exists (2)

vCenter Server generates and captures warning events when an encrypted hostd core dump exists. The message is written to the hostd logs.



The screenshot shows the vCenter Server Monitor tab for host 10.192.233.101. The left sidebar shows the navigation menu with 'Events' selected. The main table lists events, with the following data:

Description	Type	Date Time	Task	Target	User	Event Type ID
/bin/hostd crashed (1 time(s) so far) and an encrypted core file using keyid vmware:key/fqid/<VMWARE-NULL>/testcluster/14 might have been created at /var/core/hostd-zdump.000. This might have caused connections to the host to be dropped.	Warning	06/11/2019, 12:11:08 PM		10.192.233.101		esx.problem.hostd.core.d...
User dcui@127.0.0.1 lo...	Information	06/11/2019, 12:08:42 PM		10.192.233.101	dcui	vim.event.UserLoginSessi...
User dcui@127.0.0.1 lo...	Information	06/11/2019, 12:08:42 PM		10.192.233.101	dcui	vim.event.UserLoginSessi...
Alarm 'Host error' on t...	Information	06/11/2019, 12:08:16 PM		10.192.233.101		vim.event.AlarmSnmppCo...
Alarm 'Host error' on t...	Information	06/11/2019, 12:08:16 PM		10.192.233.101		vim.event.AlarmActionTrig...
UserDump: CoreDum...	Warning	06/11/2019, 12:08:09 PM		10.192.233.101		com.vmware.vc.StatelessA...
User dcui@127.0.0.1 lo...	Information	06/11/2019, 12:02:01 PM		10.192.233.101	dcui	vim.event.GeneralHostWa...
User dcui@127.0.0.1 lo...	Information	06/11/2019, 12:02:01 PM		10.192.233.101	dcui	vim.event.UserLogoutSes...

The event details for the warning are shown below the table:

Date Time: 06/11/2019, 12:11:08 PM
Type: Warning
Target: 10.192.233.101

Description: /bin/hostd crashed (1 time(s) so far) and an encrypted core file using keyid vmware:key/fqid/<VMWARE-NULL>/testcluster/14 might have been created at /var/core/hostd-zdump.000. This might have caused connections to the host to be dropped.

Related events: There are no related events.

7-22 Critical Event: Insufficient Disk Space

A critical event is posted when encryption fails because of insufficient disk space.

If disk space is low, and you try to encrypt or decrypt a powered-off VM, this event might occur. Twice the amount of disk space is required temporarily to complete encrypting or decrypting operations.

Locate the datastore where the encrypted VM resides and monitor its disk space.

The screenshot displays the vSphere Client interface for the environment 'sc2-rdops-vm03-dhcp-126-175.eng.vmware.com'. The left sidebar shows the navigation tree with 'Datacenter' expanded, containing '10.161.81.98 (Not responding)', 'testVMNext (disconnected)', and '10.192.52.82'. The main pane shows the 'Monitor' tab for the selected VM, displaying a list of events. A critical event is highlighted: 'Cryptographic operation on virtual machine [vm.VirtualMachine:vm-1005,New Virtual Machine] failed due to insufficient disk space...'. Below this, a task is listed: 'Reconfigure virtual machine' with a status of 'Insufficient disk sp...'. The bottom section shows a table of recent tasks.

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Reconfigure virtual machine	New Virtual Machine	Insufficient disk sp...		VSPHERE.LOCAL/Admini...	11 ms	06/10/2019, 4:41:15 PM	06/10/2019, 4:41:15 PM	sc2-rdops-vm03-dhcp-1...
Create virtual machine	Datacenter	Completed		VSPHERE.LOCAL/Admini...	5 ms	06/10/2019, 4:40:43 PM	06/10/2019, 4:40:46 PM	sc2-rdops-vm03-dhcp-1...

7-23 Critical Event: Key Generation Failed

A critical event is posted when the KMS fails to generate a new key for vCenter Server. The event is also captured in the vCenter Server logs.

The screenshot displays the vSphere Client interface for the environment 'sc2-rdops-vm03-dhcp-126-175.eng.vmware.com'. The left sidebar shows the navigation tree with 'Datacenter' expanded, listing 'testVMNext (disconnected)' and 'New Virtual Machine'. The main pane shows the 'Events' tab, displaying a list of events. A red box highlights the event: 'Key creation failed on KMS cluster cluster1 with error code QLC_ERR_COMMUNICATE. Check log for details.' The event is of type 'Error' and occurred on 06/10/2019 at 4:48:00 PM. Below the event list, the 'Recent Tasks' table is visible, showing two tasks: 'Reconfigure virtual machine' and 'Create virtual machine', both completed successfully.

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Reconfigure virtual machine	vmf1	Completed	Executing callbacks	VSPHERE.LOCAL\Admini...	19 ms	06/10/2019, 4:48:00 PM	06/10/2019, 4:48:01 PM	sc2-rdops-vm03-dhcp-1...
Create virtual machine	Datacenter	Completed		VSPHERE.LOCAL\Admini...	12 ms	06/10/2019, 4:47:49 PM	06/10/2019, 4:47:50 PM	sc2-rdops-vm03-dhcp-1...

Search for Key creation failed in the `/var/log/vmware/vpxd/vpxd.log` file.

Verify that the KMS is powered on and operating properly.

Verify that the KMS and vCenter Server can communicate with each other.

7-24 Critical Event: KMS Unreachable

The KMS Unreachable critical event is posted when vCenter Server detects that its connection with the KMS is down.

If the connection is down, the KMS health status alarm is triggered.

If the KMS health status alarm is triggered, take the following actions:

- Check vCenter Server logs for related messages.
- Verify that the KMS server is not down or powered off.
- Check the connection between the KMS server and vCenter Server:
 - If a proxy is required to reach your KMS, check that it is still reachable.
 - Verify that you are using the correct IP addresses or FQDNs.

7-25 Lab 6: Cloning and Decrypting VMs

Use the vSphere Client to clone an encrypted VM to an unencrypted copy:

1. Verify That the Source VM Is Encrypted
2. Clone the VM from the Storage Policy with Encryption to the Storage Policy with No Encryption
3. Verify That the Cloned VM Is Not Encrypted

7-26 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe enhancements to cloning, migration, and encryption operations
- Identify the enhancements to the VM encryption architecture
- Describe how to enable encryption mode on ESXi hosts in vSphere 7
- Recognize new VM encryption events and alarms

7-27 **Lesson 2: vSphere Trust Authority**

7-28 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe how vSphere Trust Authority secures a vSphere environment and its workloads
- Recognize use cases for vSphere Trust Authority
- Explain the general architecture of vSphere Trust Authority
- Configure vSphere Trust Authority

7-29 About vSphere Trust Authority

vSphere Trust Authority secures a vSphere environment and its workloads:

- Uses a hardware root of trust to secure the environment
- Reports on the software running on a host
- Measures or attests software using a Trusted Platform Module (TPM)

A Trust Authority administrator must configure the following parameters:

- Permitted versions of ESXi
- Trusted TPM manufacturers or devices

A trusted key provider is required. A valid attestation report is needed before access is granted to the key provider.

7-30 Use Case for vSphere Trust Authority

You can use vSphere Trust Authority to perform security tasks:

- Provide ESXi hosts with a hardware root of trust and remote attestation capabilities.
- Perform cryptographic operations on VMs with an enhanced level of encryption key management.
- Restrict encryption key management by releasing keys only to attested ESXi hosts.

In vSphere 7, you can create a trusted computing base, which consists of a secure, manageable set of ESXi hosts.

vSphere Trust Authority implements a remote attestation service for the ESXi hosts that you want to trust.

Furthermore, vSphere Trust Authority improves on Trusted Platform Module (TPM) 2.0 attestation support (added to vSphere beginning in the 6.7 release). vSphere Trust Authority implements access restrictions on encryption keys and better protects VM workload secrets.

In addition, vSphere Trust Authority allows only authorized Trust Authority administrators to configure vSphere Trust Authority services and hosts. The Trust Authority administrator can be the same user as the vSphere administrator user, or a separate user.

7-31 vSphere Trust Authority

vSphere Trust Authority improves on encryption and attestation:


- You require attestation for key access.
- vSphere Trust Authority uses a hardware root of trust.
- Key handoff is through the vSphere Trust Authority Cluster.
- You can encrypt vCenter Server when using vSphere Trust Authority

In vSphere 6.x, VM encryption and attestation work in the following ways:

- Attestation is a read-only view, and no repercussions occur when it fails.
- The attester is vCenter Server.
- Key handoff occurs through vCenter.
- You cannot safely encrypt vCenter because it forms a dependency on key access.

7-32 vSphere Trust Authority Components

vSphere Trust Authority comprises a Trust Authority Cluster and a trusted key provider. Other ESXi clusters are attested against the Trust Authority Cluster.



Several concepts apply to vSphere Trust Authority:

- The vSphere Trust Authority Cluster is a set of restricted ESXi hosts with a known good configuration.
- vSphere Trust Authority is administered by a small number of trusted administrators.
- The trusted key provider (also called key management server, or KMS) is a key provider that only the Trust Authority Cluster should know.
- The attested cluster is the Trusted Cluster. The attestation report for a cluster must be validated by the Trust Authority Cluster.

7-33 About vSphere Trust Authority Administrator


To enable vSphere Trust Authority, you must add a user to the vSphere TrustedAdmins group:

- This user becomes the Trust Authority administrator.
- You use a separate user from the vCenter Server administrator as your vSphere Trust Authority administrator.
- You must add a Trust Authority administrator for both the Trust Authority Cluster and the Trusted Cluster.
- Group membership is validated by ESXi using SAML token authentication.

7-34 ESXi Components

New services run on ESXi enabled with vSphere Trust Authority:

- Attestation Service (attestd)
- Key Provider Service (kmxd)
- Trusted infrastructure agent (kmtx)



7-35 Attestation Service

The Attestation Service on the ESXi host verifies the remote software and hardware:

- Receives an attestation report from the trusted infrastructure agent of the attesting ESXi hosts
- Verifies that a TPM device is present on the attesting ESXi hosts
- Verifies the ESXi version of the attesting ESXi hosts
- Generates a signed document describing the binary and configuration state of the attesting ESXi hosts

A Trust Authority administrator must configure:

- Which versions of ESXi to trust
- Which TPM device manufacturers to trust
- (Optional) Which specific TPM devices to trust

The Attestation Service performs the following functions:

- Generates a signed document that contains assertions describing the binary and configuration state of the remote ESXi hosts in the Trusted Cluster.
- Attests the state of the ESXi hosts using a Trusted Platform Module (TPM) 2.0 chip for software measurement and reporting. The TPM measures the software stack and configuration data and reports it to the Attestation service.
- Verifies that the software measurement signature can be attributed to a previously configured trusted TPM endorsement key (EK).
- Ensures that the software measurement matches one of a set of previously compliant ESXi images.
- Signs a SAML token that it issues to the ESXi host, providing the assertions about the identity, validity, and configuration of the ESXi host.

7-36 Key Provider Service

For the Key Provider Service to connect to a key provider, you must create a trusted key provider, configure a trust setup between the Trust Authority Cluster and the KMS.

The Key Provider Service (`kmxd`) provides a method for encapsulating encryption key sources:

- The service supports KMIP-compliant key servers.
- Additional protocols can be added in the future.
- The service allows encryption key access only if a client presents a valid attestation report.

With the Key Provider Service, vCenter Server and ESXi hosts do not require direct key management server (KMS) credentials.

In vSphere Trust Authority, an ESXi host can access an encryption key by authenticating with the Key Provider Service rather than with vCenter Server (as in vSphere 6.7).

For the Key Provider Service to connect to a KMS, the Trust Authority administrator must configure a trust setup. For most servers that are compliant with the Key Management Interoperability Protocol (KMIP), you configure a trust setup by creating client and server certificates.



7-37 Trusted Infrastructure Agent

The trusted infrastructure agent (kmsxa) communicates with a trusted Key Provider Service and Attestation Service:

- The trusted infrastructure agent retrieves encryption keys from the trusted Key Provider Service.
- The agent can be issued with encryption keys only if the ESXi host passes attestation.
- The ESXi host does not receive encryption keys from vCenter Server when the host uses a trusted key provider.



7-38 Ports Used by vSphere Trust Authority


All vSphere Trust Authority services run behind the reverse proxy and API forwarder. The trusted infrastructure agent communicates over port 443 to the vSphere Trust Authority hosts.

The following internal ports are used:

- Attestation Service: Port 7889
- Key Provider Service: Port 7888
- Trusted infrastructure agent: Port 7890


7-39 vCenter Server Topologies (1)

vSphere Trust Authority requires two or more vCenter Server systems. The Trust Authority Cluster must reside in a different vCenter Server inventory than the Trusted (attested) Cluster but can participate in Enhanced Linked Mode.



7-40 vCenter Server Topologies (2)


A single vCenter Server instance that manages both the Trust Authority Cluster and the Trusted (attested) Cluster is not supported.



7-41 vSphere Trust Authority Workflow

In the vSphere Trust Authority workflow, a Trust Authority administrator performs the following tasks:

1. Exports the identity of the hosts to be attested
2. Enables trusted host services on a Trust Authority Cluster
3. Imports the identity of the attesting hosts into the Trust Authority Cluster
4. Configures a trusted key provider
5. Exports the identity of the Trust Authority Cluster
6. Imports the Trust Authority Cluster identity into the attesting cluster
7. Configures the Trust Authority Cluster as the key provider for the attesting cluster



7-42 Enabling vSphere Trust Authority

In vSphere 7, you enable vSphere Trust Authority primarily through PowerCLI:

- PowerCLI 12.0 provides a new set of PowerCLI cmdlets.
- To run the PowerCLI cmdlets, you must be a member of the TrustedAdmins single sign-on group.



7-43 VM Encryption with vSphere Trust Authority

The user experience of VM encryption does not change when it is used with vSphere Trust Authority in vSphere 7:


- You can encrypt a VM by applying an encryption storage policy.
- When you add a virtual TPM device to a VM, the VM is encrypted.
- An attested cluster uses a trusted key provider.
- You can use a mix of a trusted key provider and a standard key provider in the same vSphere environment.



7-44 Identifying VM Encryption Key Providers

In the vSphere Client, the VM summary indicates the type of key provider that is used for encryption.

A VM that is encrypted with a trusted key provider has a slightly different icon from a VM encrypted with a standard key provider.




The `crypto-util` dictionary describe `/path/vm.vmx` does not indicate whether the VM is encrypted with a standard or trusted key provider. However, the key provider name is present in the output so you can determine the key provider type.

A VM encrypted with a standard key provider can be migrated to an attested ESXi host, but it continues to use the standard key provider.


7-45 Migrating a VM Encrypted with a Trusted Key Provider (1)

A VM that is encrypted with a trusted key provider cannot be migrated to an ESXi host that is not attested against the same vSphere Trust Authority Cluster.



7-46 Migrating a VM Encrypted with a Trusted Key Provider (2)

In the vSphere Client, a compatibility issue indicates that the destination host is not trusted and attested.



7-47 Labs

Lab: Assigning a vSphere Trust Authority Administrator

Lab: Configuring vSphere Trust Authority

Lab: Encrypting a VM with a Trusted Key Provider

7-48 Lab 7: Assigning a vSphere Trust Authority Administrator

Assign a vSphere Trust Authority Administrator role to allow a user to configure and manage vSphere Trust Authority:

1. Assign a vSphere Trust Authority Administrator

7-49 Lab 8: Enabling and Configuring vSphere Trust Authority

Enable and configure vSphere Trust Authority:

1. Export the TPM Certificate and ESXi Image Metadata
2. Export the Trusted User Principal
3. Enable vSphere Trust Authority Services
4. Import the Trusted Host Information to the Trust Authority Cluster
5. Create a Trusted Key Provider on the Trust Authority Cluster
6. Export the Trust Authority Cluster Settings
7. Import the Trust Authority Cluster Settings into the Trusted Hosts Cluster
8. Configure the Trusted Key Provider for the Trusted Hosts Cluster

7-50 Lab 9: Encrypting a VM with a Trusted Key Provider

Encrypt a VM with a trusted key provider:

1. Encrypt a VM with a Trusted Key Provider

7-51 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe how vSphere Trust Authority secures a vSphere environment and its workloads
- Recognize use cases for vSphere Trust Authority
- Explain the general architecture of vSphere Trust Authority
- Configure vSphere Trust Authority

7-52 Key Points

- In vSphere 7, you can perform encryption and decryption operations on a VM during cloning.
- You can perform cross vCenter migrations and cloning of encrypted VMs.
- vSphere Trust Authority verifies that ESXi hosts are running trusted software and releases encryption keys only to trusted ESXi hosts.
- Software is attested (measured) using a Trusted Platform Module (TPM).
- The user experience of VM encryption does not change when it is used with vSphere Trust Authority in vSphere 7.

Questions?

Module 8

Cluster Enhancements

8-2 Importance

vSphere DRS provides automated resource balancing on DRS enabled clusters. Using the Cluster Quickstart workflow, you can enable vSphere features at the cluster level in vCenter Server. Understanding how these new and enhanced features work in vSphere 7 can help you maximize the scalability of your VM resources.

8-3 Module Lessons

1. Cluster Quickstart Workflow
2. vSphere DRS Enhancements

8-4 **Lesson 1: Cluster Quickstart Workflow**


8-5 Learner Objectives

After completing this lesson, you should be able to meet the following objective:

- Recognize how to configure clusters using the vSphere Cluster Quickstart workflow

8-6 Cluster Quickstart Workflow

In the vSphere Client, you can use the Cluster Quickstart workflow to create, configure, and expand clusters.



The Cluster Quickstart workflow guides you through the process of deploying clusters. It covers every aspect of the initial configuration, such as host, network, and vSphere settings.

Cluster Quickstart reduces the amount of time it takes to configure a cluster.

8-7 Creating a Cluster

When you create a cluster, you can enable cluster services:

- vSphere DRS
- vSphere HA
- vSAN

The steps of the Cluster Quickstart workflow vary depending on the selected services.

New Cluster

Datacenter

×

Name	Cluster-01
Location	 Datacenter
 vSphere DRS	☒
 vSphere HA	☒
vSAN	☒

These services will have default settings - these can be changed later in the Cluster Quickstart workflow.

☐ Manage image setup and updates on all hosts collectively 

Requires hosts to have ESXi 7.0 or higher. Will replace Update Manager. Cannot be turned off.

CANCEL

OK

8-8 Using Cluster Quickstart Workflow to Add Hosts (1)

On the Cluster quickstart page, you click **Add** to add hosts to your cluster.

The screenshot shows the vSphere Cluster Quickstart workflow for Cluster-01. The interface includes a left-hand navigation pane with sections for Services, Configuration, Trust Authority, and vSAN. The main content area is titled 'Cluster quickstart' and contains three steps: 1. Cluster basics (marked with a green checkmark), 2. Add hosts, and 3. Configure cluster. Step 1 lists selected services: vSphere DRS, vSphere HA, and vSAN, with an 'EDIT' button below. Step 2, 'Add hosts', has an 'ADD' button highlighted with an orange border and a 'RE-VALIDATE' button. Step 3 has 'CONFIGURE' and 'RE-VALIDATE' buttons. A 'SKIP QUICKSTART' button is located in the top right corner. The top navigation bar includes tabs for Summary, Monitor, Configure (selected), Permissions, Hosts, VMs, Datastores, Networks, and Updates.

Cluster-01 ACTIONS

Summary Monitor **Configure** Permissions Hosts VMs Datastores Networks Updates

Services

- vSphere DRS
- vSphere Availability

Configuration

- Quickstart**
- General
- Security
- Licensing
- VMware EVC
- VM/Host Groups
- VM/Host Rules
- VM Overrides
- I/O Filters
- Host Options
- Host Profile

Trust Authority

- Alarm Definitions
- Scheduled Tasks

vSAN

Cluster quickstart

SKIP QUICKSTART

We have collected some common configuration tasks to make it easier to get your cluster up and running. If you prefer to configure your cluster manually, you can choose not to use this automated workflow.

1. Cluster basics

Selected services:

- vSphere DRS
- vSphere HA
- vSAN

EDIT

2. Add hosts

Add new and existing hosts to your cluster.

ADD RE-VALIDATE

3. Configure cluster

Configure network settings for vMotion and vSAN traffic, review and customize cluster services, and set up a vSAN datastore.

CONFIGURE RE-VALIDATE

8-9 Using Cluster Quickstart Workflow to Add Hosts (2)

You enter the IP address or FQDN of each host that you want to add to the cluster. From the inventory, you can also select existing hosts that are not part of a cluster.

The left screenshot shows the 'Add hosts' dialog with the 'Add new and existing hosts to your cluster' tab selected. The 'New hosts (2)' tab is active, and the 'Use the same credentials for all hosts' checkbox is checked. The table below shows the hosts being added:

IP address or FQDN	Username	Password
sa-esxi-01	root	*****
sa-esxi-02	root	*****

The right screenshot shows the 'Add hosts' dialog with the 'Add new and existing hosts to your cluster' tab selected. The 'Existing hosts (2 from 2)' tab is active, and the 'Use the same credentials for all hosts' checkbox is checked. The table below shows the existing hosts being added:

Hostname / IP Address	ESX Version	Model
sa-esxi-04.vclass1.ocli	7.0.0	VMware, Inc. VMware ESX Platform
sa-esxi-03.vclass1.ocli	7.0.0	VMware, Inc. VMware ESX Platform


By selecting the **Use the same credentials for all hosts** check box, you can streamline the process of adding hosts to the cluster.

If the system cannot validate a host, you are prompted to manually validate its certificate and accept its thumbprint in the Security Alert pop-up.

The Host summary page lists all hosts that are added to the cluster and related warnings.

8-10 Using Cluster Quickstart Workflow to Configure a Cluster

To configure the host networking settings and services on your cluster by using the Cluster Quickstart workflow, click **CONFIGURE**.



If you click **SKIP QUICKSTART**, you can manually configure your cluster by using the menus in the vSphere Client. By skipping the Cluster Quickstart workflow, you cannot restore it. Any hosts added to this cluster must be configured manually.

8-11 Configuring a Cluster: Distributed Switches

Using the Cluster Quickstart workflow, select up to three distributed switches.

When you enable vSAN, default networks are created for vSphere vMotion and for vSAN.

When creating a distributed switch, select at least one physical adapter.

Configure cluster

1 Distributed switches

2 vMotion traffic

3 Storage traffic

4 Advanced options

5 Claim disks

6 Proxy settings

7 Review

Distributed switches

Configure the distributed switches

Name	Port Groups	Uplinks	
DSwitch	USE EXIS...	2	5

Port groups

The following default port groups will be assigned to the distributed switch.

vMotion network

DSwitch ▾

DSwitch-vMotion

vSAN network

DSwitch ▾

DSwitch-vSAN

Physical adapters

One uplink port group will be created on each switch containing all the specified physical adapters.

Adapter 0 (vmnic0)

Not in use ▾

Adapter 1 (vmnic1)

DSwitch ▾

Adapter 2 (vmnic2)

DSwitch ▾

CANCEL NEXT

Alternatively, you can select the **Configure networking settings later** check box to configure the default settings only for the cluster services and to hide all options that are related to host networking. But, by selecting this option, you cannot perform the networking configuration by using the Configure cluster wizard.

8-12 Configuring a Cluster: vMotion Traffic

On the vMotion traffic page, you enter the IP address information for the vSphere vMotion traffic.

Configure cluster

1 Distributed switches

2 vMotion traffic

3 Storage traffic

4 Advanced options

5 Claim disks

6 Proxy settings

7 Review

vMotion traffic

Specify the IP addresses for the vMotion traffic

Distributed switch

DSwitch

Distributed port group name

DSwitch-vMotion

☒ Use VLAN

12

Protocol

IPv4

IPv4 Configuration

IP type

Static IPs

Each host is configured automatically based on the input below. Empty default gateway might result in a segmented network.

sa-esxi-01.vclass.local	172.20.12.51	255.255.255.0
	Default gateway	AUTOFILL
sa-esxi-02.vclass.local	172.20.12.52	255.255.255.0
	Default gateway	
sa-esxi-03.vclass.local	172.20.12.53	255.255.255.0
	Default gateway	
sa-esxi-04.vclass.local	172.20.12.54	255.255.255.0

CANCEL

BACK

NEXT

You enter the information for the first host and click **AUTOFILL** to complete the fields for the following hosts.

8-13 Configuring a Cluster: Storage Traffic

On the Storage traffic page, you enter the IP address information for the storage traffic.

Configure cluster

1 Distributed switches

2 vMotion traffic

3 Storage traffic

4 Advanced options

5 Claim disks

6 Proxy settings

7 Review

Storage traffic

Specify the IP addresses for the vSAN traffic

Distributed switch

DSwitch

Distributed port group name

DSwitch-VSAN

☒ Use VLAN

14

Protocol

IPv4

IPv4 Configuration

IP type

Static IPs

Each host is configured automatically based on the input below.

sa-esxi-01.vclass.local	172.20.14.51	255.255.255.0	AUTOFILL
sa-esxi-02.vclass.local	172.20.14.52	255.255.255.0	
sa-esxi-03.vclass.local	172.20.14.53	255.255.255.0	
sa-esxi-04.vclass.local	172.20.14.54	255.255.255.0	

CANCEL

BACK

NEXT

8-14 Configuring a Cluster: Advanced Options

On the Advanced options page, you get different settings depending on the cluster services that are enabled:

- High Availability (optional)
- Distributed Resource Scheduler (optional)
- Host Options
- Enhanced vMotion Compatibility

Configure cluster

- 1 Distributed switches
- 2 vMotion traffic
- 3 Storage traffic
- 4 Advanced options**
- 5 Claim disks
- 6 Proxy settings
- 7 Review

Advanced options ✕

Customize the cluster settings.

▼ vSphere HA	
Host Failure Monitoring	☒
Virtual Machine Monitoring	☐
Admission Control	☒
Host failure to tolerate	1 ▼ The maximum is one less than number of hosts in cluster.
▼ vSphere DRS	
Automation level	Fully Automated ▼
Migration threshold	3 ▼ ⓘ
▼ vSAN Options	
Deployment type	Single site cluster ▼
Data-At-Rest Encryption	☐
KMS Cluster	▼
Deduplication and compression	☐ Requires all-flash configuration
Fault domains	☐

CANCEL BACK NEXT

8-15 Configuring a Cluster: Claim Disks

On the Claim disks page, you can claim cache and capacity disks for vSAN on each host. You manually create disks groups after the cluster is created.

Configure cluster

1 Distributed switches

2 vMotion traffic

3 Storage traffic

4 Advanced options

5 Claim disks

6 Proxy settings

7 Review

Claim disks

Select disks to contribute to the vSAN datastore.

Claim disks on hosts for cache and capacity. Non-empty disks will be deleted.

Claimed capacity40.00 GB

Claimed cache20.00 GB

Unclaimed storage0.00 B

CLAIM DISK FOR

DISK DRIVE TYPE

Group by: Disk model/size

Disk Model/Serial Number	Claim For	Drive Type	Disk Distribution/
VMware Virtual disk, 5.00 ...	Cache tier	Flash	1 disk on 4 hosts
Local VMware Disk (mpx....	Cache tier	Flash	sa-esxi-01.v...
Local VMware Disk (mpx....	Cache tier	Flash	sa-esxi-02.v...
Local VMware Disk (mpx....	Cache tier	Flash	sa-esxi-04.v...
Local VMware Disk (mpx....	Cache tier	Flash	sa-esxi-03.v...
VMware Virtual disk, 10.0...	Capacity tier	HDD	1 disk on 4 hosts
Local VMware Disk (mpx....	Capacity tier	HDD	sa-esxi-01.v...
Local VMware Disk (mpx....	Capacity tier	HDD	sa-esxi-02.v...
Local VMware Disk (mpx....	Capacity tier	HDD	sa-esxi-04.v...
Local VMware Disk (mpx....	Capacity tier	HDD	sa-esxi-03.v...

10 items

Configuration correct.

CANCEL

BACK

NEXT

8-16 Configuring a Cluster: Proxy Settings

To send data for the Customer Experience Improvement Program (CEIP), configure a proxy server.

Configure cluster

1 Distributed switches

2 vMotion traffic

3 Storage traffic

4 Advanced options

5 Claim disks

6 Proxy settings

7 Review

Proxy settings

×

Configure proxy to establish connection to send data for CEIP. vSAN Support Insight requires to allow outbound traffic to https://vcsa.vmware.com:443/ph/api/* and http://www.vmware.com:80/*

☐ Configure the Proxy server if your system uses one

Host name: *

Port: *

User name:

Password:


CANCEL

BACK

NEXT

8-17 Additional Cluster Configuration

After you complete the Cluster Quickstart workflow, you can configure additional settings, review and customize cluster services, and set up a vSAN datastore.




The Cluster Quickstart workflow uses the vSAN health services to validate the configuration and to help you correct configuration issues. It performs checks and makes recommendations according to VMware best practices to help ensure the highest levels of stability and performance.

8-18 Using Quickstart to Expand Clusters

You can quickly expand an existing cluster while maintaining consistency and adhering to VMware Validated Design recommendations.

For example, you can add hosts to a vSAN cluster, keep the hosts uniformly configured, and avoid disrupting any ongoing operations:

- Apply the existing networking settings.
- Claim disks on each new host.
- Place the new hosts into fault domains.



All hosts in a cluster should have the same or a similar hardware configuration. The Cluster Quickstart workflow simplifies the process of keeping HCI clusters consistent.

8-19 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objective:

- Recognize how to configure clusters using the vSphere Cluster Quickstart workflow

8-20 **Lesson 2: vSphere DRS Enhancements**

8-21 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize improvements to vSphere DRS
- Identify improvements to vSphere ESX Agent Manager
- Describe scalable shares
- Configure vSphere DirectPath I/O

8-22 DRS in vSphere 6.x

In vSphere 6.x, the DRS load-balancing algorithm runs every 5 minutes to study the host load balance in the cluster.

When a cluster becomes imbalanced, vSphere DRS makes recommendations or migrates VMs based on the aggregated cluster capacity:

- Identifies VM migrations that improve the host load balance
- Uses vSphere vMotion to migrate VMs between the hosts in the cluster

8-23 vSphere 6.x DRS Summary

The **Summary** tab of a cluster provides an overview of the configuration and status of DRS. It also shows the current level of host load balance in the cluster.

vSphere DRS

Balanced

Migration automation level:

Fully Automated

Migration threshold:

Apply priority 1, priority 2, and priority 3 recommendations.

Power management automation level:

Off

DRS recommendations:

0

DRS faults:

0

8-24 DRS in vSphere 7

In vSphere 7, the DRS load-balancing algorithm runs every minute. It focuses on the VM DRS score to make recommendations or to migrate VMs:

- Uses VM metrics to calculate the VM DRS score and cluster DRS score
- Identifies VM migrations that improve the VM DRS score
- Uses vSphere vMotion to migrate VMs between the hosts in the cluster

By performing migrations, DRS ensures that VMs receive the required resources.


8-25 VM DRS Score

The VM DRS score is a new metric that measures the execution efficiency of a virtual machine.

VM DRS score considerations:

- VM resource requirements fulfillment
- ESXi host resource capacity

DRS migrates VMs to the host that provides the highest score and ensures fair resource allocation to all VMs.



The VM DRS score roughly represents the percentage of time that a VM has its resource requirements satisfied over a sample period. The scores are filed into one of five buckets. These buckets are 20 percent ranges. A score of 80-100% indicates little to no resource contention.

DRS evaluates the current performance of a VM. Then, it runs what-if scenarios on compatible hosts to evaluate the cost of the VM migration and the potential performance. If the migration helps the VM run more efficiently, DRS migrates the VM to the new host.


Because of the new focus on the VM DRS Score, more migrations are likely after DRS is enabled. The cluster must be normalized to ensure that VMs have the required resources regardless of the individual host that they are placed on.

The Cluster DRS and VM DRS scores replace the balance deviation level in the cluster **Summary** tab.

8-26 Cluster DRS Score

The Cluster DRS score is a new metric that averages the individual scores of VMs.


DRS focuses on maximizing the VM execution efficiency rather than on the host load balance in the cluster.



8-27 DRS Settings and Score List


On the vSphere 7 DRS portlet, you can access links for the DRS parameters and score list:

- Select **VIEW DRS SETTINGS** to view the DRS parameters and their values.
- Select **VIEW ALL VMS** to access the new VM DRS Score list.



8-28 VM DRS Score List

In vSphere 7, DRS provides a listing of the VM DRS scores and more detailed metrics for all cluster VMs.

The screenshot shows the vSphere 7 interface for a cluster named 'SA-Cluster'. The 'Monitor' tab is selected, and the left-hand navigation pane is expanded to 'vSphere DRS', where 'VM DRS Score' is highlighted with a red rectangle. The main table displays a list of VMs with their DRS scores, active CPU, and used CPU. The table has columns for Name, DRS Score (sorted descending), Active CPU, and Used CPU. The VMs are listed in descending order of DRS score, with most at 94% and a few at 92% and 91%.

Name	DRS Score ↓	Active CPU	Used CPU
sa-vm-wln-110	94%	20 MHz	20 MHz
sa-vm-wln-114	94%	20 MHz	20 MHz
sa-vm-wln-123	94%	20 MHz	20 MHz
sa-vm-wln-112	94%	41 MHz	41 MHz
sa-vm-wln-22241	94%	41 MHz	20 MHz
sa-vm-wln-16	94%	20 MHz	20 MHz
SA-VM-Win-01	94%	41 MHz	20 MHz
sa-vm-wln-22214	94%	20 MHz	20 MHz
sa-vm-wln-22242	94%	20 MHz	20 MHz
New Virtual Machine	94%	20 MHz	20 MHz
sa-vm-wln-117	94%	20 MHz	20 MHz
sa-vm-wln-02	92%	20 MHz	20 MHz
sa-vm-wln-22250	92%	41 MHz	20 MHz
sa-vm-wln-22224	92%	41 MHz	41 MHz
sa-vm-wln-2226	92%	20 MHz	20 MHz
sa-vm-wln-2225	92%	20 MHz	20 MHz
sa-vm-wln-22234	92%	41 MHz	41 MHz
sa-vm-wln-22212	92%	41 MHz	20 MHz
sa-vm-wln-22230	92%	20 MHz	20 MHz

A VM DRS score is computed from an individual VM's CPU, memory, and network metrics. You can use these metrics to gauge the goodness or wellness of the VM. The metrics are useful in determining which hosts are experiencing resource contention and of what type.

8-29 Cluster DRS Score in Performance Charts

In vSphere, DRS includes a new metric in the performance charts. The new metric is the cluster DRS score.

This metric shows the cluster DRS score over the time span in the advanced performance chart.

Chart Options

SA-Cluster

×

Chart options: --Select option-- Save Options As... Delete Options

Chart Metrics

CPU

Cluster services

Memory

Virtual machine operations

Select counters for this chart:

	Rollups	Units	Internal Name	Stat Type	Description
☒ Counters					
☒ DRS Score	Latest	%	clusterDrsScore	Absolute	DRS score of the cluster
☐ Effective CPU resources	Average	MHz	effectivecpu	Rate	Total available CPU resources o...
☐ Effective memory resourc...	Average	MB	effectivemem	Absolute	Total amount of machine memo...

Timespan:

Last day

Last: 5

Hour(s)

From: 01/08/2020

12:03:17

To: 01/09/2020

12:03:17

(time is in ISO 8601 format)

Chart Type:

Line Graph

Select object for this chart:

☒ Target Objects


☒ SA-Cluster

CANCEL

OK

8-30 Cluster DRS Score Chart

A cluster DRS score uses performance charts.



8-31 Agent VMs

An agent VM is a virtual machine that performs a specific function for the virtual infrastructure.


For example, an agent can act as a firewall or gather information about the virtual machines on the host.

Using the ESX Agent Manager (EAM) service, vCenter Server can deploy and manage agent VMs.

8-32 Agent VM Management in vSphere 6.x

In vSphere 6.x, agent VMs are pinned to an ESXi host.


vSphere DRS does not migrate agent VMs.



8-33 Agent VM Management in vSphere 7

In vSphere 7, agent VMs are pinned to an ESXi cluster.


vSphere DRS can migrate agent VMs between ESXi hosts in the cluster.



If an ESXi host fails, vSphere HA always starts agent VMs first. This restart priority is nonconfigurable.

8-34 Resource Pool Shares

A resource pool is a logical abstraction of hierarchically managed CPU and memory resources. Shares specify the resource pool priority relative to other resource pools.



Resource pools can be organized hierarchically. The root resource pool is the topmost resource pool. The root resource pool includes the sum for all CPUs (in megahertz) and the sum of all the installed RAM (in megabytes) available in the compute environment (standalone host or cluster).


A child resource pool is used to allocate resources from the parent resource pool for the child's consumers. A child resource pool cannot exceed the capacity of the parent resource pool. By creating a child pool, you reserve resources from the parent pool. In the example, the root resource pool is Cluster-A. Cluster-A has two child resource pools: the Engineering pool and the Finance pool.

Shares specify the relative priority or Importance of a resource pool. In the example, you want the finance VMs to have twice as much processor time as the engineering VMs. You can configure the Finance pool with twice as many CPU shares as the Engineering pool.

You can also specify shares per VM and distribute the resource allocation within the resource pool. VM shares are relative only to other powered-on VMs within the same resource pool. In the example, all VMs have an equal number of shares. However, the Finance VMs are entitled to about 67% of the cluster processor time. Each Finance VM gets about 33% of the entire cluster CPU resources, or twice as much as for engineering VMs.

8-35 Diluted Shares

In vSphere 6.x, when you power on more VMs in a resource pool, you can dilute the resource allocation for each VM in the pool.




When you power on additional VMs in a resource pool, the resource allocation per VM is recalculated based on the value configured for shares. In the example, two more VMs are powered on in the Finance pool. All VMs in the cluster have an equal number of shares. Therefore, the resources are divided by four.

Because the finance VMs are entitled to about 68% of the cluster processor time, each finance VM gets about 17% of the entire cluster CPU resources. The value of shares within the Finance pool is diluted and is no longer twice as much as the engineering VMs.

8-36 About Scalable Shares

In vSphere 7, you can enable scalable shares to dynamically scale up and down the shares of sibling resource pools and maintain the resource allocation ratio.



When you enable scalable shares on a resource pool, all child resource pools maintain the original ratio of shares configured in the parent pool.

In the example, you enable scalable shares on Cluster-A to dynamically adjust the resource allocation of child resource pools. All VMs have an equal number of shares, so the shares configuration of the child resource pools adjusts to give every finance VM twice as much as the engineering VMs. The shares configuration of the child resource pools adjusts to 80% for the Finance pool, and 20% for the Engineering pool. Each finance VM gets 20% of the Cluster-A processor time, and the engineering VMs get 10%.

8-37 Enabling Scalable Shares

You can enable scalable shares at the cluster or resource pool level. The resource allocation of child resource pools becomes scalable.

Edit Cluster SettingsCluster-A

vSphere DRS

Automation

Additional Options

Power Management

Advanced Options

VM Distribution

☐ Enforce a more even distribution of virtual machines across hosts in the cluster for availability (may see a degradation of DRS when utilizing this setting).

CPU Over-Commitment ⓘ

☐ Enable

Over-commitment ratio: 0 :1 (vCPU:pCPU)

Scalable Shares

☒ Enable scalable shares for the resource pools on this cluster

CANCEL

OK

New Resource PoolCluster-A

Name

Finance Pool

Scale Descendant's Shares ⓘ

☒ Yes, make them scalable

CPU

Shares

High

8000

Reservation

0

MHz

Max reservation: 5,712 MHz

Reservation Type

☒ Expandable

Limit

Unlimited

MHz

Max limit: 5,712 MHz

Memory

Shares

Normal

163840

Reservation

0

MB

Max reservation: 9,559 MB

Reservation Type

☒ Expandable

Limit

Unlimited

MB

Max limit: 9,559 MB

CANCEL

OK


8-38 vSphere DirectPath I/O

Using vSphere DirectPath I/O, VMs can directly access the physical PCI and PCIe devices connected to the ESXi host.

In earlier versions of vSphere, VMs identify the PCI devices by using the bus slot:

- The VM is bound to the host hardware.
- DRS does not migrate the VM using vSphere vMotion.
- vSphere HA does not restart the VM on a different host.


You configure PCI devices on the host to make them available for passthrough to a VM. With PCI passthrough, VMs can directly access high-performance graphics or sound cards, bypassing the VMkernel.



8-39 Dynamic DirectPath I/O

In vSphere 7, with dynamic vSphere DirectPath I/O, VMs can identify PCI devices by using their vendor and model names:

- A pool of PCI devices that are available in the cluster can be assigned to the VM.
- At power-on, DRS performs initial placement of the VM on any host that can provide the identified PCI device.
- You cannot migrate the VM using vSphere vMotion.
- vSphere HA can restart the VM on any host that can provide the identified PCI device.



8-40 Adding a Hardware Label to PCI Devices

You can add a hardware label to a passthrough-enabled PCI device. Labels help to easily identify devices from a list.

The screenshot shows the vSphere Client interface for a host named 10.160.145.108. The left sidebar shows the navigation tree with 'Hardware' > 'PCI Devices' selected. The main area displays the 'Edit PCI Device Availability' dialog box for the host 10.160.145.108. The dialog contains a table of PCI devices with columns for ID, Status, Vendor Name, Device Name, and Hardware Label. A red box highlights the 'Hardware Label' column, which is currently empty for all devices. Another red box highlights the 'CONFIGURE PASSTHROUGH' button in the top right corner of the main window.


ID	Status	Vendor Name	Device Name	Hardware Label
0000:04:00.0	Available	Intel Corporation	Virtual 82574L Gigabit Ethernet Controller	--
0000:00:0F.0	Unavailable	VMware	SVGA II Adapter	--
0000:18:00.0	Available	Intel Corporation	Virtual 82574L Gigabit Ethernet Controller	--
0000:02:00.0	Unavailable	VMware	USB1.1 UHCI Controller	--
0000:02:01.0	Unavailable	VMware	USB2 EHCI Controller	--
0000:03:00.0	Unavailable	VMware Inc.	PVSCSI SCSI Controller	--
0000:0B:00.0	Unavailable	Intel Corporation	Virtual 82574L Gigabit Ethernet Controller	--
0000:13:00.0	Available	Intel Corporation	Virtual 82574L Gigabit Ethernet Controller	red

8-41 Adding Dynamic vSphere DirectPath I/O Devices to a VM

You can add multiple dynamic vSphere DirectPath I/O devices to a VM:

- PCI pass-through devices
- Shared PCI pass-through devices, for example, NVIDIA GRID GPUs

Full memory reservation is applied to the VMs with pass-through devices.



To add a dynamic vSphere DirectPath I/O device to a VM, go to Edit Settings for the VM and follow these steps:

1. Click **ADD NEW DEVICE** and select **PCI Device**.
2. For New PCI device, select **Dynamic DirectPath IO** and click **SELECT HARDWARE**.

A list of devices that can be attached to the VM appears.

3. Select one or more devices from the list.

8-42 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Recognize improvements to vSphere DRS
- Identify improvements to vSphere ESX Agent Manager
- Describe scalable shares
- Configure vSphere DirectPath I/O

8-43 Key Points

- The Cluster Quickstart workflow simplifies the configuration of clusters and helps ensure consistency across all the hosts in the cluster.
- In vSphere 7, DRS is VM-focused and balances the host load based on a VM DRS score.
- ESX Agent Manager can deploy agent VMs and pin them to the required hosts.
- In vSphere 7, you can enable scalable shares on a resource pool to dynamically maintain the resource allocation ratio.
- With dynamic vSphere DirectPath I/O, DRS can perform initial placement of a VM on any host that can provide a PCI device defined for the VM.

Questions?

Module 9

vSphere Lifecycle Management

9-2 Importance

Managing the life cycle of vSphere involves keeping vCenter Server and ESXi hosts up to date and integrated with other VMware and third-party solutions. To achieve these goals, you must understand how to use the new features provided by vSphere Lifecycle Manager, namely, cluster-level management of ESXi hosts and the vCenter Server Update Planner.

9-3 Module Lessons

1. vCenter Server Update Planner
2. Overview of vSphere Lifecycle Manager
3. Working with vSphere Lifecycle Manager

9-4 **Lesson 1: vCenter Server Update Planner**

9-5 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe features of vCenter Server Update Planner
- Run vCenter Server upgrade prechecks and interoperability reports
- Export prechecks and interoperability report results

9-6 Overview of vCenter Server Update Planner

In vSphere 7, you can use the Update Planner feature to plan updates to vCenter Server and other VMware products registered with it.


The Update Planner can perform the following tasks:

- Retrieve information about VMware products associated with vCenter Server.
- List available vCenter Server updates and upgrades.
- Create interoperability reports.
- Perform a precheck to verify that vCenter Server meets the update or upgrade requirements.

9-7 Update Planner Requirements

The Update Planner feature is available for vCenter Server 7.0 or later.

You must join the VMware Customer Experience Improvement Program (CEIP) to generate an interoperability or precheck report.



sa-vcsa-02.vclass.local | ACTIONS ▾

Summary **Monitor** Configure Permissions Datacenters Hosts & Clusters VMs ...

Issues and Alarms ▾
All Issues
Triggered Alarms

Tasks and Events ▾
Tasks
Events
Sessions
Security

vCenter Server ▾
Interoperability

Cloud Native Storage ▾
Container Volumes
Skyline Health

Interoperability ⓘ

⚠ This feature only works when vCenter Server Appliance on sa-vcsa-02.vclass.local has joined in Customer Experience Improvement Program. To enable this feature, please join VMware's Customer Experience Improvement Program. If you choose to participate in VMware's Customer Experience Improvement Program, it will take some time until your choice is replicated across all vCenter Server Appliances that are part of Embedded Linked Mode. If you see this error after you have joined, please retry running the operation again after several minutes.
[Learn more and enable CEIP](#)

MODIFY PRODUCT LIST EXPORT ↻

Product	Current Version	Compatible Version(s)	Release Notes
			
0 items			


When generating reports, if the Customer Experience Improvement Program (CEIP) is not accepted, a prompt describing CEIP appears. If discarded, reports are not generated.

9-8 Generating Update Planner Reports

When a new vCenter Server version is available, the new version appears on the **Updates** tab of the vSphere Client.

You can generate different reports:

- Product interoperability
- Update precheck



When new vCenter Server updates are released, the vSphere Client shows a notification in the **Summary** tab. Clicking the notification directs you to the **Updates** tab.

The **Updates** tab has a new Update Planner page. This page shows a list of vCenter Server versions that you can select.

Details include release date, version, build, and other information about each vCenter Server version available.

The Type column tells you if the release item is an update, an upgrade, or a patch. All three types appear on the same page.

If multiple versions appear, the recommended version is preselected.

After selecting a vCenter Server version from the list, you can generate product interoperability reports and preupdate reports.

9-9 Exporting Update Planner Reports

You can export reports in CSV format and use the report as a guide to prepare for an update.

Both product interoperability and precheck reports can be exported.

Product Interoperability ⓘ

+ ADD PRODUCT

EXPORT

↺

Product	Current Version	Compatible Version (s)	Release Notes
✓ vSphere Replication	<version>	7.0 ▾	Link
✓ vRealize Orchestrator	<version>	7.0 ▾	Link
⚠ ESXI 123	<version>	7.0 ▾	Link
⚠ ESXI 123	<version>	7.0 ▾	Link
✓ NSX for vSphere	<version>	7.0 ▾	Link
✓ vRealize Automation	<version>	7.0 ▾	Link
✓ Site Recovery Manager	<version>	7.0 ▾	Link

1-10 of 23 Items 1 2 3

9-10 Upgrading vCenter Server

When updating and patching vCenter Server, use the virtual appliance management interface (VAMI). When upgrading vCenter Server, run the vCenter Server Appliance installer.

The screenshot shows the vCenter Server Management interface. The top navigation bar includes the VMware logo, 'vCenter Server Management', the date and time 'Wed 02-12-2020 03:34 PM UTC', language 'English', and links for 'Help' and 'Actions'. A left sidebar contains navigation links: Summary, Monitor, Access, Networking, Firewall, Time, Services, Update (highlighted), Administration, Syslog, and Backup. The main content area is titled 'Current version details' and shows 'Product Type' as 'vCenter Server with an embedded Platform Services Controller' and 'Version' as '7.0.0.10000'. Below this, the 'Available updates' tab is selected, showing 'Last Checked Feb 12, 2020, 10:33:44 AM' and buttons for 'SETTINGS' and 'CHECK UPDATES'. A table header shows columns for Version, Type, Release Date, and Reboot Required, but the main area displays a funnel icon and the message 'No applicable update found.'

The screenshot shows the vCenter Server 7.0 Installer interface. The top bar includes the VMware logo, 'vCenter Server 7.0 Installer', and a language dropdown 'English'. The main area features four large cards with icons and text: 'Install' (Install a new vCenter Server), 'Upgrade' (Upgrade an existing vCenter Server Appliance or Upgrade and Converge an existing vCenter Server Appliance with an external Platform Services Controller), 'Migrate' (Migrate from an existing Windows vCenter Server or Migrate and Converge an existing Windows vCenter Server with an external Platform Services Controller), and 'Restore' (Restore from a previous vCenter Server Backup).

9-11 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe features of vCenter Server Update Planner
- Run vCenter Server upgrade prechecks and interoperability reports
- Export prechecks and interoperability report results

9-12 **Lesson 2: Overview of vSphere Lifecycle Manager**

9-13 Learner Objectives


After completing this lesson, you should be able to meet the following objectives:

- Describe vSphere Lifecycle Manager
- Compare vSphere Lifecycle Manager to vSphere Update Manager
- Identify the components of an ESXi image

9-14 Introduction to vSphere Lifecycle Manager

In vSphere 7, vSphere Lifecycle Manager enhances the functionality that vSphere Update Manager provided in previous vSphere releases. It adds features for managing the ESXi host life cycle at the cluster level.

vSphere Lifecycle Manager is a service that runs in vCenter Server. After deploying vCenter Server Appliance, you can access the vSphere Lifecycle Manager user interface in the HTML5-based vSphere Client.



9-15 vSphere Update Manager and vSphere Lifecycle Manager: Functionalities


Using baselines or images, vSphere Lifecycle Manager standardizes ESXi clusters, components, vendor add-ons, and drivers and firmware.

Functionality	vSphere Update Manager	vSphere Lifecycle Manager
Upgrade and patch ESXi hosts	✓	✓
Install and update third-party software on ESXi hosts	✓	✓
Manage VMware Tools updates	✓	✓
Manage virtual machine hardware upgrades	✓	✓
Standardize the use of a single image across an entire cluster		✓
Install the required ESXi version across an entire cluster		✓
Install and update third-party software and firmware across an entire cluster		✓
Manage firmware and driver updates	*	✓

*Using vSphere Update Manager, you can install and update drivers on ESXi hosts. Device firmware is supplied directly by the hardware vendor.

9-16 Elements of ESXi Images


An ESXi image consists of several elements. To maintain consistency, you apply a single ESXi image to all hosts in a cluster.



9-17 About ESXi Base Image

With every vSphere release, VMware provides you with a base image that contains the ESXi software fixes and enhancements.

The base image is the collection of components that is necessary to start an ESXi host.



9-18 About Components

A component is the logical grouping of one or more VIBs that encapsulates a functionality in ESXi.


Starting with vSphere 7, the component is the smallest unit that is used by vSphere Lifecycle Manager to install VMware and third-party software on ESXi hosts.

Adding components to a cluster image is optional.

Components are the basic packaging for VIBs and additional metadata that provides the name and version of the component.

Like a bulletin, a component might contain one or multiple VIBs. However, unlike the bulletin, the component is a logical grouping of VIBs. On installation, it provides you with a visible feature. vSphere HA, for example, is provided as a component for clusters managed by an image.

Starting with vSphere 7, you upgrade components, not VIBs.




9-19 About Vendor Add-Ons

You can customize an ESXi release by adding a vendor add-on to an ESXi base image.

When combined with an ESXi base image, the add-on can add, update, or remove components that are part of the ESXi base image.

Adding a vendor add-on to a cluster image is optional.




9-20 About Firmware and Driver Add-Ons

The firmware and drivers add-on is a vendor-provided add-on.

It contains the components that encapsulate firmware update packages for a specific server type.

The firmware and drivers add-on might also contain the necessary drivers.

To add a firmware and drivers add-on to your image, you must first install the Hardware Support Manager plug-in for the respective family of servers.



9-21 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe vSphere Lifecycle Manager
- Compare vSphere Lifecycle Manager to vSphere Update Manager
- Identify the components of an ESXi image

9-22 **Lesson 3: Working with vSphere Lifecycle Manager**

9-23 Learner Objectives


After completing this lesson, you should be able to meet the following objectives:

- Import ESXi updates in the vSphere Client
- Define a cluster image using vSphere Lifecycle Manager
- Validate ESXi host compliance against a cluster image
- Update ESXi hosts using vSphere Lifecycle Manager

9-24 vSphere Lifecycle Manager Home View

In the vSphere Lifecycle Manager home view, you configure and administer the vSphere Lifecycle Manager instance that runs on your vCenter Server system.

From the drop-down menu at the top of the Lifecycle Manager pane, you can select the vCenter Server System that you want to manage.



You do not require special privileges to access the vSphere Lifecycle Manager home view.

To access the vSphere Lifecycle Manager home view in the vSphere Client, select **Menu > Lifecycle Manager**.

9-25 Image Depot

The landing page for the vSphere Lifecycle Manager home view is the **Image Depot** tab.

In the **Image Depot** tab, you can view details about downloaded ESXi elements:

- ESXi versions
- Vendor add-ons
- Components

The screenshot displays the vSphere Lifecycle Manager (vLCM) Image Depot interface. The top navigation bar includes 'Image Depot', 'Updates', 'Imported ISOs', 'Baselines', and 'Settings'. The 'Image Depot' tab is active, showing a list of ESXi versions. A detailed view for 'ESXi Version 7.0 GA - 15843807' is open, displaying the release date (03/16/2020), category (Enhancement), and a list of components.

Name	Version
ESXi	7.0 GA - 15843807

ESXi
Version 7.0 GA - 15843807
03/16/2020
Category: Enhancement

The general availability release of VMware ESXi Server 7.0.0 brings whole new levels of virtualization performance to datacenters and enterprises.

Components

Broadcom NetXtreme-E Network and ROCE/RDMA Drivers	216.0.0.0-1vmw
Broadcom Emulex Connectivity Division fcoe driver for FCoE adapters	12.0.1500.0-1vmw
Broadcom Emulex Connectivity Division NVMe over Fabric driver for FC adapters	12.4.293.2-3vmw
Broadcom IMA plugin for Emulex OneConnect iSCSI Driver	12.0.1200.0-2vmw
Emulex OneConnect iSCSI Driver	12.0.1200.0-1vmw
Broadcom Emulex Connectivity Division lpfc driver for FC adapters	12.4.293.3-5vmw
Emulex OneConnect Network Driver	12.0.1250.0-5vmw
Broadcom Emulex OneConnect Network Driver	11.4.62.0-1vmw
Broadcom 12Gbps SAS/PCIe MegaRAID Driver	7.712.50.00-1vmw

When you select a downloaded file, the details appear to the right:


- When you select an ESXi version, the details include the version name, build number, category, and description, and the list of components that make up the base image.
- When you select a vendor add-on, the details include the add-on name, version, vendor name, release date, category, and the list of added or removed components.
- When you select a component, the details include the component name, version, publisher, release date, category, severity, and contents (VIBs).

9-26 Importing Updates

vSphere Lifecycle Manager downloads software from online depots that you can configure and synchronize.

You can also use ESXi updates from an offline bundle:

- From the **Actions** drop-down menu, select **Import Updates**.
- Enter a URL or browse for a ZIP file that contains an ESXi image.



9-27 Patch Settings

By default, vSphere Lifecycle Manager is configured to download patch metadata automatically from the VMware repository.

Select **Settings > Patch Setup** to change the patch download source or add a URL to configure a custom download source.


The screenshot shows the vSphere Lifecycle Manager interface. The top navigation bar includes 'Lifecycle Manager', a dropdown for 'sa-vcsa-01.vclass.local', and an 'ACTIONS' dropdown. Below this is a tab bar with 'Image Depot', 'Updates', 'Imported ISOs', 'Baselines', and 'Settings'. The left sidebar has a tree view with 'Administration' (containing 'Patch Downloads' and 'Patch Setup') and 'Host Remediation' (containing 'Images', 'Baselines', and 'VMs'). The 'Patch Setup' page is active, showing a title 'Patches downloaded from the internet' and a 'CHANGE DOWNLOAD SOURCE' button. Below the title are links for 'NEW', 'EDIT', 'DELETE', 'ENABLE', and 'DISABLE'. A table lists four download sources, all from VMware, with columns for 'Download Source', 'Enabled', 'Connectivity Status', 'Source', 'Component', and 'Description'. At the bottom, there is an 'EXPORT' button and a count of '4 Download Sources'.

	Download Source	Enabled	Connectivity Status	Source	Component	Description
☐	https://hostupdate.vmware.com/software/...	Yes	Connected	VMware	Host	Partner provided Addons for ESXi
☐	https://hostupdate.vmware.com/software/...	Yes	Connected	VMware	Host	Download vSphere ESXi and ESX patches
☐	https://hostupdate.vmware.com/software/...	Yes	Connected	VMware	Host	VMware Certified Async Drivers for ESXi
☐	https://hostupdate.vmware.com/software/...	Yes	Connected	VMware	Host	VMware Async Releases for VM-tools on ESXi

9-28 Using Images to Perform ESXi Host Life Cycle Operations

After all ESXi hosts in a cluster are upgraded to vSphere 7, you can convert their lifecycle management from baselines to images.

You set up a single image and apply it to all hosts in a cluster. This step ensures cluster-wide host image homogeneity.



To set up a cluster image, the following requirements apply:

- All ESXi hosts must be version 7 or later.
- All ESXi hosts must be statefully installed.
- Ideally, all physical hosts should be from the same vendor.

9-29 Creating an ESXi Image for a New Cluster

When creating a cluster, you can create a corresponding cluster image.

New Cluster | SA-Datacenter

Name	New Cluster
Location	SA-Datacenter
vSphere DRS	☒
vSphere HA	☒
vSAN	☒

These services will have default settings - these can be changed later in the Cluster Quickstart workflow.

☒ **Manage all hosts in the cluster with a single image**

Image setup	
ESXi Version	7.0 GA - 15843807
Vendor Addon (optional)	None

The cluster image can be further customized later.

CANCEL **OK**

When creating a cluster, you can create a corresponding cluster image:

1. Create a cluster.
2. Select the **Manage image setup and updates on all hosts collectively** check box.
3. Define the ESXi version for your cluster image.
4. (Optional) Select vendor add-ons for the host.

Only add-ons that are compatible with the selected vSphere version appear in the drop-down menu.

After your cluster is created, add ESXi hosts to it.

The Create New Cluster wizard introduces a switch for enabling vSphere Lifecycle Manager and selecting elements for the desired cluster image.


You can further customize the image in the cluster update settings.

9-30 Customizing Cluster Images

You can edit the cluster image by updating the ESXi image version, or by changing, adding, or removing components and add-ons.

Before saving the image specification, you can validate it:


- Ensures completeness of the image
- Verifies that the image has no missing component dependencies
- Confirms that components do not conflict with one another



The screenshot shows the 'Cluster' configuration page with the 'Updates' tab selected. On the left, a sidebar lists 'Hosts' with a dropdown arrow, and under it, 'Image' is selected. Other options in the sidebar are 'Hardware Compatibility', 'VMware Tools', and 'VM Hardware'. The main content area is titled 'Image' and contains the text: 'Hosts in this cluster are managed collectively. This image below will be applied to all hosts in this cluster.' Below this text is a table with the following information:

ESXi Version	7.0 GA - 15843807
Vendor Addon	None
Firmware and Drivers Addon	None
Components	No additional components Show details

At the top right of the 'Image' section, there are two buttons: 'EDIT' and '...'.



The screenshot shows the 'Cluster' configuration page with the 'Updates' tab selected. On the left, a sidebar lists 'Hosts' with a dropdown arrow, and under it, 'Image' is selected. Other options in the sidebar are 'Hardware Compatibility', 'VMware Tools', and 'VM Hardware'. The main content area is titled 'Edit Image' and contains the text: 'Select the version of ESXi and other components that you want for the hosts in this cluster. The same image will be applied consistently to all these hosts.' Below this text is a form with the following fields:

ESXi Version	7.0 GA - 15843807 (released 03/16/2020)
Vendor Addon	SELECT (optional)
Firmware and Drivers Addon	SELECT (optional)
Components	No additional components Show details

At the bottom of the form, there are three buttons: 'SAVE', 'VALIDATE', and 'CANCEL'.

9-31 Checking Image Compliance

After you define a valid image, you can perform a compliance check to compare that image with the image that runs on the ESXi hosts in your cluster.

You can check the image compliance at the level of various vCenter Server objects:

- At the host level for a specific ESXi host
- At the cluster level for all ESXi hosts in the cluster
- At the data center level for all clusters and hosts in the data center
- At the vCenter Server level for all data centers, clusters, and ESXi hosts in the vCenter Server inventory.

The screenshot displays the 'Image Compliance' section in vCenter. At the top, there is a 'CHECK COMPLIANCE' button highlighted with an orange box. Below it, a status message indicates '1 of 2 hosts is out of compliance with the image'. Two buttons, 'REMEDIATE ALL' and 'RUN PRE-CHECK', are visible. On the left, a list of hosts shows 'sa-esxi-03.vclass.local' with a warning icon. The main panel shows details for this host, including a warning that the host will be rebooted during remediation. Below this, there are two tables for 'Software compliance'.

Image Compliance CHECK COMPLIANCE ...

Last checked on 04/08/2020, 4:53:45 PM (0 days ago)

⚠ 1 of 2 hosts is out of compliance with the image

REMEDIATE ALL RUN PRE-CHECK

Hosts

- sa-esxi-03.vclass.local

1 host

sa-esxi-03.vclass.local ✕

⚠ Host is out of compliance with the image

⚠ The host will be rebooted during remediation.

Software compliance Show Only drift comparison ▾

Image	Host Version	Image Version
ESXi Version	7.0 GA - 15525992	7.0 GA - 15843807

Components	Host Version	Image Version
VMware-VIB-Test-Certs	7.0.0-1.0.15531524	None ⓘ

Components per page 4 ▾ 1 component

The status of a host can be unknown, compliant, out of compliance, or not compatible with the image:

- A host status is unknown before you check compliance.
- A compliant host is one that has the same ESXi image defined for the cluster and with no standalone VIBs or differing components.
- If the host is out of compliance, a message about the impact of remediation appears. In the example, the host must be rebooted as part of the remediation. Another impact that might be reported is the requirement that the host enters maintenance mode.
- A host is not compatible if it runs an image version that is later than the desired cluster image version, or if the host does not meet the installation requirements for the vSphere build.

9-32 Running a Remediation Precheck

To ensure that the cluster health is optimal and that no problems occur during the remediation of the cluster, you can perform a remediation precheck.

To run a remediation precheck, you take the following steps:

- In the vSphere Client, click **Host and Clusters** and select a cluster that is managed by an image.
- In the **Updates** tab, click **Image**.
- On the Image Compliance card, click **RUN PRE-CHECK**.

The screenshot shows the 'Image Compliance' section in the vSphere Client. At the top, it says 'Last checked on 04/08/2020, 4:53:45 PM (0 days ago)' and '1 of 2 hosts is out of compliance with the image'. There are two buttons: 'REMEDiate ALL' and 'RUN PRE-CHECK', with the latter highlighted by an orange box. Below this is a list of hosts, with 'sa-esxi-03.vclass.local' selected. A detailed view for this host is shown on the right, indicating it is out of compliance and will be rebooted during remediation. It also displays a table of software compliance details.

Image Compliance CHECK COMPLIANCE ...

Last checked on 04/08/2020, 4:53:45 PM (0 days ago)

⚠ 1 of 2 hosts is out of compliance with the image

REMEDiate ALL **RUN PRE-CHECK**

Hosts

- sa-esxi-03.vclass.local

1 host

sa-esxi-03.vclass.local ×

⚠ Host is out of compliance with the image

⚠ The host will be rebooted during remediation.

Software compliance Show Only drift comparison

Image	Host Version	Image Version
ESXi Version	7.0 GA - 15525992	7.0 GA - 15843807

Components	Host Version	Image Version
VMware-VIB-Test-Certs	7.0.0-1.0.15531524	None ⓘ

Components per page **4** 1 component

9-33 Hardware Compatibility

The hardware compatibility check verifies the underlying hardware of the ESXi host in the cluster against the vSAN Hardware Compatibility List (HCL).

If your vCenter Server instance does not have Internet access, you can manually download a copy of the vSAN HCL database file.

The screenshot shows the vCenter interface for a cluster. The left sidebar has a 'Hosts' dropdown menu with options: Image, Hardware Compatibility (selected), VMware Tools, and VM Hardware. The main panel is titled 'Hardware Compatibility' and includes a sub-header 'Discover and resolve potential hardware compatibility issues.' Below this is a blue information box stating: 'Compatibility is checked only for vSAN storage controllers, and not with full VMware Compatibility Guide.' The 'Compatibility check results' section shows the check was performed on 04/08/2020 at 5:15:52 PM, with a 'Run checks' link. A blue warning box indicates: 'The image has 2 compatibility issues. Review them before remediating the cluster.' Below this is a table of incompatible devices:

Devices incompatible with ESXi
NVMe SSD Controller SM961/PM961 present on 2 hosts
NVMe SSD Controller SM981/PM981/PM983 present on 4 hosts


At the bottom, a blue information box states: 'No compatible devices were found'.

9-34 Standalone VIBs

When you convert a cluster to use vSphere Lifecycle Manager, ESXi hosts are scanned. During this scan, a VIB that is not part of a component is identified as standalone, and a warning appears.

Before updating ESXi hosts, you can import or ignore standalone VIBs:

- Import a component that contains the VIB and add it to the cluster image.
- Ignore the warning and allow the update process to remove the VIB from the host.



A warning about a standalone VIB does not block a conversion. If you continue to update ESXi, the VIB is uninstalled from the host as part of the conversion.

You cannot include standalone VIBs in a cluster image.

9-35 Remediating a Cluster Against a Single Image

When you remediate a cluster that you manage with an image, vSphere Lifecycle Manager applies the following image components to the ESXi hosts:

- ESXi image version
- Optional: vendor add-on
- Optional: firmware and drivers add-on
- Optional: user specified components

Image Compliance CHECK COMPLIANCE ...

Last checked on 04/08/2020, 4:53:45 PM (0 days ago)

⚠ 1 of 2 hosts is out of compliance with the image

REMEDiate ALL RUN PRE-CHECK

Hosts

- sa-esxi-03.vclass.local

1 host

sa-esxi-03.vclass.local ×

⚠ Host is out of compliance with the image

⚠ The host will be rebooted during remediation.

Software compliance Show Only drift comparison ▼

Image	Host Version	Image Version
ESXi Version	7.0 GA - 15525992	7.0 GA - 15843807

Components	Host Version	Image Version
VMware-VIB-Test-Certs	7.0.0-1.0.15531524	None ⓘ

Components per page 4 ▼ 1 component

Remediation makes the selected hosts compliant with the desired image.

You can remediate a single ESXi host or an entire cluster, or simply precheck hosts without updating them.

During each step of a remediation process, ESXi Health Perspectives determines the readiness of the host to enter or exit maintenance mode or be rebooted.

You can also click **RUN PRE-CHECK** to precheck hosts without updating them.

9-36 Reviewing Remediation Impact

vSphere Lifecycle Manager performs a precheck before each remediation.

Review Remediation Impact

Impact summary

Applicable remediation settings

End User License Agreement

Impact to specific hosts

sa-esxi-04.vclass.local

sa-esxi-03.vclass.local

Impact summary

- 2 host(s) are non-compliant with the image.
- 1 host(s) will be live remediated.
- 1 host(s) will be rebooted.

Notes

VM states honor remediation settings
VMs may be powered off, suspended or migrated to other hosts based on the applicable remediation settings.

Pre-check will be run again as a part of the remediation
Pre-check will be run again as a part of the remediation process. This is to ensure that no new issues have arisen on the cluster or hosts since the last pre-check (if any) that prevent remediation.

Hosts are remediated one at a time
Hosts will be remediated one at a time, so hosts will not reboot/go into maintenance mode simultaneously.

Order of host remediation is determined at runtime
Hosts will be remediated in an order determined at runtime. Hence that order may not correspond to the order in which they appear here.

☒ I accept the terms of the [end user license agreement](#)

[EXPORT IMPACT DETAILS](#)


CLOSE

START REMEDIATION

9-37 Recommended Images

Using vSphere Lifecycle Manager, you can check for recommended images for a cluster that you manage with an image.

vSphere Lifecycle Manager checks for compatibility across the image components. This process ensures that the recommended image fulfills all software dependencies.



You check for image recommendations on demand and per cluster. You can check for recommendations for different clusters at the same time. When recommendation checks run concurrently with other checks, with compatibility scans, and with remediation operations, the checks are queued to run one at a time.

If you have never checked recommendations for the cluster, the **View recommended images** option is dimmed.

After you select **Check for recommended images**, the results for that cluster are generated.

The task Checking for recommended images is visible to all user sessions and cannot be canceled.

When the check completes, you can select **View recommended images**.

9-38 Viewing Recommended Images

To view recommended images for a cluster:

1. Click the **Updates** tab for the cluster.
2. Select **Actions** under Cluster Image.
3. Select **View recommended images**.
4. Select an image and click **CONTINUE**.

vSphere shows the recommended images for clusters in the following categories:

- Latest release within the current series
- Latest major release

ACTIONS ▾

- Edit
- Import
- Export
- Check for recommended images
- View recommended images**

Recommended Images ×

Generated at <timestamp>

🔔 Valid only if no new updates were added to the depot or hardware was changed since then.

<Show correct text about validated images here> We have pre-built these two valid images for this cluster. If you use one of these images you will not have any firmware or hardware compatibility issues. If you do not see images for more recent versions, that means they are not applicable to this cluster. [View less](#)

Images Not Recommended	Explanation
8.0U2	Hardware not compatible for this version.
7.0U6	This update was only for Cisco hosts.

	CURRENT IMAGE	LATEST IN CURRENT SERIES	LATEST AND GREATEST
	7.0	7.0U5	8.0U1
Select image	☐	☐	☒
ESXi Version (Release date)	7.0 (January 2019)	7.0U5 (December 2019)	8.0U1 (July 2020)
Vendor Addon	Dell Addon v2.5	Dell Addon v2.5	Dell Addon v2.5
Firmware Bundle	None	None	None
Components			

[VIEW ALL COMPONENTS](#) CANCEL CONTINUE

When you view recommended images, vSphere shows the following types of images:

- **CURRENT IMAGE:** The image specification that is running on the cluster.
- **LATEST IN CURRENT SERIES:** If available, a later version within the same release series appears. For example, if the cluster is running vSphere 7.0 and vSphere 7.1 is released, an image based on vSphere 7.1 appears.
- **LATEST AND GREATEST:** If available, a later version in a later major release. For example, if the cluster is running vSphere 7.0 or 7.1 and vSphere 8.0 is released, an image based on vSphere 8.0 appears.

vSphere might show one recommendation, two, or no recommendations at all:

- If the latest release within the current series is the same as the latest major version released, only one result appears.
- If the two releases are different, two recommendations appear.
- If the current image is the same as the latest release, no recommendations appear.

9-39 Selecting a Recommended Image

You can select a recommended image and then validate and save it as the desired cluster image.

The screenshot shows the vSphere Client interface. On the left, a tree view shows the hierarchy: vSphere > DC 1 > DC 2 > Cluster 1 > New Cluster. The main panel is titled 'New Cluster' and has tabs for Summary, Monitor, Configure, Permissions, Hosts, VMs, Datastores, Networks, and Cluster Image (which is selected). The 'Edit Image' section is active, showing 'Select the version of ESXi and other components.' Below this, there's a dropdown for 'ESXi Version' set to '8.0U1' (build number: 878789, released 7 Jan 2018) and a 'Vendor Addon' section with 'Dell Addon v2.5'. A 'Components' section shows '4 manually added components'. Below this is a table with 4 components, each with a name, version, and notes. At the bottom are 'SAVE', 'CANCEL', and 'VALIDATE' buttons.

Component Name	Version	Notes
additional-component-1	2.3	Manually added
additional-component-2	7.2	Manually added
additional-component-3	2.3	Manually added
additional-component-4	5.6	Manually added

You can use a recommended image as a starting point to customize the cluster image. When you select a recommended image, the Edit Image workflow appears.

You can perform these actions:

- Add or remove image components.
- Validate and save the image.
- Scan the cluster for compatibility.
- Remediate the cluster.

9-40 Esxcli Support for vSphere Lifecycle Manager Images

In vSphere 7, new command options are introduced for `esxcli software`.

```
[root@sa-esxi-01:~] esxcli software
Usage: esxcli software {cmd} [cmd options]
```

Available Namespaces:

<code>component</code>	Apply, remove, list or display individual Components
<code>sources</code>	Query depot contents for Components, VIBs and image profiles
<code>vib</code>	Install, update, remove, or display individual VIB packages
<code>acceptance</code>	Retrieve and set the host acceptance level setting
<code>addon</code>	Display Addons
<code>baseimage</code>	Display base images
<code>profile</code>	Display, install, update or validates image profiles

Available Commands:

<code>apply</code>	Applies a complete image with a software spec that specifies base image, addon and components to install on the host.

```
[root@sa-esxi-01:~]
```

With changes to ESXCLI commands in vSphere 7, you can interact with the software specification of ESXi hosts:

- The existing `esxcli software sources` command has new options to manage the components, add-ons, and base images in a depot.
- `esxcli software baseimage` and `esxcli software addon` are new namespaces to display the base image and add-on installed on the host.

9-41 Lab 10: Updating ESXi Hosts with vSphere Lifecycle Manager

Import updates from an offline bundle and configure clusters with vSphere Lifecycle Manager:

1. Import Updates
2. Create a Cluster with vSphere Lifecycle Manager Enabled
3. Add ESXi Hosts to the Cluster
4. Update ESXi Hosts with vSphere Lifecycle Manager
5. Manage an Existing Cluster with a Cluster Image

9-42 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Import ESXi updates in the vSphere Client
- Define a cluster image using vSphere Lifecycle Manager
- Validate ESXi host compliance against a cluster image
- Update ESXi hosts using vSphere Lifecycle Manager

9-43 Key Points

- With Update Planner, you can perform prechecks to verify that your vCenter Server meets the update or upgrade requirements.
- Using vSphere Lifecycle Manager, you can manage ESXi hosts using baselines or manage a cluster of ESXi hosts using images.
- vSphere Lifecycle Manager centralizes version management for clusters, ESXi, drivers, and firmware.
- To ensure that your environment runs the latest verified software, you can check for a recommended image by VMware.

Questions?

Module 10

Introduction to vSphere with Kubernetes

10-2 Importance

vSphere with Kubernetes introduces native container support into vSphere. To use the features and capabilities of vSphere with Kubernetes, you must understand container and Kubernetes constructs and fundamentals.

10-3 Module Lessons

1. Introduction to Containers
2. Introduction to Kubernetes
3. Introduction to vSphere with Kubernetes

10-4 **Lesson 1: Introduction to Containers**

10-5 Learner Objectives


After completing this lesson, you should be able to meet the following objectives:

- Differentiate between containers and virtual machines
- Identify the parts of a container system
- Recognize the steps in a basic Docker workflow

10-6 Virtual Machines and Containers (1)

VMs provide virtual hardware that the guest OS uses to run applications. Multiple applications run on a single VM but they are logically separated and isolated.

With containers, developers take a streamlined base OS file system and layer on only the required binaries and libraries that the application depends on.



With virtualization, multiple physical machines can be consolidated into a single physical machine that runs multiple VMs. Each VM provides virtual hardware that the guest OS uses to run applications. Multiple applications run on a single VM but these applications are still logically separated and isolated.

A concern about VMs is that they are hundreds of megabytes to gigabytes in size and contain many binaries and libraries that are not relevant to the main application running on them.

With containers, developers take a streamlined base OS file system and layer on only the required binaries and libraries that the application depends on. When a container is run as a process on the container host OS, the container can see its dependencies and base OS packages. The container is isolated from all other processes on the container host OS. The container processes are the only processes that run on a minimal system.

From the container host OS perspective, the container is another process that is running, but it has a restricted view of the file system and potentially restricted CPU and memory.

10-7 Virtual Machines and Containers (2)

Virtual machines and containers differ in several ways.

Virtual Machines	Containers
Encapsulation of an entire operating system	Encapsulation of an application and dependent binaries or libraries
Scheduled by the hypervisor	Scheduled by the container host OS
Run on the hypervisor	Run on the container host OS
Starting a VM means starting an operating system (seconds to minutes)	Starting a container means starting the application process (milliseconds to seconds)

Containers are the ideal technology for microservices because the goals of containers (lightweight, easily packaged, can run anywhere) align with the goals and benefits of the microservices architecture.

Operators get modularized application components that are small and can fit into existing resources.

Developers can focus on the logic of modularized application components, knowing that the infrastructure is reliable and supports the scalability of modules.

10-8 About Container Hosts

Container hosts can be of the following types:

- Standard OS with a container engine installed:
 - Ubuntu with a container engine
- OS that is developed specifically with containers in mind:
 - Photon
 - CoreOS
- Virtual machine or physical machine:

Among the many benefits of using VMs are easy management and scalability.

In vSphere with Kubernetes, the container hosts are Photon-based VMs.


Administrators provide container hosts, which are the base structure that developers use to run their containers. A robust microservices system includes more deliverables, many of which are built using containers.

For developers to focus on providing services to customers, operations must provide a reliable container host infrastructure.

10-9 Container Workflow Example

A container workflow follows these steps:

1. Build an image from the source code and dependencies.
2. Send the image to the image registry.
3. Take the image from the image registry.
4. Run the image as a container.



In the example workflow, an image is created from source code files and other dependencies. A container can be stored in a registry, and the image can be pulled into a container host where it is used to spawn containers.

With this type of workflow, developers write code once. The code works across many different environments.


In the Dockerfile, a key specification is the name of the base OS image that is required for the container. A container can be deployed only on hosts running the OS that matches the one in the Dockerfile.

10-10 About Container Engines

A container engine is a control plane that is installed on each container host. The control plane manages the containers on that host.

Container engines perform several functions:

- Build container images from source code (for example, Dockerfile).
Alternatively, load container images from a repository.
- Create running containers based on a container image.
- Commit a running container to an image.
- Save an image and push it to a repository.
- Stop and remove containers.
- Suspend and restart containers.
- Report container status.



Docker is the most commonly used container engine.

The container engine runs as a daemon process on the container host OS. When a user requests that a container is run, the container engine gets the container image from an image registry (or locally, if already downloaded) and runs the container as a process.

10-11 About Dockerfile

Dockerfile is a plain text file that declares how to create an image.

Dockerfile is similar to the source code of an image.

You can use the BUILD command to create an image from Dockerfile

Sample Dockerfile:

```
1 FROM ubuntu
2 RUN apt-get install -y nginx php
3 COPY index.php /var/www/html/index.php
```

10-12 Container Images

Container images are comparable to a VM template. They have the following characteristics:

- They contain application code and application dependencies.
- They are built using a layered file system and can be made up of one or more layers.
- Docker images are built using a Dockerfile.
- Each line in a Dockerfile represents a layer in a container image.

Docker uses a union file system to combine multiple layers of an image into a single image. These layers (also called intermediate images) are generated when the commands in the Dockerfile are run during the Docker image build.

The intermediate images are layered together to form a new single coherent directory of files that makes up the image on the container host.

To analyze a Docker image, run these commands:

- `docker image inspect`
- `docker image history`

```
Layer 3 - Application code
/var/www/html/index.php
```

```
Layer 2 - Application dependencies
/usr/bin/php
/usr/sbin/nginx
/usr/lib/php/20170718/pdo.so
/usr/lib/php/20170718/json.so
/etc/php/7.2/fpm/php-fpm.conf
/etc/php/7.2/fpm/php.ini
/etc/nginx/mime.types
/etc/nginx/nginx.conf
/var/log/nginx/access.log
/var/log/nginx/error.log
```

```
Layer 1 - Base OS
/bin/bash
/bin/cat
/bin/chmod
/bin/chown
/bin/cp
/bin/echo
/bin/grep
/bin/tar
/etc/hosts
/etc/fstab
/etc/passwd
/etc/group
/etc/hostname
/etc/networks
/etc/nsswitch.conf
/etc/resolv.conf
/etc/security/limits.conf
/etc/sysctl.conf
/lib/x86_64-linux-gnu/ld-2.27.so
/lib/x86_64-linux-gnu/ld-linux-x86-64.so.2
/lib/x86_64-linux-gnu/libc-2.27.so
/lib/x86_64-linux-gnu/libc.so.6
```

10-13 Image Registry

Container images are stored in a central image registry:

- The image registry maintains multiple versions of container images.
- Container engines can take images from an image registry to run them.

VMware Harbor is a registry that can be installed on premises. It provides the following features:

- Image signing
- Vulnerability scanning



For more information about the Docker Hub library for container images, see <https://hub.docker.com/>.

<https://hub.docker.com/> For more information about the Harbor registry server project, see <https://goharbor.io/>.

10-14 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Differentiate between containers and virtual machines
- Identify the parts of a container system
- Recognize the steps in a basic Docker workflow

10-15 **Lesson 2: Introduction to Kubernetes**

10-16 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Explain the Importance of Kubernetes
- Recognize the basic architecture of Kubernetes
- Describe a basic Kubernetes workflow

10-17 Problems Solved by Kubernetes


With Docker, containers are managed on a single container host. Managing multiple containers across multiple container hosts creates many problems:

- Managing large numbers of containers
- Restarting failed containers
- Scaling containers to meet capacity
- Networking and load balancing

Kubernetes provides an orchestration layer to solve these issues.

10-18 Kubernetes Architecture

Each component in the Kubernetes architecture focuses on increasing resource utilization, reliability, and consistency.



The Kubernetes architecture includes the following components:

- **kubectl**: The command-line interface (CLI) to the Kubernetes API
- **Kubernetes master VM**:
 - **API server**: The entry point to the Kubernetes platform
 - **etcd**: The key or value datastore that stores the configuration and current state of the system
 - **Controller manager**: Runs various controllers that watch the API for changes and takes appropriate action
 - **Scheduler**: Balances pods across nodes
- **Kubernetes node**:
 - **kubelet**: Watches the API for work assigned to it

- Kubectl: A command-line tool for controlling Kubernetes clusters
- kube-proxy: Configures networking rules to route traffic to containers
- Docker: Runs pods when requested by kubelet
- Pod: A group of one or more containers (such as Docker containers), with shared storage and network, and a specification for how to run the containers
- Node: The smallest unit of computing resources in Kubernetes

Kubernetes automates many key operational responsibilities, providing the developer with a reliable environment.

10-19 About Manifests

In Kubernetes, manifest files declare the desired state of objects, including the container images that should be running.

Manifests have the following properties:

- YAML format
- Declarative configuration
- Desired API primitives

Kubernetes manages the creation of the requested primitives.

Manifest files are the new shared language between the developer and the operator. They act as the single source of truth for the configuration of a set of application components (containers).

For more information about YAML, see *Yet Another Markup Language* at <https://yaml.org/spec/history/2001-05-26.html>.

<https://yaml.org/spec/history/2001-05-26.html> With declarative configuration, you can tell Kubernetes what you want, and it knows what to do.

10-20 About Pods

A pod is a set of one or more tightly coupled containers.

It is the smallest unit of work in Kubernetes.

Containers in a pod start and stop together.

Example: Run a pod.

- The pod name is nginx-demo.
- The pod is assigned a label
- It runs image nginx with version 1.7.9.
- It should listen on port 80.

For more information about Kubernetes pods, see <https://kubernetes.io/>.

The Kubernetes architecture has many objects and other conceptual primitives. This topic is an overview of a few fundamental objects and primitives .

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx-demo
  labels:
    app: nginx
spec:
  containers:
    - name: nginx
      image: nginx: 1.7.9
      ports:
        - containerPort: 80
```

10-21 About Deployments

A deployment is the most commonly used primitive:

- Declares whether pods can be upgraded and whether they can be patched without disrupting services
- Provides rolling updates by creating ReplicaSets and destroying old ReplicaSets
- Enables a new version of an image to be deployed without downtime

Example: Deploy a deployment.

- The deployment name is nginx.
- Three replicas are expected to be running.
- The deployment runs image nginx with version 1.7.9.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx
spec:
  selector:
    matchLabels:
      app: nginx
  replicas: 3
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx: 1.7.9
          ports:
            - containerPort: 80
```


10-22 About ReplicaSets

A ReplicaSet declares how the functionality of a pod is made scalable and resilient through redundancy.

The ReplicaSet ensures that a specified number of pods is kept running.

Example: Deploy a ReplicaSet.

- The ReplicaSet name is nginx-replica-demo.
- Two replicas are expected to be running.
- The ReplicaSet applies to pods with the label nginx.

```
apiVersion: apps/v1
kind: ReplicaSet
metadata:
  name: nginx-replica-demo
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx: 1.7.9
          ports:
            - containerPort: 80
```

10-23 About Services

A service describes how pods discover and communicate with each other and external networks.

A service exposes a deployment as a single IP address.

Example: Deploy a service.

- The service name is nginx-service.
- The service type is LoadBalancer.
- It should listen on port 80.

A service IP address is exposed through one of these ServiceTypes:

- ClusterIP: Exposes the service on a cluster-internal IP address
- NodePort: Exposes the service externally as a static port
- LoadBalancer: Exposes the service externally using a cloud provider's load balancer
- ExternalName: Maps the service to a CNAME record

```
apiVersion: v1
kind: Service
metadata:
  name: nginx-service
spec:
  type: LoadBalancer
  selector:
    app: nginx
  ports:
    - port: 80
      protocol: TCP
      targetPort: 80
```

10-24 About Network Policies

A network policy is a specification of how groups of pods are allowed to communicate with each other and other network endpoints.

The policy declares rules for ingress and egress traffic.


Example: Deploy a network policy.

- The policy name is network-policy.
- It allows all ingress traffic from 172.20.10.0/24.
- It allows all egress traffic to 10.0.0.0/8.

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: network-policy
  namespace: namespace-01
spec:
  podSelector:
    matchLabels:
      app: nginx
  policyTypes:
    - Ingress
    - Egress
  ingress:
    - from:
        - ipBlock:
            cidr: 172.20.10.0/24
      - podSelector:
          matchLabels:
            app: nginx
      ports:
        - protocol: TCP
          port: 80
  egress:
    - to:
        - ipBlock:
            cidr: 10.0.0.0/8
      ports:
        - protocol: TCP
          port: 80
```

10-25 Kubernetes Workflow Example

The example workflow shows how a developer might use Docker to build an image locally, store the image in a registry, and use Kubernetes to manage the running of the container.



The example workflow includes the following steps:

1. Build an image from source code and dependencies.
2. Send the image to the image registry.
3. Tell Kubernetes to use the image to run a pod.
4. Scheduler assigns the pod to a node.
5. Kubelet accepts the pod.
6. Docker takes the image from the image registry.
7. Docker starts the container process inside a pod.

The operational intent is to secure containers, maintain the communication between pods, and ensure that all the code in images runs.

The container is managed by the Kubernetes-specific requirements in the manifest for maintaining a robust, resilient, and scalable application by the worker kubelet. The container

is also managed by the requirements of the manifest in master etcd. With these details, the developer and operator teams can work together to declare what is needed and, in the process, fulfill customers' needs.

Kubernetes supports a fast, dynamic, and continuous process for meeting customer needs.

10-26 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Explain the Importance of Kubernetes
- Recognize the basic architecture of Kubernetes
- Describe a basic Kubernetes workflow

10-27 **Lesson 3: Introduction to vSphere with Kubernetes**


10-28 Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe the purpose of vSphere with Kubernetes and how it fits into the VMware Tanzu portfolio
- Identify the capabilities of vSphere with Kubernetes
- Recognize the key benefits of vSphere with Kubernetes
- Describe the vSphere with Kubernetes Supervisor Cluster
- Identify the components that make up the vSphere with Kubernetes Supervisor Cluster
- Describe the Tanzu Kubernetes Grid service for vSphere

10-29 vSphere with Kubernetes and VMware Cloud Foundation 4

vSphere with Kubernetes is available through VMware Cloud Foundation 4.



At initial launch, VMware Cloud Foundation is the release vehicle for vSphere with Kubernetes.

Hybrid infrastructure services are as follows:

- vSphere Pod Service: Developers can run containers natively and securely on vSphere without managing VMs or Kubernetes clusters.
- Network Service: Developers can define virtual routers, load balancers, and firewall rules for use with their application.
- Registry Service: Developers can store and manage Docker and OCI images using Harbor. Harbor is an open-source container image registry that secures images with role-based access control.
- Storage Service: vCenter Server storage policies and devices can be consumed as Kubernetes storage classes and used as persistent disks for use with containers, Kubernetes clusters, and VMs.

Tanzu Runtime services include Tanzu Kubernetes Grid service. With it, developers can manage consistent, compliant, and conformant Kubernetes clusters, which are Tanzu Kubernetes clusters.

VMware Cloud Foundation 4 does not support upgrades from previous VMware Cloud Foundation releases.

10-30 vSphere with Kubernetes and VMware Tanzu (1)

Introducing

VMware Tanzu

Build Modern Apps

Run Enterprise Kubernetes

Manage Kubernetes for Developers AND IT



10-31 vSphere with Kubernetes and VMware Tanzu (2)

vSphere with Kubernetes integrates and embeds Kubernetes directly into vSphere and is the foundation of the VMware Tanzu portfolio for running Kubernetes:

- Unites vSphere and Kubernetes
- Extends vSphere for all modern applications
- Supports collaboration between development and IT operations (DevOps)



Build

- Bitnami: Package management system for web applications, solution stacks, and virtual machines
- Pivotal: Cloud native application development

Run


- vSphere with Kubernetes: Runs Kubernetes workloads natively on vSphere. Enables self-provisioning of Tanzu Kubernetes clusters running on vSphere with Kubernetes.
- Tanzu Kubernetes Grid: Standalone offering of Tanzu Kubernetes Grid that can be provisioned on vSphere 6.7 U3 and later and on Amazon Web Services.
- Tanzu Kubernetes Grid+: Adds customer reliability engineering to the Tanzu Kubernetes Grid entitlement.

Manage

- Tanzu Mission Control: Manages multiple Tanzu Kubernetes clusters from one view, no matter where the Tanzu Kubernetes Grid clusters are running (vSphere, AWS, and more).

10-32 About vSphere with Kubernetes

vSphere with Kubernetes transforms vSphere into a native Kubernetes platform.



A cluster that is enabled with vSphere with Kubernetes is called a Supervisor Cluster.

The Supervisor Cluster runs on top of an SDDC layer that includes the following elements:

- ESXi for compute
- NSX-T Data Center for networking
- vSAN or other shared storage solution as the shared storage for vSphere native pods and VMs that run inside the Supervisor Cluster


After a Supervisor Cluster is created, you can create namespaces in the Supervisor Cluster that are called supervisor namespaces.

DevOps engineers can run workloads consisting of containers running inside vSphere native pods and create upstream Kubernetes clusters by using the Tanzu Kubernetes Grid service for vSphere.

10-33 About Supervisor Clusters (1)

In vSphere with Kubernetes, a Supervisor Cluster is a special Kubernetes cluster that uses ESXi as the worker nodes instead of Linux.


In a Supervisor Cluster, each ESXi host runs an instance of Spherelet, which is a special implementation of kubelet, directly on ESXi and not as a VM.



10-34 About Supervisor Clusters (2)

The Supervisor Cluster includes a set of control plane VMs that provide API access to the Supervisor Cluster.

The control plane VM serves as an endpoint for Kubernetes users to manage their namespaces using the familiar kubectl command-line interface.



10-35 About Control Plane VMs (1)

The control plane VMs are equivalent to Kubernetes master nodes.

These control plane VMs are created and run on the hosts that are part of the Supervisor Cluster:

- ESX Agent Manager (EAM) manages the deployment of the control plane VMs.
- DRS determines the placement of these VMs and migrates them when needed.
- VM anti-affinity rules are automatically created for the control plane VMs.


The control plane VMs provide the management and developer interface to the Supervisor Cluster:

- An API endpoint
- Infrastructure services and pods for the Supervisor Cluster
- Lifecycle management of the control plane VMs

Anti-affinity of the control plane VMs is done using a new method similar to compute policies in VMware Cloud on AWS. This method is not available in the vSphere Client. Control plane anti-affinity does not appear in the vSphere DRS configuration settings.

10-36 About Control Plane VMs (2)


The control plane VMs run infrastructure services and pods for the Supervisor Cluster, similar to Kubernetes master nodes.



10-37 Control Plane Networking

The control plane VMs use management, corporate, and cluster networks:

- vSphere Client management is accessed from the management network using a floating IP.
- etcd communication between each control plane VM uses the management network.
- Developers access the infrastructure using a virtual IP.
- The control plane communicates with pods using a cluster network.



The floating IP follows the `etcd` leader, which ensures that, under partition, traffic is routed to the functional portion of the cluster.

The floating IP is not always assigned. During a leader election, or during loss of `etcd` quorum, no control plane possesses the floating IP, and `wcp` (vCenter Server Appliance) and `Spherelets` (ESX) cannot connect to the `apiserver`.

All components retry their connections indefinitely and resume normal function when the `etcd` cluster regains quorum.

10-38 About Spherelet

Spherelet is a kubelet that is ported natively to ESXi. It allows the ESXi host to become part of a Kubernetes cluster. Spherelet performs the following functions:

- Communicates with the control plane VMs
- Manages node configuration
- Starts vSphere pods
- Monitors vSphere pods

10-39 About CRX

Container Runtime Executive (CRX) is built into ESXi. It has the following properties:

- Paravirtualized Photon Linux kernel
- Same hardware virtualization, boundaries, and isolation as VMs


Container Runtime Executive (CRX) is similar to a VM from the perspective of hostd and vCenter Server. CRX includes a heavily paravirtualized Linux kernel that works with the hypervisor.

CRX uses the same hardware virtualization techniques as VMs, and it has a VM boundary around it.

CRX uses a direct boot technique, which allows the Linux guest of CRX to initiate the main init process without passing through kernel initialization. In this way, vSphere native pods can boot nearly as fast as containers.

VMX: Virtual machine runtime

VMM: Virtual machine Monitor




10-40 About vSphere Pods

A vSphere pod is a VM with a small footprint that runs one or more Linux containers.

With vSphere pods, workloads have the following capabilities:


- Strong isolation from a Linux kernel based on Photon OS
- Resource management using DRS
- High performance from the same level of resource isolation as VMs
- Open Container Initiative (OCI) compatible
- Equivalent to a Kubernetes container host



10-41 vSphere Pods Network and Storage


vSphere pods use different types of storage depending on the objects that are stored. The types of storage are ephemeral virtual machine disks (VMDKs), persistent volume VMDKs, and container image VMDKs:

- Storage policies for container image and ephemeral disks are defined at the cluster level.
- Storage policies for persistent volumes are defined at the namespace level.
- Networking for vSphere pods uses the topology provided by NSX.



10-42 Comparing vSphere Pods with Kubernetes

vSphere with Kubernetes architecture differs slightly from Kubernetes architecture. The main difference is that ESXi becomes a Kubernetes node.




In traditional Kubernetes, the pods run inside the Kubernetes worker nodes (Linux VMs). In vSphere with Kubernetes, the vSphere pods run on ESXi similar to a VM.

10-43 About Namespaces

Namespaces are a Kubernetes construct used to divide cluster resources:

- Namespaces support multiuser access.
- Resources are controlled using resource quotas.
- The vSphere administrator can control quotas and access from the vSphere Client.



A namespace sets the resource boundaries where vSphere native pods and Kubernetes clusters, created by using the Kubernetes service, can run. When initially created, the namespace has unlimited resources in the Supervisor Cluster. You can set limits for CPU, memory, storage, and the number of Kubernetes objects that can run in the namespace. A resource pool is created per each namespace in vSphere. Storage limitations are represented as storage quotas in Kubernetes.

To provide access to namespaces to DevOps engineers, you assign permission to users or user groups that are available in an identity source that is associated with vCenter Single Sign-On.

After a namespace is created and configured with resource and object limits, and permissions and storage policies, DevOps engineers can access the namespace to run Kubernetes workloads and create Kubernetes clusters by using the Kubernetes service.

10-44 Container Registry

Containers need a repository to store their images.

vSphere with Kubernetes integrates the Harbor image registry into vSphere.

Each namespace in the Supervisor Cluster is assigned its own project in the Harbor registry.

Compute-01

ACTIONS

SummaryMonitorConfigurePermissionsHostsVMsNamespacesDatastoresNetworksUpdates

Services >

Configuration >

Trust Authority >

Alarm Definitions

Scheduled Tasks

Namespaces >

General

Network

Storage

Certificates

Image Registry

vSAN >

Image Registry

Embedded Harbor registry

DISABLE

Health

Running

Storage

WP Storage Policy

Storage Usage

58 MB Used200 GB Capacity

Link to Harbor UI

https://172.18.248.4

Garbage collection policy

Scheduled for every Saturday at 2:00 AM


Root certificate

Download SSL Root Certificate

When enabled, Harbor is deployed as a series of vSphere pods to the Supervisor Cluster in its own dedicated, system-managed, namespace.

10-45 Tanzu Kubernetes Grid Service for vSphere

Tanzu Kubernetes Grid service for vSphere enables a Tanzu Kubernetes cluster that runs inside VMs in a namespace in the Supervisor Cluster. With this service, developers can control their own isolated Kubernetes cluster.




A Tanzu Kubernetes Grid cluster can be compared to running traditional Kubernetes clusters using VMs and a container host OS. The Tanzu Kubernetes Grid service for vSphere automates the deployment of such a cluster. A Tanzu Kubernetes cluster object can consist only of its control plane and worker VMs. vSphere pods cannot be deployed into a Tanzu Kubernetes Grid cluster object.

10-46 vSphere with Kubernetes Objects

vSphere with Kubernetes objects are visible in the vSphere Client.

Objects visible in the Hosts and Clusters view include:

1. Namespace
2. vSphere pods
3. Tanzu Kubernetes cluster
4. Harbor registry
5. Control plane VMs



10-47 Review of Learner Objectives

After completing this lesson, you should be able to meet the following objectives:

- Describe the purpose of vSphere with Kubernetes and how it fits into the VMware Tanzu portfolio
- Identify the capabilities of vSphere with Kubernetes
- Recognize the key benefits of vSphere with Kubernetes
- Describe the vSphere with Kubernetes Supervisor Cluster
- Identify the components that make up the vSphere with Kubernetes Supervisor Cluster
- Describe the Tanzu Kubernetes Grid service for vSphere

10-48 Key Points

- Containers are an encapsulation of an application and its dependent binaries or libraries.
- Docker is the main container virtualization standard that is used with Kubernetes.
- Kubernetes is an open-source container orchestration platform.
- vSphere with Kubernetes transforms vSphere into a native Kubernetes platform.
- Using VMware Tanzu, you can build, manage, and run Kubernetes workloads.

Questions?