VMware vSphere: Install, Configure, Manage

Lecture Manual
ESXi 7 and vCenter Server 7

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

https://t.me/learningnets

VMware vSphere:
Install, Configure, Manage
Lecture Manual
ESXi 7 and vCenter Server 7
Part Number EDU-EN-VSICM7-LECT (4/2020)

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, Enhanced vMotion™ Compatibility, vCenter Linked Mode, VMware ESX®, VMware ESXi™, VMware Horizon®, VMware Horizon® 7, VMware Horizon® 7 on VMware Cloud™ on AWS, VMware Horizon® View™, VMware Host Client™, VMware Photon™, VMware PowerCLI™, VMware Remote Console™, VMware Site Recovery Manager™, VMware Site Recovery™, VMware Skyline Advisor™, VMware Skyline™, VMware Tools™, VMware vCenter Server®, VMware vCenter Server® High Availability, VMware vCenter® Lifecycle Manager™, VMware vCenter® Server Appliance™, VMware vCenter® Single Sign-On, VMware Verify™, VMware View®, VMware vRealize®, VMware vRealize® Log Insight™, VMware vRealize® Log Insight™ for vCenter™, 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® Orchestrator™, VMware vRealize® Suite Lifecycle Manager™, VMware vSAN™, VMware vSphere®, VMware vSphere® API, VMware vSphere® API for Storage Awareness™, VMware vSphere® Client™, VMware vSphere® Command-Line Interface, VMware vSphere® DirectPath I/O™, VMware vSphere® Distributed Power Management™, VMware vSphere® Distributed Resource Scheduler™, VMware vSphere® ESXi™ Dump Collector, VMware vSphere® ESXi™ Shell, VMware vSphere® Fault Tolerance, VMware vSphere® High Availability, VMware vSphere® Replication™, VMware vSphere® Standard Edition™, VMware vSphere® Storage APIs - Array Integration, VMware vSphere® Storage APIs - Data Protection, VMware vSphere® Storage vMotion®, VMware vSphere® Virtual Symmetric Multiprocessing, VMware vSphere® Virtual Volumes™, VMware vSphere® VMFS, and VMware vSphere® vMotion® 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....2
1-3 Importance....3
1-4 Learner Objectives (1)....4
1-5 Learner Objectives (2)....5
1-6 Course Outline....6
1-7 Typographical Conventions....7
1-8 References (1)....8
1-9 References (2)....9
1-10 VMware Online Resources....10
1-11 VMware Education Overview....11
1-12 VMware Certification Overview....12
1-13 VMware Badge Overview....13
1-14 Virtual Beans: Introduction....14

Module 2 Introduction to vSphere and the Software-Defined Data Center 15

2-2 Importance....16
2-3 Module Lessons....17
2-4 Virtual Beans: Data Center....18
2-5 Lesson 1: Overview of vSphere and Virtual Machines....19
2-6 Learner Objectives....20
2-7 Terminology (1)....21
2-8 Terminology (2)....22
2-9 About Virtual Machines....23
2-10 Benefits of Using Virtual Machines....24
2-11 Types of Virtualization....26

https://t.me/learningnets

2-12 About the Software-Defined Data Center 27
2-13 vSphere and Cloud Computing 29
2-14 About VMware Skyline 31
2-15 VMware Skyline Family 32
2-16 Review of Learner Objectives 34
2-17 Lesson 2: vSphere Virtualization of Resources 35
2-18 Learner Objectives 36
2-19 Virtual Machine: Guest and Consumer of ESXi Host 37
2-20 Physical and Virtual Architecture 38
2-21 Physical Resource Sharing 39
2-22 CPU Virtualization 41
2-23 Physical and Virtualized Host Memory Usage 42
2-24 Physical and Virtual Networking 43
2-25 Physical File Systems and Datastores 45
2-26 GPU Virtualization 47
2-27 Review of Learner Objectives 48
2-28 Lesson 3: vSphere User Interfaces 49
2-29 Learner Objectives 50
2-30 vSphere User Interfaces 51
2-31 About VMware Host Client 52
2-32 About vSphere Client 53
2-33 About PowerCLI and ESXCLI 54
2-34 Lab 1: Accessing the Lab Environment 55
2-35 Review of Learner Objectives 56
2-36 Lesson 4: Overview of ESXi 57
2-37 Learner Objectives 58
2-38 About ESXi 59
2-39 Configuring an ESXi Host 61
2-40 Configuring an ESXi Host: Root Access 62
2-41 Configuring an ESXi Host: Management Network 63
2-42 Configuring an ESXi Host: Other Settings 64

https://t.me/learningnets

2-43 Controlling Remote Access to an ESXi Host 65
2-44 Managing User Accounts: Best Practices 66
2-45 ESXi Host as an NTP Client 67
2-46 Demonstration: Installing and Configuring ESXi Hosts 68
2-47 Lab 2: Configuring an ESXi Host 69
2-48 Review of Learner Objectives 70
2-49 Virtual Beans: Data Center 71
2-50 Key Points 72

Module 3 Virtual Machines 73

3-2 Importance 74
3-3 Module Lessons 75
3-4 Virtual Beans: Virtualizing Workloads 76
3-5 Lesson 1: Creating Virtual Machines 77
3-6 Learner Objectives 78
3-7 About Provisioning Virtual Machines 79
3-8 Creating VMs with the New Virtual Machine Wizard (1) 81
3-9 Creating VMs with the New Virtual Machine Wizard (2) 82
3-10 New Virtual Machine Wizard Settings 84
3-11 Installing the Guest Operating System 85
3-12 Deploying OVF Templates 86
3-13 About VMware Tools 87
3-14 Installing VMware Tools 89
3-15 Downloading VMware Tools 90
3-16 Labs 91
3-17 Lab 3: Creating a Virtual Machine 92
3-18 Lab 4: Installing VMware Tools 93
3-19 Review of Learner Objectives 94
3-20 Lesson 2: Virtual Machine Hardware Deep Dive 95
3-21 Learner Objectives 96
3-22 Virtual Machine Encapsulation 97

https://t.me/learningnets

3-23 About Virtual Machine Files 98
3-24 About VM Virtual Hardware 100
3-25 Virtual Hardware Versions 102
3-26 About CPU and Memory 103
3-27 About Virtual Storage 105
3-28 About Thick-Provisioned Virtual Disks 107
3-29 About Thin-Provisioned Virtual Disks 108
3-30 Thick-Provisioned and Thin-Provisioned Disks 109
3-31 About Virtual Networks 110
3-32 About Virtual Network Adapters 111
3-33 Other Virtual Devices 114
3-34 About the Virtual Machine Console 115
3-35 Lab 5: Adding Virtual Hardware 116
3-36 Review of Learner Objectives 117
3-37 Lesson 3: Introduction to Containers 118
3-38 Learner Objectives 119
3-39 Traditional Application Development 120
3-40 Modern Application Development 122
3-41 Benefits of Microservices and Containerization 123
3-42 Container Terminology 124
3-43 About Containers 125
3-44 Rise of Containers 126
3-45 About Container Hosts 127
3-46 Containers at Runtime 128
3-47 About Container Engines 129
3-48 Virtual Machines and Containers (1) 130
3-49 Virtual Machines and Containers (2) 131
3-50 About Kubernetes 132
3-51 Challenges of Running Kubernetes in Production 134
3-52 Architecting with Common Application Requirements 135
3-53 Review of Learner Objectives 136

https://t.me/learningnets

3-54 Virtual Beans: Virtualizing Workloads 137
3-55 Key Points 138

Module 4 vCenter Server 139

4-2 Importance 140
4-3 Module Lessons 141
4-4 Virtual Beans: vCenter Server Requirements 142
4-5 Lesson 1: Centralized Management with vCenter Server 143
4-6 Learner Objectives 144
4-7 About the vCenter Server Management Platform 145
4-8 About vCenter Server Appliance 146
4-9 vCenter Server Services 147
4-10 vCenter Server Architecture 148
4-11 About vCenter Single Sign-On 149
4-12 About Enhanced Linked Mode 150
4-13 ESXi and vCenter Server Communication 151
4-14 vCenter Server Appliance Scalability 153
4-15 Review of Learner Objectives 154
4-16 Lesson 2: Deploying vCenter Server Appliance 155
4-17 Learner Objectives 156
4-18 Preparing for vCenter Server Appliance Deployment 157
4-19 vCenter Server Appliance Native GUI Installer 158
4-20 vCenter Server Appliance Installation 159
4-21 vCenter Server Appliance Installation: Stage 1 160
4-22 vCenter Server Appliance Installation: Stage 2 161
4-23 Getting Started with vCenter Server 162
4-24 Configuring vCenter Server Using the vSphere Client 163
4-25 vCenter Server Appliance Management Interface 164
4-26 vCenter Server Appliance Multihoming 165
4-27 Demonstration: Deploying vCenter Server Appliance 166
4-28 Review of Learner Objectives 167

https://t.me/learningnets

4-29 Lesson 3: vSphere Licensing 168
4-30 Learner Objectives....169
4-31 vSphere Licensing Overview....170
4-32 vSphere License Service 171
4-33 Adding License Keys to vCenter Server 172
4-34 Assigning a License to a vSphere Component....173
4-35 Viewing Licensed Features 174
4-36 Lab 6: Adding vSphere Licenses....175
4-37 Review of Learner Objectives....176
4-38 Lesson 4: Managing the vCenter Server Inventory 177
4-39 Learner Objectives....178
4-40 vSphere Client Shortcuts Page 179
4-41 Using the Navigation Pane 180
4-42 vCenter Server Views for Hosts, Clusters, VMs, and Templates 181
4-43 vCenter Server Views for Storage and Networks....182
4-44 Viewing Object Information 183
4-45 About Data Center Objects....184
4-46 Organizing Inventory Objects into Folders....185
4-47 Adding a Data Center and Organizational Objects to vCenter Server....187
4-48 Adding ESXi Hosts to vCenter Server....188
4-49 Creating Custom Tags for Inventory Objects....189
4-50 Labs....190
4-51 Lab 7: Creating and Managing the vCenter Server Inventory....191
4-52 Lab 8: Configuring Active Directory: Joining a Domain 192
4-53 Review of Learner Objectives....193
4-54 Lesson 5: vCenter Server Roles and Permissions 194
4-55 Learner Objectives....195
4-56 About vCenter Server Permissions....196
4-57 About Roles 197
4-58 About Objects....199
4-59 Adding Permissions to the vCenter Server Inventory 200

https://t.me/learningnets

4-60 Viewing Roles and User Assignments....201
4-61 Applying Permissions: Scenario 1....202
4-62 Applying Permissions: Scenario 2....203
4-63 Activity: Applying Group Permissions (1)....204
4-64 Activity: Applying Group Permissions (2)....205
4-65 Applying Permissions: Scenario 3....206
4-66 Applying Permissions: Scenario 4....207
4-67 Creating a Role....208
4-68 About Global Permissions....209
4-69 Labs....210
4-70 Lab 9: Configuring Active Directory: Adding an Identity Source....211
4-71 Lab 10: Users, Groups, and Permissions....212
4-72 Review of Learner Objectives....213
4-73 Lesson 6: Backing Up and Restoring vCenter Server Appliance....214
4-74 Learner Objectives....215
4-75 Virtual Beans: vCenter Server Operations....216
4-76 About vCenter Server Backup and Restore....217
4-77 Methods for vCenter Server Appliance Backup and Restore....218
4-78 File-Based Backup of vCenter Server Appliance....219
4-79 File-Based Restore of vCenter Server Appliance....220
4-80 Scheduling Backups....221
4-81 Viewing the Backup Schedule....222
4-82 Demonstration: Backing Up and Restoring a vCenter Server Appliance Instance....223
4-83 Review of Learner Objectives....224
4-84 Lesson 7: Monitoring vCenter Server and Its Inventory....225
4-85 Learner Objectives....226
4-86 vCenter Server Events....227
4-87 About Log Levels....228
4-88 Setting Log Levels....229
4-89 Forwarding vCenter Server Appliance Log Files to a Remote Host....230

https://t.me/learningnets

4-90 vCenter Server Database Health 231
4-91 Monitoring vCenter Server Appliance 232
4-92 Monitoring vCenter Server Appliance Services 233
4-93 Monthly Patch Updates for vCenter Server Appliance 234
4-94 Review of Learner Objectives 235
4-95 Lesson 8: vCenter Server High Availability 236
4-96 Learner Objectives 237
4-97 Importance of Keeping vCenter Server Highly Available 238
4-98 About vCenter Server High Availability 239
4-99 Scenario: Active Node Failure 240
4-100 Scenario: Passive Node Failure 241
4-101 Scenario: Witness Node Failure 242
4-102 Benefits of vCenter Server High Availability 243
4-103 vCenter Server High Availability Requirements 244
4-104 Demonstration: Configuring vCenter Server High Availability 245
4-105 Review of Learner Objectives 246
4-106 Virtual Beans: vCenter Server Maintenance and Operations 247
4-107 Key Points 248

Module 5 Configuring and Managing Virtual Networks 249

5-2 Importance 250
5-3 Module Lessons 251
5-4 Virtual Beans: Networking Requirements 252
5-5 Lesson 1: Introduction to vSphere Standard Switches 253
5-6 Learner Objectives 254
5-7 About Virtual Switches 255
5-8 Types of Virtual Switch Connections 256
5-9 Virtual Switch Connection Examples 257
5-10 About VLANs 258
5-11 Types of Virtual Switches 260
5-12 Adding ESXi Networking 261

https://t.me/learningnets

5-13 Viewing the Configuration of Standard Switches....262
5-14 Network Adapter Properties....263
5-15 Distributed Switch Architecture....264
5-16 Standard and Distributed Switches: Shared Features....265
5-17 Additional Features of Distributed Switches....266
5-18 Lab 11: Using Standard Switches....267
5-19 Review of Learner Objectives....268
5-20 Lesson 2: Configuring Standard Switch Policies....269
5-21 Learner Objectives....270
5-22 Network Switch and Port Policies....271
5-23 Configuring Security Policies....272
5-24 Traffic-Shaping Policies....274
5-25 Configuring Traffic Shaping....275
5-26 NIC Teaming and Failover Policies....277
5-27 Load-Balancing Method: Originating Virtual Port ID....279
5-28 Load-Balancing Method: Source MAC Hash....281
5-29 Load-Balancing Method: Source and Destination IP Hash....283
5-30 Detecting and Handling Network Failure....285
5-31 Physical Network Considerations....287
5-32 Review of Learner Objectives....288
5-33 Virtual Beans: Networking Requirements....289
5-34 Key Points....290

Module 6 Configuring and Managing Virtual Storage 291

6-2 Importance....292
6-3 Module Lessons....293
6-4 Virtual Beans: Storage....294
6-5 Lesson 1: Storage Concepts....295
6-6 Learner Objectives....296
6-7 About Datastores....297
6-8 Storage Overview....298

https://t.me/learningnets

6-9 Storage Protocol Overview....300
6-10 About VMFS....302
6-11 About NFS....304
6-12 About vSAN....305
6-13 About vSphere Virtual Volumes....306
6-14 About Raw Device Mapping....307
6-15 Physical Storage Considerations....308
6-16 Review of Learner Objectives....309
6-17 Lesson 2: Fibre Channel Storage....310
6-18 Learner Objectives....311
6-19 About Fibre Channel....312
6-20 Fibre Channel SAN Components....313
6-21 Fibre Channel Addressing and Access Control....315
6-22 Multipathing with Fibre Channel....317
6-23 FCoE Adapters....319
6-24 Configuring Software FCoE: Creating VMkernel Ports....320
6-25 Configuring Software FCoE: Activating Software FCoE Adapters....321
6-26 Review of Learner Objectives....322
6-27 Lesson 3: iSCSI Storage....323
6-28 Learner Objectives....324
6-29 iSCSI Components....325
6-30 iSCSI Addressing....327
6-31 Storage Device Naming Conventions....329
6-32 iSCSI Adapters....330
6-33 ESXi Network Configuration for IP Storage....332
6-34 Activating the Software iSCSI Adapter....333
6-35 Discovering iSCSI Targets....334
6-36 iSCSI Security: CHAP....335
6-37 Multipathing with iSCSI Storage....337
6-38 Binding VMkernel Ports with the iSCSI Initiator....338
6-39 Lab 12: Accessing iSCSI Storage....339

https://t.me/learningnets

6-40 Review of Learner Objectives 340
6-41 Lesson 4: VMFS Datastores 341
6-42 Learner Objectives 342
6-43 Creating a VMFS Datastore 343
6-44 Browsing Datastore Contents 344
6-45 About VMFS Datastores 345
6-46 Managing Overcommitted Datastores 346
6-47 Increasing the Size of VMFS Datastores 347
6-48 Datastore Maintenance Mode 348
6-49 Deleting or Unmounting a VMFS Datastore 349
6-50 Multipathing Algorithms 351
6-51 Configuring Storage Load Balancing 352
6-52 Lab 13: Managing VMFS Datastores 354
6-53 Review of Learner Objectives 355
6-54 Lesson 5: NFS Datastores 356
6-55 Learner Objectives 357
6-56 NFS Components 358
6-57 NFS 3 and NFS 4.1 359
6-58 NFS Version Compatibility with Other vSphere Technologies 360
6-59 Configuring NFS Datastores 362
6-60 Configuring ESXi Host Authentication and NFS Kerberos Credentials 363
6-61 Configuring the NFS Datastore to Use Kerberos 365
6-62 Unmounting an NFS Datastore 366
6-63 Multipathing and NFS Storage 367
6-64 Enabling Multipathing for NFS 4.1 369
6-65 Lab 14: Accessing NFS Storage 370
6-66 Review of Learner Objectives 371
6-67 Lesson 6: vSAN Datastores 372
6-68 Learner Objectives 373
6-69 About vSAN Datastores 374
6-70 Disk Groups 375

https://t.me/learningnets

6-71 vSAN Hardware Requirements 376
6-72 Viewing the vSAN Datastore Summary 378
6-73 Objects in vSAN Datastores 379
6-74 VM Storage Policies 380
6-75 Viewing VM Settings for vSAN Information 381
6-76 Lab 15: Using a vSAN Datastore 382
6-77 Review of Learner Objectives 383
6-78 Virtual Beans: Storage 384
6-79 Activity: Using vSAN Storage at Virtual Beans (1) 385
6-80 Activity: Using vSAN Storage at Virtual Beans (2) 386
6-81 Key Points 387

Module 7 Virtual Machine Management 389

7-2 Importance 390
7-3 Module Lessons 391
7-4 Virtual Beans: VM Management 392
7-5 Lesson 1: Creating Templates and Clones 393
7-6 Learner Objectives 394
7-7 About Templates 395
7-8 Creating a Template: Clone VM to Template 396
7-9 Creating a Template: Convert VM to Template 397
7-10 Creating a Template: Clone a Template 398
7-11 Updating Templates 399
7-12 Deploying VMs from a Template 400
7-13 Cloning Virtual Machines 401
7-14 Guest Operating System Customization 402
7-15 About Customization Specifications 403
7-16 Customizing the Guest Operating System 404
7-17 About Instant Clones 405
7-18 Use Cases for Instant Clones 406
7-19 Lab 16: Using VM Templates: Creating Templates and Deploying VMs 407

https://t.me/learningnets

7-20 Review of Learner Objectives....408
7-21 Lesson 2: Working with Content Libraries....409
7-22 Learner Objectives....410
7-23 About Content Libraries....411
7-24 Benefits of Content Libraries....412
7-25 Types of Content Libraries....413
7-26 Adding VM Templates to a Content Library....415
7-27 Deploying VMs from Templates in a Content Library....416
7-28 Lab 17: Using Content Libraries....417
7-29 Review of Learner Objectives....418
7-30 Lesson 3: Modifying Virtual Machines....419
7-31 Learner Objectives....420
7-32 Modifying Virtual Machine Settings....421
7-33 Hot-Pluggable Devices....423
7-34 Dynamically Increasing Virtual Disk Size....425
7-35 Inflating Thin-Provisioned Disks....426
7-36 VM Options: General Settings....427
7-37 VM Options: VMware Tools Settings....428
7-38 VM Options: VM Boot Settings....429
7-39 Removing VMs....431
7-40 Lab 18: Modifying Virtual Machines....432
7-41 Review of Learner Objectives....433
7-42 Lesson 4: Migrating VMs with vSphere vMotion....434
7-43 Learner Objectives....435
7-44 About VM Migration....436
7-45 About vSphere vMotion....437
7-46 Enabling vSphere vMotion....438
7-47 vSphere vMotion Migration Workflow....439
7-48 VM Requirements for vSphere vMotion Migration....441
7-49 Host Requirements for vSphere vMotion Migration (1)....442
7-50 Host Requirements for vSphere vMotion Migration (2)....443

https://t.me/learningnets

7-51 Checking vSphere vMotion Errors 444
7-52 Encrypted vSphere vMotion 445
7-53 Cross vCenter Migrations 446
7-54 Cross vCenter Migration Requirements 447
7-55 Network Checks for Cross vCenter Migrations 448
7-56 VMkernel Networking Layer and TCP/IP Stacks 449
7-57 vSphere vMotion TCP/IP Stacks 451
7-58 Long-Distance vSphere vMotion Migration 452
7-59 Networking Prerequisites for Long-Distance vSphere vMotion 453
7-60 Lab 19: vSphere vMotion Migrations 454
7-61 Review of Learner Objectives 455
7-62 Lesson 5: Enhanced vMotion Compatibility 456
7-63 Learner Objectives 457
7-64 CPU Constraints on vSphere vMotion Migration 458
7-65 About Enhanced vMotion Compatibility 459
7-66 Enhanced vMotion Compatibility Cluster Requirements 461
7-67 Enabling EVC Mode on an Existing Cluster 462
7-68 Changing the EVC Mode for a Cluster 463
7-69 Virtual Machine EVC Mode 464
7-70 Review of Learner Objectives 465
7-71 Lesson 6: Migrating VMs with vSphere Storage vMotion 466
7-72 Learner Objectives 467
7-73 About vSphere Storage vMotion 468
7-74 vSphere Storage vMotion In Action 469
7-75 Identifying Storage Arrays That Support vSphere Storage APIs - Array Integration 471
7-76 vSphere Storage vMotion Guidelines and Limitations 472
7-77 Changing Both Compute Resource and Storage During Migration (1) 473
7-78 Changing Both Compute Resource and Storage During Migration (2) 474
7-79 Lab 20: vSphere Storage vMotion Migrations 475
7-80 Review of Learner Objectives 476

https://t.me/learningnets

7-81 Lesson 7: Creating Virtual Machine Snapshots 477
7-82 Learner Objectives....478
7-83 VM Snapshots....479
7-84 Taking Snapshots....480
7-85 Types of Snapshots....481
7-86 VM Snapshot Files 483
7-87 VM Snapshot Files Example (1) 485
7-88 VM Snapshot Files Example (2) 486
7-89 VM Snapshot Files Example (3) 487
7-90 Managing Snapshots 488
7-91 Deleting VM Snapshots (1)....490
7-92 Deleting VM Snapshots (2)....491
7-93 Deleting VM Snapshots (3)....492
7-94 Deleting All VM Snapshots 493
7-95 About Snapshot Consolidation....494
7-96 Discovering When to Consolidate Snapshots....495
7-97 Consolidating Snapshots 496
7-98 Lab 21: Working with Snapshots....497
7-99 Review of Learner Objectives....498
7-100 Lesson 8: vSphere Replication and Backup 499
7-101 Learner Objectives....500
7-102 About vSphere Replication....501
7-103 About the vSphere Replication Appliance....502
7-104 Replication Functions 504
7-105 Deploying the vSphere Replication Appliance 505
7-106 Configuring vSphere Replication for a Single VM....506
7-107 Configuring Recovery Point Objective and Point in Time Instances 507
7-108 Recovering Replicated VMs....508
7-109 Backup and Restore Solution for VMs....510
7-110 vSphere Storage APIs - Data Protection: Offloaded Backup Processing 511
7-111 vSphere Storage APIs - Data Protection: Changed-Block Tracking 513

https://t.me/learningnets

7-112 Review of Learner Objectives 514
7-113 Activity: Virtual Beans VM Management (1) 515
7-114 Activity: Virtual Beans VM Management (2) 516
7-115 Activity: Virtual Beans VM Management (3) 517
7-116 Key Points 518

Module 8 Resource Management and Monitoring 519

8-2 Importance 520
8-3 Module Lessons 521
8-4 Virtual Beans: Resource Management and Monitoring 522
8-5 Lesson 1: Virtual CPU and Memory Concepts 523
8-6 Learner Objectives 524
8-7 Memory Virtualization Basics 525
8-8 VM Memory Overcommitment 526
8-9 Memory Overcommit Techniques 528
8-10 Configuring Multicore VMs 530
8-11 About Hyperthreading 532
8-12 CPU Load Balancing 533
8-13 Review of Learner Objectives 534
8-14 Lesson 2: Resource Controls 535
8-15 Learner Objectives 536
8-16 Reservations, Limits, and Shares 537
8-17 Resource Allocation Reservations: RAM 538
8-18 Resource Allocation Reservations: CPU 539
8-19 Resource Allocation Limits 540
8-20 Resource Allocation Shares 541
8-21 Resource Shares Example (1) 542
8-22 Resource Shares Example (2) 543
8-23 Resource Shares Example (3) 544
8-24 Resource Shares Example (4) 545
8-25 Defining Resource Allocation Settings for a VM 546

https://t.me/learningnets

8-26 Viewing VM Resource Allocation Settings....547
8-27 Lab 22: Controlling VM Resources....548
8-28 Review of Learner Objectives....549
8-29 Lesson 3: Resource Monitoring Tools....550
8-30 Learner Objectives....551
8-31 Performance-Tuning Methodology....552
8-32 Resource-Monitoring Tools....553
8-33 Guest Operating System Monitoring Tools....554
8-34 Using Perfmon to Monitor VM Resources....555
8-35 Using esxtop to Monitor VM Resources....556
8-36 Monitoring Inventory Objects with Performance Charts....557
8-37 Working with Overview Performance Charts....558
8-38 Working with Advanced Performance Charts....559
8-39 Chart Options: Real-Time and Historical....560
8-40 Chart Types: Bar and Pie....562
8-41 Chart Types: Line....563
8-42 Chart Types: Stacked....564
8-43 Chart Types: Stacked Per VM....565
8-44 Saving Charts....566
8-45 About Objects and Counters....567
8-46 About Statistics Types....568
8-47 About Rollup....569
8-48 Review of Learner Objectives....571
8-49 Lesson 4: Monitoring Resource Use....572
8-50 Learner Objectives....573
8-51 Interpreting Data from Tools....574
8-52 CPU-Constrained VMs (1)....575
8-53 CPU-Constrained VMs (2)....577
8-54 Memory-Constrained VMs (1)....578
8-55 Memory-Constrained VMs (2)....579
8-56 Memory-Constrained Hosts....580

https://t.me/learningnets

8-57 Disk-Constrained VMs 581
8-58 Monitoring Disk Latency 582
8-59 Network-Constrained VMs 583
8-60 Lab 23: Monitoring Virtual Machine Performance 584
8-61 Review of Learner Objectives 585
8-62 Lesson 5: Using Alarms 586
8-63 Learner Objectives 587
8-64 About Alarms 588
8-65 Predefined Alarms (1) 589
8-66 Predefined Alarms (2) 590
8-67 Creating a Custom Alarm 591
8-68 Defining the Alarm Target Type 592
8-69 Defining the Alarm Rule: Trigger (1) 593
8-70 Defining the Alarm Rule: Trigger (2) 594
8-71 Defining the Alarm Rule: Setting the Notification 595
8-72 Defining the Alarm Reset Rules 596
8-73 Enabling the Alarm 597
8-74 Configuring vCenter Server Notifications 598
8-75 Lab 24: Using Alarms 599
8-76 Review of Learner Objectives 600
8-77 Activity: Virtual Beans Resource Monitoring (1) 601
8-78 Activity: Virtual Beans Resource Management and Monitoring (2) 602
8-79 Key Points 603

Module 9 vSphere Clusters 605

9-2 Importance 606
9-3 Module Lessons 607
9-4 Virtual Beans: vSphere Clusters 608
9-5 Lesson 1: vSphere Clusters Overview 609
9-6 Learner Objectives 610
9-7 About vSphere Clusters 611

https://t.me/learningnets

9-8 Creating a vSphere Cluster and Enabling Cluster Features 612
9-9 Configuring the Cluster Using Quickstart 613
9-10 Configuring the Cluster Manually 615
9-11 Adding a Host to a Cluster 616
9-12 Viewing Cluster Summary Information 617
9-13 Monitoring Cluster Resources 618
9-14 Review of Learner Objectives 619
9-15 Lesson 2: vSphere DRS 620
9-16 Learner Objectives 621
9-17 About vSphere DRS 622
9-18 vSphere DRS: VM Focused 623
9-19 About the VM DRS Score 624
9-20 VM DRS Score List 625
9-21 Viewing VM DRS Scores Using Performance Charts (1) 626
9-22 Viewing VM DRS Scores Using Performance Charts (2) 627
9-23 Viewing vSphere DRS Settings 628
9-24 vSphere DRS Settings: Automation Level 629
9-25 vSphere DRS Settings: Migration Threshold 630
9-26 vSphere DRS Settings: Predictive DRS 632
9-27 vSphere DRS Settings: VM Swap File Location 633
9-28 vSphere DRS Settings: VM Affinity 634
9-29 vSphere DRS Settings: DRS Groups 635
9-30 vSphere DRS Settings: VM-Host Affinity Rules 636
9-31 VM-Host Affinity Preferential Rules 637
9-32 VM-Host Affinity Required Rules 638
9-33 vSphere DRS Settings: VM-Level Automation 639
9-34 vSphere DRS Cluster Requirements 640
9-35 Viewing vSphere DRS Cluster Resource Utilization 641
9-36 Viewing vSphere DRS Recommendations 642
9-37 Maintenance Mode and Standby Mode 643
9-38 Removing a Host from the vSphere DRS Cluster 644

https://t.me/learningnets

9-39 vSphere DRS and Dynamic DirectPath I/O 645
9-40 Adding a Dynamic DirectPath I/O Device to a VM 646
9-41 Lab 25: Implementing vSphere DRS Clusters 647
9-42 Review of Learner Objectives 648
9-43 Lesson 3: Introduction to vSphere HA 649
9-44 Learner Objectives 650
9-45 Protection at Every Level 651
9-46 About vSphere HA 653
9-47 vSphere HA Scenario: ESXi Host Failure 654
9-48 vSphere HA Scenario: Guest Operating System Failure 655
9-49 vSphere HA Scenario: Application Failure 656
9-50 vSphere HA Scenario: Datastore Accessibility Failures 657
9-51 vSphere HA Scenario: Protecting VMs Against Network Isolation 659
9-52 Importance of Redundant Heartbeat Networks 660
9-53 Redundancy Using NIC Teaming 661
9-54 Redundancy Using Additional Networks 662
9-55 Review of Learner Objectives 663
9-56 Lesson 4: vSphere HA Architecture 664
9-57 Learner Objectives 665
9-58 vSphere HA Architecture: Agent Communication 666
9-59 vSphere HA Architecture: Network Heartbeats 669
9-60 vSphere HA Architecture: Datastore Heartbeats 670
9-61 vSphere HA Failure Scenarios 671
9-62 Failed Subordinate Hosts 672
9-63 Failed Master Hosts 674
9-64 Isolated Hosts 675
9-65 VM Storage Failures 676
9-66 Protecting Against Storage Failures with VMCP 677
9-67 vSphere HA Design Considerations 678
9-68 Review of Learner Objectives 679
9-69 Lesson 5: Configuring vSphere HA 680

https://t.me/learningnets

9-70 Learner Objectives....681
9-71 vSphere HA Prerequisites....682
9-72 Configuring vSphere HA Settings....683
9-73 vSphere HA Settings: Failures and Responses....684
9-74 vSphere HA Settings: VM Monitoring....686
9-75 vSphere HA Settings: Heartbeat Datastores....687
9-76 vSphere HA Settings: Admission Control....688
9-77 Example: Admission Control Using Cluster Resources Percentage....690
9-78 Example: Admission Control Using Slots (1)....691
9-79 Example: Admission Control Using Slots (2)....692
9-80 vSphere HA Settings: Performance Degradation VMs Tolerate....693
9-81 vSphere HA Setting: Default VM Restart Priority....695
9-82 vSphere HA Settings: Advanced Options....696
9-83 vSphere HA Settings: VM-Level Settings....697
9-84 About vSphere HA Orchestrated Restart....698
9-85 VM Dependencies in Orchestrated Restart (1)....699
9-86 VM Dependencies in Orchestrated Restart (2)....700
9-87 Network Configuration and Maintenance....701
9-88 Monitoring vSphere HA Cluster Status....702
9-89 Using vSphere HA with vSphere DRS....703
9-90 Lab 26: Using vSphere HA....704
9-91 Review of Learner Objectives....705
9-92 Lesson 6: Introduction to vSphere Fault Tolerance....706
9-93 Learner Objectives....707
9-94 About vSphere Fault Tolerance....708
9-95 vSphere Fault Tolerance Features....709
9-96 vSphere Fault Tolerance with vSphere HA and vSphere DRS....710
9-97 Redundant VMDK Files....711
9-98 vSphere Fault Tolerance Checkpoint....712
9-99 vSphere Fault Tolerance: Precopy....713
9-100 vSphere Fault Tolerance Fast Checkpointing....714

https://t.me/learningnets

9-101 vSphere Fault Tolerance Shared Files....715
9-102 Enabling vSphere Fault Tolerance on a VM....716
9-103 Review of Learner Objectives....717
9-104 Activity: Virtual Beans Clusters (1)....718
9-105 Activity: Virtual Beans Clusters (2)....719
9-106 Key Points....720

Module 10 vSphere Lifecycle Management 721

10-2 Importance....722
10-3 Module Lessons....723
10-4 Virtual Beans: Lifecycle Management....724
10-5 Lesson 1: vCenter Server Update Planner....725
10-6 Learner Objectives....726
10-7 Overview of vCenter Server Update Planner....727
10-8 Update Planner Requirements....728
10-9 Update Planner View in the vSphere Client....729
10-10 Interoperability View in vSphere Client....730
10-11 Exporting Report Results....731
10-12 Managing the vCenter Server Life Cycle....732
10-13 Review of Learner Objectives....733
10-14 Lesson 2: Overview of vSphere Lifecycle Manager....734
10-15 Learner Objectives....735
10-16 Introduction to vSphere Lifecycle Manager....736
10-17 Baselines and Images....737
10-18 vSphere Lifecycle Manager Home View....738
10-19 Patch Settings....739
10-20 vSphere Lifecycle Manager Integration with vSphere DRS....740
10-21 Review of Learner Objectives....741
10-22 Lesson 3: Working with Baselines....742
10-23 Learner Objectives....743
10-24 Baselines and Baseline Groups....744

https://t.me/learningnets

10-25 Creating and Editing Patch or Extension Baselines 745
10-26 Creating a Baseline 746
10-27 Creating a Baseline: Name and Description 747
10-28 Creating a Baseline: Select Patches Automatically....748
10-29 Creating a Baseline: Select Patches Manually....749
10-30 Updating Your Host or Cluster with Baselines 750
10-31 Remediation Precheck....751
10-32 Remediating Hosts....752
10-33 Review of Learner Objectives....753
10-34 Lesson 4: Working with Images....754
10-35 Learner Objectives....755
10-36 Elements of ESXi Images....756
10-37 Image Depots....758
10-38 Importing Updates 759
10-39 Using Images to Perform ESXi Host Life Cycle Operations....760
10-40 Creating an ESXi Image for a New Cluster....761
10-41 Checking Image Compliance....762
10-42 Running a Remediation Precheck....763
10-43 Hardware Compatibility 764
10-44 Standalone VIBs....765
10-45 Remediating a Cluster Against an Image....766
10-46 Reviewing Remediation Impact....767
10-47 Recommended Images....768
10-48 Viewing Recommended Images 769
10-49 Selecting a Recommended Image 771
10-50 Customizing Cluster Images 772
10-51 Lab 27: Using vSphere Lifecycle Manager....773
10-52 Review of Learner Objectives....774
10-53 Lesson 5: Managing the Life Cycle of VMware Tools and VM Hardware....775
10-54 Learner Objectives....776
10-55 Keeping VMware Tools Up To Date....777

https://t.me/learningnets

10-56 Upgrading VMware Tools (1) 778
10-57 Upgrading VMware Tools (2) 779
10-58 Keeping VM Hardware Up To Date 780
10-59 Upgrading VM Hardware (1) 781
10-60 Upgrading VM Hardware (2) 782
10-61 Review of Learner Objectives 783
10-62 Virtual Beans: Conclusion 784
10-63 Key Points 785

https://t.me/learningnets

Module 1

Course Introduction

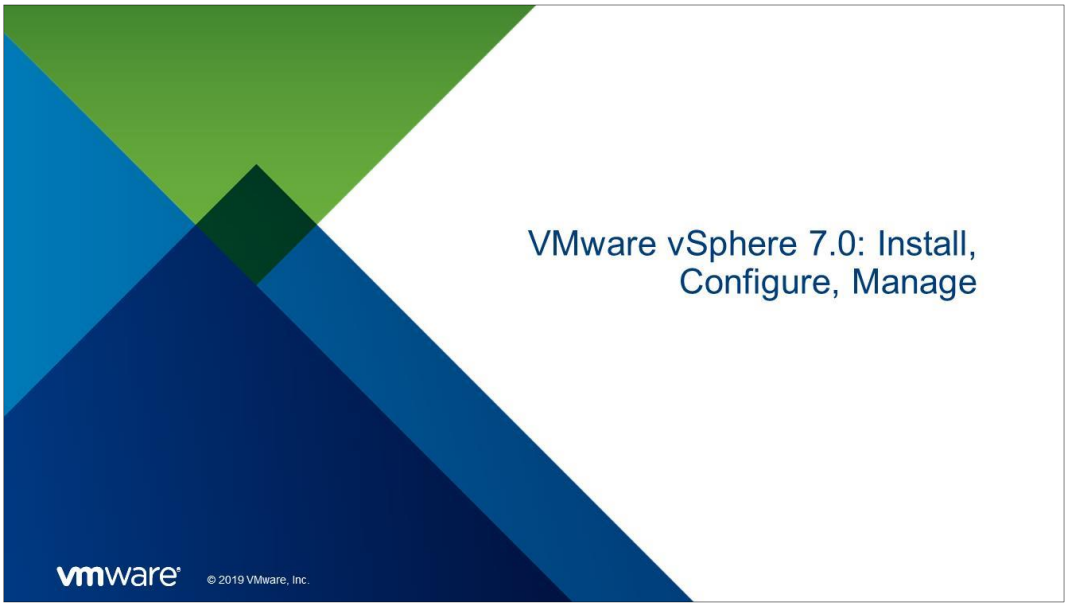

https://t.me/learningnets

1-2 Course Introduction

https://t.me/learningnets

1-3 Importance

As a vSphere administrator, you require knowledge about vSphere components and resources and how they work together in your environment. You also require practical skills in installing, deploying, and managing these components and resources. By developing your knowledge and skills, you can build and run a highly scalable vSphere virtual infrastructure.

https://t.me/learningnets

1-4 Learner Objectives (1)

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

- Install and configure ESXi hosts
- Deploy and configure vCenter Server Appliance
- Use the vSphere Client to manage the vCenter Server inventory and the vCenter Server configuration
- Create virtual networks with vSphere standard switches
- Describe the storage technologies supported by vSphere
- Configure virtual storage using iSCSI and NFS storage
- Create and manage VMFS datastores
- Use the vSphere Client to create virtual machines, templates, clones, and snapshots
- Create a content library for deploying virtual machines

https://t.me/learningnets

1-5 Learner Objectives (2)

- Manage virtual machine resource use
- Migrate virtual machines with vSphere vMotion and vSphere Storage vMotion
- Create and manage a vSphere cluster that is enabled with vSphere HA and vSphere DRS
- Use vSphere Lifecycle Manager to perform upgrades to ESXi hosts and virtual machines

https://t.me/learningnets

1-6 Course Outline

1. Course Introduction
2. Introduction to vSphere and the Software-Defined Data Center
3. Virtual Machines
4. vCenter Server
5. Configuring and Managing Virtual Networks
6. Configuring and Managing Virtual Storage
7. Virtual Machine Management
8. Resource Management and Monitoring
9. vSphere Clusters
10. vSphere Lifecycle Management

https://t.me/learningnets

1-7 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 `esxtop` command. • ... found in the `var/log/messages` file.
`Monospace Bold`	Identifies user inputs: • Enter **`ipconfig /release`**.
Boldface	Identifies user interface controls: • Click the **Configuration** tab.
Italic	Identifies book titles: • *vSphere Virtual Machine Administration*
`< >`	Indicates placeholder variables: • <ESXi_host_name> • ... the `Settings/<Your_Name>.txt` file

https://t.me/learningnets

1-8 References (1)

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

https://t.me/learningnets

1-9 References (2)

Title	Location
vSphere Security	https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-52188148-C579-4F6A-8335-CFBCE0DD2167.html
vSphere Resource Management	https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.resmgmt.doc/GUID-98BD5A8A-260A-494F-BAAE-74781F5C4B87.html
VMware Compatibility Guide	https://vmware.com/resources/compatibility
VMware Configuration Maximums	https://configmax.vmware.com

https://t.me/learningnets

1-10 VMware Online Resources

Documentation for vSphere: 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.

https://t.me/learningnets

1-11 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.

https://t.me/learningnets

1-12 VMware Certification Overview

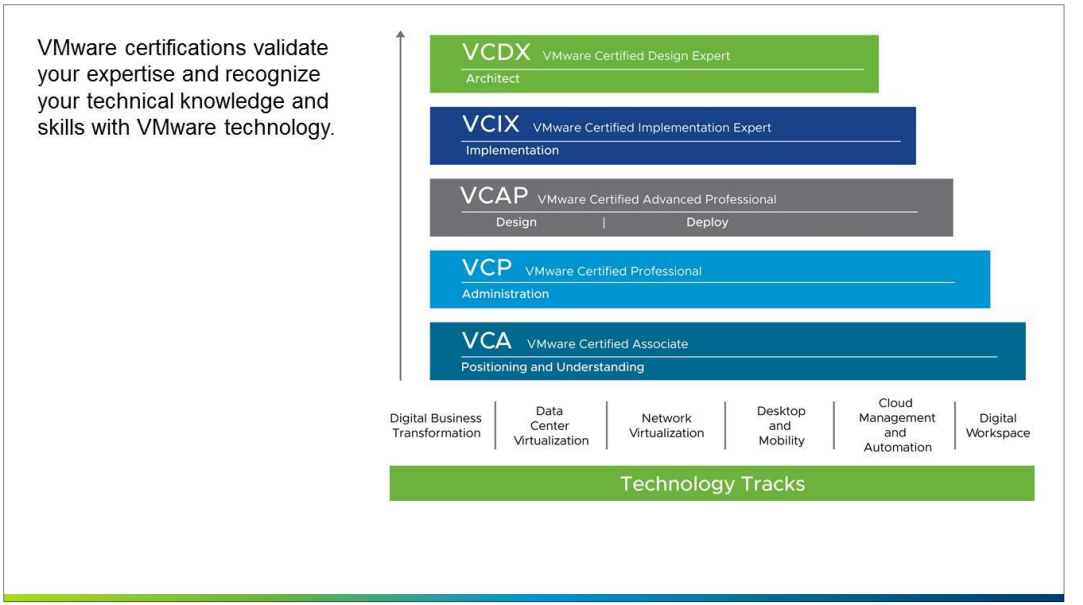

VMware certification sets the standards for IT professionals who work with VMware technology. Certifications are grouped into technology tracks. Each track offers one or more levels of certification (up to five levels).

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

https://t.me/learningnets

1-13 VMware Badge Overview

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.

https://t.me/learningnets

1-14 Virtual Beans: Introduction

Virtual Beans is a coffee company that owns a chain of cyber cafés. Each café sells coffee drinks, snacks, and packaged coffee beans. Each café is also equipped with a variety of video games and high-speed Internet access.

Virtual Beans has an online store (vmbeans.com) where you can purchase coffee beans and various accessories.

Virtual Beans is a fast-growing company. After much success over the years, it went from a single, small café to a company that owns a chain of cafés spanning multiple cities. The online web store is also a huge success.

Virtual Beans recently purchased vSphere 7 for its data center.

You work as a system administrator at Virtual Beans. You will be part of the IT team in charge of deploying vSphere 7 in the data center. You are new to vSphere, but you have two years experience working for Virtual Beans.

https://t.me/learningnets

Module 2

Introduction to vSphere and the Software-Defined Data Center

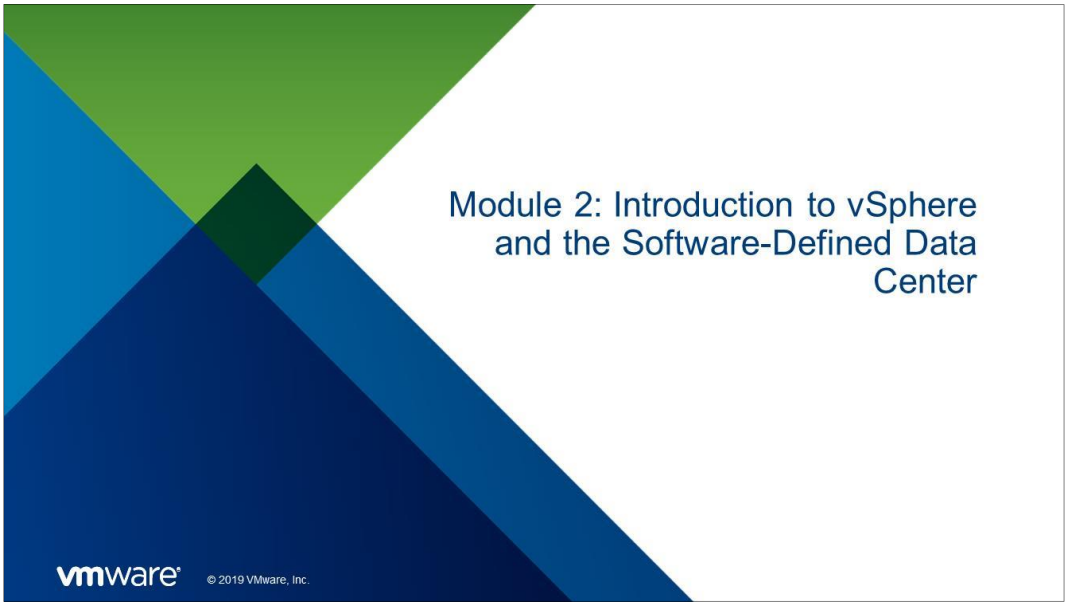

https://t.me/learningnets

2-2 Importance

As a vSphere administrator, you must be familiar with the components on which vSphere is based. You must also understand the following concepts:

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

https://t.me/learningnets

2-3 Module Lessons

1. Overview of vSphere and Virtual Machines
2. vSphere Virtualization of Resources
3. vSphere User Interfaces
4. Overview of ESXi

https://t.me/learningnets

2-4 Virtual Beans: Data Center

Virtual Beans has a data center at its company headquarters. The company's goals are as follows:

- Use the latest version of vSphere.
- Create a cost-effective, leading-edge data center.
- Create a secure, scalable, high-performing, and highly available infrastructure.
- Create a vSphere infrastructure that follows VMware best practices.
- Open a second data center to serve as a backup site to the primary data center and to host new applications.

As a Virtual Beans administrator, you must decide how to implement these goals. But first, you must understand how a vSphere data center works.

https://t.me/learningnets

2-5 Lesson 1: Overview of vSphere and Virtual Machines

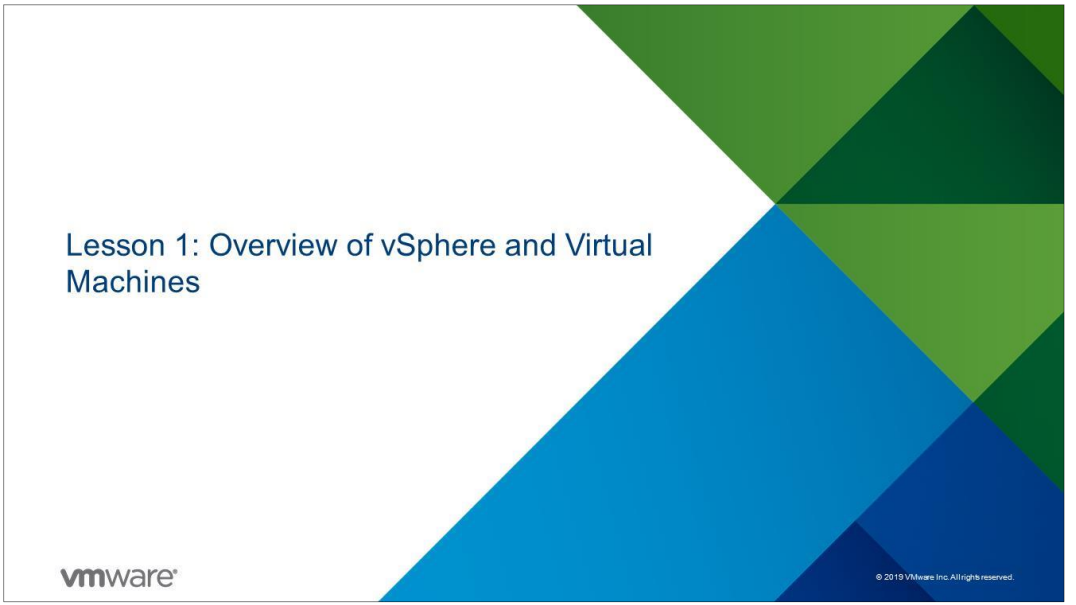

https://t.me/learningnets

2-6 Learner Objectives

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

- Explain basic virtualization concepts
- Describe how vSphere fits into the software-defined data center and the cloud infrastructure
- Describe how to proactively manage your vSphere environment

https://t.me/learningnets

2-7 Terminology (1)

Virtualization is associated with several key concepts, products, and features.

Term	Definition	Examples
Operating system	Software designed to allocate physical resources to applications	Microsoft Windows, Linux
Application	Software that runs on an operating system, consuming physical resources	Microsoft Office, Chrome
Virtual machine	Specialized application that abstracts hardware resources into software	
Guest	The operating system that runs in a VM (also called the guest operating system)	Microsoft Windows, Linux
Hypervisor	Specialized operating system designed to run VMs	ESXi, Workstation, Fusion
Host	Physical computer that provides resources to the ESXi hypervisor	

https://t.me/learningnets

2-8 Terminology (2)

Term	Definition
vSphere	Server virtualization product of VMware that combines the ESXi hypervisor and the vCenter Server management platform
Cluster	Group of ESXi hosts whose resources are shared by VMs
vSphere vMotion	Feature that supports the migration of powered-on VMs from host to host without service interruption
vSphere HA	Cluster feature that protects against host hardware failures by restarting VMs on hosts that are running normally
vSphere DRS	Cluster feature that uses vSphere vMotion to place VMs on hosts and ensure that each VM receives the resources that it needs

https://t.me/learningnets

2-9 About Virtual Machines

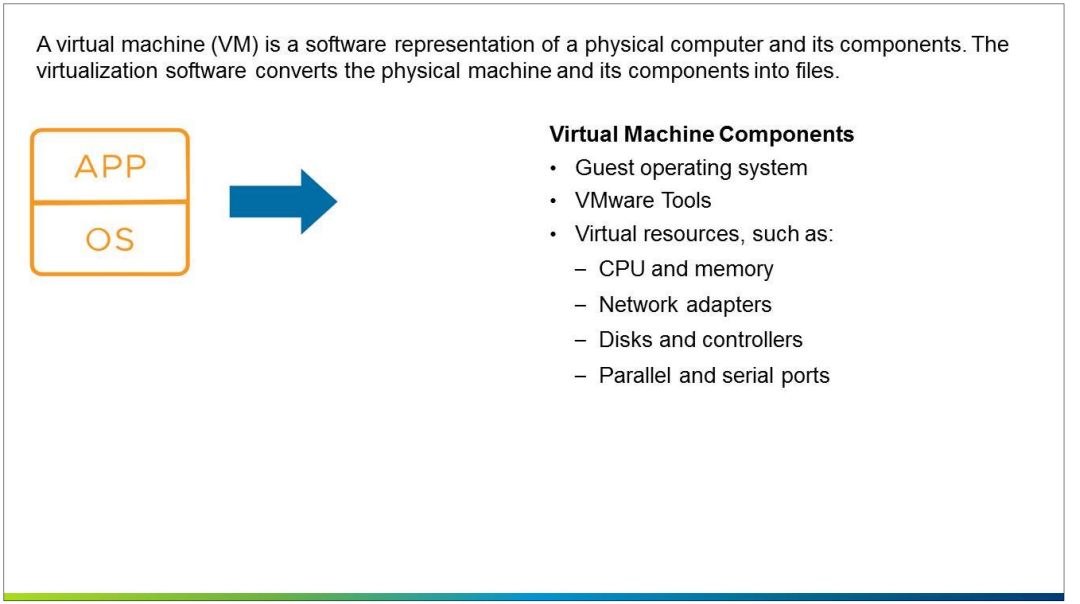

A virtual machine (VM) includes a set of specification and configuration files and is supported by the physical resources of a host. Every VM has virtual devices that provide the same functionality as physical hardware but are more portable, more secure, and easier to manage.

VMs typically include an operating system, applications, VMware Tools, and both virtual resources and hardware that you manage in much the same way as you manage a physical computer.

VMware Tools is a bundle of drivers. Using these drivers, the guest operating system can interact efficiently with the guest hardware. VMware Tools adds extra functionality so that ESXi can better manage the VM's use of physical hardware.

https://t.me/learningnets

2-10 Benefits of Using Virtual Machines

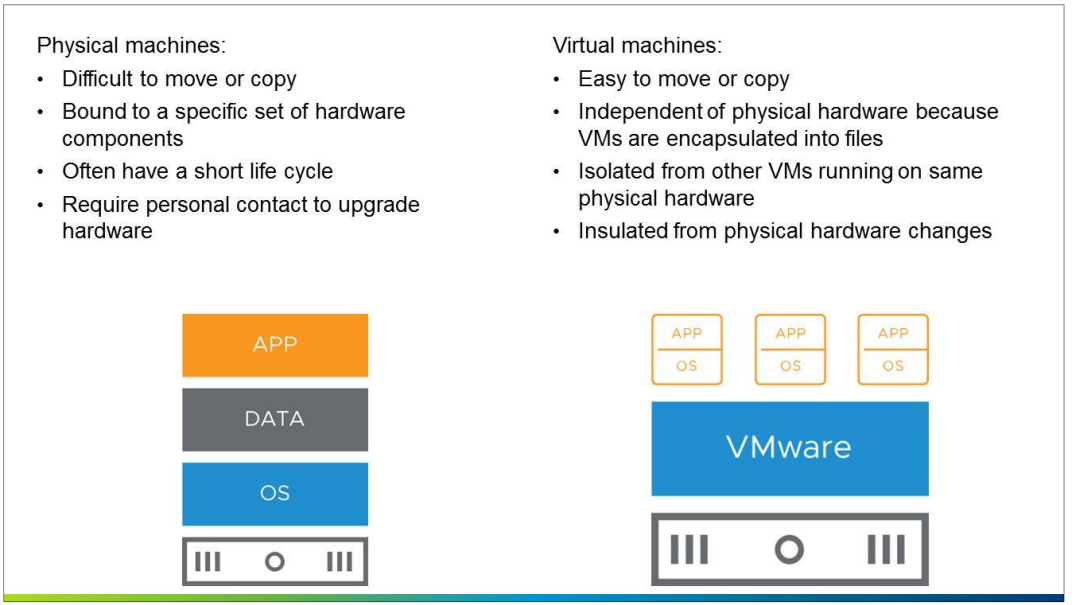

In a physical machine, the operating system (for example, Windows or Linux) is installed directly on the hardware. The operating system requires specific device drivers to support specific hardware. If the computer is upgraded with new hardware, new device drivers are required.

If applications interface directly with hardware drivers, an upgrade to the hardware, drivers, or both can have significant repercussions if incompatibilities exist. Because of these potential repercussions, hands-on technical support personnel must test hardware upgrades against a wide variety of application suites and operating systems. Such testing costs time and money.

Virtualizing these systems saves on such costs because VMs are 100 percent software.

Multiple VMs are isolated from one another. You can have a database server and an email server running on the same physical computer. The isolation between the VMs means that software-dependency conflicts are not a problem. Even users with system administrator privileges on a VM's guest operating system cannot breach this layer of isolation to access another VM. These users must explicitly be granted access by the ESXi system administrator. As a result of VM

https://t.me/learningnets

isolation, if a guest operating system running in a VM fails, other VMs on the same host are unaffected and continue to run.

A guest operating system failure does not affect access and performance:

- Users can still access the other VMs.
- The operational VMs can access the resources that they need.
- The other VMs can still perform.

With VMs, you can consolidate your physical servers and make more efficient use of your hardware. Because a VM is a set of files, features that are not available or not as efficient on physical architectures are available to you, for example:

- You can rapidly and consistently provision VMs.
- With VMs, you can use live migration, fault tolerance, high availability, and disaster recovery scenarios to increase uptime and reduce recovery time from failures.
- You can use multitenancy to mix VMs into specialized configurations, such as a DMZ.

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

https://t.me/learningnets

2-11 Types of Virtualization

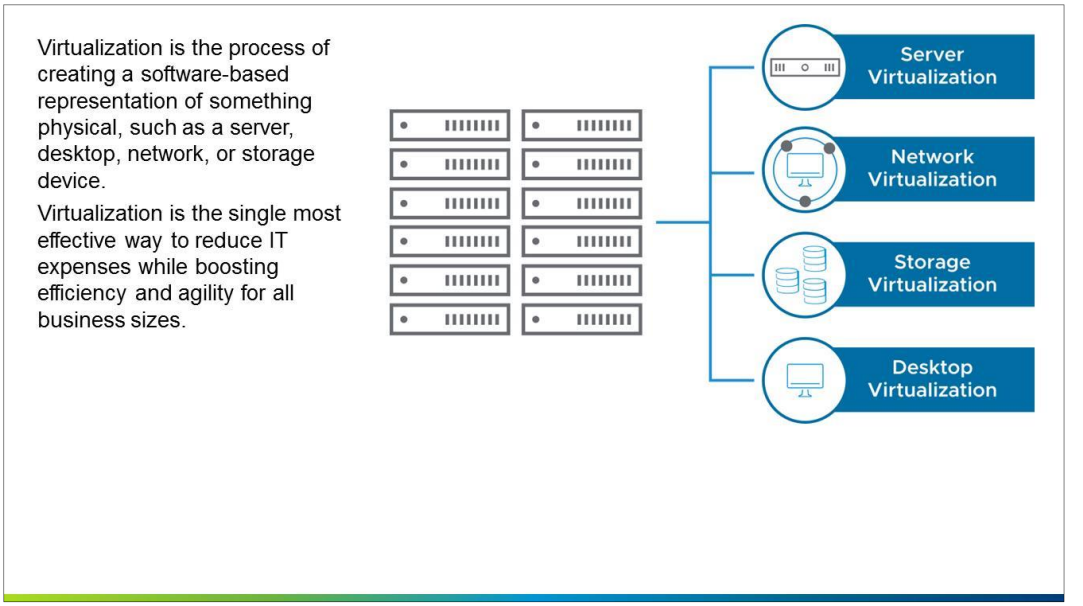

Server virtualization addresses inefficiencies by allowing multiple operating systems to run on a single physical server as VMs, each with access to the underlying server's computing resources.

Network virtualization is the complete reproduction of a physical network in software. Applications run on the virtual network exactly as if on a physical network.

Storage virtualization is the process of creating a software-based representation of network storage devices into what appears to be a single unit.

By deploying desktops as a managed service, you can respond more quickly to changing needs and opportunities.

https://t.me/learningnets

2-12 About the Software-Defined Data Center

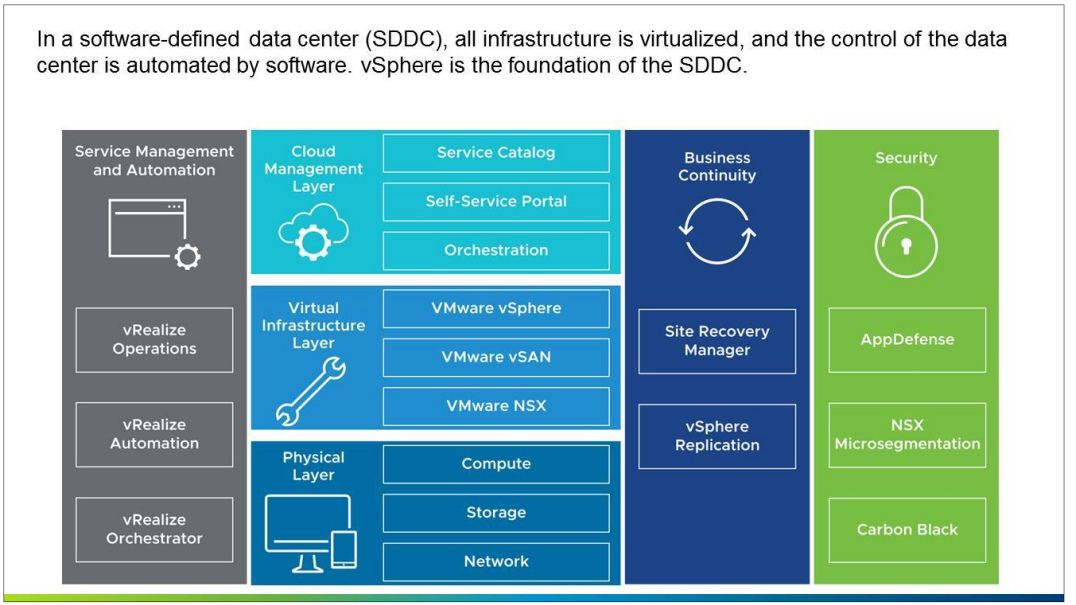

A software-defined virtual data center (SDDC) is deployed with isolated computing, storage, networking, and security resources that are faster than the traditional, hardware-based data center.

All the resources (CPU, memory, disk, and network) of a software-defined data center are abstracted into files. This abstraction brings the benefits of virtualization at all levels of the infrastructure, independent of the physical infrastructure.

An SDDC can include the following components:

- Service management and automation: Use service management and automation to track and analyze the operation of multiple data sources in the multiregion SDDC. Deploy vRealize Operations Manager and vRealize Log Insight across multiple nodes for continued availability and increased log ingestion rates.
- Cloud management layer: This layer includes the service catalog, which houses the facilities to be deployed. The cloud management layer also includes orchestration, which provides the

https://t.me/learningnets

workflows to deploy catalog items, and the self-service portal for end users to access and use the SDDC.

- Virtual infrastructure layer: This layer establishes a robust virtualized environment that all other solutions integrate with. The virtual infrastructure layer includes the virtualization platform for the hypervisor, pools of resources, and virtualization control. Additional processes and technologies build on the infrastructure to support Infrastructure as a Service (IaaS) and Platform as a Service (PaaS).
- Physical layer: The lowest layer of the solution includes compute, storage, and network components.
- Security: Customers use this layer of the platform to meet demanding compliance requirements for virtualized workloads and to manage business risk.

https://t.me/learningnets

2-13 vSphere and Cloud Computing

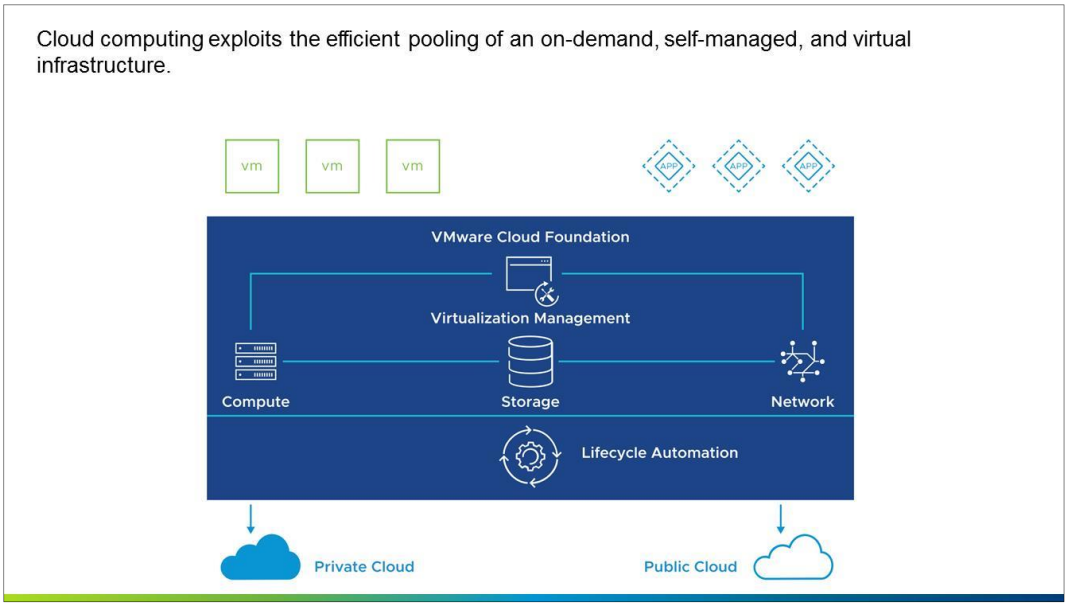

As defined by the National Institute of Standards and Technology (NIST), cloud computing is a model for the ubiquitous, convenient, and on-demand network access to a shared pool of configurable computing resources.

For example, networks, servers, storage, applications, and services can be rapidly provisioned and released with minimal management effort or little service provider interaction.

vSphere is the foundation for the technology that supports shared and configurable resource pools. vSphere abstracts the physical resources of the data center to separate the workload from the physical hardware. A software user interface can provide the framework for managing and maintaining this abstraction and allocation.

VMware Cloud Foundation is the unified SDDC platform that bundles vSphere (ESXi and vCenter Server), vSAN, and NSX into a natively integrated stack to deliver enterprise-ready cloud infrastructure. VMware Cloud Foundation discovers the hardware, installs the VMware stack (ESXi, vCenter Server, vSAN, and NSX), manages updates, and performs lifecycle management. VMware Cloud Foundation can be self-deployed on compatible hardware or preloaded by partners

https://t.me/learningnets

and can be used in both private and public clouds (VMware Cloud on AWS or VMware cloud providers).

Use cases:

- Cloud infrastructure: Exploit the high performance, availability, and scalability of the SDDC to run mission-critical applications such as databases, web applications, and virtual desktop infrastructure (VDI).
- IT automation: Automate infrastructure and application delivery with self-service capabilities.
- VDI: Provide a complete solution for VDI deployment at scale. It simplifies the planning and design with standardized and tested solutions fully optimized for VDI workloads.
- Hybrid cloud: Build a hybrid cloud with a common infrastructure and a consistent operational model, connecting your on-premises and off-premises data center that is compatible, stretched, and distributed.

To find out more about VMware cloud computing, go to http://www.vmware.com/cloud-computing/overview.html.

https://t.me/learningnets

2-14 About VMware Skyline

VMware Skyline is a proactive support technology that provides predictive analysis and proactive recommendations to help you 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. 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.

https://t.me/learningnets

2-15 VMware Skyline Family

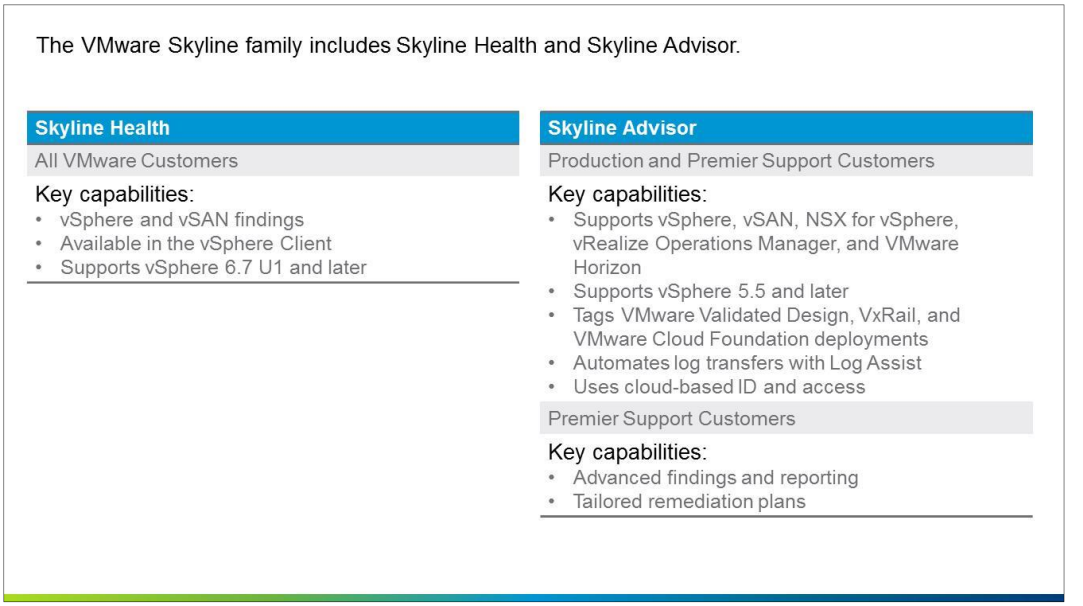

With Basic Support, you can access Skyline findings and recommendations for vSphere and vSAN by using Skyline Health in the vSphere Client (version 6.7 and later).

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

https://t.me/learningnets

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

- An advanced set of proactive 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 the following services:
  - Designated support team
  - Direct access to senior-level technical support engineers
  - Assistance with multivendor troubleshooting
  - 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/kb/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.

https://t.me/learningnets

2-16 Review of Learner Objectives

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

- Explain basic virtualization concepts
- Describe how vSphere fits into the software-defined data center and the cloud infrastructure
- Describe how to proactively manage your vSphere environment

https://t.me/learningnets

https://t.me/learningnets

2-18 Learner Objectives

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

- Explain how vSphere interacts with CPUs, memory, networks, and storage

https://t.me/learningnets

2-19 Virtual Machine: Guest and Consumer of ESXi Host

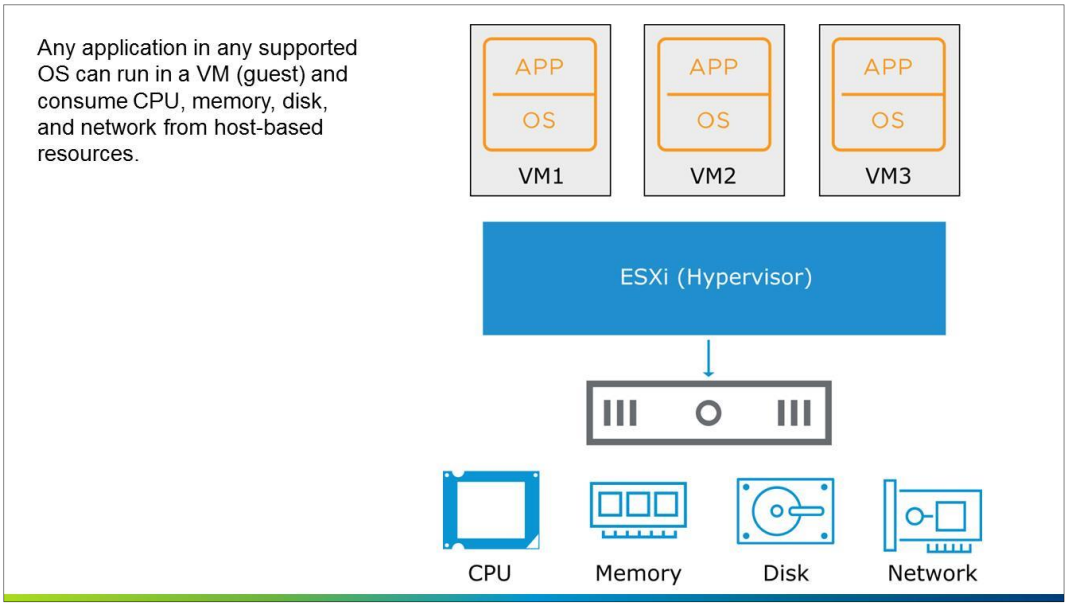

A virtual machine is an abstraction in software of a physical machine. A VM turns components into files that act like physical components.

For the list of all supported operating systems, see *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

https://t.me/learningnets

2-20 Physical and Virtual Architecture

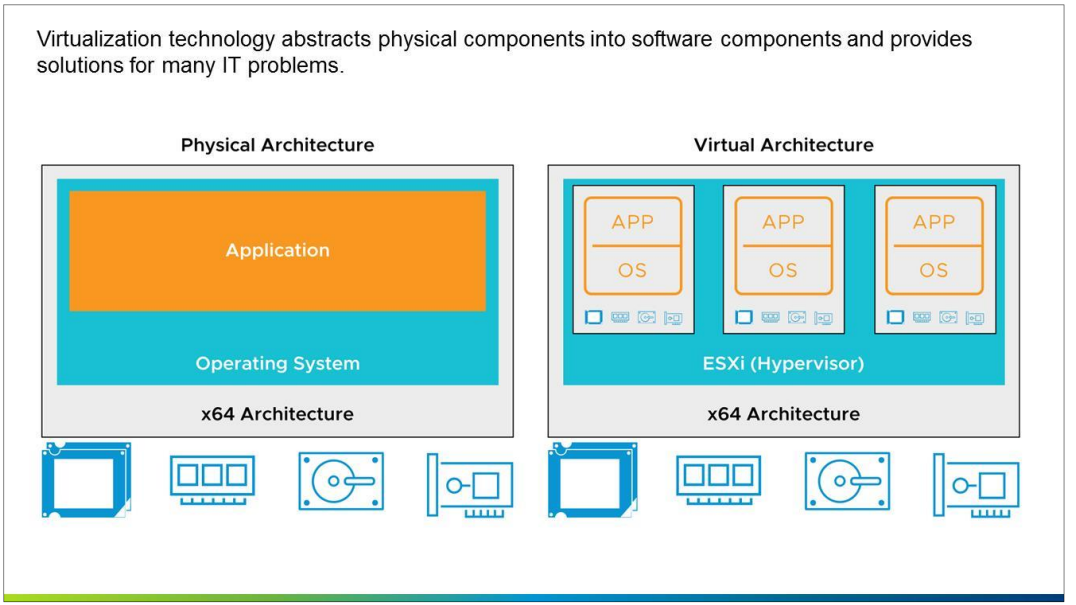

You can use virtualization to consolidate and run multiple workloads as VMs on a single computer.

The slide shows the differences between a virtualized and a nonvirtualized host.

In traditional architectures, the operating system interacts directly with the installed hardware. The operating system schedules processes to run, allocates memory to applications, sends and receives data on network interfaces, and both reads from and writes to attached storage devices.

In comparison, a virtualized host interacts with the installed hardware through a thin layer of software called the virtualization layer or hypervisor.

The hypervisor provides physical hardware resources dynamically to VMs as needed to support the operation of the VMs. With the hypervisor, VMs can operate with a degree of independence from the underlying physical hardware. For example, a VM can be moved from one physical host to another. In addition, its virtual disks can be moved from one type of storage to another without affecting the functioning of the VM.

https://t.me/learningnets

2-21 Physical Resource Sharing

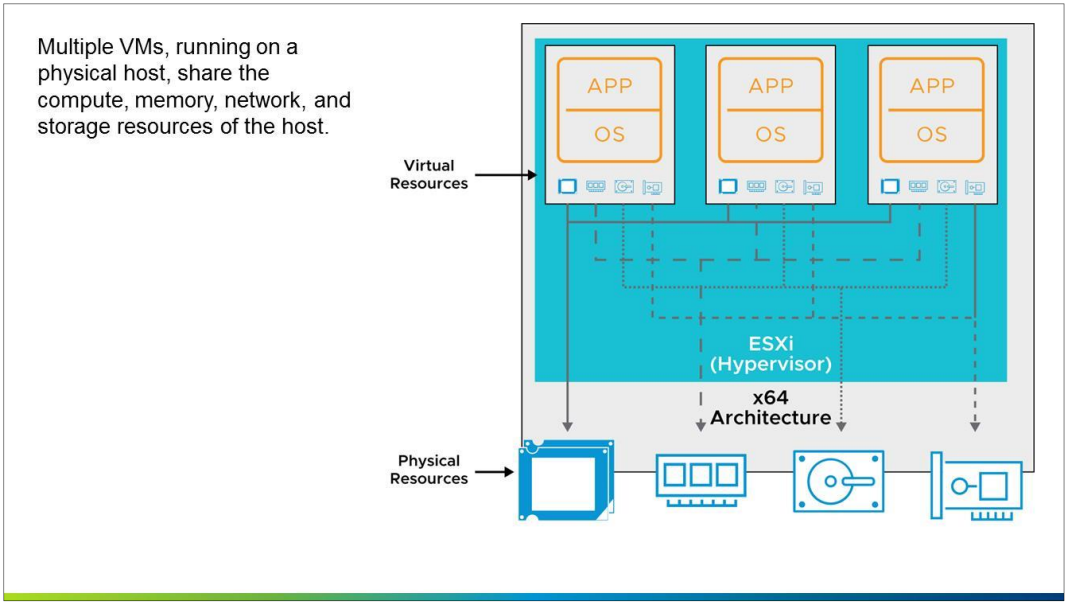

With virtualization, you can run multiple VMs on a single physical host, with each VM sharing the resources of one physical computer across multiple environments. VMs share access to CPUs and are scheduled to run by the hypervisor.

In addition, VMs are assigned their own region of memory to use and share access to the physical network cards and disk controllers. Different VMs can run different operating systems and applications on the same physical computer.

When multiple VMs run on an ESXi host, each VM is allocated a portion of the physical resources. The hypervisor schedules VMs like a traditional operating system allocates memory and schedules applications. These VMs run on various CPUs. The ESXi hypervisor can also overcommit memory. Memory is overcommitted when your VMs can use more virtual RAM than the physical RAM that is available on the host

VMs, like applications, use network and disk bandwidth. However, VMs are managed with elaborate control mechanisms to manage how much access is available for each VM. With the

https://t.me/learningnets

default resource allocation settings, all VMs associated with the same ESXi host receive an equal share of available resources.

https://t.me/learningnets

2-22 CPU Virtualization

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

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

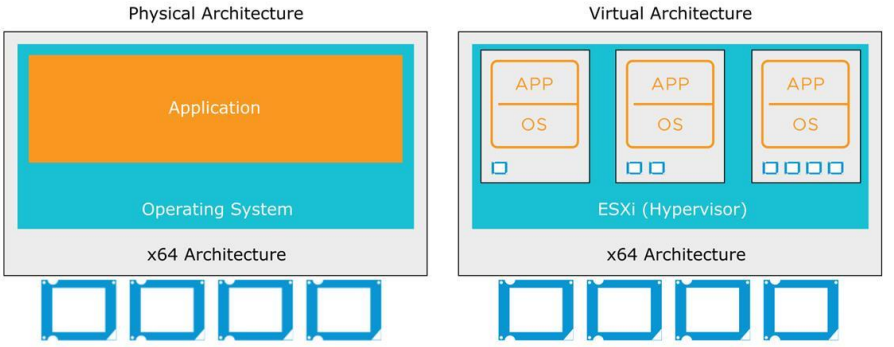

The virtualization layer runs instructions only when needed to make VMs operate as if they were running directly on a physical machine. CPU virtualization is not emulation. With a software emulator, programs can run on a computer system other than the one for which they were originally written.

Emulation provides portability but might negatively affect performance. CPU virtualization is not emulation because the supported guest operating systems are designed for x64 processors. Using the hypervisor the operating systems can run natively on the hosts' physical x64 processors.

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

https://t.me/learningnets

2-23 Physical and Virtualized Host Memory Usage

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

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

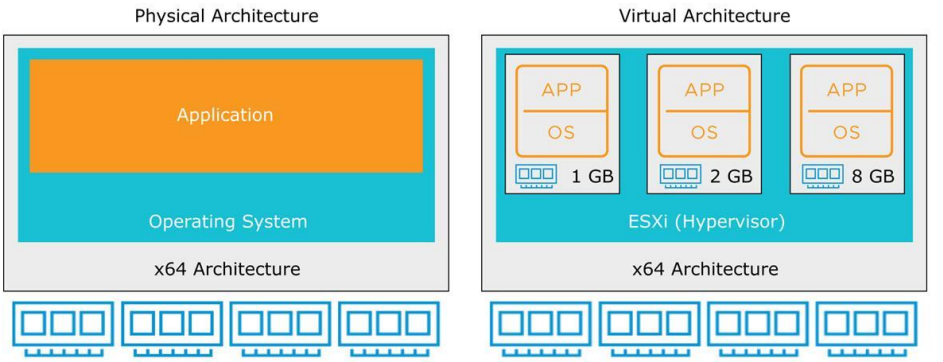

When an application starts, it uses the interfaces provided by the operating system to allocate or release virtual memory pages during the execution. Virtual memory is a decades-old technique used in most general-purpose operating systems. Operating systems use virtual memory to present more memory to applications than they physically have access to. Almost all modern processors have hardware to support virtual memory.

Virtual memory creates a uniform virtual address space for applications. With the operating system and hardware, virtual memory can handle the address translation between the virtual address space and the physical address space. This technique adapts the execution environment to support large address spaces, process protection, file mapping, and swapping in modern computer systems.

In a virtualized environment, the VMware virtualization layer creates a contiguous addressable memory space for the VM when it is started. The allocated memory space is configured when the VM is created and has the same properties as the virtual address space. With this configuration, the hypervisor can run multiple VMs simultaneously while protecting the memory of each VM from being accessed by others.

https://t.me/learningnets

2-24 Physical and Virtual Networking

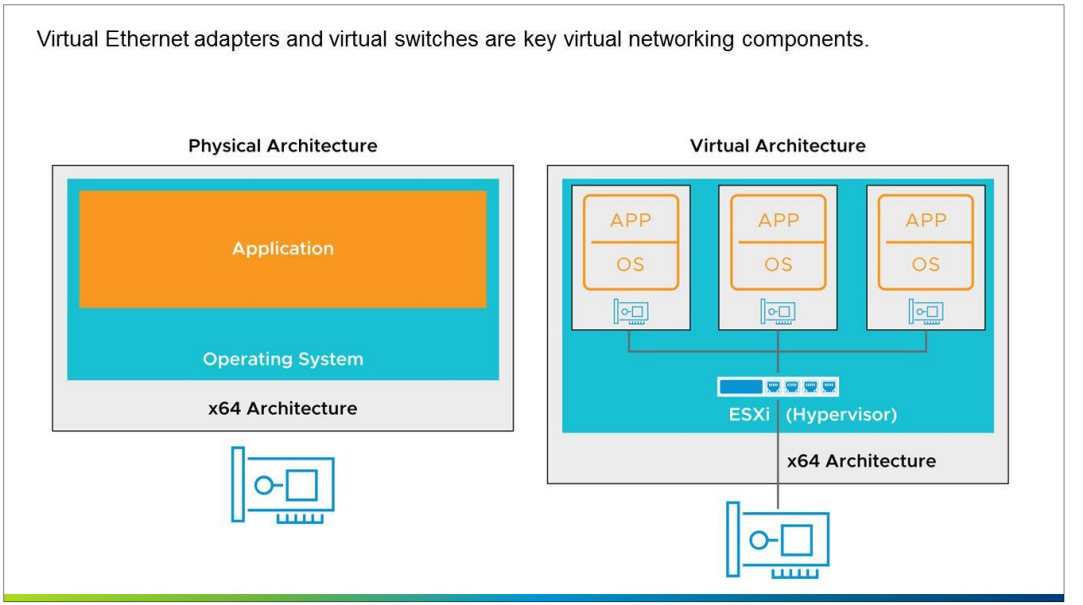

A VM can be configured with one or more virtual Ethernet adapters. VMs use virtual switches on the same ESXi host to communicate with one another by using the same protocols that are used over physical switches, without the need for additional hardware.

Virtual switches also support VLANs that are compatible with standard VLAN implementations from other networking equipment vendors. With VMware virtual networking, you can link local VMs together and link local VMs to the external network through a virtual switch.

A virtual switch, like a physical Ethernet switch, forwards frames at the data link layer. An ESXi host might contain multiple virtual switches. The virtual switch connects to the external network through outbound Ethernet adapters, called vmnics. The virtual switch can bind multiple vmnics together, like NIC teaming on a traditional server, offering greater availability and bandwidth to the VMs using the virtual switch.

Virtual switches are similar to modern physical Ethernet switches in many ways. Like a physical switch, each virtual switch is isolated and has its own forwarding table. So every destination that

https://t.me/learningnets

the switch looks up can match only ports on the same virtual switch where the frame originated. This feature improves security, making it difficult for hackers to break virtual switch isolation.

Virtual switches also support VLAN segmentation at the port level, so that each port can be configured as an access or trunk port, providing access to either single or multiple VLANs.

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

https://t.me/learningnets

2-25 Physical File Systems and Datastores

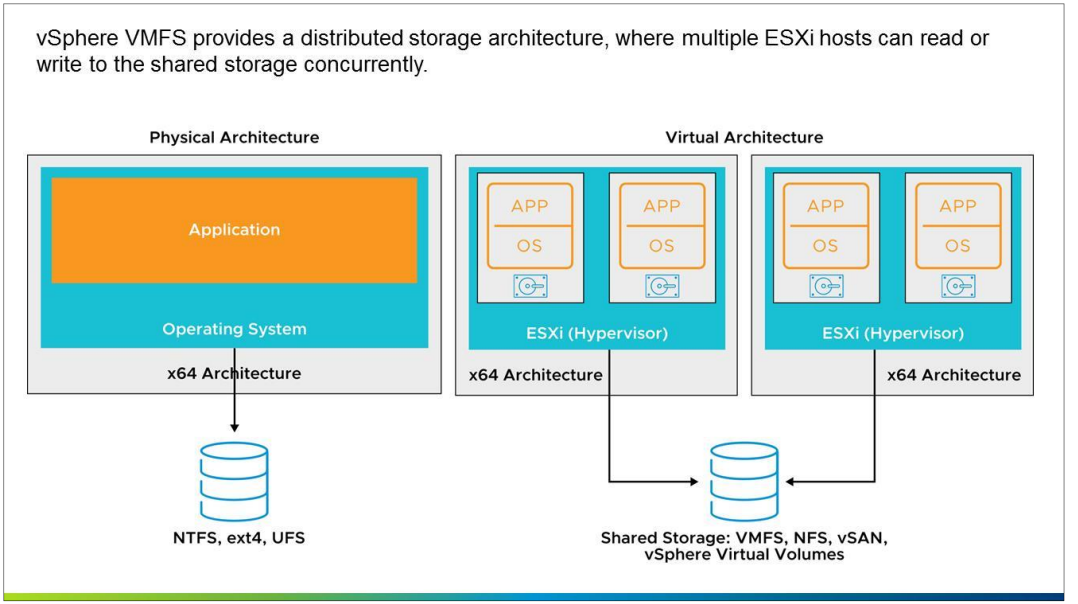

To store virtual disks, ESXi uses datastores, which are logical containers that hide the specifics of physical storage from VMs and provide a uniform model for storing VM files. Datastores that you deploy on block storage devices use the VMFS format, a special high-performance file system format that is optimized for storing virtual machines.

VMFS is designed, constructed, and optimized for a virtualized environment. It is a high-performance cluster file system designed for virtual machines. It functions in the following ways:

- Uses distributed journaling of its file system metadata changes for fast and resilient recovery if a hardware failure occurs
- Increases resource usage by providing multiple VMs with shared access to a consolidated pool of clustered storage
- Is the foundation of distributed infrastructure services, such as live migration of VMs and VM files, dynamically balanced workloads across available compute resources, automated restart of VMs, and fault tolerance

https://t.me/learningnets

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

With the distributed locking methods, VMFS forges the link between the VM and the underlying storage resources. VMs can use the unique capabilities of VMFS to join a cluster seamlessly, with no management overhead.

https://t.me/learningnets

2-26 GPU Virtualization

GPU graphics devices optimize complex graphics operations. These operations can run at high performance without overloading the CPU.
Virtual GPUs can be added to VMs for the following use cases:

- Rich 2D and 3D graphics
- VMware Horizon virtual desktops
- Graphics-intensive applications, such as those used by architects and engineers
- Server applications for massively parallel tasks, such as scientific computation applications

You can configure VMs with up to four vGPU devices to cover use cases requiring multiple GPU accelerators.

VMware supports AMD and NVIDIA graphics cards.

GPUs can be used by developers of server applications. Although servers do not usually have monitors, GPU support is important and relevant to server virtualization.

https://t.me/learningnets

2-27 Review of Learner Objectives

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

- Explain how vSphere interacts with CPUs, memory, networks, and storage

https://t.me/learningnets

https://t.me/learningnets

2-29 Learner Objectives

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

- Recognize the user interfaces for accessing the vCenter Server system and ESXi hosts

https://t.me/learningnets

2-30 vSphere User Interfaces

You can use the vSphere Client, PowerCLI, VMware Host Client, and ESXCLI to interact with the vSphere environment.

For information on ports and protocols, see http://ports.vmware.com.

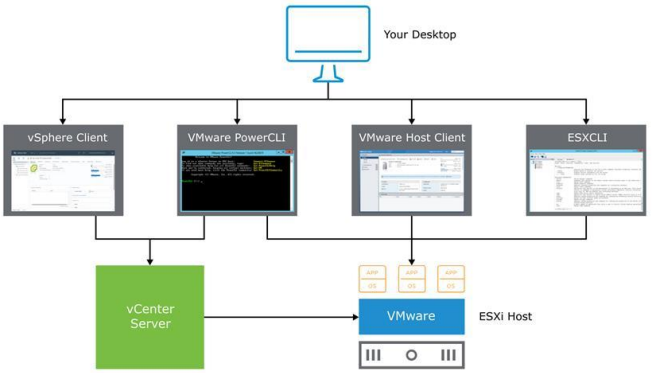

VMware Host Client provides direct management of individual ESXi hosts. VMware Host Client is generally used only when management through vCenter Server is not possible.

With the vSphere Client, an HTML5-based client, you can manage vCenter Server Appliance and the vCenter Server object inventory.

VMware Host Client and the vSphere Client provide the following benefits:

- Clean, modern UI
- No browser plug-ins to install or manage
- Integrated into vCenter Server and ESXi

https://t.me/learningnets

2-31 About VMware Host Client

VMware Host Client is an HTML5-based user interface that you can use to manage individual ESXi hosts directly when vCenter Server is unavailable.

VMware Host Client is served from ESXi, and you access it from a supported browser at https://ESXi_FQDN_or_IP_Address/ui.

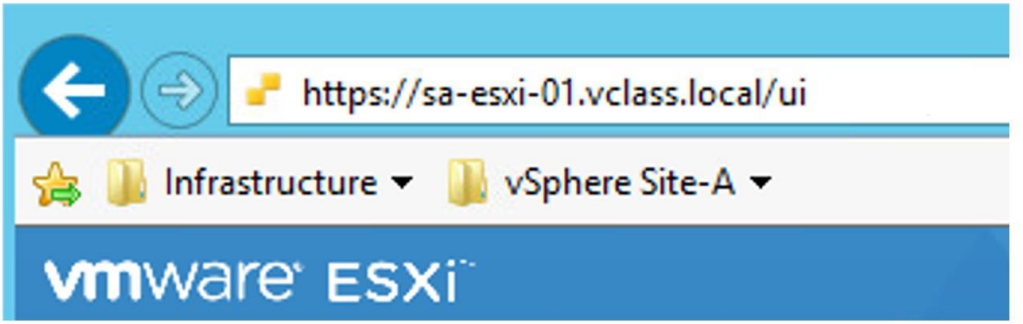

VMware ESXi in the upper-left corner of the banner on the VMware Host Client interface helps you to differentiate VMware Host Client from other clients.

https://t.me/learningnets

2-32 About vSphere Client

The vSphere Client is an HTML5-based client. You manage the vSphere environment with the vSphere Client by connecting to vCenter Server Appliance.

You access the vSphere Client from a supported browser at https://vCenter_Server_Appliance_FQDN_or_IP_Address/ui.

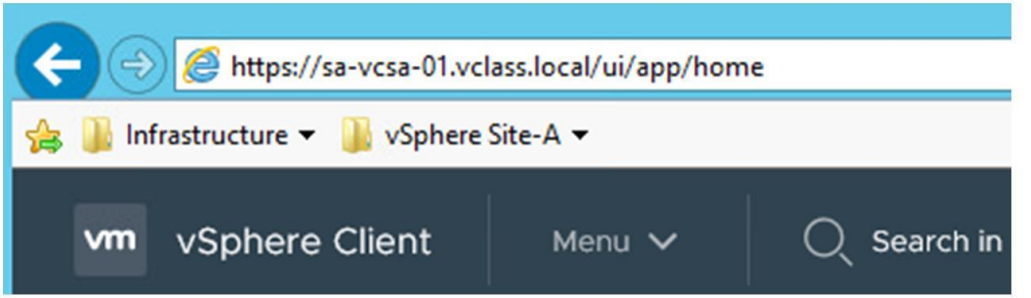

vSphere Client, which in the upper-left corner of the banner on the vSphere Client interface, helps you differentiate vSphere Client from other clients.

When you use https://vCenter_Server_Appliance_FQDN_or_IP_Address/ui to access the vSphere Client, the URL internally redirects to port 9443 on your vCenter Server system.

With the vSphere Client, you can manage vCenter Server Appliance through a web browser, and Adobe Flex does not have to be enabled in the browser.

https://t.me/learningnets

2-33 About PowerCLI and ESXCLI

PowerCLI is a command-line and scripting tool that is built on Windows PowerShell:

- Provides a PowerShell interface to vSphere API
- Provides more than 700 cmdlets for managing and automating vSphere

The ESXCLI tool allows for remote management of ESXi hosts by using the ESXCLI command set:

- ESXCLI can be downloaded from the VMware {code} page at https://code.vmware.com/web/tool/7.0/esxcli.
- ESXCLI commands can be run against a vCenter Server system and target any ESXi system.

You can install ESXCLI on a Windows or Linux system. You can run ESXCLI commands from the Windows or Linux system to manage ESXi systems.

For more information about ESXCLI, see https://code.vmware.com/web/tool/7.0/esxcli.

For more information about PowerCLI, see https://code.vmware.com/web/tool/12.0.0/vmware-powercli.

https://t.me/learningnets

2-34 Lab 1: Accessing the Lab Environment

Log in to the student desktop and access the vSphere Client and VMware Host Client:

1. Access the Student Desktop
2. Log In to an ESXi Host with VMware Host Client
3. Log In to vCenter Server with the vSphere Client

https://t.me/learningnets

2-35 Review of Learner Objectives

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

- Recognize the user interfaces for accessing the vCenter Server system and ESXi hosts

https://t.me/learningnets

https://t.me/learningnets

2-37 Learner Objectives

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

- Describe the ESXi host architecture
- Navigate the Direct Console User Interface (DCUI) to configure an ESXi host
- Recognize user account best practices
- Install an ESXi host
- Configure ESXi host settings

https://t.me/learningnets

2-38 About ESXi

ESXi is a hypervisor that you can buy with vSphere or get in a free, downloadable version. ESXi has the following features:

- High security:
  - Host-based firewall
  - Memory hardening
  - Kernel module integrity
  - Trusted Platform Module (TPM 2.0)
  - UEFI secure boot
  - Encrypted core dumps
- Small disk footprint
- Quick boot for faster patching and upgrades
- Installable on hard disks, SAN LUNs, SSD, USB devices, SD cards, SATADOM, and diskless hosts

To ensure that your physical servers are supported by ESXi 7.0, check *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

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

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

- Host-based firewall: To minimize the risk of an attack through the management interface, ESXi includes a firewall between the management interface and the network.
- Memory hardening: The ESXi kernel, user-mode applications, and executable components, such as drivers and libraries, are located at random, nonpredictable memory addresses. Combined with the nonexecutable memory protections made available by microprocessors, memory hardening provides protection that makes it difficult for malicious code to use memory exploits to take advantage of vulnerabilities.

https://t.me/learningnets

- Kernel module integrity: Digital signing ensures the integrity and authenticity of modules, drivers, and applications as they are loaded by the VMkernel.
- Trusted Platform Module: TPM is a hardware element that creates a trusted platform. This element affirms that the boot process and all drivers loaded are genuine.
- UEFI secure boot: This feature is for systems that support UEFI secure boot firmware, which contains a digital certificate that the VMware infrastructure bundles (VIBs) chain to. At boot time, a verifier is started before other processes to check the VIB's chain to the certificate in the firmware.
- Lockdown modes: This vSphere feature disables login and API functions from being executed directly on an ESXi host.
- ESXi Quick Boot: With this feature, ESXi can reboot without reinitializing the physical server BIOS. Quick Boot reduces remediation time during host patch or host upgrade operations. Quick Boot is enabled by default on supported hardware.

https://t.me/learningnets

2-39 Configuring an ESXi Host

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

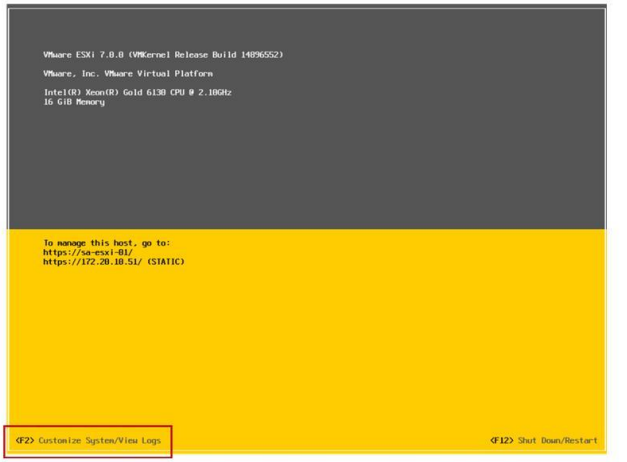

You use the Direct Console User Interface (DCUI) to configure certain settings for ESXi hosts. The DCUI is a low-level configuration and management interface, accessible through the console of the server, that is used primarily for initial basic configuration. You press F2 to start customizing system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

https://t.me/learningnets

2-40 Configuring an ESXi Host: Root Access

Administrators use the DCUI to configure root access settings:

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

The administrative user name for the ESXi host is root. The root password must be configured during the ESXi installation process.

https://t.me/learningnets

2-41 Configuring an ESXi Host: Management Network

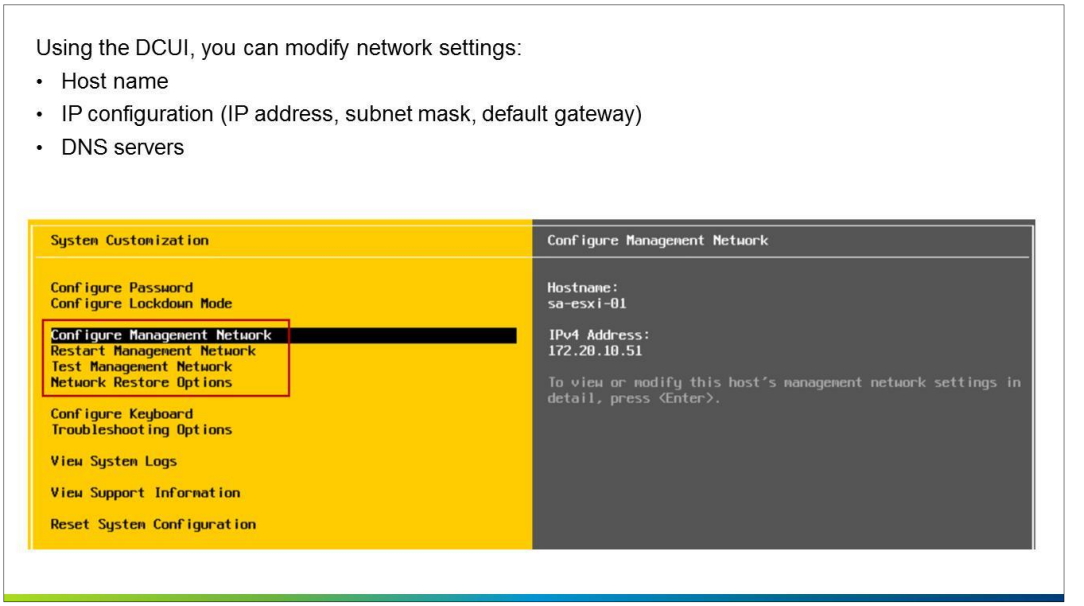

You must set up your IP address before your ESXi host is operational. By default, a DHCP-assigned address is configured for the ESXi host. To change or configure basic network settings, you use the DCUI.

In addition to changing IP settings, you perform the following tasks from the DCUI:

- Configure VLAN settings.
- Configure IPv6 addressing.
- Set custom DNS suffixes.
- Restart the management network (without rebooting the system).
- Test the management network (using ping and DNS requests).
- Disable a management network.

https://t.me/learningnets

2-42 Configuring an ESXi Host: Other Settings

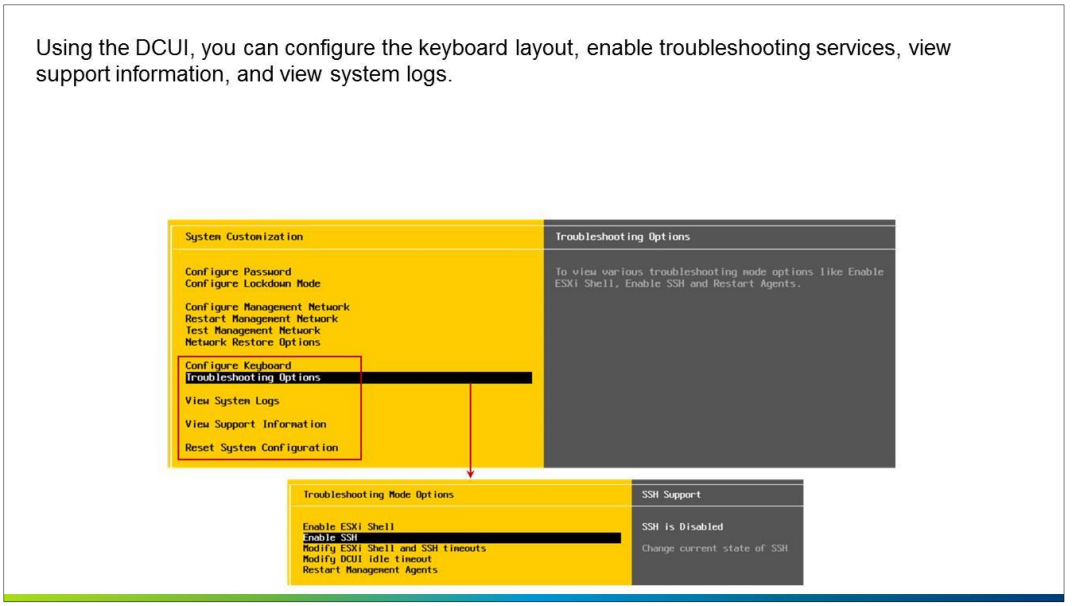

From the DCUI, you can change the keyboard layout, view support information, such as the host's license serial number, and view system logs. The default keyboard layout is U.S. English.

You can use the troubleshooting options, which are disabled by default, to enable or disable troubleshooting services:

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

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

By selecting the **Reset System Configuration** option, you can reset the system configuration to its software defaults and remove custom extensions or packages that you added to the host.

https://t.me/learningnets

2-43 Controlling Remote Access to an ESXi Host

You can use the vSphere Client to customize essential security settings that control remote access to an ESXi host:

- The ESXi firewall is enabled by default:
  - The firewall blocks incoming and outgoing traffic, except for the traffic that is enabled in the host's firewall settings.
- Services, such as the NTP client and the SSH client, can be managed by the administrator.
- Lockdown mode prevents remote users from logging in to the host directly. The host is accessible only through the DCUI or vCenter Server.

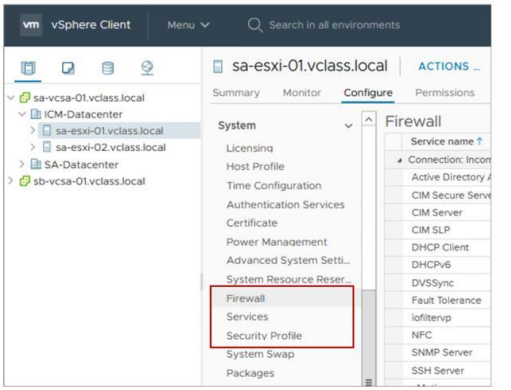

An ESXi host includes a firewall as part of the default installation. On ESXi hosts, remote clients are typically prevented from accessing services on the host. Similarly, local clients are typically prevented from accessing services on remote hosts.

To ensure the integrity of the host, few ports are open by default. To provide or prevent access to certain services or clients, you must modify the properties of the firewall.

You can configure firewall settings for incoming and outgoing connections for a service or a management agent. For some services, you can manage service details.

For example, you can use the **Start**, **Stop**, or **Restart** buttons to change the status of a service temporarily. Alternatively, you can change the startup policy so that the service starts with the host or with port use. For some services, you can explicitly specify IP addresses from which connections are allowed.

https://t.me/learningnets

2-44 Managing User Accounts: Best Practices

When assigning user accounts to access ESXi hosts or vCenter Server systems, ensure that you follow these security guidelines:

- Strictly control root privileges to ESXi hosts.
- Create strong root account passwords that have at least eight characters. Use special characters, case changes, and numbers. Change passwords periodically.
- Manage ESXi hosts centrally through the vCenter Server system by using the appropriate vSphere client.
- Minimize the use of local users on ESXi hosts:
  - Add the ESXi hosts to Active Directory and add the relevant administrator users to the ESX Admins domain group. Users in the ESX Admins domain group have root privileges on ESXi hosts, by default.
  - If local users are created, manage them centrally using the `esxcli` command in the vSphere CLI.

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

Whenever possible, use the vSphere Client to log in to the vCenter Server system and manage your ESXi hosts. In some unusual circumstances, for example, when the vCenter Server system is down, you use VMware Host Client to connect directly to the ESXi host.

Although you can log in to your ESXi host through the vSphere CLI or through vSphere ESXi Shell, these access methods should be reserved for troubleshooting or configuration that cannot be accomplished by using VMware Host Client.

If a host must be managed directly, avoid creating local users on the host. If possible, join the host to a Windows domain and log in with domain credentials instead.

https://t.me/learningnets

2-45 ESXi Host as an NTP Client

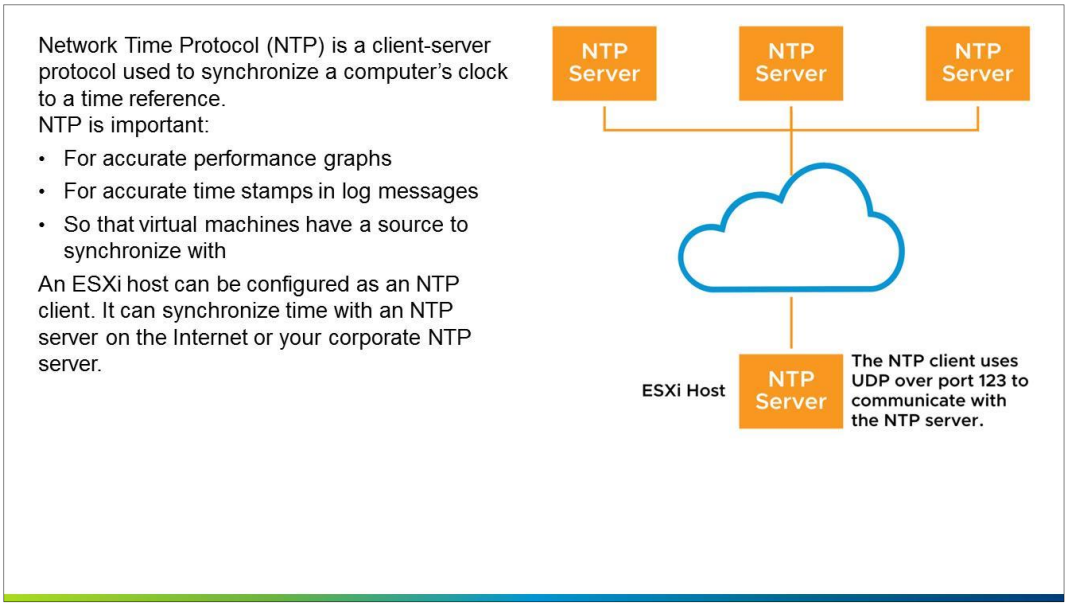

Network Time Protocol (NTP) is an Internet standard protocol that is used to synchronize computer clock times in a network. The benefits of synchronizing an ESXi host's time include:

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

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

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

For more information about timekeeping, see VMware knowledge base article 1318 at http://kb.vmware.com/kb/1318.

https://t.me/learningnets

2-46 Demonstration: Installing and Configuring ESXi Hosts

Your instructor will run a demonstration.

https://t.me/learningnets

2-47 Lab 2: Configuring an ESXi Host

Use VMware Host Client to configure an ESXi host:

1. Add an ESXi Host to Active Directory
2. Log In to the ESXi Host as an Active Directory User
3. Enable the SSH and vSphere ESXi Shell Services

https://t.me/learningnets

2-48 Review of Learner Objectives

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

- Describe the ESXi host architecture
- Navigate the Direct Console User Interface (DCUI) to configure an ESXi host
- Recognize user account best practices
- Install an ESXi host
- Configure ESXi host settings

https://t.me/learningnets

2-49 Virtual Beans: Data Center

As a Virtual Beans administrator, you now understand essential vSphere terminology. Your initial takeaways about vSphere are as follows:

- vSphere is the starting point for building a software-defined data center.
- ESXi hosts are highly secure platforms on which Virtual Beans applications run.
- Check the VMware Compatibility Guide to ensure that your physical servers support ESXi 7.0.

https://t.me/learningnets

2-50 Key Points

- Virtual machines are hardware independent.
- VMs share the physical resources of the ESXi host on which they reside.
- vSphere abstracts CPU, memory, storage, and networking for VM use.
- The ESXi hypervisor runs directly on the host.

Questions?

https://t.me/learningnets

Module 3

Virtual Machines

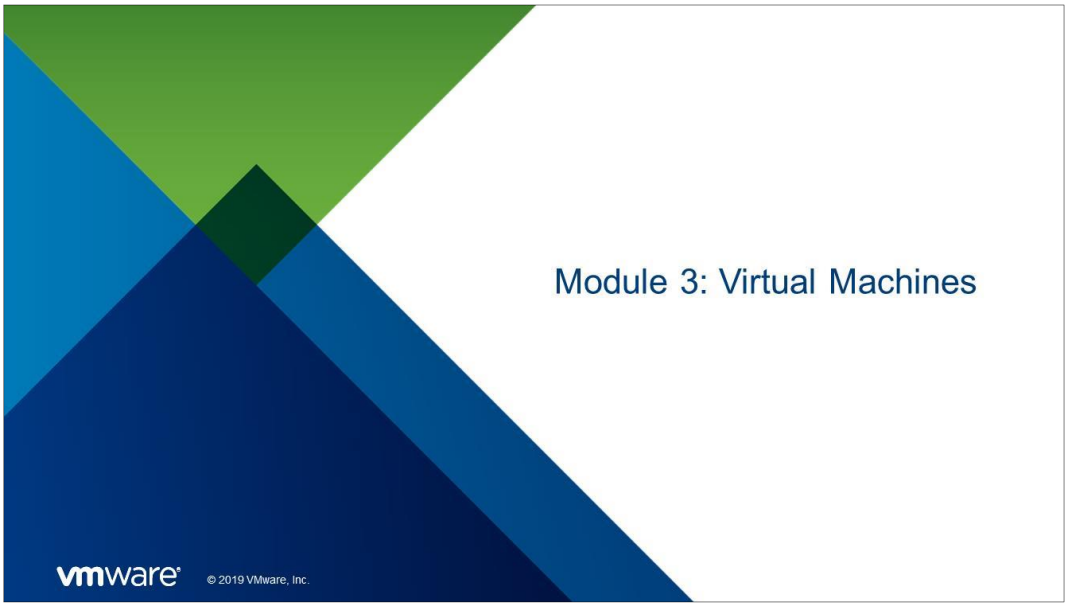

https://t.me/learningnets

3-2 Importance

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

https://t.me/learningnets

3-3 Module Lessons

1. Creating Virtual Machines
2. Virtual Machine Hardware Deep Dive
3. Introduction to Containers

https://t.me/learningnets

3-4 Virtual Beans: Virtualizing Workloads

Virtual Beans uses internally developed applications that run in an environment with Windows and Linux systems.
Virtual Beans requires that 100% of their workloads be virtualized:

- Business-critical applications
- Nonbusiness-critical applications

In addition, Virtual Beans application developers are creating and testing a new order-fulfillment system based on container technology.

As a Virtual Beans administrator, you must familiarize yourself with the components of a virtual machine and the virtual devices that are supported. You also want to learn about containers because future applications will use this technology.

https://t.me/learningnets

3-5 Lesson 1: Creating Virtual Machines

https://t.me/learningnets

3-6 Learner Objectives

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

- Create and provision a virtual machine
- Describe how to import a virtual appliance OVF template
- Explain the importance of VMware Tools
- Install VMware Tools

https://t.me/learningnets

3-7 About Provisioning Virtual Machines

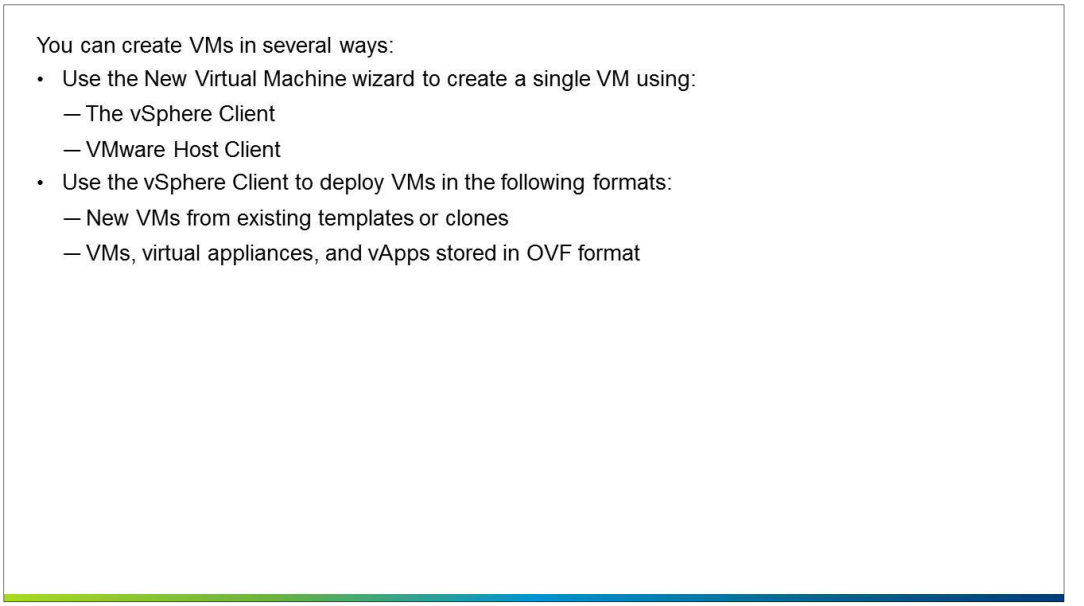

The optimal method for provisioning VMs for your environment depends on factors such as the size and type of your infrastructure and the goals that you want to achieve.

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

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

In addition to using the vSphere Client, you can also use VMware Host Client to create a VM by using OVF files. However, several limitations apply when you use VMware Host Client for this

https://t.me/learningnets

deployment method. For information about OVF and OVA limitations for the VMware Host Client, see *vSphere Single Host Management - VMware Host Client* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.hostclient.doc/GUID-509C12B2-32F2-4928-B81B-DE87C7B2A5F6.html.

https://t.me/learningnets

3-8 Creating VMs with the New Virtual Machine Wizard (1)

You can use the New Virtual Machine wizard in the vSphere Client to create a VM.

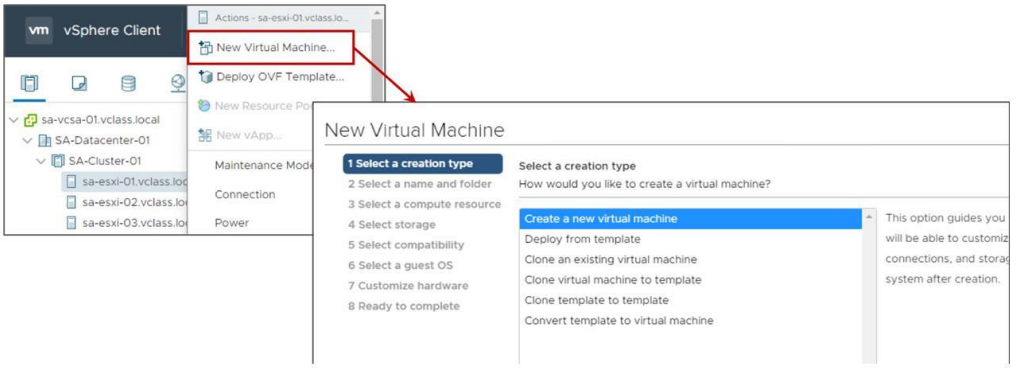

https://t.me/learningnets

3-9 Creating VMs with the New Virtual Machine Wizard (2)

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

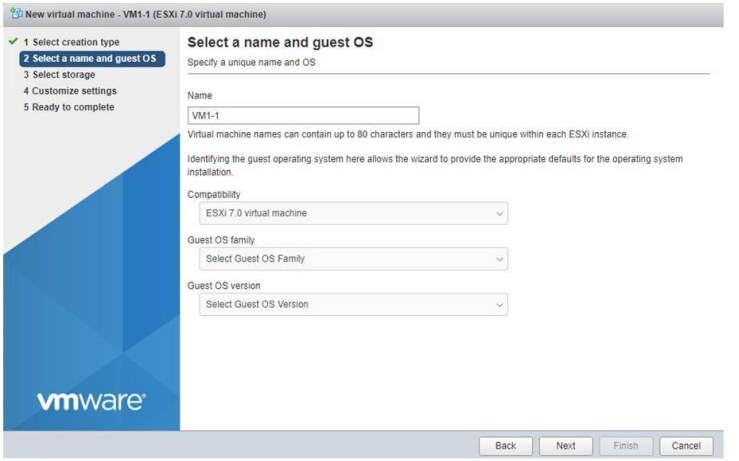

The New Virtual Machine wizard prompts you for standard information:

- The VM name

  If using the vSphere Client, you can also specify the folder in which to place the VM.

- The resource on which the VM runs

  If using VMware Host Client, you create the VM on the host that you are logged in to.

  If using the vSphere Client, you can specify a host, a cluster, a vApp, or a resource pool. The VM can access the resources of the selected object.

- The datastore on which to store the VM's files

  Each datastore might have a different size, speed, availability, and other properties. The available datastores are accessible from the destination resource that you select.

- The guest operating system to be installed into the VM

https://t.me/learningnets

- The number of NICs, the network to connect to, and the network adapter type
- Virtual disk provisioning choice

https://t.me/learningnets

3-10 New Virtual Machine Wizard Settings

VM configuration settings are based on prior choices that you made about the operating system.

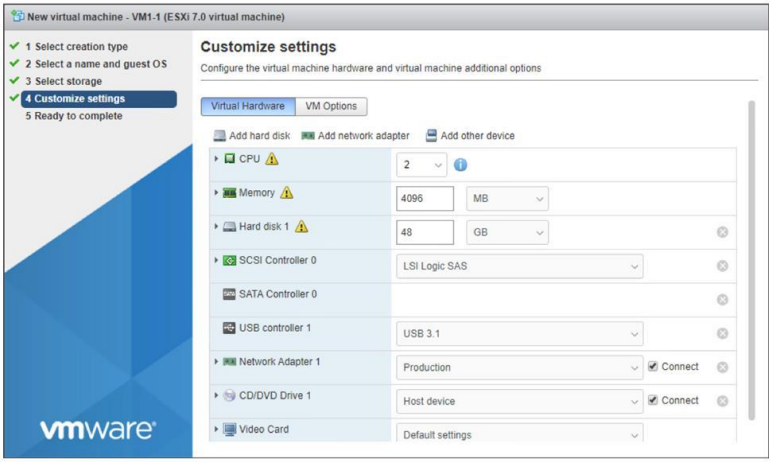

https://t.me/learningnets

3-11 Installing the Guest Operating System

Installing a guest operating system in your VM is similar to installing it on a physical computer.

To install the guest operating system, you interact with the VM through the VM console. Using the vSphere Client, you can attach a CD, DVD, or ISO image containing the installation image to the virtual CD/DVD drive.

On the slide, the Windows Server 2008 guest operating system is being installed. You can use the vSphere Client to install a guest operating system. You can also install a guest operating system from an ISO image or a CD. Installing from an ISO image is typically faster and more convenient than a CD installation.

For more information about installing guest operating systems, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html.

For more about the supported guest operating systems, see *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

https://t.me/learningnets

3-12 Deploying OVF Templates

You can deploy any VM or virtual appliance stored in OVF format. Virtual appliances are preconfigured VMs:

- They are usually designed for a single purpose.
- They are also available from VMware Solution Exchange.

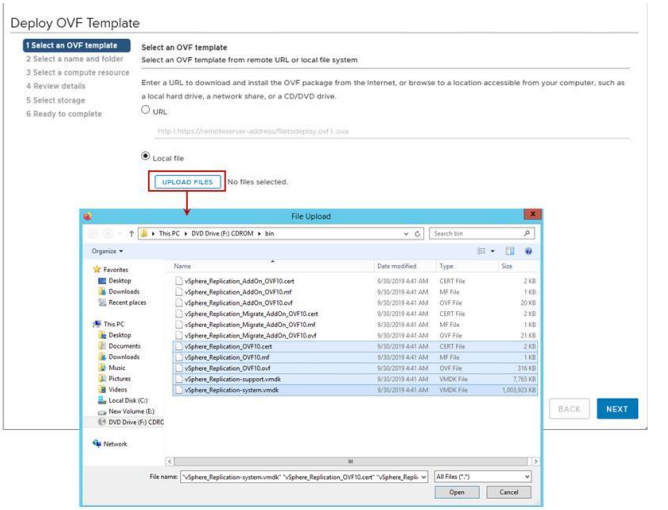

A virtual appliance is a preconfigured VM that typically includes a preinstalled guest operating system and other software. A virtual appliance is usually designed for a specific purpose, for example, to provide a secure web browser, a firewall, or a backup and recovery utility.

A virtual appliance can be added or imported to your vCenter Server system inventory or ESXi inventory. Virtual appliances can be imported from websites such as the VMware Virtual Appliance Marketplace at https://marketplace.vmware.com/vsx/.

Virtual appliances are deployed as OVF templates. OVF is a platform-independent, efficient, extensible, and open packaging and distribution format for VMs. OVF files are compressed, resulting in faster downloads. The vSphere Client validates an OVF file before importing it and ensures that it is compatible with the intended destination server. If the appliance is incompatible with the selected host, you cannot import it.

https://t.me/learningnets

3-13 About VMware Tools

VMware Tools is a set of features that enhance the performance of a VM's guest operating system. Benefits and features include:

- Device drivers
  - SVGA display
  - VMXNET/VMXNET3
  - Balloon driver for memory management
  - Sync driver for quiescing I/O
- Increased graphics performance
- Improved mouse performance
- Guest OS heartbeat service
- Time synchronization
- Ability to shut down the VM

VMware Tools improves management of the VM by replacing generic operating system drivers with VMware drivers tuned for virtual hardware. You install VMware Tools into the guest operating system. When you install VMware Tools, you install these items:

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

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

- Faster graphics performance and Windows Aero on operating systems that support Aero
- Shared folders between host and guest file systems

https://t.me/learningnets

- Copying and pasting text, graphics, and files between the virtual machine and the host or client desktop
- Scripting that helps automate guest operating system operations

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

https://t.me/learningnets

3-14 Installing VMware Tools

Ensure that you select the correct version of VMware Tools for your guest operating system.

To find out which VMware Tools ISO images are bundled with vSphere 7, see the vSphere 7 Release Notes.

The method for installing VMware Tools depends on the guest operating system type.

Guest Operating System Type	VMware Tools Installation Method
Microsoft Windows	Install from `windows.iso` for Vista and later guests.
Linux	To allow vSphere to manage the installation, install from `linux.iso`.
	To allow the guest OS to manage the installation, use `open-vm-tools`, available in various Linux package management systems, such as `yum`, `apt`, or `rpm`.
MacOS	Install from `darwin.iso` for Mac OS X versions 10.11 and later.
Solaris	Install from `solaris.iso` for Solaris versions 10 and later.

For more information about using Open VM tools, see *VMware Tools User Guide* at https://docs.vmware.com/en/VMware-Tools/index.html.

https://t.me/learningnets

3-15 Downloading VMware Tools

You can download a specific version of VMware Tools from the VMware vSphere product download page.

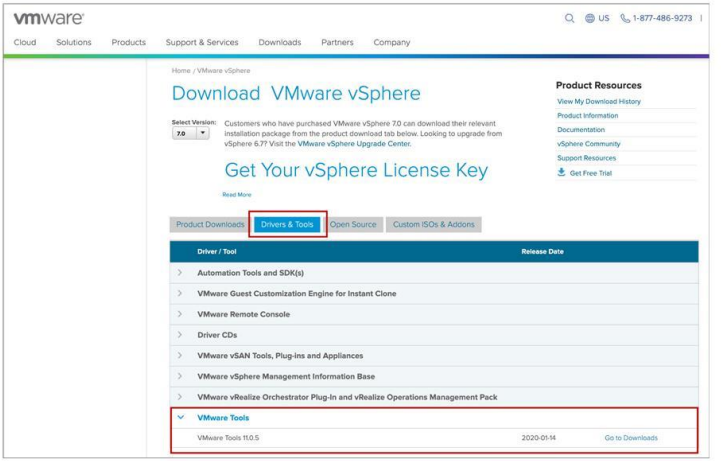

https://t.me/learningnets

3-16 Labs

Lab: Creating a Virtual Machine
Lab: Installing VMware Tools

https://t.me/learningnets

3-17 Lab 3: Creating a Virtual Machine

Use VMware Host Client to create and delete a virtual machine:

1. Create a Virtual Machine
2. Delete the Virtual Machine

https://t.me/learningnets

3-18 Lab 4: Installing VMware Tools

Use VMware Host Client to install VMware Tools into an existing Windows VM:

1. Power On and Open a Console to the VM
2. Install VMware Tools

https://t.me/learningnets

3-19 Review of Learner Objectives

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

- Create and provision a virtual machine
- Describe how to import a virtual appliance OVF template
- Explain the importance of VMware Tools
- Install VMware Tools

https://t.me/learningnets

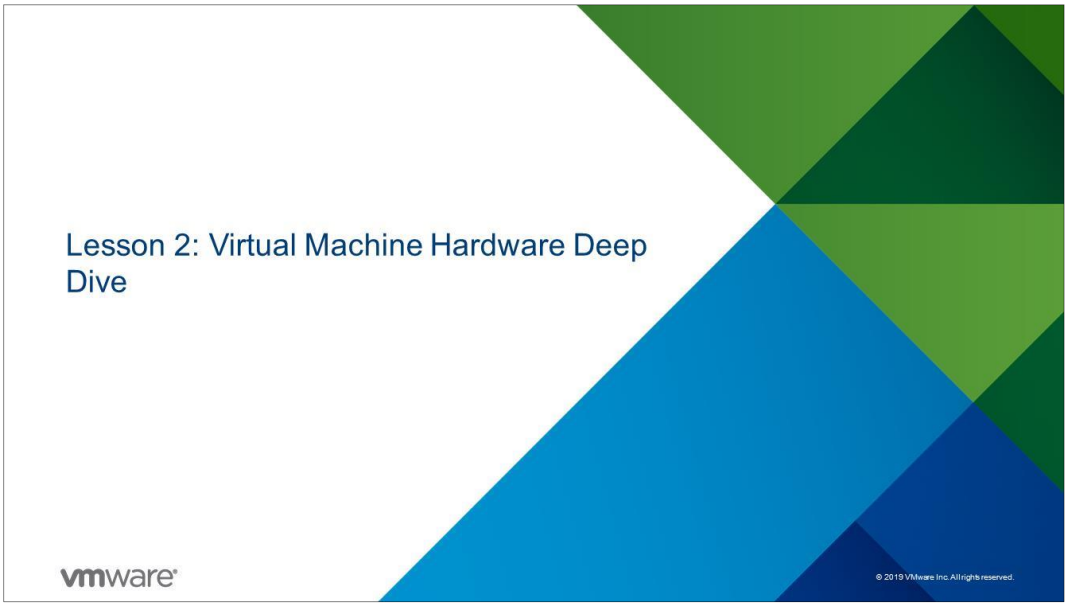

https://t.me/learningnets

3-21 Learner Objectives

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

- Identify the files that make up a VM
- Compare VM hardware versions
- Recognize the components of a VM
- Navigate the vSphere Client and examine VM settings
- Identify methods for accessing a VM console
- Identify virtual network adapters, including the enhanced VMXNET3
- Distinguish between types of virtual disk provisioning

https://t.me/learningnets

3-22 Virtual Machine Encapsulation

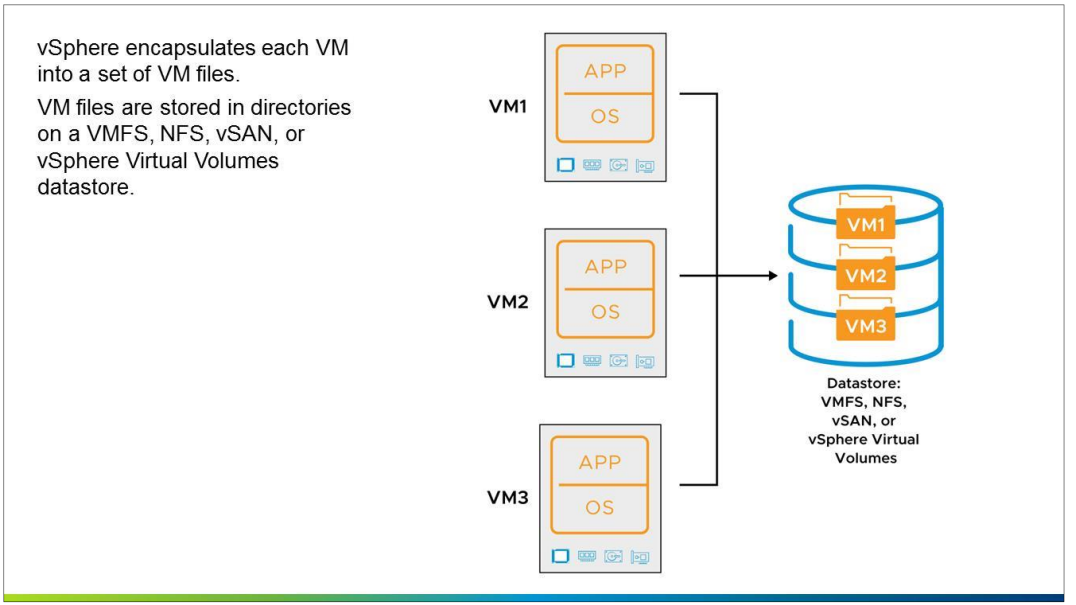

vSphere encapsulates each VM into a few files or objects, making VMs easier to manage and migrate. The files and objects for each VM are stored in a separate folder on a datastore.

https://t.me/learningnets

3-23 About Virtual Machine Files

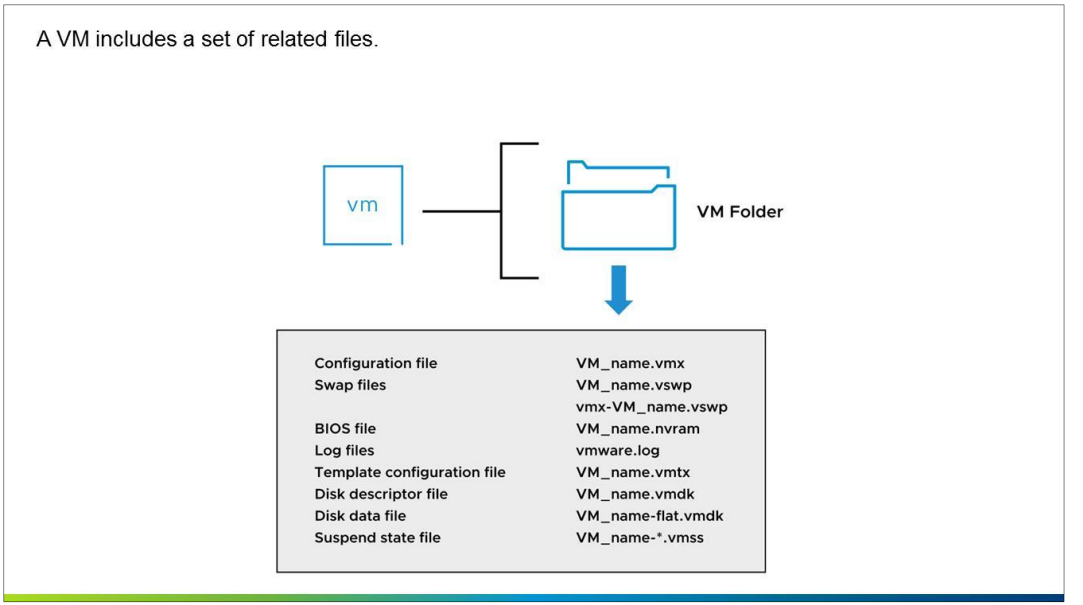

The slide lists some of the files that make up a VM. Except for the log files, the name of each file starts with the VM's name <VM_name>. A VM consists of the following files:

- A configuration file (.vmx).
- Swap files (.vswp) used to reclaim memory during periods of contention.
- A file containing the VM's BIOS settings (.nvram).
- A VM's current log file (.log) and a set of files used to archive old log entries (-#.log).

  In addition to the current log file, vmware.log, up to six archive log files are maintained at one time. For example, -1.log to -6.log might exist at first.

  The next time an archive log file is created, for example, when the VM is powered off and powered back on, the following actions occur: The -6.log is deleted, the -5.log is

https://t.me/learningnets

recalled to `-6.log`, and so on. Finally, the previous `vmware.log` is recalled to the `-1.log`.

- One or more virtual disk files. The first virtual disk has files `VM_name.vmdk` and `VM_name-flat.vmdk`.

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

- If the VM is converted to a template, a VM template configuration file (`.vmtx`) replaces the VM configuration file (`.vmx`). A VM template is a master copy of the VM.

The list of files shown on the slide is not comprehensive. For a complete list of all the types of VM files, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html.

https://t.me/learningnets

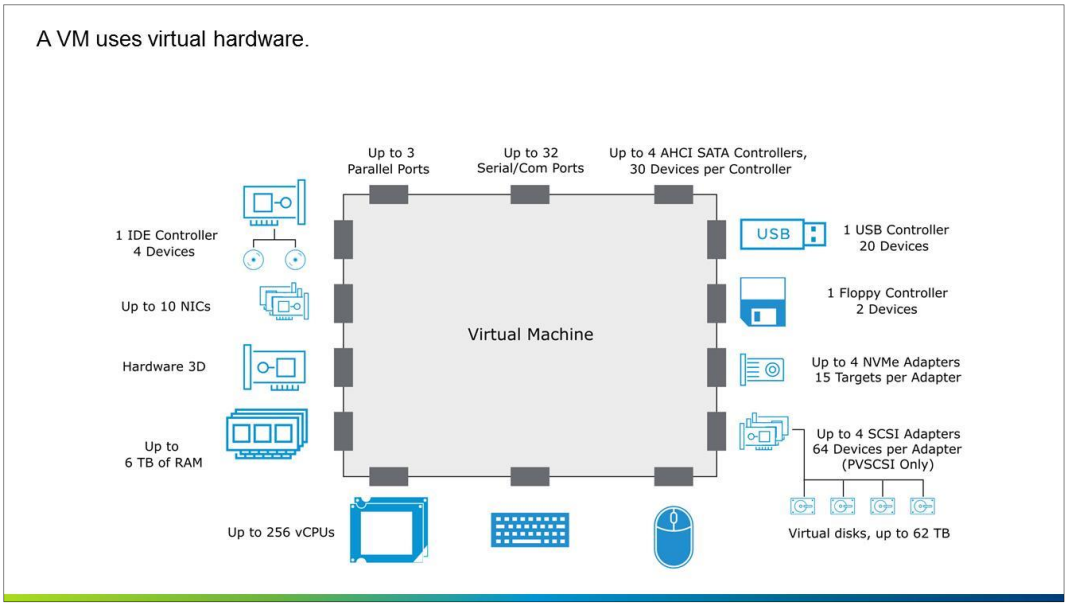

Each guest OS sees ordinary hardware devices. The guest OS does not know that these devices are virtual. All VMs have uniform hardware, except for a few variations that the system administrator can apply. Uniform hardware makes VMs portable across VMware virtualization platforms.

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

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

https://t.me/learningnets

You can add up to 16 PCI vSphere DirectPath I/O devices to a VM. The devices must be reserved for PCI passthrough on the host on which the VM runs. Snapshots are not supported with vSphere DirectPath I/O pass-through devices.

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

The Virtual Machine Communication Interface (VMCI) is an infrastructure that provides a high-speed communication channel between a VM and the hypervisor. You cannot add or remove VMCI devices.

The VMCI SDK facilitates the development of applications that use the VMCI infrastructure. Without VMCI, VMs communicate with the host using the network layer. Using the network layer adds overhead to the communication. With VMCI, communication overhead is minimal and tasks that require communication can be optimized. VMCI can go up to nearly 10 Gbit/s with 128 K sized queue pairs.

The following types of communication are available:

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

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

For more information about virtual hardware, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html.

https://t.me/learningnets

3-25 Virtual Hardware Versions

The virtual hardware version, or VM compatibility level, determines the operating system functions that a VM supports.

Do not use a later version that is not supported by the VMware product.

Compatibility	Virtual Hardware Version
ESXi 7.0	17
ESXi 6.7 U2 and later	15
ESXi 6.7 and later	14
ESXi 6.5 and later	13
ESXi 6.0 and later	11

Virtual hardware versions 12 and 16 are specific to Workstation and Fusion Pro.

Each release of a VMware product has a corresponding VM hardware version included. The table shows the latest hardware version that each ESXi version supports. Each VM compatibility level supports at least five major or minor vSphere releases.

For a complete list of virtual machine configuration maximums, see VMware Configuration Maximums at https://configmax.vmware.com.

https://t.me/learningnets

3-26 About CPU and Memory

You can add, change, or configure CPU and memory resources to improve VM performance. The maximum number of virtual CPUs (vCPUs) that you can assign to a VM depends on the following factors:

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

A VM running on an ESXi 7.0 host can have up to 256 vCPUs.

The maximum memory size of a VM depends on the VM's compatibility setting.

The maximum memory size of a VM with ESXi 7.0 compatibility running on ESXi 7.0 is 6 TB.

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

You can use the multicore vCPU feature to control the number of cores per virtual socket in a VM. With this capability, operating systems with socket restrictions can use more of the host CPU's cores, increasing overall performance.

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

You can set most of the memory parameters during VM creation or after the guest operating system is installed. Some actions require that you power off the VM before changing the settings.

The memory resource settings for a VM determine how much of the host's memory is allocated to the VM.

https://t.me/learningnets

The virtual hardware memory size determines how much memory is available to applications that run in the VM. A VM cannot benefit from more memory resources than its configured virtual hardware memory size.

ESXi hosts limit the memory resource use to the maximum amount useful for the VM so that you can accept the default of unlimited memory resources. You can reconfigure the amount of memory allocated to a VM to enhance performance. Maximum memory size for a VM depends on the VM's compatibility setting.

https://t.me/learningnets

3-27 About Virtual Storage

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

- BusLogic Parallel
- LSI Logic Parallel
- LSI Logic SAS
- VMware Paravirtual SCSI
- AHCI SATA controller
- Virtual NVMe

Storage adapters provide connectivity for your ESXi host to a specific storage unit or network.

ESXi supports different classes of adapters, including SCSI, iSCSI, RAID, Fibre Channel, Fibre Channel over Ethernet (FCoE), and Ethernet. ESXi accesses the adapters directly through device drivers in the VMkernel:

- BusLogic Parallel: The latest Mylex (BusLogic) BT/KT-958 compatible host bus adapter.
- LSI Logic Parallel: The LSI Logic LSI53C10xx Ultra320 SCSI I/O controller is supported.
- LSI Logic SAS: The LSI Logic SAS adapter has a serial interface.
- VMware Paravirtual SCSI: A high-performance storage adapter that can provide greater throughput and lower CPU use.

https://t.me/learningnets

- AHCI SATA controller: Provides access to virtual disks and CD/DVD devices. The SATA virtual controller appears to a VM as an AHCI SATA controller. AHCI SATA is available only for VMs with ESXi 5.5 and later compatibility.
- Virtual NVMe: NVMe is an Intel specification for attaching and accessing flash storage devices to the PCI Express bus. NVMe is an alternative to existing block-based server storage I/O access protocols.

https://t.me/learningnets

3-28 About Thick-Provisioned Virtual Disks

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

VM disks consume all the capacity, as defined at creation, regardless of the amount of data in the guest operating system file system. Thick-provisioned disk types are eager zeroed or lazy zeroed:

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

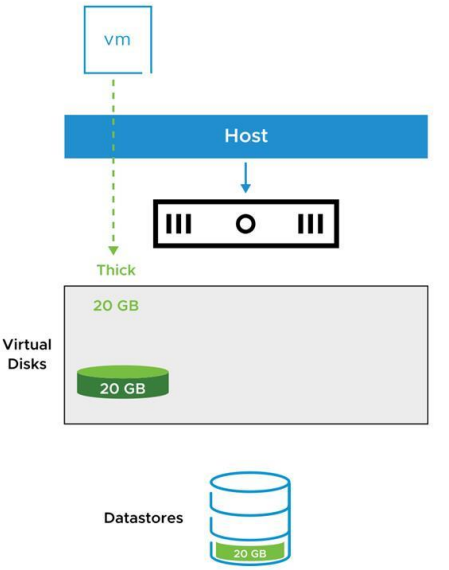

In a lazy-zeroed thick-provisioned disk, space required for the virtual disk is allocated during creation. Data remaining on the physical device is not erased during creation. Later, the data is zeroed out on demand on first write from the VM. This disk type is the default.

In an eager-zeroed thick-provisioned disk, the space required for the virtual disk is allocated during creation. Data remaining on the physical device is zeroed out when the disk is created.

https://t.me/learningnets

3-29 About Thin-Provisioned Virtual Disks

With thin provisioning, VMs use storage space as needed:

- Virtual disks consume only the capacity needed to hold the current files.
- The VM always sees the full allocated disk size.

Run the `unmap` command to reclaim unused space from the array.

Reporting and alerts help manage allocations and capacity.

You can mix thick and thin formats.
More efficient use of storage, for example:

- Provisioned space for virtual disks: 140 GB
- Available datastore capacity: 100 GB
- Used datastore capacity: 80 GB

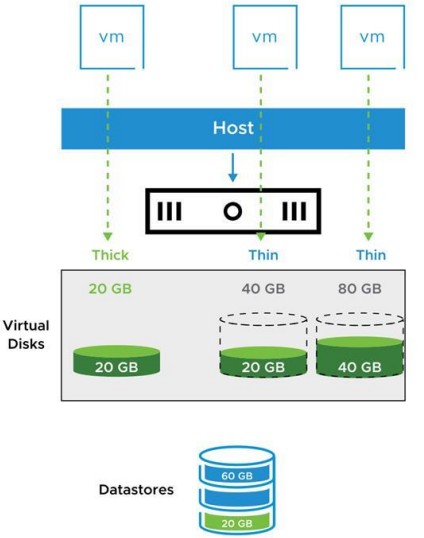

A thin-provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can expand to the maximum capacity allocated to it.

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

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

https://t.me/learningnets

3-30 Thick-Provisioned and Thin-Provisioned Disks

Virtual disk options differ in terms of creation time, block allocation, layout, and zeroing out of allocated file blocks.

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

https://t.me/learningnets

3-31 About Virtual Networks

VMs and physical machines communicate through a virtual network.
When you configure networking for a VM, you select or change the following settings:

- Network adapter type
- Port group to connect to
- Network connection state
- Whether to connect to the network when the VM powers on

For more information about virtual networks, see *vSphere Networking* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.networking.doc/GUID-35B40B0B-0C13-43B2-BC85-18C9C91BE2D4.html.

https://t.me/learningnets

3-32 About Virtual Network Adapters

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

Network Adapter Type	Description
E1000-E1000E	Emulated version of an Intel Gigabit Ethernet NIC, with drivers available in most newer guest operating systems.
VMXNET3	Available only with VMware Tools.
Flexible	Can function as either a Vlance or VMXNET adapter.
SR-IOV pass-through	Allows VM and physical adapter to exchange data without using the VMkernel as an intermediary.
vSphere DirectPath I/O	Allows VM access to physical PCI network functions on platforms with an I/O memory management unit.
PVRDMA	Paravirtualized device that provides improved virtual device performance. It provides an RDMA-like interface for vSphere guests.

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

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

https://t.me/learningnets

The following NIC types are supported:

- E1000E: Emulated version of the Intel 82574 Gigabit Ethernet NIC. E1000E is the default adapter for Windows 8 and Windows Server 2012.
- E1000: Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in most newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later.
- Flexible: Identifies itself as a Vlance adapter when a VM starts, but initializes itself and functions as either a Vlance or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlance adapter to the higher performance VMXNET adapter.
- Vlance: Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in 32-bit legacy guest operating systems. A VM configured with this network adapter can use its network immediately.
- VMXNET2 (Enhanced): Based on the VMXNET adapter but provides high-performance features commonly used on modern networks, such as jumbo frames and hardware offloads. VMXNET2 (Enhanced) is available only for some guest operating systems on ESX/ESXi 3.5 and later. It is not supported for ESXi 6.7 and later.
- VMXNET3: A paravirtualized NIC designed for performance. VMXNET3 offers all the features available in VMXNET2 and adds several new features, such as multiqueue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery.
- SR-IOV pass-through: Representation of a virtual function on a physical NIC with SR-IOV support. This adapter type is suitable for VMs that require more CPU resources or where latency might cause failure. If VMs are sensitive to network delay, SR-IOV can provide direct access to the virtual functions of supported physical NICs, bypassing the virtual switches and reducing overhead.

  SR-IOV pass-through is available in ESXi 6.0 and later for Red Hat Enterprise Linux 6 and later, and Windows Server 2008 R2 with SP2. An operating system release might contain a default virtual function driver for certain NICs. For others, you must download and install it from a location provided by the NIC or host vendor.

https://t.me/learningnets

- vSphere DirectPath I/O allows a guest operating system on a VM to directly access physical PCI and PCIe devices connected to a host. Pass-through devices help your environment use resources efficiently and improve performance. You can configure a pass-through PCI device on a VM by using the vSphere Client. VMs configured with vSphere DirectPath I/O do not have the following features:
  - Hot adding and removing of virtual devices
  - Suspend and resume
  - Record and replay
  - Fault tolerance
  - High availability
  - vSphere DRS: Limited availability

    The VM can be part of a cluster but cannot migrate across hosts.
  - Snapshots.
- With PVRDMA, multiple guests can access the RDMA device by using verbs API, an industry-standard interface. A set of these verbs was implemented to expose an RDMA-capable guest device (PVRDMA) to applications. The applications can use the PVRDMA guest driver to communicate with the underlying physical device. PVRDMA supports RDMA, providing the following functions:
  - OS bypass
  - Zero-copy
  - Low latency and high bandwidth
  - Less power use and faster data access

https://t.me/learningnets

3-33 Other Virtual Devices

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

- CD/DVD drive: For connecting to a CD, DVD, or ISO image.
- USB 3.0 and 3.1: Supported with host-connected and client-connected devices.
- Floppy drive: For connecting a VM to a floppy drive or a floppy image.
- Generic SCSI devices: A VM can be connected to additional SCSI adapters.
- vGPUs: A VM can use GPUs on the physical host for high-computation activities.

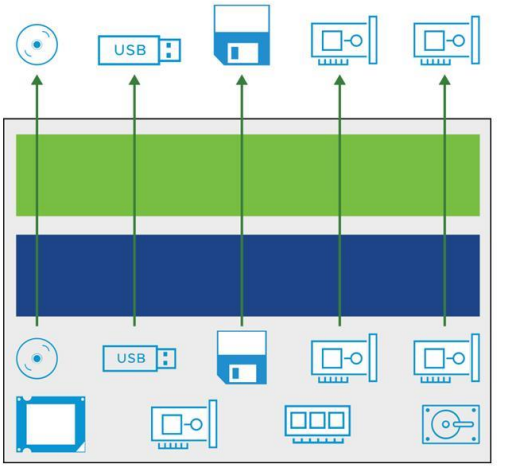

Virtual CPU (vCPU) and virtual memory are the minimum required virtual hardware. Having a virtual hard disk, virtual NICs, and other virtual devices make the VM more useful.

For information about adding virtual devices to a VM, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html.

https://t.me/learningnets

3-34 About the Virtual Machine Console

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

You can use the standalone VMware Remote Console Application (VMRC) to connect to client devices.

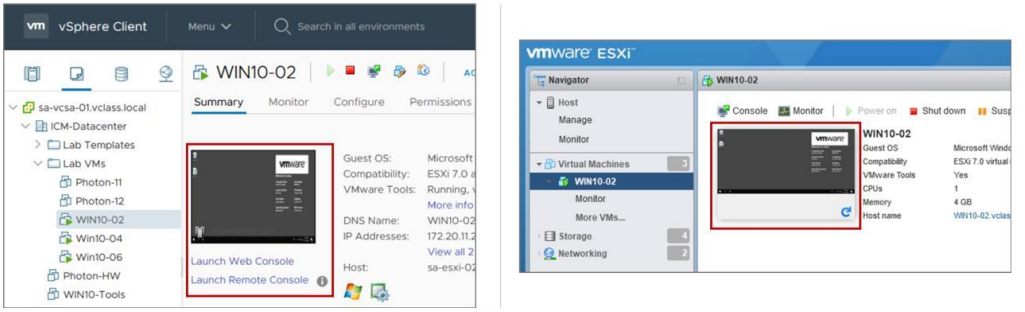

You can open the VM console from the vSphere Client.

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

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

https://t.me/learningnets

3-35 Lab 5: Adding Virtual Hardware

Use VMware Host Client to examine a virtual machine's configuration and add virtual hardware to the virtual machine:

1. Examine a Virtual Machine's Configuration
2. Add Virtual Hard Disks to the Virtual Machine
3. Compare Thin-Provisioned and Thick-Provisioned Disks

https://t.me/learningnets

3-36 Review of Learner Objectives

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

- Identify the files that make up a VM
- Compare VM hardware versions
- Recognize the components of a VM
- Navigate the vSphere Client and examine VM settings
- Identify methods for accessing a VM console
- Identify virtual network adapters, including the enhanced VMXNET3
- Distinguish between types of virtual disk provisioning

https://t.me/learningnets

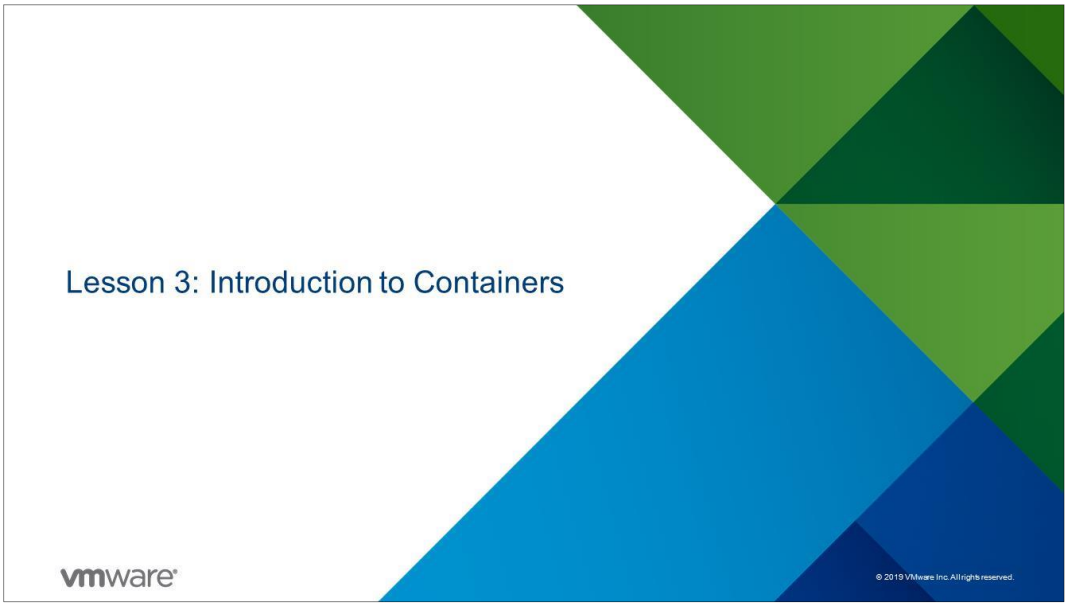

https://t.me/learningnets

3-38 Learner Objectives

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

- Describe the benefits and use cases for containers
- Identify the parts of a container system
- Differentiate between containers and virtual machines

https://t.me/learningnets

3-39 Traditional Application Development

In data centers, traditional applications are enhanced with modern application capabilities and models. But traditional application development is different from modern application development.

Traditional Application Development	
Waterfall development	Development cycles are relatively long (several months). Requirements can change over time.
Handover to the operations team	The operations team is responsible for the code in production. Training is not provided.
Monolithic applications	All functionality is handled by a single large application. Typically, monolithic applications can only be scaled vertically, not horizontally.
Separate development, staging, and production environments	Inconsistencies between environments occur. Issues in production do not occur during staging.

Waterfall development: Waterfall development cycles take from 6 to 12 months to deliver a product. Because this cycle is relatively long in the context of software development, requirements are at risk of changing. In addition, initial requirements might be misunderstood, but this misunderstanding might be realized only at the end of the project.

Handover to the operations team: When a product is ready for production, it is handed over to the operations team. The operations team deploys and manages the software from that point. Without proper training and documentation, the team can find it difficult to skill up and effectively manage the software.

Monolithic applications: Traditional applications are developed to run as a single large monolithic process. Large does not refer to the lines of code but to the large number of functionalities and responsibilities. Typically, traditional applications are deployed to a single VM using manual processes. And they are not typically designed to be scalable. The only option is to increase CPU, disk, and memory to achieve higher performance.

Separate environments: Developers start developing on their workstations. Eventually, code moves to testing, staging, and production environments. Each environment is manually

https://t.me/learningnets

configured, resulting in a relatively large amount of effort in all identical environments. Each environment creates different software libraries, packages, and configurations. This variation causes issues for developers who must determine why the application works in one environment but not in the others.

https://t.me/learningnets

3-40 Modern Application Development

Modern application development is transforming modern business.

Modern Application Development	
Typically use microservices style architectures.	Monolithic applications are broken into many smaller standalone modular functions or services that make it easier for developers to be innovative when producing and changing code.
Minimize time to market.	Streamline the process of deploying new code into a staging environment for testing. Identify and address bugs almost immediately. Quickly deploy small, incremental changes in the production environment and easily withdraw if problems arise.
Deliver updates and features quickly.	Minimize the time it takes to build, test, and release new features.
Increase product quality and avoid risk.	Automate tests, get user feedback, and improve software iteratively.
Fewer resource requirements and more productivity.	Apply continuous development and continuous integration in small iterations to reduce labor.

https://t.me/learningnets

3-41 Benefits of Microservices and Containerization

Containers are an ideal technology for supporting microservices because the goals of containers (lightweight, easily packaged, can run anywhere) align well with the goals of a microservices architecture.

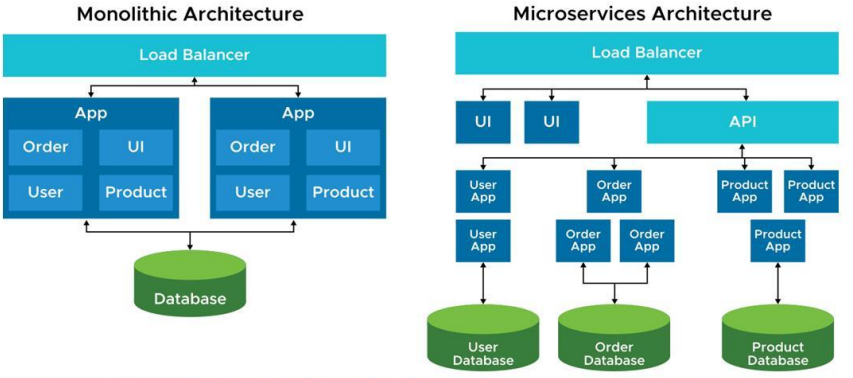

Applications that run on cloud-based environments are designed with failure in mind. They are built to be resilient, to tolerate network or database outages, and to degrade gracefully.

Typically, cloud-native applications use microservice-based architectures. The term micro does not correlate to lines of code. It refers to functionality and responsibility.

Each microservice should be responsible for specific parts of the system.

In the example, the application is broken into multiple services, including a UI and user, order, and product services. Each service has its own database. With this architecture, each service can be scaled independently. For example, during busy times, the order service might need to be scaled to handle high throughput.

The Twelve-Factor App principles describe characteristics of microservice and cloud-native applications.

https://t.me/learningnets

3-42 Container Terminology

Several terms and concepts apply to containers.

Term	Definition
Container	An application packaged with dependencies
Container engine	A runtime engine that manages the containers
Docker	The most recognized runtime engine for container support, and it is often used as a synonym for many aspects of container technologies
Container host	A virtual machine or physical machine on which the containers and container engine run
Kubernetes	Google-developed orchestration for containers

https://t.me/learningnets

3-43 About Containers

A container is an encapsulation of an application and dependent binaries and libraries. The application is decoupled from the operating system and becomes a serverless function.

Among the reasons that containers were popularized by software developers are:

- They make coding easier, locally and anywhere.
- You can deploy and test applications quickly in a staging environment. No operating system or load is required.

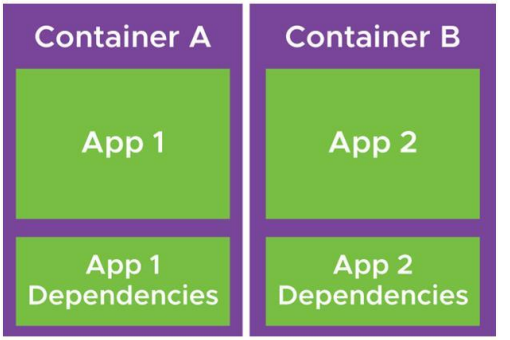

https://t.me/learningnets

3-44 Rise of Containers

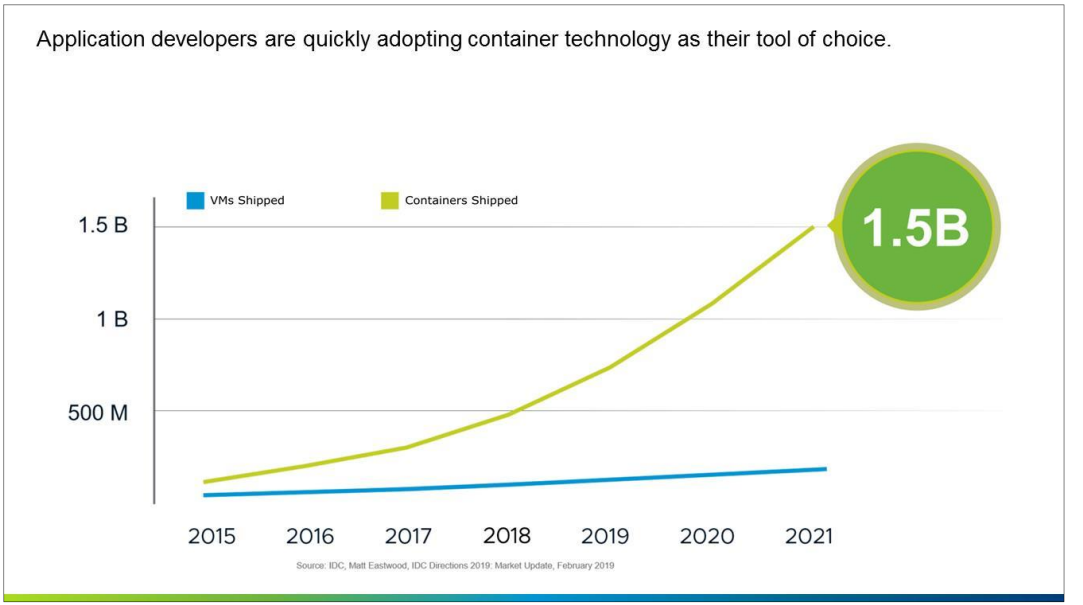

Containers are a new format of virtualized workload. They require CPU, memory, network, security, and storage.

Containers satisfy developers' need for speed by removing dependencies on underlying operating systems:

- Change the paradigm on security by using a discard and restart approach to patching and upgrades.
- Use structured tooling to fully automate updates of application logic running inside.
- Provide an easy user experience for developers that is infrastructure-agnostic (meaning that it can run on any cloud).

The opportunities containers present are many, given the infrastructure and operational complexity that they offer.

https://t.me/learningnets

3-45 About Container Hosts

The container host runs the operating system on which the containers run.

Container hosts can be of the following types:

- Standard OS with a container engine installed:
  - Ubuntu with Docker
- OS developed specifically with containers in mind:
  - Photon OS
  - Fedora CoreOS
- Virtual machine or physical machine:
  - Among the many benefits of using VMs are easy management and scalability.

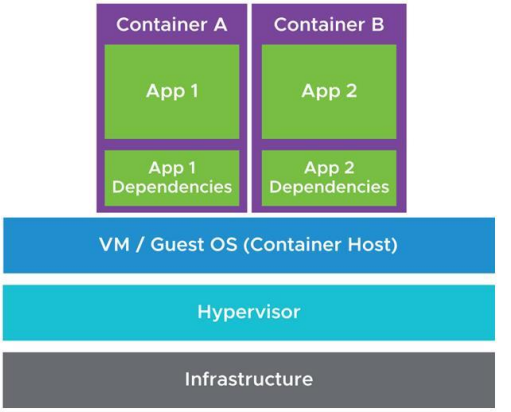

Administrators provide container hosts, which are the base structure that developers use to run their containers. A robust microservices system includes more deliverables, many of which are built using containers.

For developers to focus on providing services to customers, operations must provide a reliable container host infrastructure.

In vSphere with Kubernetes, the container hosts are Photon-based VMs.

https://t.me/learningnets

3-46 Containers at Runtime

Containers have the following characteristics:

- A container can run on any container host with the same operating system kernel that is specified by that container.
- A running container is accessed using its FQDN or its unique IP address.
- Each container can access only its own resources in the shared environment.

  When you log into a container using a remote terminal (such as SSH), you see no indication that other containers are running on the same container host.

https://t.me/learningnets

3-47 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.

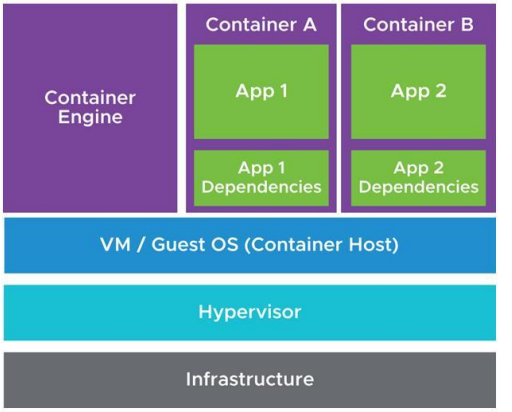

Docker is the most commonly used container engine.

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.

https://t.me/learningnets

3-48 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.

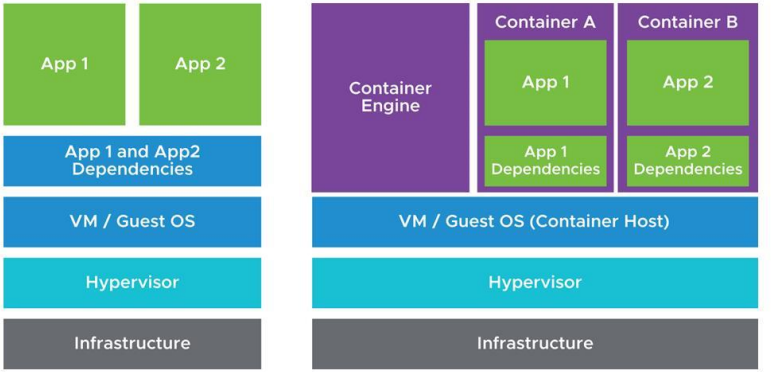

With virtualization, multiple physical machines can be consolidated into a single physical machine that runs multiple VMs. Each VM provides virtual hardware that the guest OS uses to run applications. Multiple applications run on a single VM but these applications are still logically separated and isolated.

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.

https://t.me/learningnets

3-49 Virtual Machines and Containers (2)

VMs and containers work in different 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.

https://t.me/learningnets

3-50 About Kubernetes

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 problems.

Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications.

Kubernetes automates many key operational responsibilities, providing the developer with a reliable environment.

Kubernetes performs the following functions:

- Groups containers that make up an application into logical units for easy management and discovery
- Automatically places containers based on their resource requirements
- Restarts failed containers, replaces and reschedules containers when hosts fail, and stops containers that do not respond to your user-defined health check
- Progressively rolls out changes to your application, ensuring that it does not stop all your instances at the same time and enabling zero downtime
- Allocates IP addresses, mounts the storage system of your choice, load balances, and generally looks after the containers

https://t.me/learningnets

Kubernetes manages containers across multiple container hosts, similar to how vCenter Server manages all ESXi hosts in a cluster. Running Docker without Kubernetes is like running ESXi hosts without vCenter Server to manage them.

https://t.me/learningnets

3-51 Challenges of Running Kubernetes in Production

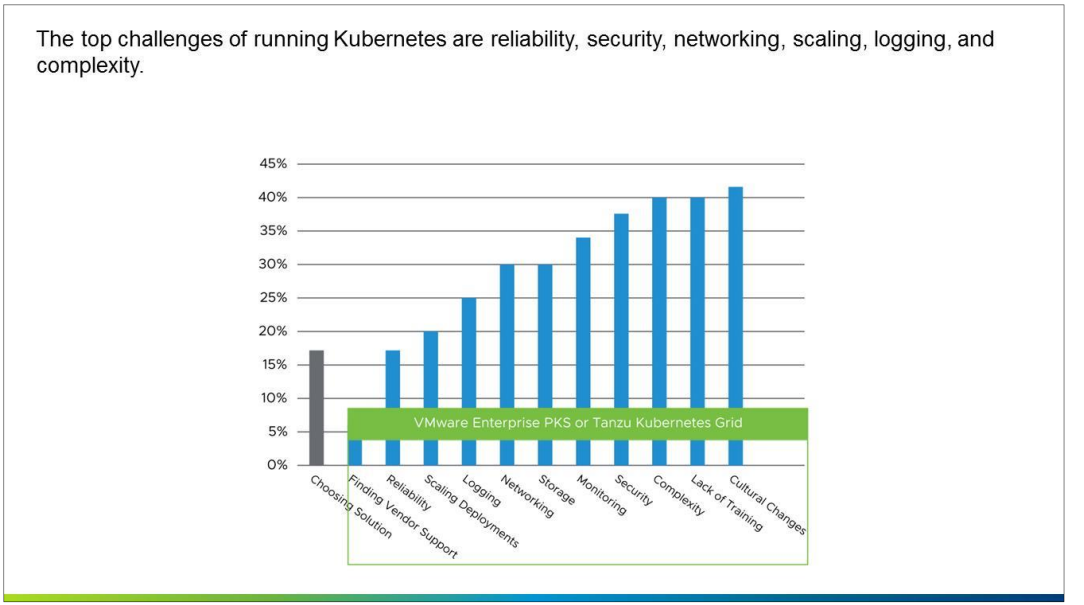

Kubernetes orchestrates containers that support the application. However, running Kubernetes in production is not easy, especially for operations teams. The top challenges of running Kubernetes are related to reliability, security, networking, scaling, logging, and complexity. How do you monitor Kubernetes and the underlying infrastructure? How do you build a reliable platform to deploy your applications? How do you handle the complexity that this layer of abstraction introduces?

For years, VMware has helped to solve these types of problems for IT. VMware can offer its expertise and solutions in this area.

https://t.me/learningnets

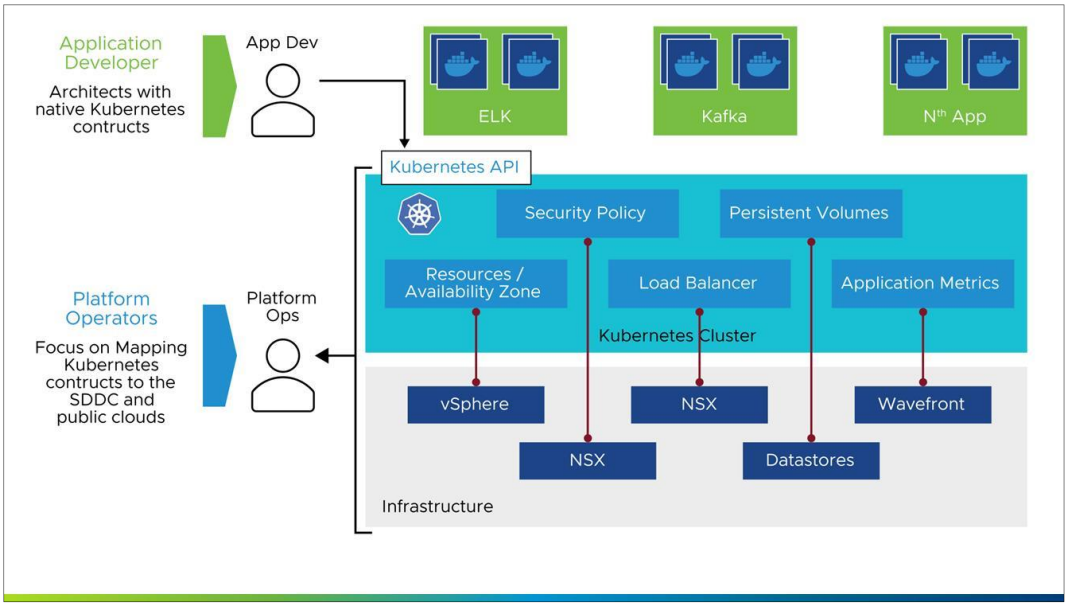

Application developers prefer using Kubernetes rather than programming to the infrastructure. For example, an application developer must build an ELK stack. The developer prefers to deal with the Kubernetes API. The developer wants to use the resources, load balancer, and all the primitives that Kubernetes constructs, rather than worry about the underlying infrastructure.

But the infrastructure is still there. It must be mapped for Kubernetes to use it. Usually, that mapping is done by a platform operator so the developer can use the Kubernetes constructs.

The slide shows how the mapping is done with the VMware software-defined data center (SDDC). The resources and availability zones map to vSphere clusters, security policy and load-balancing map to NSX, persistent volumes map to vSphere datastores and metrics map to Wavefront. Each of these items provides value.

https://t.me/learningnets

3-53 Review of Learner Objectives

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

- Describe the benefits and use cases for containers
- Identify the parts of a container system
- Differentiate between containers and virtual machines

https://t.me/learningnets

3-54 Virtual Beans: Virtualizing Workloads

As a Virtual Beans administrator, you want to start creating VMs with different configurations and testing your applications.
Your key takeaways are:

- The VMware Compatibility Guide can help you determine what versions of Windows and Linux guest operating systems are supported in ESXi 7.0.
- Virtual machines support a wide selection of virtual hardware devices, for example, vGPUs and NVME adapters.
- vSphere provides the underlying infrastructure on which containers and Kubernetes run.

https://t.me/learningnets

3-55 Key Points

- A VM is a set of files that are encapsulated into a folder and placed on a datastore.
- VMs can be provisioned using the vSphere Client and VMware Host Client.
- VMware Tools increases the overall performance of the VM's guest operating system.
- The virtual hardware version, or VM compatibility level, determines the operating system functions that a VM supports.
- Containers are the ideal technology for microservices because the goals of containers align with the goals and benefits of the microservices architecture.

Questions?

https://t.me/learningnets

Module 4

vCenter Server

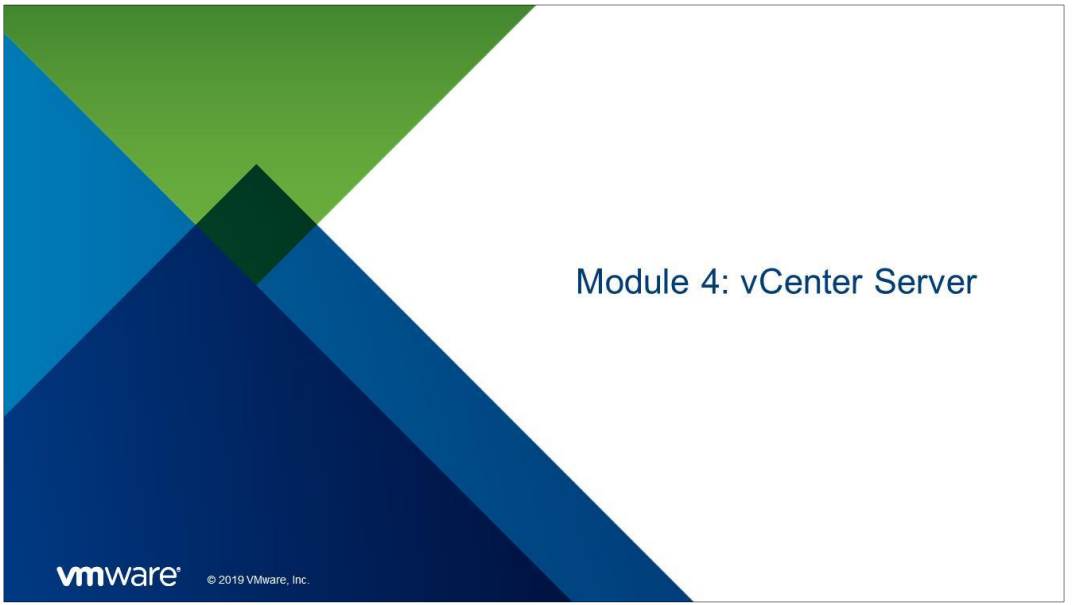

https://t.me/learningnets

4-2 Importance

vCenter Server helps you centrally manage multiple ESXi hosts and their virtual machines. If you do not properly deploy, configure, and manage vCenter Server Appliance, your environment might experience reduced administrative efficiency or ESXi host and virtual machine downtime.

https://t.me/learningnets

4-3 Module Lessons

1. Centralized Management with vCenter Server
2. Deploying vCenter Server Appliance
3. vSphere Licensing
4. Managing the vCenter Server Inventory
5. vCenter Server Roles and Permissions
6. Backing Up and Restoring vCenter Server Appliance
7. Monitoring vCenter Server Appliance
8. vCenter Server High Availability

https://t.me/learningnets

4-4 Virtual Beans: vCenter Server Requirements

Virtual Beans has the following requirements for vCenter Server (the management platform):

- Use Active Directory for user and group authentication.
- Restrict user access to vCenter Server:
  - System administrators: Administrative rights to the entire vSphere environment
  - Operators: Provision VMs and monitor, and troubleshoot vSphere problems
- When the new data center comes online, manage both data centers from a centralized management console.

As a Virtual Beans administrator, you are responsible for installing and configuring vCenter Server, and setting up user access.

https://t.me/learningnets

4-5 Lesson 1: Centralized Management with vCenter Server

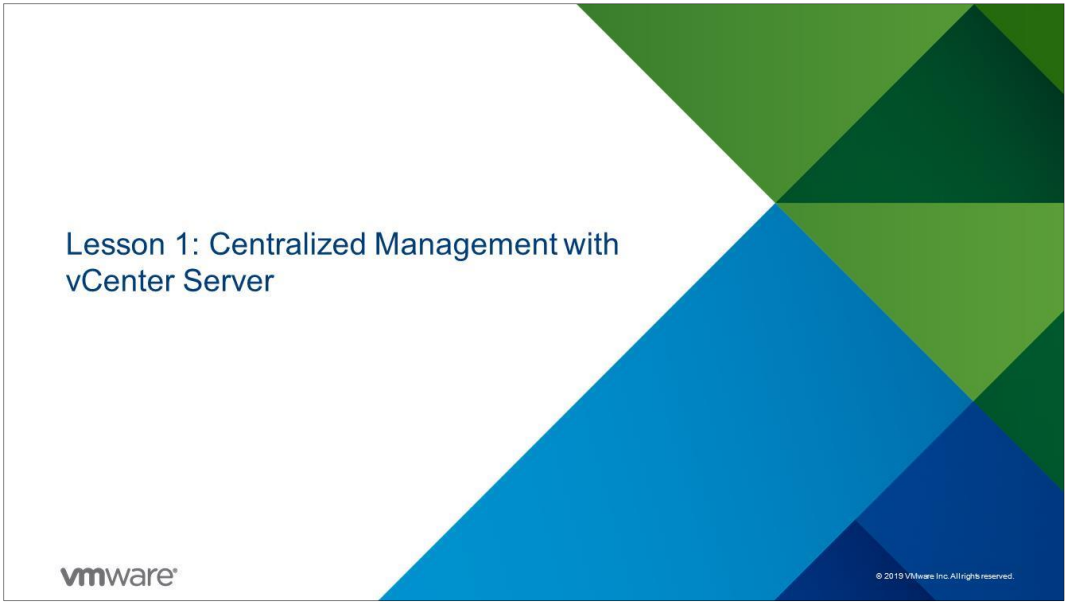

https://t.me/learningnets

4-6 Learner Objectives

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

- Describe the vCenter Server architecture
- Recognize how ESXi hosts communicate with vCenter Server
- Identify vCenter Server services

https://t.me/learningnets

4-7 About the vCenter Server Management Platform

vCenter Server acts as a central administration point for ESXi hosts and virtual machines that are connected in a network:

- Directs the actions of VMs and hosts
- Runs on a Linux-based appliance

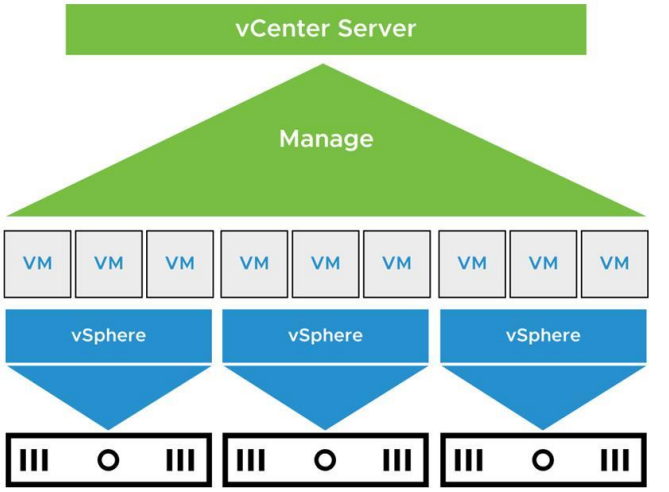

With vCenter Server, you can pool and manage the resources of multiple hosts.

You can deploy vCenter Server Appliance on an ESXi host in your infrastructure. vCenter Server Appliance is a preconfigured Linux-based virtual machine that is optimized for running vCenter Server and the vCenter Server components.

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

https://t.me/learningnets

4-8 About vCenter Server Appliance

vCenter Server Appliance is a prepackaged Linux-based VM that is optimized for running vCenter Server and associated services.

The vCenter Server Appliance package contains the following software:

- Photon
- PostgreSQL database
- vCenter Server services

During deployment, you can select the vCenter Server Appliance size for your vSphere environment and the storage size for your database requirements.

vCenter Server is a service that runs in vCenter Server Appliance. vCenter Server acts as a central administrator for ESXi hosts that are connected in a network.

https://t.me/learningnets

4-9 vCenter Server Services

vCenter Server services include:

- vCenter Server
- vSphere Client
- vCenter Single Sign-On
- License service
- vCenter Lookup Service
- VMware Certificate Authority
- Content Library
- vSphere ESXi Dump Collector

When you deploy vCenter Server Appliance, all these services are included.

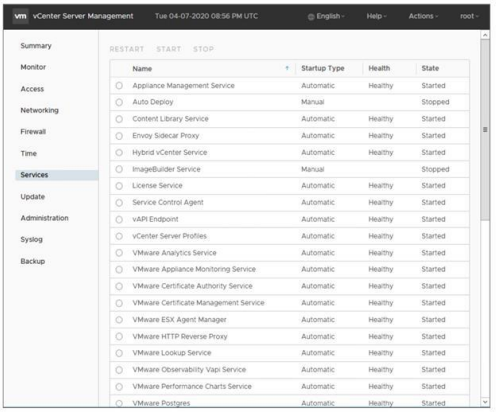

Although installation of vCenter Server services is not optional, administrators can choose whether to use their functionalities.

https://t.me/learningnets

4-10 vCenter Server Architecture

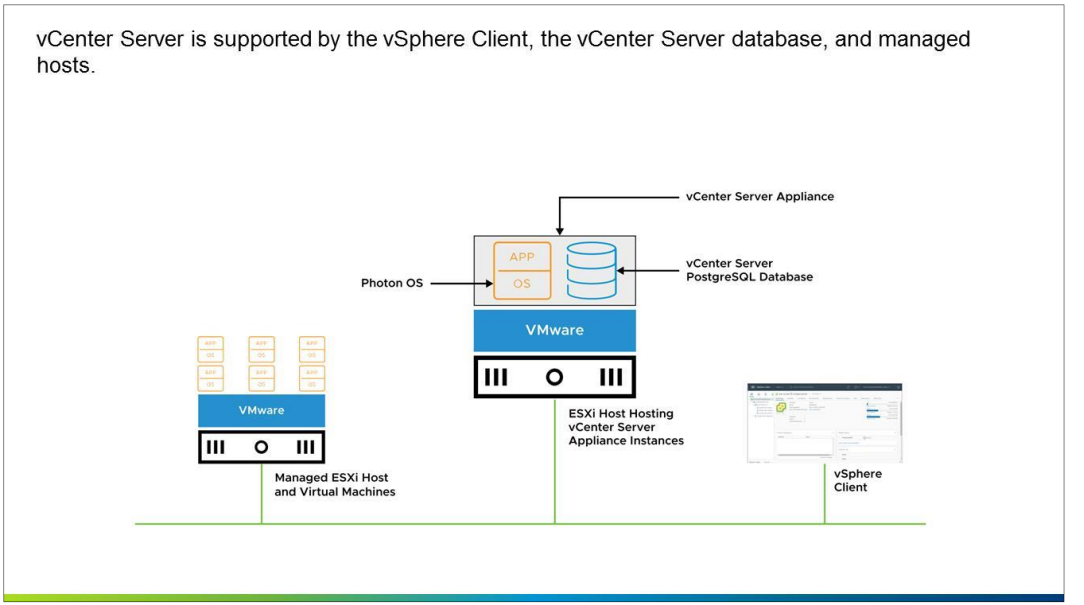

The vCenter Server architecture relies on the following components:

- vSphere Client: You use this client to connect to vCenter Server so that you can manage your ESXi hosts centrally. When an ESXi host is managed by vCenter Server, you should always use vCenter Server and the vSphere Client to manage that host.
- vCenter Server database: The vCenter Server database is the most important component. The database stores inventory items, security roles, resource pools, performance data, and other critical information for vCenter Server.
- Managed hosts: You can use vCenter Server to manage ESXi hosts and the VMs that run on them.

https://t.me/learningnets

4-11 About vCenter Single Sign-On

vCenter Single Sign-On provides authentication across multiple vSphere components through a secure token mechanism:

1. User logs in to the vSphere Client.
2. vCenter Single Sign-On authenticates credentials against a directory service (for example, Active Directory).
3. A SAML token is sent back to the user's browser.
4. The SAML token is sent to vCenter Server, and the user is granted access.

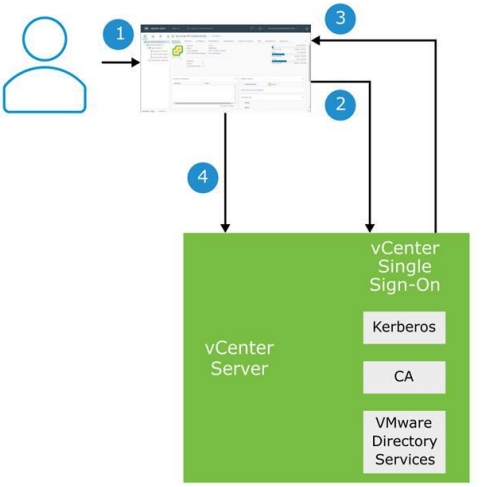

https://t.me/learningnets

4-12 About Enhanced Linked Mode

With Enhanced Linked Mode, you can log in to a single instance of vCenter Server and manage the inventories of all the vCenter Server systems in the group:

- Up to 15 vCenter Server instances can be linked in one vCenter Single Sign-On domain.
- An Enhanced Linked Mode group can be created only during the deployment of vCenter Server Appliance.

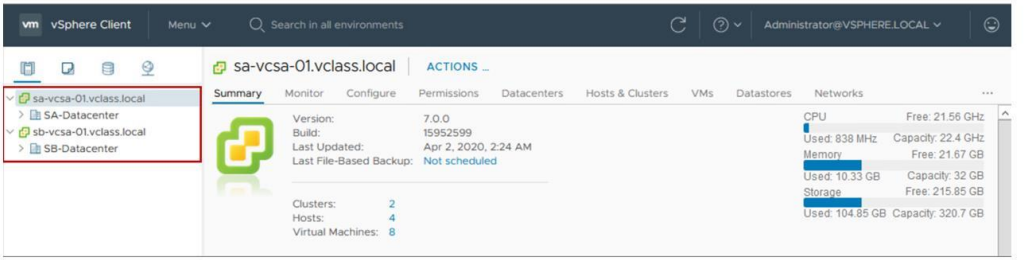

You cannot create an Enhanced Linked Mode group after you deploy vCenter Server Appliance.

Enhanced Linked Mode provides the following features:

- You can log in to all linked vCenter Server instances simultaneously with a single user name and password.
- You can view and search the inventories of all linked vCenter Server instances in the vSphere Client.
- Roles, permission, licenses, tags, and policies are replicated across linked vCenter Server instances.

To join vCenter Server instances in Enhanced Linked Mode, connect the vCenter Server instances to the same vCenter Single Sign-On domain.

Enhanced Linked Mode requires the vCenter Server Standard licensing level. This mode is not supported with vCenter Server Foundation or vCenter Server for Essentials.

https://t.me/learningnets

4-13 ESXi and vCenter Server Communication

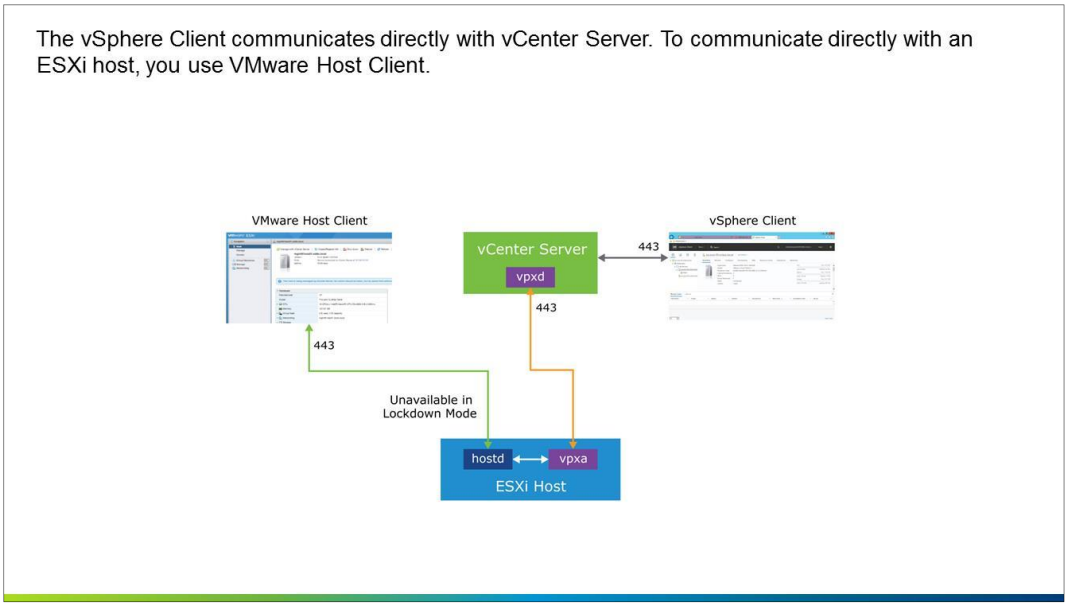

vCenter Server provides direct access to the ESXi host through a vCenter Server agent called virtual provisioning X agent (vpxa). The vpxa process is automatically installed on the host and started when the host is added to the vCenter Server inventory. The vCenter Server service (vpxd) communicates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

Clients that communicate directly with the host, and bypass vCenter Server, converse with hostd. The hostd process runs directly on the ESXi host and manages most of the operations on the ESXi host. The hostd process is aware of all VMs that are registered on the ESXi host, the storage volumes visible to the ESXi host, and the status of all VMs.

Most commands or operations come from vCenter Server through vpxa. Examples include creating, migrating, and powering on virtual machines. Acting as an intermediary between the vpxd process, which runs on vCenter Server, and the hostd process, vpxa relays the tasks to perform on the host.

When you are logged in to the vCenter Server system through the vSphere Client, vCenter Server passes commands to the ESXi host through the vpxa.

https://t.me/learningnets

The vCenter Server database is also updated. If you use VMware Host Client to communicate directly with an ESXi host, communications go directly to the hostd process and the vCenter Server database is not updated.

https://t.me/learningnets

4-14 vCenter Server Appliance Scalability

Metric	vCenter Server Appliance 7.0
Hosts per vCenter Server instance	2,500
Powered-on VMs per vCenter Server instance	40,000
Registered VMs per vCenter Server instance	45,000
Hosts per cluster	64
VMs per cluster	8,000

https://t.me/learningnets

4-15 Review of Learner Objectives

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

- Describe the vCenter Server architecture
- Recognize how ESXi hosts communicate with vCenter Server
- Identify vCenter Server services

https://t.me/learningnets

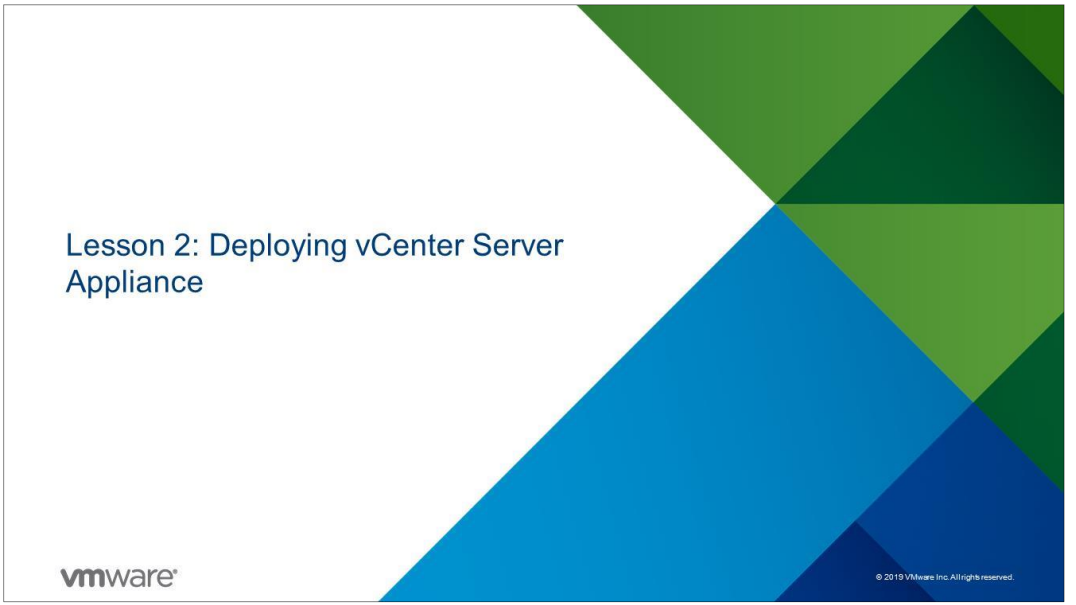

https://t.me/learningnets

4-17 Learner Objectives

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

- Deploy vCenter Server Appliance into an infrastructure
- Configure vCenter Server settings

https://t.me/learningnets

4-18 Preparing for vCenter Server Appliance Deployment

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

- Verify that all vCenter Server Appliance system requirements are met.
- Get the fully qualified domain name (FQDN) or the static IP of the host machine on which you install vCenter Server Appliance.
- Ensure that clocks on all VMs in the vSphere network are synchronized.

For more information, see *VMware ESXi Installation and Setup* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.esxi.install.doc/GUID-B2F01BF5-078A-4C7E-B505-5DFFED0B8C38.html.

https://t.me/learningnets

4-19 vCenter Server Appliance Native GUI Installer

The GUI installer has several features:

- With the GUI installer, you can perform an interactive deployment of vCenter Server Appliance.
- The GUI installer is a native application for Windows, Linux, and macOS.
- The installer has no dependency on browsers or plug-ins.
- It performs validations and prechecks during the deployment.

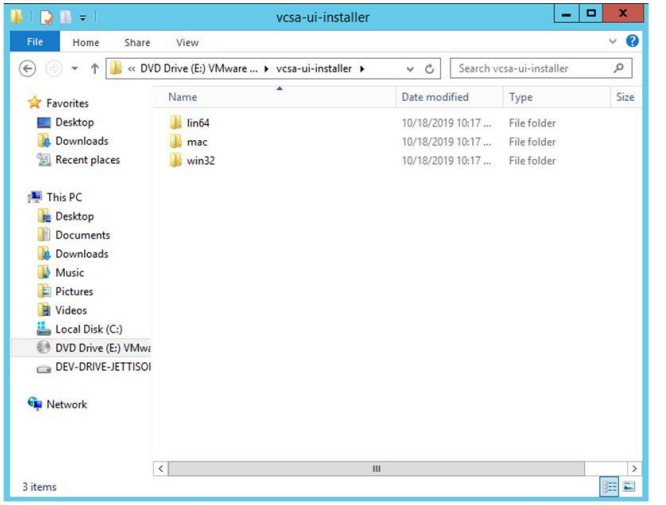

The GUI installer performs validations and prechecks during the deployment phase to ensure that no mistakes are made and that a compatible environment is created.

https://t.me/learningnets

4-20 vCenter Server Appliance Installation

The vCenter Server Appliance installation is a two-stage process:

- Stage 1: Deployment of OVF
- Stage 2: Configuration

The deployment can be fully automated by using JSON templates with the CLI installer on Windows, Linux, or macOS.

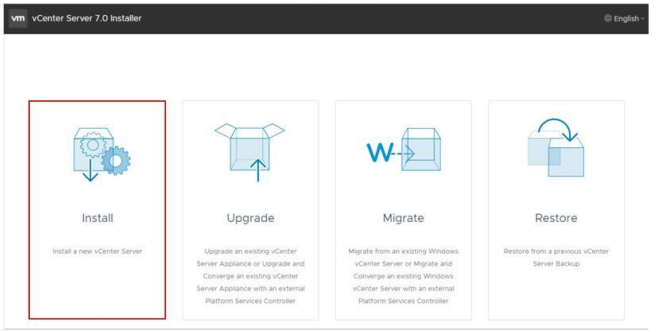

The **Install** option installs a new vCenter Server Appliance.

The **Upgrade** option upgrades an existing vCenter Server Appliance instance, or upgrades and converges an existing vCenter Server Appliance instance with external Platform Services Controller.

The **Migrate** option migrates from an existing Windows vCenter Server instance, or migrates and converges an existing Windows vCenter Server instance with external Platform Services Controller.

The **Restore** option restores from a previous vCenter Server Appliance backup.

https://t.me/learningnets

4-21 vCenter Server Appliance Installation: Stage 1

Stage 1 begins with the UI phase:

- Accept the EULA.
- Connect to the target ESXi host or vCenter Server system.
- Define the vCenter Server Appliance name and root password.
- Select compute size, storage size, and datastore location (thin disk).
- Define networking settings.

Stage 1 continues with the deployment phase:

- OVF is deployed to the ESXi host.
- Disks and networking are configured.

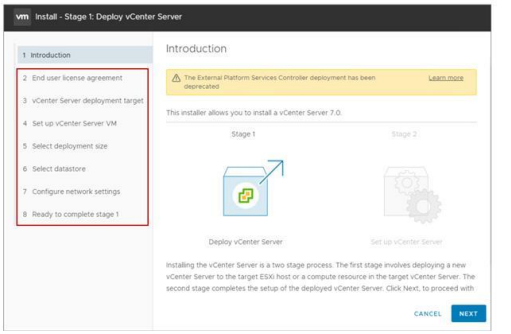

https://t.me/learningnets

4-22 vCenter Server Appliance Installation: Stage 2

Stage 2 is the configuration phase:

- Configure time synchronization mode and SSH access.
- Create a vCenter Single Sign-On domain or join an existing SSO domain.
- Join the Customer Experience Improvement Program (CEIP).

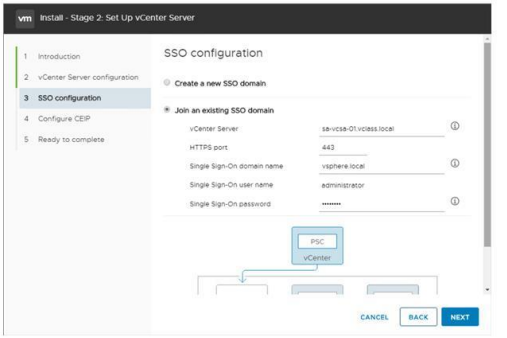

In stage 2, you configure whether to use the ESXi host or NTP servers as the time synchronization source. You can also enable SSH access. SSH access is disabled by default.

https://t.me/learningnets

After you deploy vCenter Server Appliance, use the vSphere Client to log in and manage your vCenter Server inventory: https://vCenter_Server_FQDN_or_IP_address/ui.

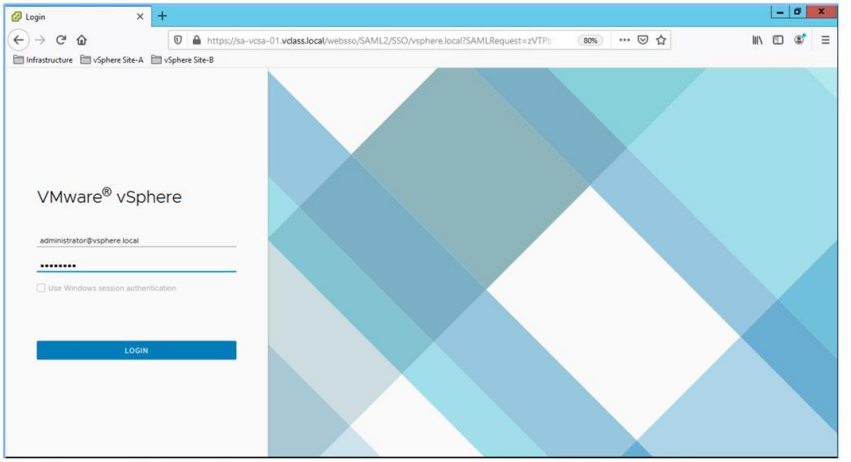

https://t.me/learningnets

Using the vSphere Client, you can configure vCenter Server, including settings such as licensing, statistics collection, and logging.

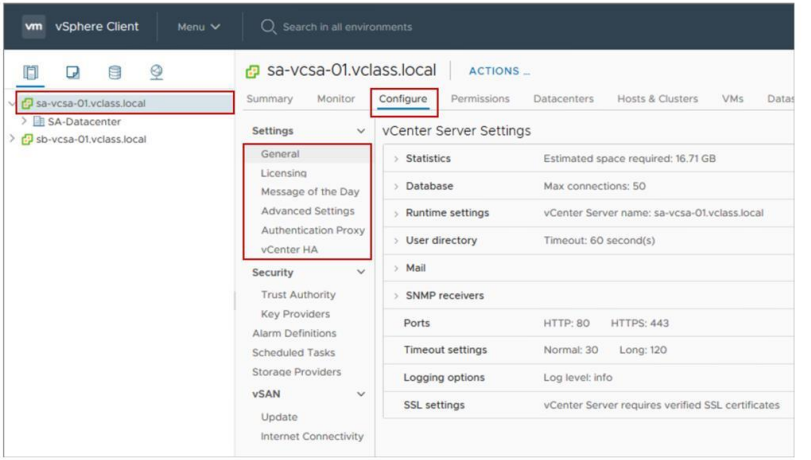

To access the vCenter Server system settings by using the vSphere Client, select the vCenter Server system in the navigation pane, click the **Configure** tab, and expand **Settings**.

https://t.me/learningnets

4-25 vCenter Server Appliance Management Interface

Using the vCenter Server Appliance Management Interface (VAMI), you can configure and monitor your vCenter Server Appliance instance.
Tasks include:

- Monitoring resource use by the appliance
- Backing up the appliance
- Monitoring vCenter Server services
- Adding additional network adapters

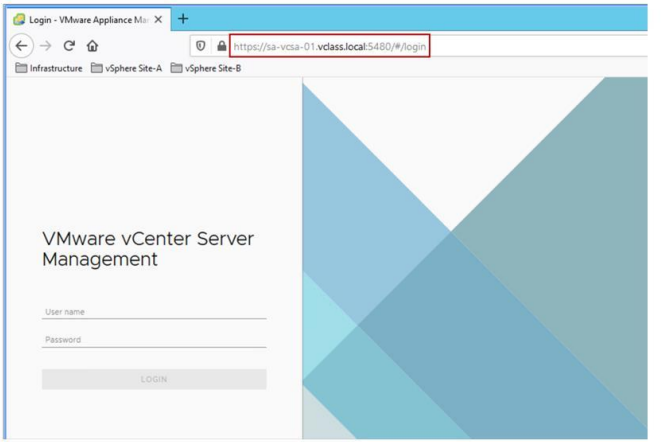

The vCenter Server Appliance Management Interface is an HTML client designed to configure and monitor vCenter Server Appliance.

The vCenter Server Appliance Management Interface connects directly to port 5480. Use the URL https://FQDN_or_IP_address:5480.

https://t.me/learningnets

4-26 vCenter Server Appliance Multihoming

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

For example, vCenter Server High Availability requires a second NIC for its private network.

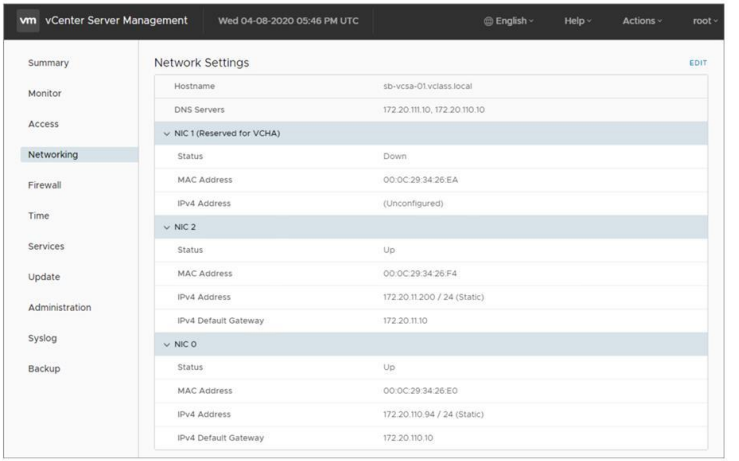

A maximum of four NICs are supported for multihoming. All four multihoming-supported NIC configurations are preserved during upgrade, backup, and restore processes.

https://t.me/learningnets

4-27 Demonstration: Deploying vCenter Server Appliance

Your instructor will run a demonstration.

https://t.me/learningnets

4-28 Review of Learner Objectives

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

- Deploy vCenter Server Appliance into an infrastructure
- Configure vCenter Server settings

https://t.me/learningnets

https://t.me/learningnets

4-30 Learner Objectives

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

- View licensed features for vCenter Server or an ESXi host
- Add license keys to vCenter Server

https://t.me/learningnets

4-31 vSphere Licensing Overview

Licensing vSphere components is a two-step process:

1. Add a license to the vCenter License Service.
2. Assign the license to the ESXi hosts, vCenter Server Appliance instances, and other vSphere components.

https://t.me/learningnets

4-32 vSphere License Service

The License Service runs on vCenter Server Appliance.
The License Service performs the following functions:

- Provides centralized license management
- Provides an inventory of vSphere licenses
- Manages the license assignments for products that integrate with vSphere, such as Site Recovery Manager.

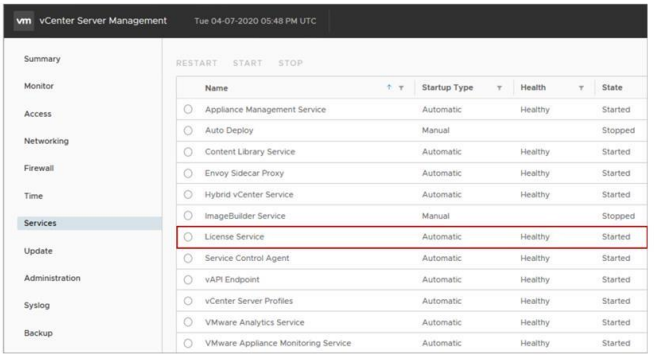

The License Service manages the license assignments for ESXi hosts, vCenter Server systems, and clusters with vSAN enabled.

You can monitor the health and status of the License Service by using the vCenter Appliance Management Interface.

https://t.me/learningnets

4-33 Adding License Keys to vCenter Server

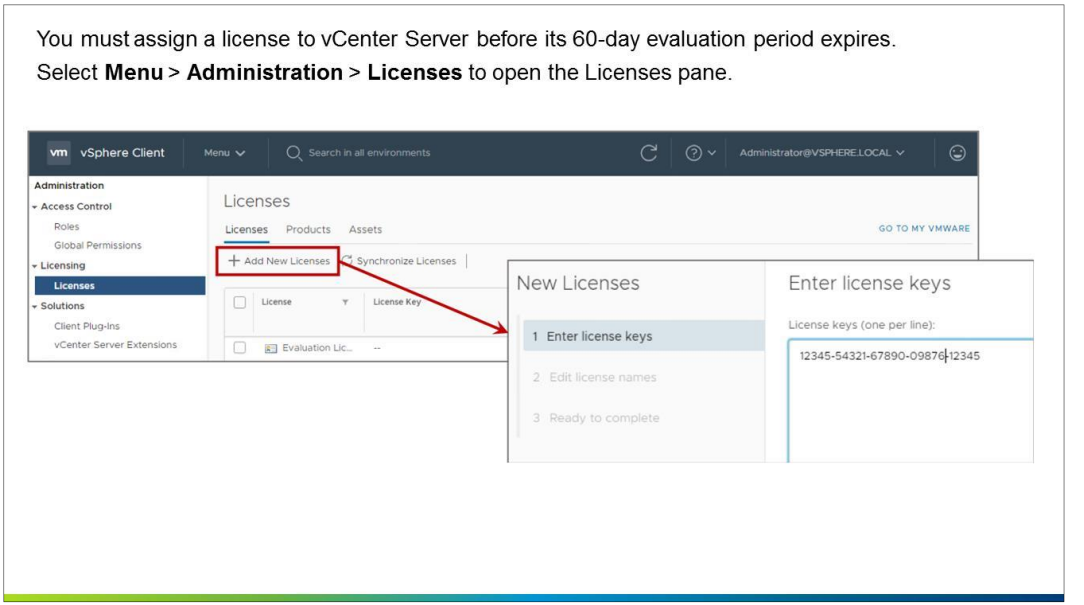

In the vSphere environment, license reporting and management are centralized. All product and feature licenses are encapsulated in 25-character license keys that you can manage and monitor from vCenter Server.

You can view license information by product, license key, or asset:

- Product: A license to use a vSphere software component or feature, for example, evaluation mode or vSphere Enterprise Plus.
- License key: The serial number that corresponds to a product.
- Asset: A machine on which a product is installed. For an asset to run certain software legally, the asset must be licensed.

https://t.me/learningnets

4-34 Assigning a License to a vSphere Component

You can assign a license to an asset, such as vCenter Server.

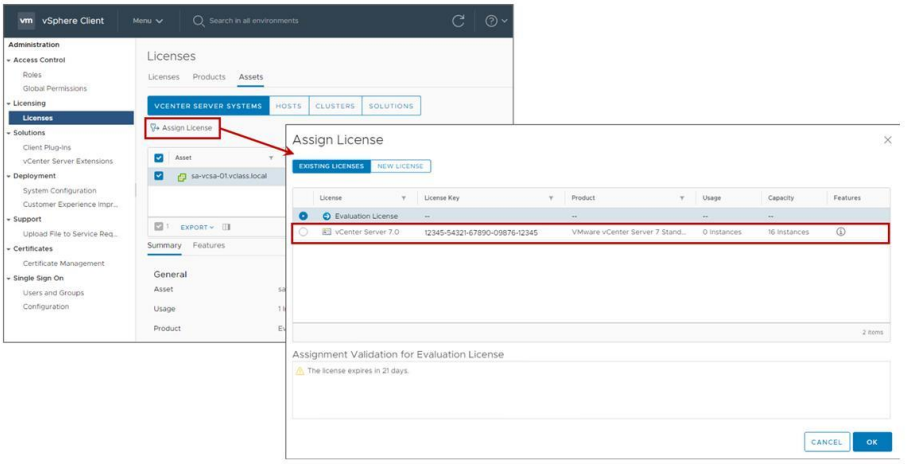

https://t.me/learningnets

4-35 Viewing Licensed Features

You assign valid license keys to your ESXi hosts and vCenter Server instance using the Licensing pane. This pane shows the type of license and available features.

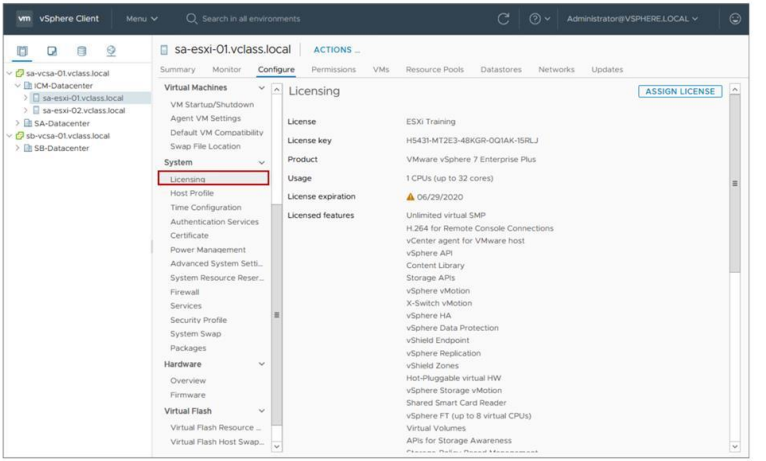

Before purchasing and activating licenses for ESXi and vCenter Server, you can install the software and run it in evaluation mode. Evaluation mode is intended for demonstrating the software or evaluating its features. During the evaluation period, the software is operational.

The evaluation period is 60 days from the time of installation. During this period, the software notifies you of the time remaining until expiration. The 60-day evaluation period cannot be paused or restarted. After the evaluation period expires, you can no longer perform some operations in vCenter Server and ESXi. For example, you cannot power on or reset your virtual machines. In addition, all hosts are disconnected from the vCenter Server system. To continue to have full use of ESXi and vCenter Server operations, you must acquire license keys.

https://t.me/learningnets

4-36 Lab 6: Adding vSphere Licenses

Use the vSphere Client to add vSphere licenses to vCenter Server and assign a license to vCenter Server:

1. Add vSphere Licenses to vCenter Server
2. Assign a License to the vCenter Server Instance

https://t.me/learningnets

4-37 Review of Learner Objectives

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

- View licensed features for vCenter Server or an ESXi host
- Add license keys to vCenter Server

https://t.me/learningnets

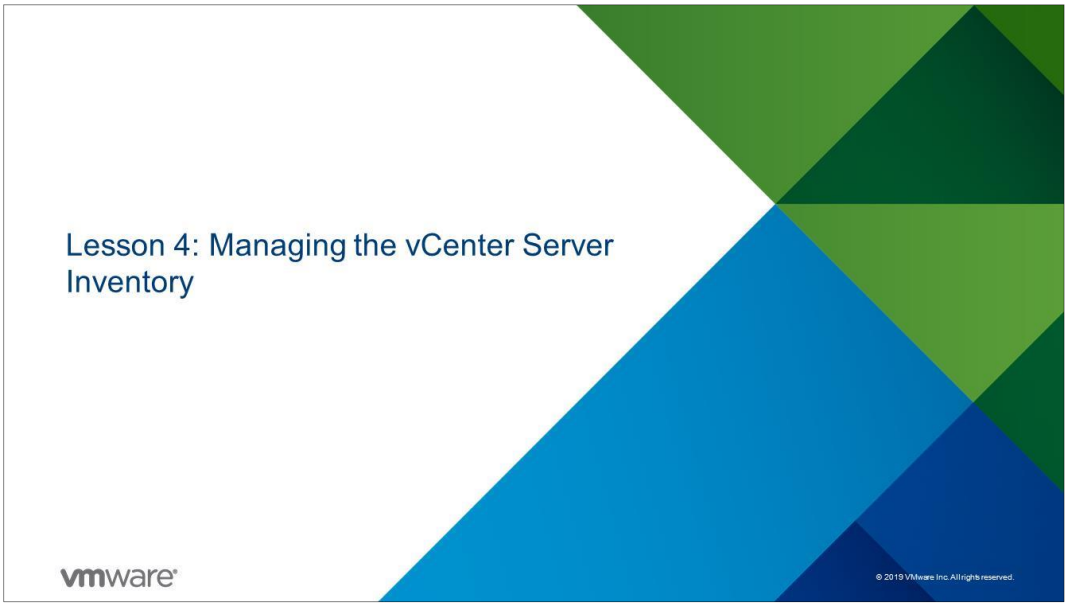

https://t.me/learningnets

4-39 Learner Objectives

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

- Use the vSphere Client to manage the vCenter Server inventory
- Create and organize vCenter Server inventory objects
- Add data center and organizational objects to vCenter Server
- Add hosts to vCenter Server
- Recognize how to create custom inventory tags for inventory objects

https://t.me/learningnets

4-40 vSphere Client Shortcuts Page

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

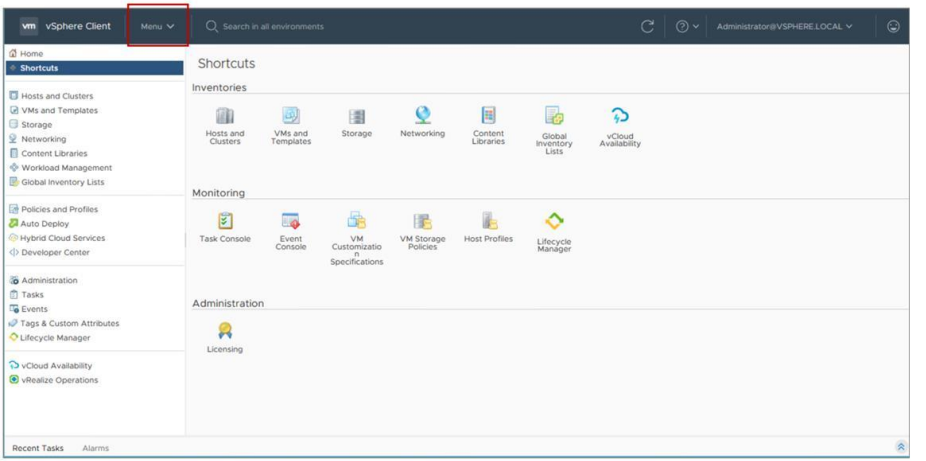

Select **Menu** > **Shortcuts**. The Shortcuts page has a navigation pane on the left and Inventories, Monitoring, and Administration panes on the right.

https://t.me/learningnets

4-41 Using the Navigation Pane

You can use the navigation pane to browse and select objects in the vCenter Server inventory.

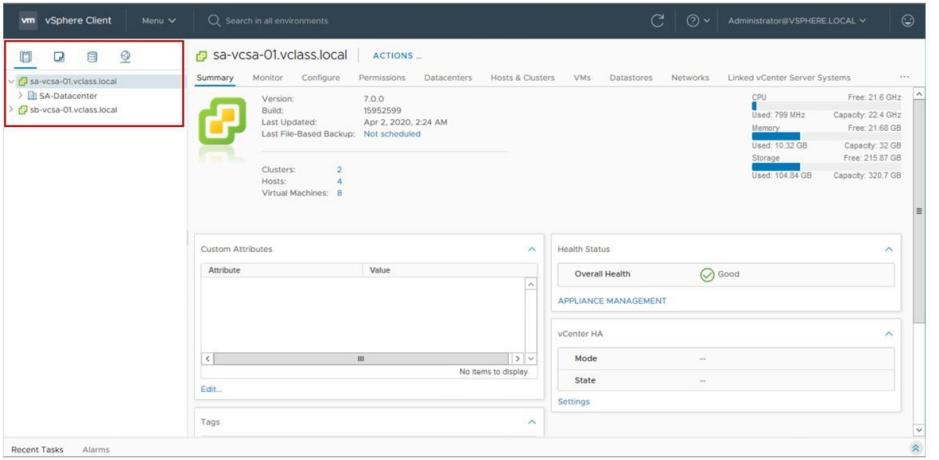

https://t.me/learningnets

4-42 vCenter Server Views for Hosts, Clusters, VMs, and Templates

Host and cluster objects are shown in one view, and VM and template objects are displayed in another view.

Host and Clusters Inventory View

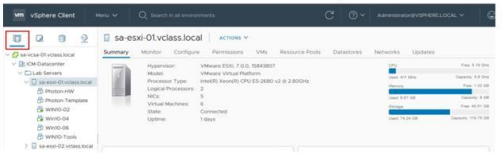

VMs and Templates Inventory View

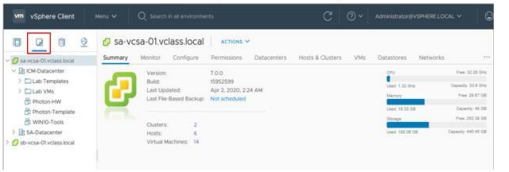

The Hosts and Clusters inventory view shows all host and cluster objects in a data center. You can further organize the hosts and clusters into folders.

The VMs and Templates inventory view shows all VM and template objects in a data center. You can also organize the VMs and templates into folders.

https://t.me/learningnets

4-43 vCenter Server Views for Storage and Networks

The Storage inventory view shows all the details for datastores in the data center. The Networking inventory view shows all standard switches and distributed switches.

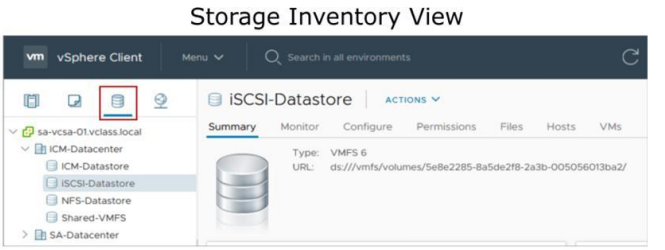

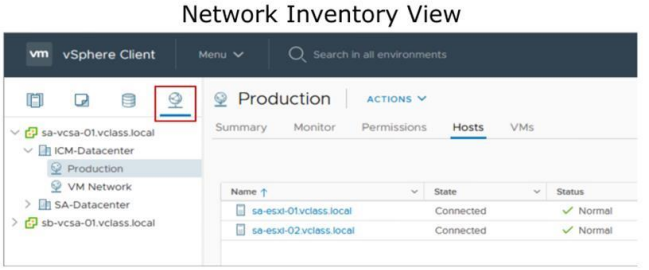

As with the other inventory views, you can organize your datastore and network objects into folders.

https://t.me/learningnets

4-44 Viewing Object Information

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

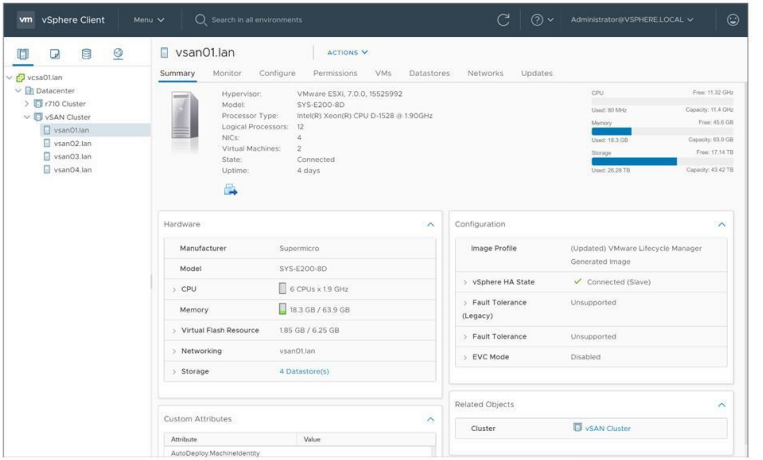

https://t.me/learningnets

4-45 About Data Center Objects

A virtual data center is a logical organization of all the inventory objects required to complete a fully functional environment for operating VMs:

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

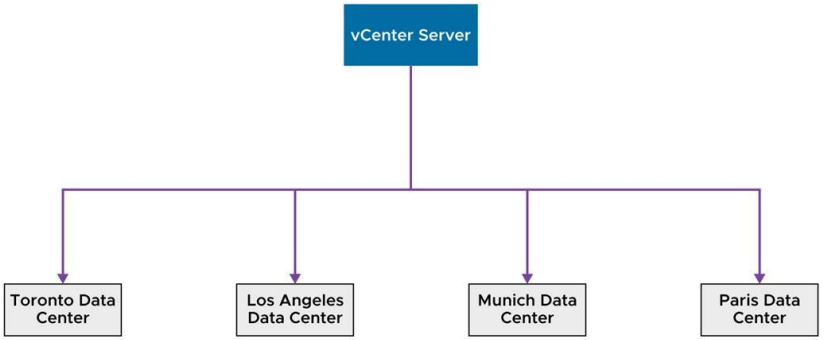

You might create a data center object for each data center geographical location. Or, you might create a data center object for each organizational unit in your enterprise.

You might create some data centers for high-performance environments and other data centers for less demanding VMs.

https://t.me/learningnets

4-46 Organizing Inventory Objects into Folders

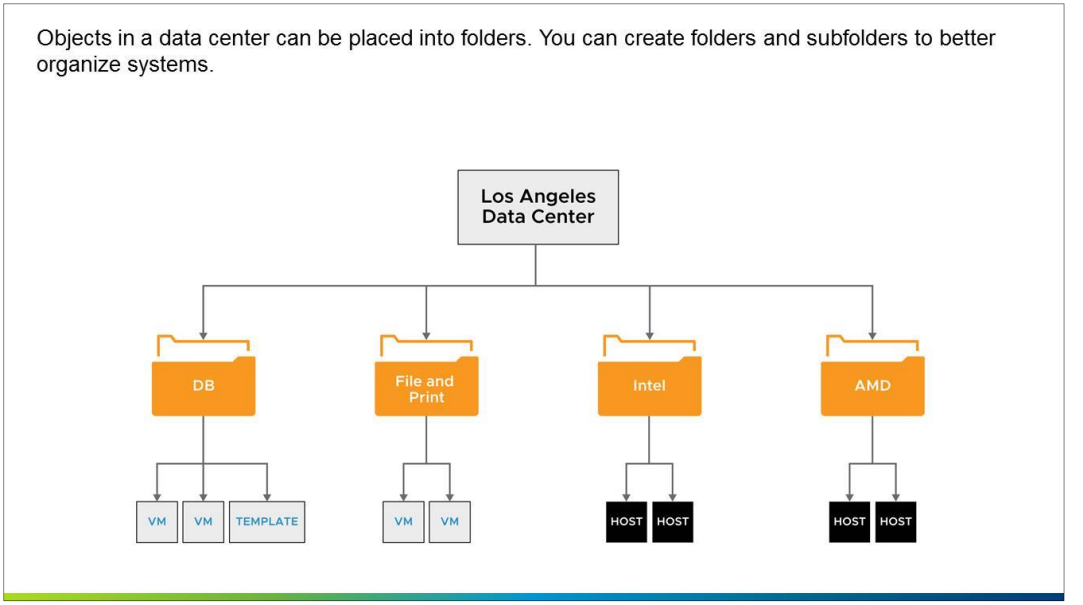

You plan the setup of your virtual environment depending on your requirements.

A large vSphere implementation might contain several virtual data centers with a complex arrangement of hosts, clusters, resource pools, and networks. It might include multiple vCenter Server systems.

Smaller implementations might require a single virtual data center with a less complex topology.

Regardless of the scale of your virtual environment, consider how the VMs that it supports are used and administered.

Populating and organizing your inventory involves the following tasks:

- Creating data centers
- Creating clusters to consolidate the resources of multiple hosts and VMs
- Adding hosts to the clusters or to the data centers

https://t.me/learningnets

- Organizing inventory objects in folders
- Setting up networking by using vSphere standard switches or vSphere distributed switches
- Configuring storage systems and creating datastore inventory objects to provide logical containers for storage devices in your inventory

https://t.me/learningnets

4-47 Adding a Data Center and Organizational Objects to vCenter Server

You can add a data center, a host, a cluster, and folders to vCenter Server.
You can use folders to group objects of the same type for easier management.

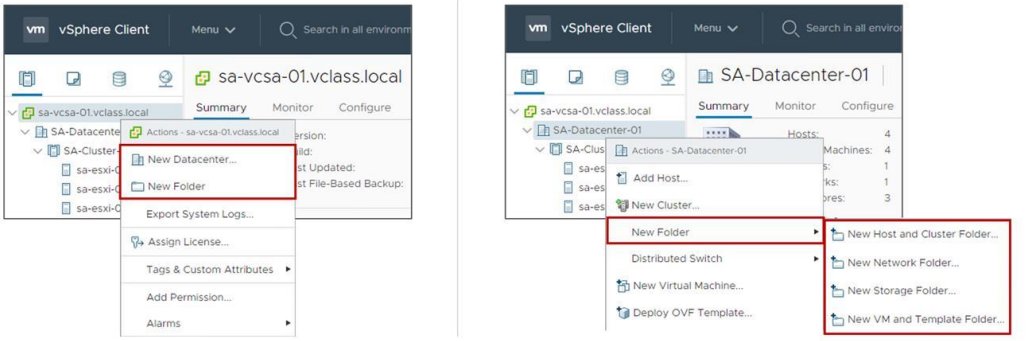

https://t.me/learningnets

4-48 Adding ESXi Hosts to vCenter Server

You can add ESXi hosts to vCenter Server using the vSphere Client.

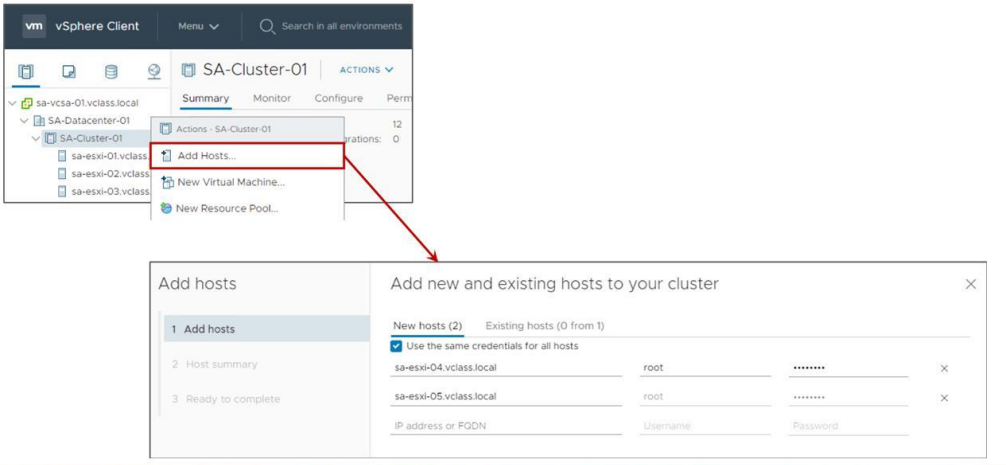

https://t.me/learningnets

4-49 Creating Custom Tags for Inventory Objects

You can use tags to attach metadata to objects in the vCenter Server inventory. Tags help make these objects more sortable.

You can associate a set of objects of the same type by searching for objectives by a given tag.

You can use tags to group and manage VMs, clusters, and datastores, for example:

- Tag VMs that run production workloads.
- Tag VMs based on their guest operating system.

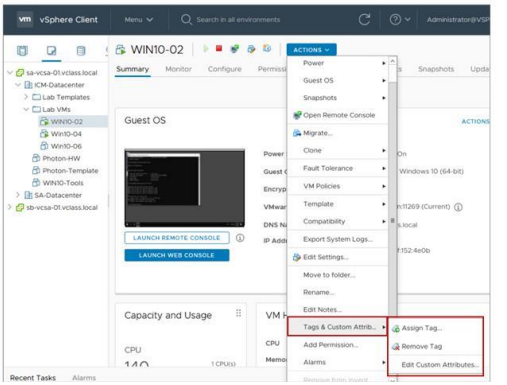

https://t.me/learningnets

Labs

Lab: Creating and Managing the vCenter Server Inventory
Lab: Configuring Active Directory: Joining a Domain

https://t.me/learningnets

4-51 Lab 7: Creating and Managing the vCenter Server Inventory

Use the vSphere Client to create and configure objects in the vCenter Server inventory:

1. Create a Data Center Object
2. Add ESXi Hosts to the Inventory
3. View Information About the ESXi Hosts
4. Configure the ESXi Hosts as NTP Clients
5. Create a Folder for the ESXi Hosts
6. Create Folders for VMs and VM Templates

https://t.me/learningnets

4-52 Lab 8: Configuring Active Directory: Joining a Domain

Join vCenter Server to the vclass.local domain:

1. Join vCenter Server to the vclass.local Domain

https://t.me/learningnets

4-53 Review of Learner Objectives

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

- Use the vSphere Client to manage the vCenter Server inventory
- Create and organize vCenter Server inventory objects
- Add data center and organizational objects to vCenter Server
- Add hosts to vCenter Server
- Recognize how to create custom inventory tags for inventory objects

https://t.me/learningnets

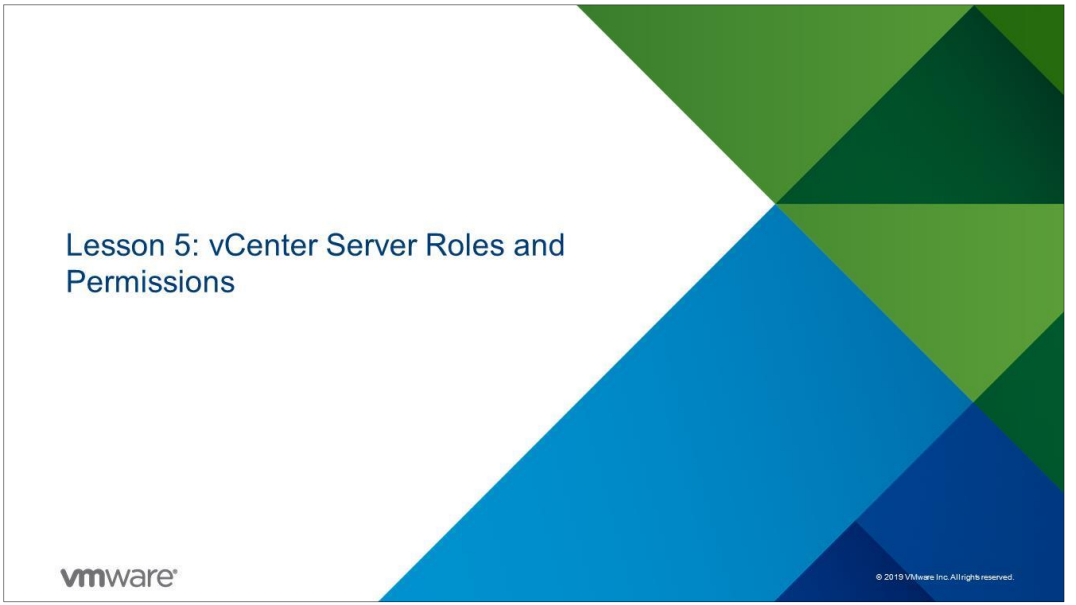

https://t.me/learningnets

4-55 Learner Objectives

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

- Define the term permission in the context of vCenter Server
- Describe the rules for applying permissions
- Create a custom role
- Create a permission

https://t.me/learningnets

4-56 About vCenter Server Permissions

Using the access control system, the vCenter Server administrator can define user privileges to access objects in the inventory.
The following concepts are important:

- Privilege: An action that can be performed
- Object: The target of the action
- User or group: Indication of who can perform the action
- Role: A set of privileges
- Permission: Gives one user or group a role (set of privileges) for the selected object

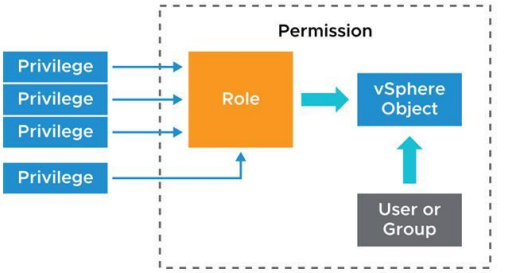

The authorization to perform tasks in vCenter Server is governed by an access control system. Through this system, the vCenter Server administrator can specify in detail which users or groups can perform which tasks on which objects.

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

By assigning a different role to a group of users on different objects, you control the tasks that those users can perform in your vSphere environment. For example, to allow a group to configure memory for the host, select that host and add a permission that grants a role to that group that includes the Host.Configuration.Memory Configuration privilege.

https://t.me/learningnets

4-57 About Roles

Privileges are grouped into roles:

- A privilege allows access to a specific task and is grouped with other privileges related to it.
- Roles allow users to perform tasks.

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

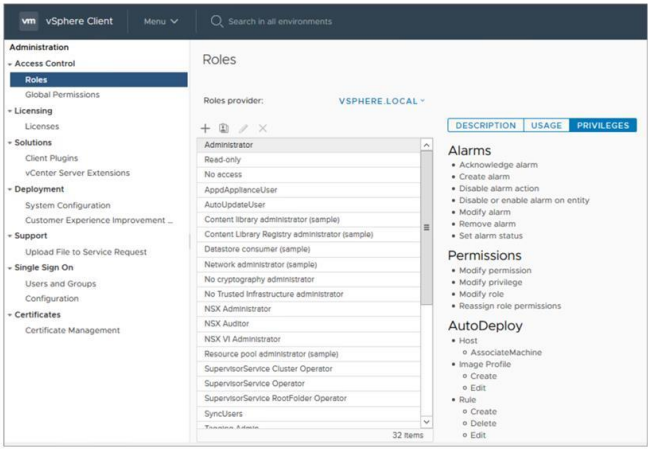

A role is a set of one or more privileges. For example, the Virtual Machine Power User sample role consists of several privileges in categories such as Datastore and Global. A role is assigned to a user or group and determines the level of access of that user or group.

You cannot change the privileges associated with the system roles:

- Administrator role: Users with this role for an object may view and perform all actions on the object.
- Read-only role: Users with this role for an object may view the state of the object and details about the object.

https://t.me/learningnets

- No access role: Users with this role for an object may not view or change the object in any way.
- No cryptography administrator role: Users with this role for an object have the same privileges as users with the Administrator role, except for privileges in the Cryptographic operations category.

All roles are independent of each other. Hierarchy or inheritance between roles does not apply.

https://t.me/learningnets

4-58 About Objects

Objects are entities on which actions are performed. Objects include data centers, folders, clusters, hosts, datastores, networks, and virtual machines.

All objects have a **Permissions** tab. The **Permissions** tab shows which user or group and role are associated with the selected object.

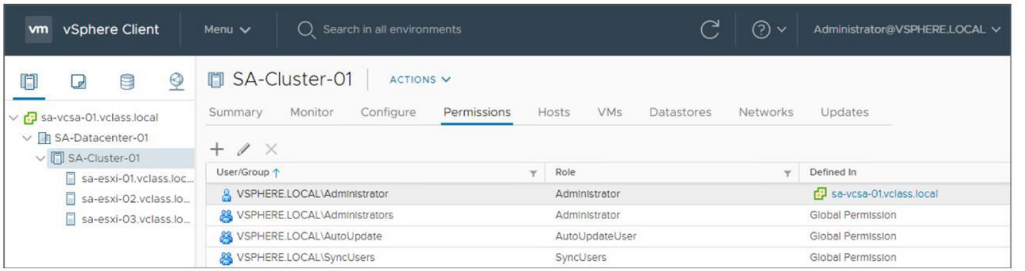

https://t.me/learningnets

4-59 Adding Permissions to the vCenter Server Inventory

To add a permission:

1. Select an object.
2. Select a user or group from a domain.
3. Select a role.
4. Propagate the permission to the child objects.

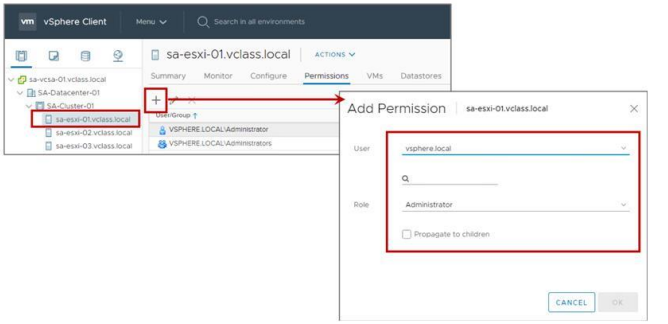

You can assign permissions to objects at different levels of the hierarchy. For example, you can assign permissions to a host object or to a folder object that includes all host objects. You can also assign permissions to a global root object to apply the permissions to all objects in all solutions.

For information about hierarchical inheritance of permissions and global permissions, see *vSphere Security* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-52188148-C579-4F6A-8335-CFBCE0DD2167.html

https://t.me/learningnets

4-60 Viewing Roles and User Assignments

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

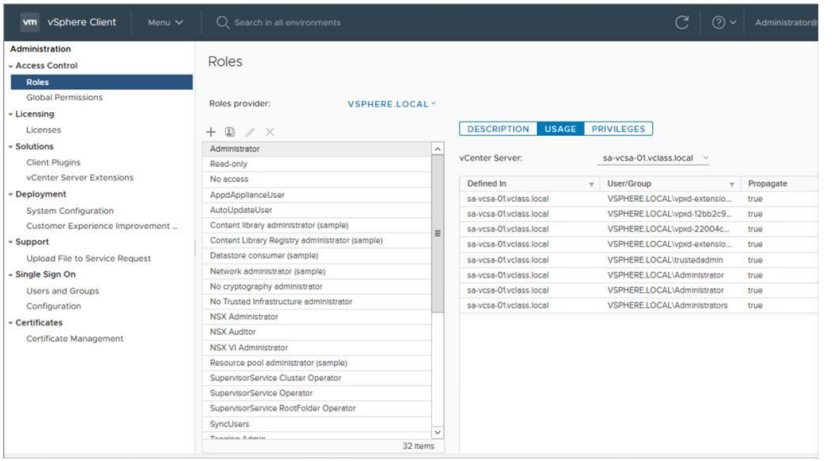

You can view all the objects to which a role is assigned and all the users or groups who are granted the role.

To view information about a role, click **Usage** in the Roles pane and select a role from the Roles list. The information provided to the right shows each object to which the role is assigned and the users and groups who were granted the role.

https://t.me/learningnets

4-61 Applying Permissions: Scenario 1

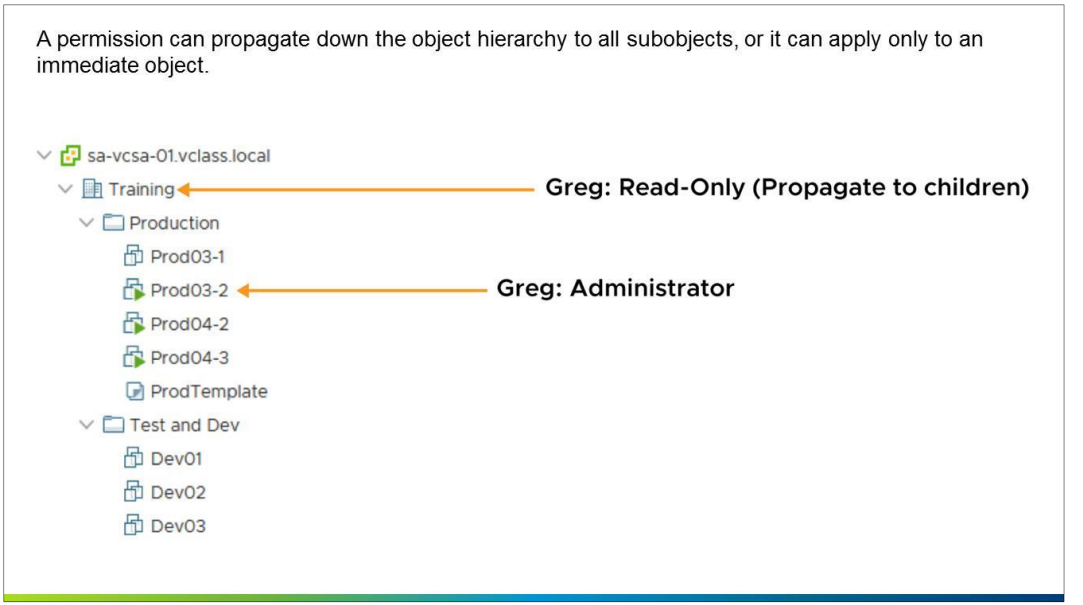

In addition to specifying whether permissions propagate downward, you can override permissions set at a higher level by explicitly setting different permissions for a lower-level object.

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

https://t.me/learningnets

4-62 Applying Permissions: Scenario 2

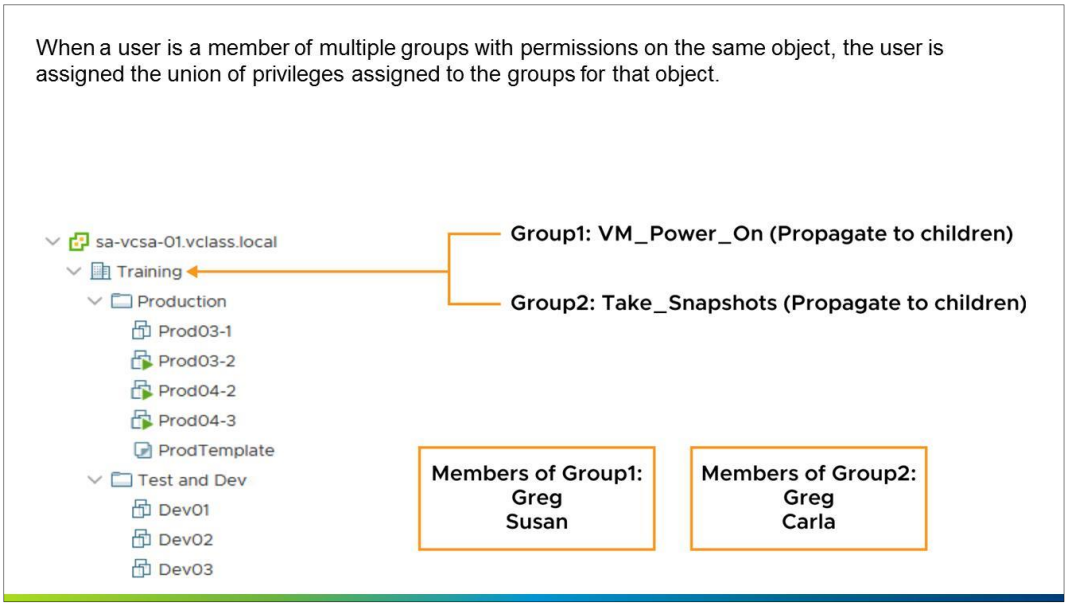

On the slide, Group1 is assigned the VM_Power_On role, a custom role that contains only one privilege: the ability to power on a VM. Group2 is assigned the Take_Snapshots role, another custom role that contains the privileges to create and remove snapshots. Both roles propagate to the child objects.

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

https://t.me/learningnets

4-63 Activity: Applying Group Permissions (1)

If Group1 has the Administrator role and Group2 has the No Access role, what permissions does Greg have?

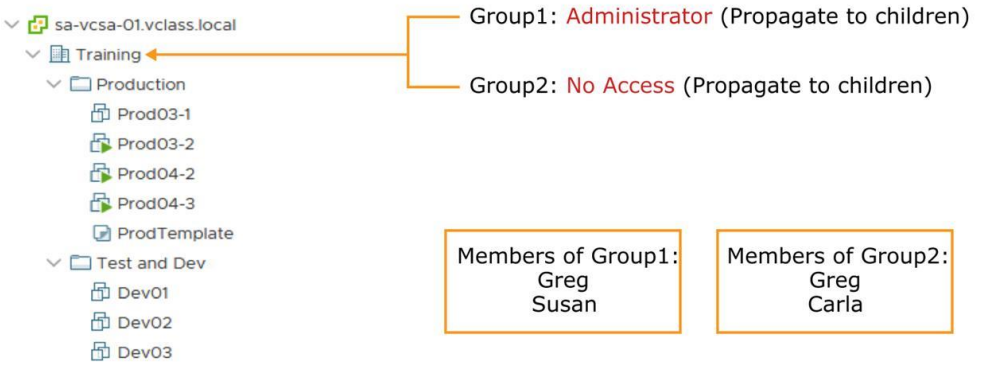

https://t.me/learningnets

4-64 Activity: Applying Group Permissions (2)

Greg has Administrator privileges.
Greg is assigned the union of privileges assigned to Group1 and Group2.

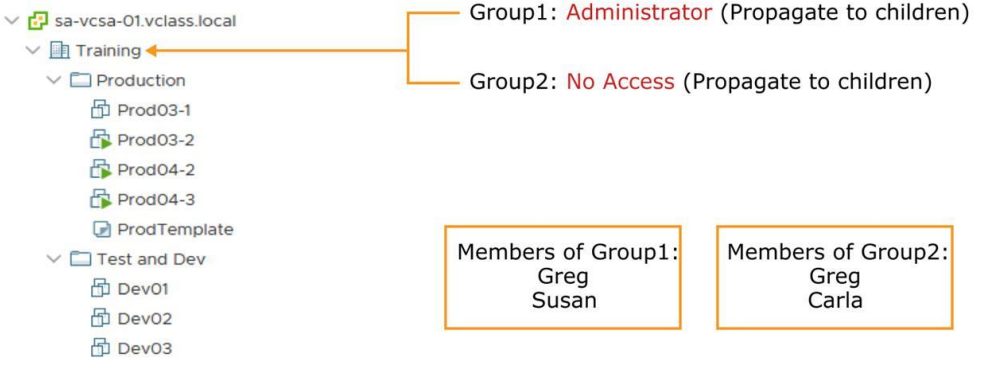

https://t.me/learningnets

4-65 Applying Permissions: Scenario 3

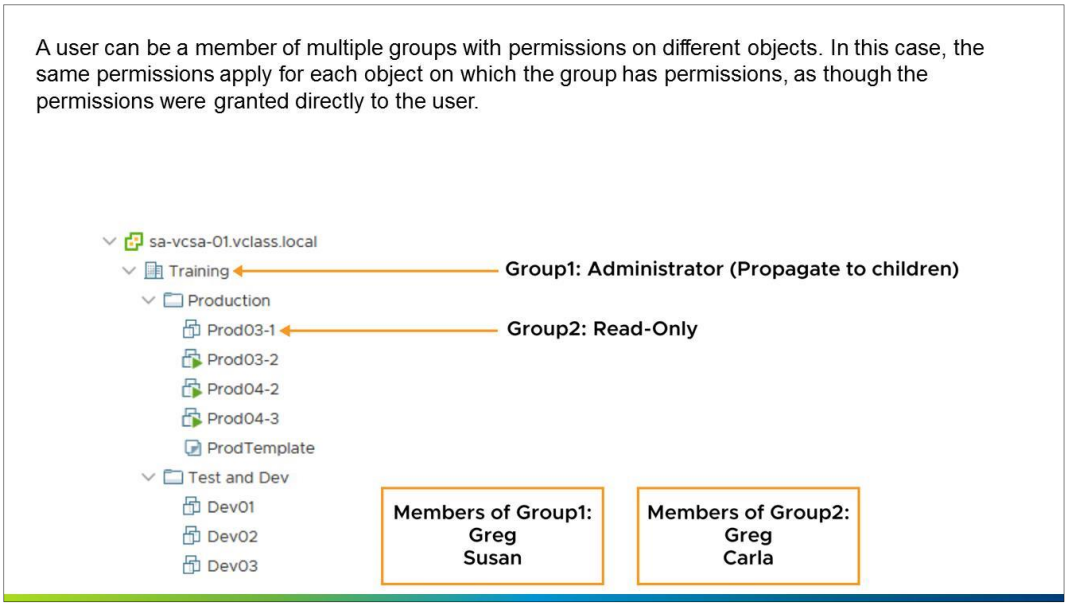

You can override permissions set for a higher-level object by explicitly setting different permissions for a lower-level object.

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

Because Greg is a member of both Group1 and Group2, he gets administrator privileges on the entire Training data center (the higher-level object), except for the VM called Prod03-1 (the lower-level object). For this VM, he gets read-only access.

https://t.me/learningnets

4-66 Applying Permissions: Scenario 4

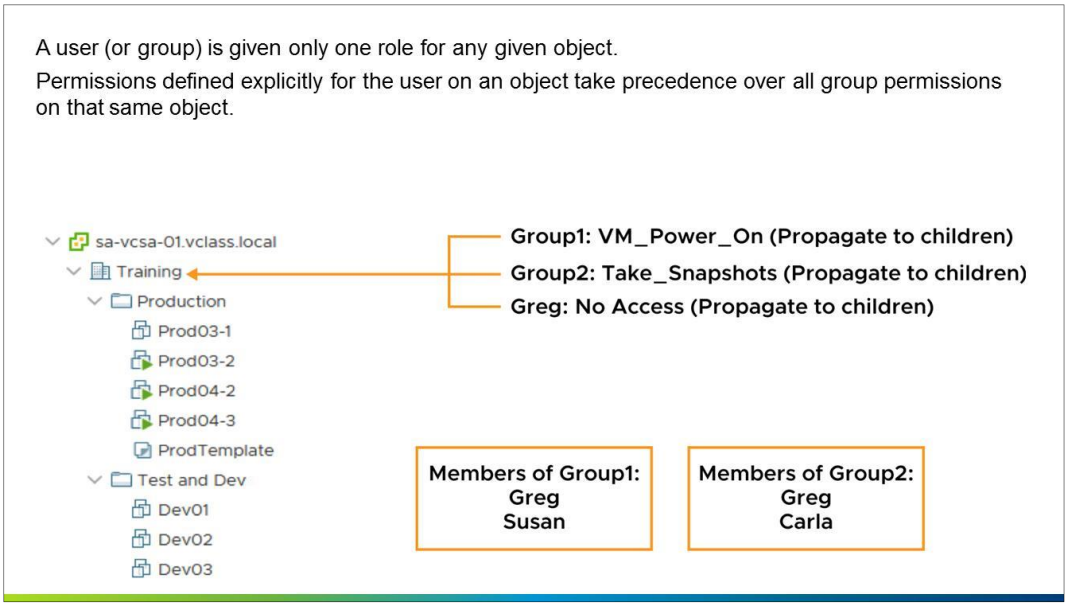

On the slide, three permissions are assigned to the Training data center:

- Group1 is assigned the VM_Power_On role.
- Group2 is assigned the Take_Snapshots role.
- Greg is assigned the No Access role.

Greg is a member of both Group1 and Group2. Assume that propagation to child objects is enabled on all roles. Although Greg is a member of both Group1 and Group2, he gets the No Access privilege to the Training data center and all objects under it. Greg gets the No Access privilege because explicit user permissions on an object take precedence over all group permissions on that same object.

https://t.me/learningnets

4-67 Creating a Role

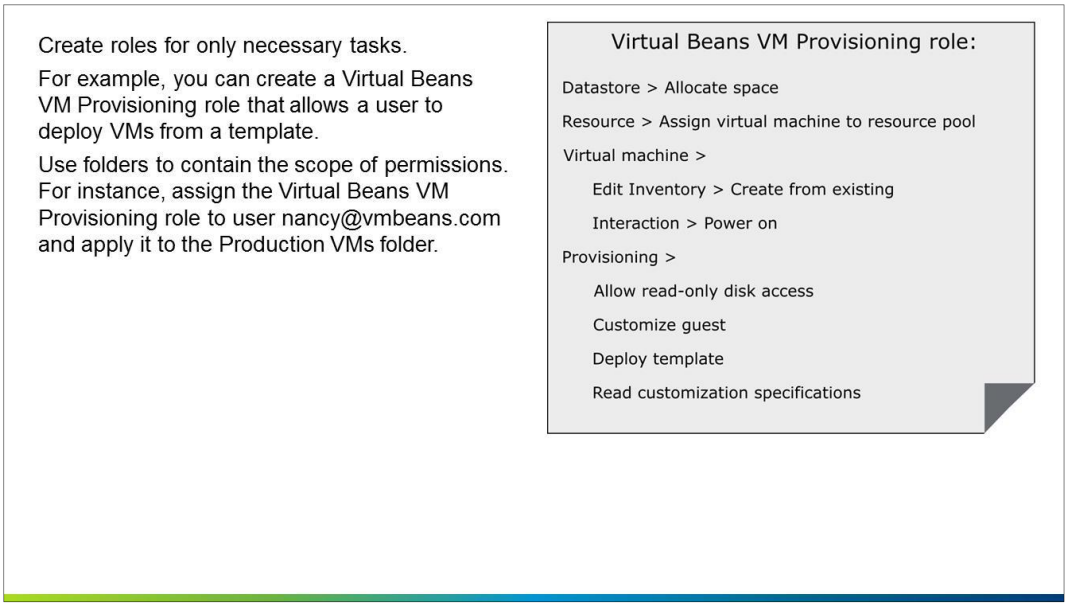

The Virtual Beans VM Provisioning role is one of many examples of roles that can be created.

Define a role using the smallest number of privileges possible to maximize security and control over your environment. Give the roles names that explicitly indicate what each role allows, to make its purpose clear.

https://t.me/learningnets

4-68 About Global Permissions

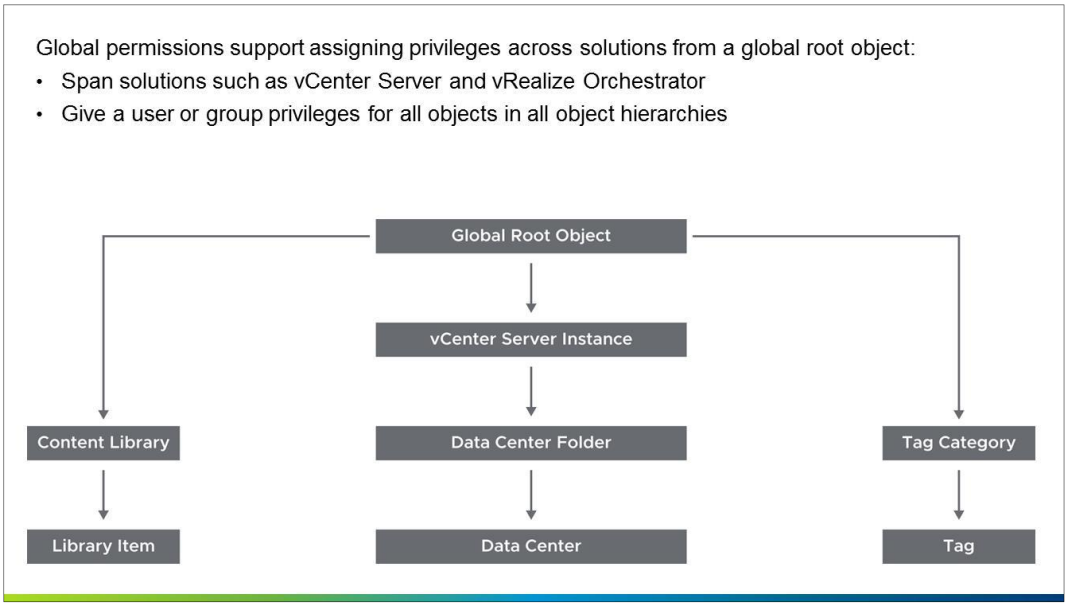

Often, you apply a permission to a vCenter Server inventory object such as an ESXi host or a VM. When you apply a permission, you specify that a user or group has a set of privileges, called a role, on the object.

Global permissions give a user or group privileges to view or manage all objects in each of the inventory hierarchies in your deployment. The example on the slide shows that the global root object has permissions over all vCenter Server objects, including content libraries, vCenter Server instances, and tags. Global permissions allow access across vCenter Server instances. vCenter Server permissions, however, are effective only on objects in a particular vCenter Server instance.

https://t.me/learningnets

Labs

Lab: Configuring Active Directory: Adding an Identity Source
Lab: Users, Groups, and Permissions

https://t.me/learningnets

4-70 Lab 9: Configuring Active Directory: Adding an Identity Source

Add vclass.local as an identity source:

1. Add vclass.local as an Identity Source

https://t.me/learningnets

4-71 Lab 10: Users, Groups, and Permissions

Assign roles and permissions so that an Active Directory user can perform functions in vCenter Server:

1. View Active Directory Users
2. Assign Object Permission to an Active Directory User
3. Assign Root-Level Global Permission to an Active Directory User
4. Log In as an Active Directory User
5. Use an Active Directory User to Create a Virtual Machine

https://t.me/learningnets

4-72 Review of Learner Objectives

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

- Define the term permission in the context of vCenter Server
- Describe the rules for applying permissions
- Create a custom role
- Create a permission

https://t.me/learningnets

Lesson 6: Backing Up and Restoring vCenter Server Appliance

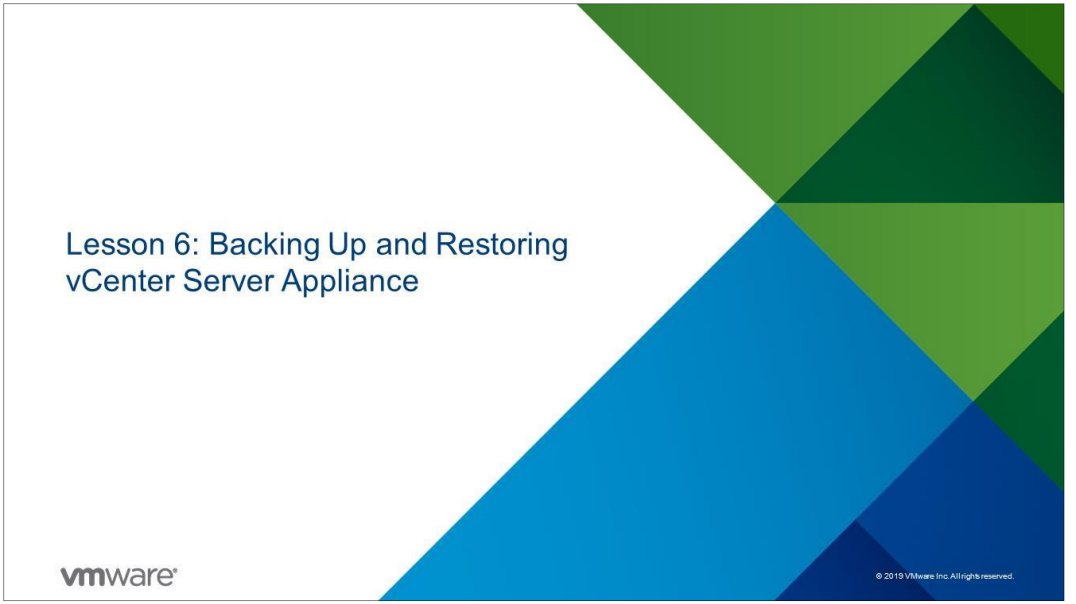

https://t.me/learningnets

4-74 Learner Objectives

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

- Create a vCenter Server Appliance backup schedule
- Restore vCenter Server Appliance from backup

https://t.me/learningnets

4-75 Virtual Beans: vCenter Server Operations

As a Virtual Beans administrator, you are responsible for the maintenance and daily operation of vCenter Server.
To align with Virtual Beans policies, you perform the following tasks:

- Back up vCenter Server data monthly.
- Make vCenter Server highly available:
  - 99.99 percent available (downtime per year of 52.56 minutes)
- Monitor vCenter Server performance to avoid potential problems in the infrastructure.

https://t.me/learningnets

4-76 About vCenter Server Backup and Restore

vCenter Server backup and restore operations protect data. These operations work in the following ways:

- Use the vCenter Server Appliance Management Interface
- Remove dependency on third-party backup solutions
- Support FTPS, HTTP, HTTPS, SFTP, FTP, NFS, and SMB
- Include the option to encrypt backups
- Restore directly from the vCenter Server Appliance GUI installer
- Restore a vCenter Server instance to a brand-new appliance

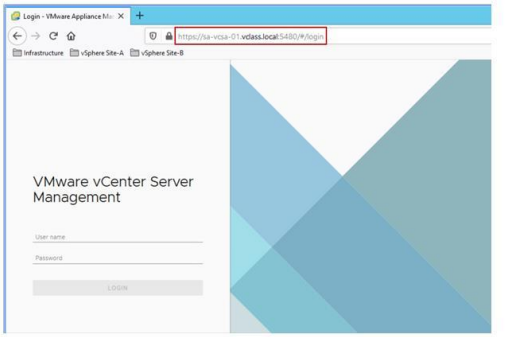

The vCenter Server Appliance Management Interface supports backing up key parts of the appliance. You can protect vCenter Server data and minimize the time required to restore data center operations.

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

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

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

https://t.me/learningnets

4-77 Methods for vCenter Server Appliance Backup and Restore

You can use different methods to back up and restore vCenter Server Appliance:

- File-based backup and restore:
  - Use the vCenter Server Appliance Management Interface to create a file-based backup.
  - Restore the backup through the GUI installer of the appliance.
  - Schedule the file-based backup and restore.
- Image-based backup and restore:
  - Use vSphere Storage APIs - Data Protection with a third-party backup product to perform centralized, efficient, off-host, LAN-free backups.

https://t.me/learningnets

4-78 File-Based Backup of vCenter Server Appliance

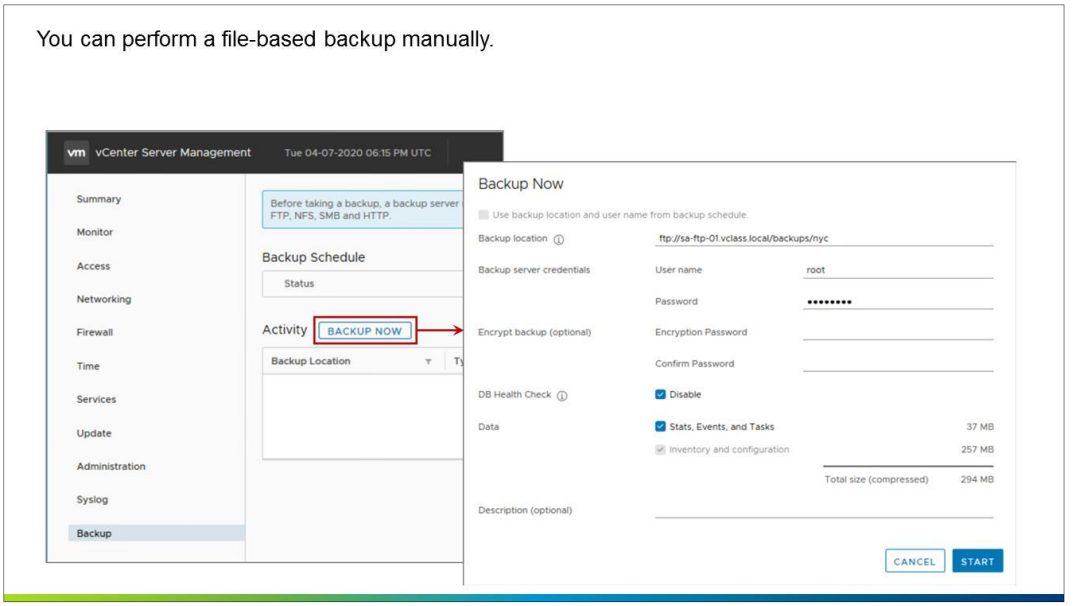

You use the vCenter Server Appliance Management Interface to perform a file-based backup of the vCenter Server core configuration, inventory, and historical data of your choice. The backed-up data is streamed over the selected protocol to a remote system. The backup is not stored on vCenter Server Appliance.

When specifying the backup location, use the following syntax: `protocol:<server-address<:port-number>/folder/subfolder`.

https://t.me/learningnets

4-79 File-Based Restore of vCenter Server Appliance

Use the vCenter Server Appliance GUI installer to restore a vCenter Server Appliance to an ESXi host or a vCenter Server instance.
The restore procedure occurs in stages:

1. A new appliance is deployed.
2. The newly deployed vCenter Server Appliance is populated with the data stored in the file-based backup.

When you use the file-based restore method, reconciliation is automatically performed.

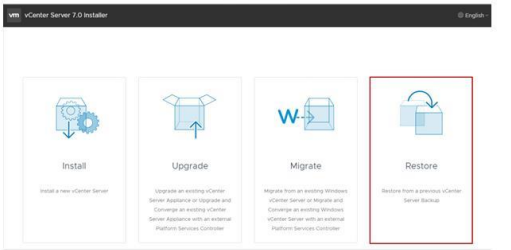

You can perform a file-based restore only for a vCenter Server Appliance instance that you previously backed up by using the vCenter Server Appliance Management Interface. You can perform the restore operation by using the GUI installer of vCenter Server Appliance. The process consists of deploying a new vCenter Server Appliance instance and copying the data from the file-based backup to the new appliance.

You can also perform a restore operation by deploying a new vCenter Server Appliance instance and using the vCenter Server Appliance Management Interface to copy the data from the file-based backup to the new appliance.

https://t.me/learningnets

4-80 Scheduling Backups

You can schedule automatic file-based backups.
The backup scheduler supports:

- A retention policy to keep all backups or a defined number of backups
- Daily, weekly, or custom schedule

Failed backups trigger an alarm in the vSphere Client.

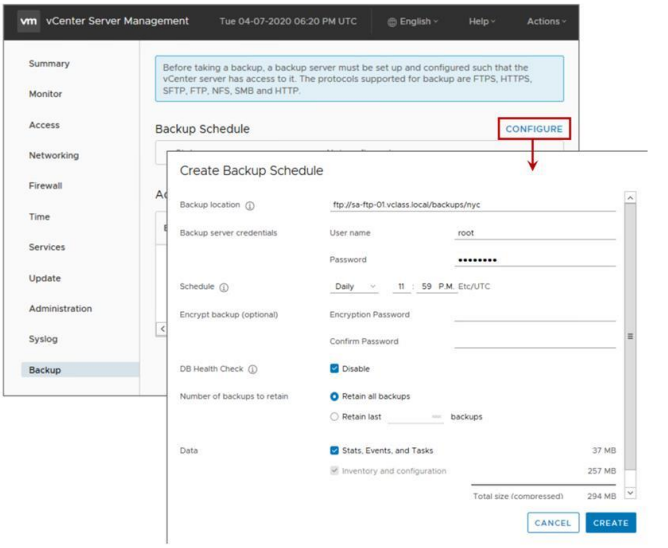

You can set up a file-based backup schedule to perform periodic backups:

- The schedule can be set up with information about the backup location, recurrence, and retention for the backups.
- You can set up only one schedule at a time.

https://t.me/learningnets

4-81 Viewing the Backup Schedule

You can view the existing defined backup schedule from the vCenter Server Appliance Management Interface.

The backup schedule can be edited, disabled, or deleted.

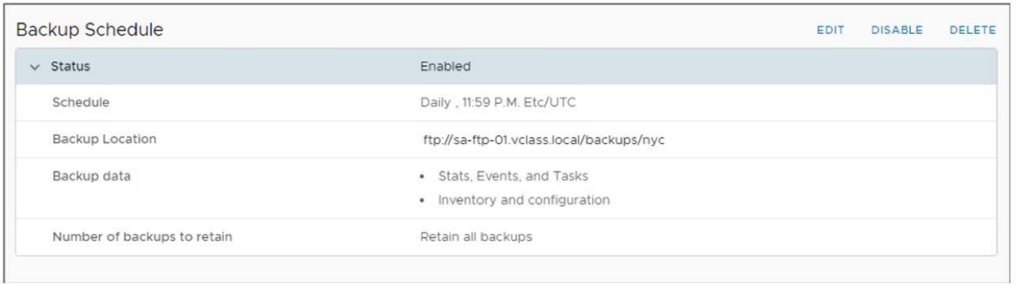

https://t.me/learningnets

4-82 Demonstration: Backing Up and Restoring a vCenter Server Appliance Instance

Your instructor will run a demonstration.

https://t.me/learningnets

4-83 Review of Learner Objectives

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

- Create a vCenter Server Appliance backup schedule
- Restore vCenter Server Appliance from backup

https://t.me/learningnets

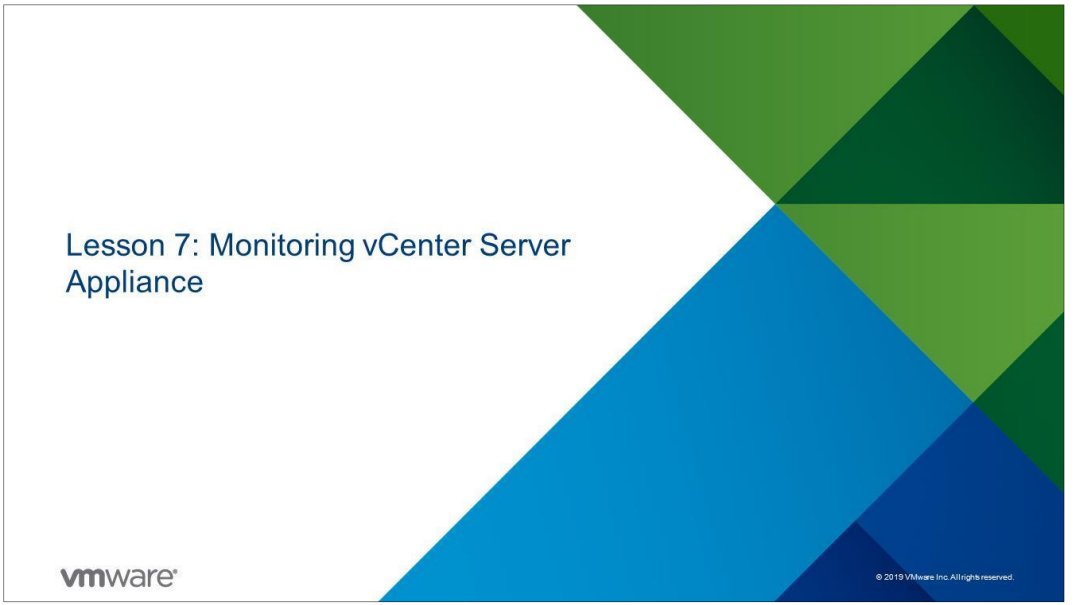

https://t.me/learningnets

4-85 Learner Objectives

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

- View vCenter Server logs and events
- Manage vCenter Server services
- Monitor vCenter Server Appliance for service and disk space usage
- Use vSphere alarms for resource exhaustion and service failures

https://t.me/learningnets

4-86 vCenter Server Events

The vCenter Server events and audit trails allow selectable retention periods in increments of 30 days:

- User-action information includes the user's account and specific event details.
- All actions are reported, including file ID, file path, source of operation, operation name, and date and time of operation.
- Events and alarms are displayed to alert the user to changes in the vCenter Server service health or when a service fails.

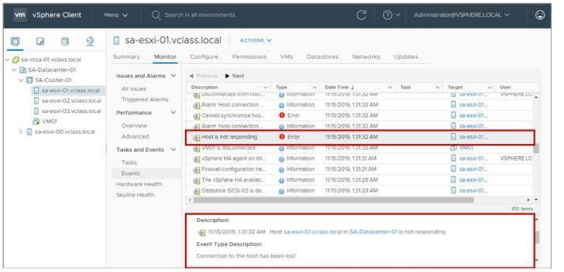

https://t.me/learningnets

4-87 About Log Levels

You can set log levels to control the quantity and type of information logged.
Examples of when to set log levels:

- When troubleshooting complex issues, set the log level to verbose or trivia. Troubleshoot and set it back to info.
- For controlling the amount of information being stored in the log files.

Option	Description
None	Turns off logging
Error (errors only)	Displays only error log entries
Warning (errors and warnings)	Displays warning and error log entries
Info (normal logging)	Displays information, error, and warning log entries
Verbose	Displays information, error, warning, and verbose log entries
Trivia (extended verbose)	Displays information, error, warning, verbose, and trivia log entries

Changes to the logging settings take effect immediately. You do not have to restart the vCenter Server system.

https://t.me/learningnets

4-88 Setting Log Levels

You can configure the amount of detail that vCenter Server collects in log files:

- You can edit the log levels in the vSphere Client.
- More verbose logging requires more space on your vCenter Server system.

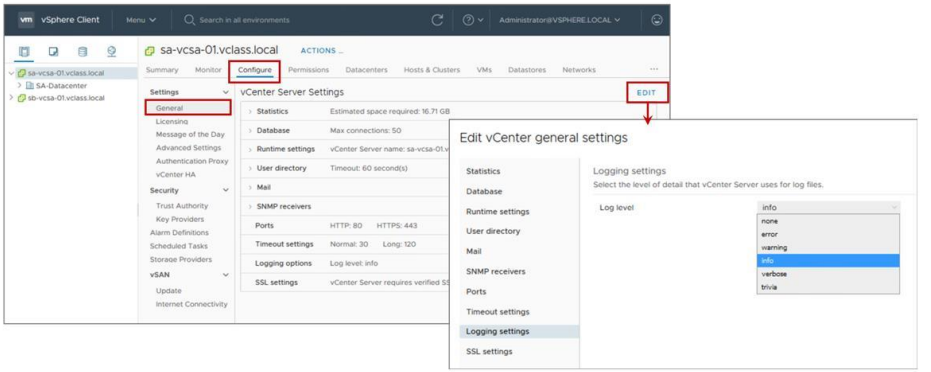

To configure logging levels, follow these steps:

1. In the vSphere Client, select the vCenter Server instance in the navigation pane.
2. Click the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **EDIT**.
5. Under Edit vCenter general settings, select **Logging settings** in the left pane.
6. Select an option from the **Log level** drop-down menu.

https://t.me/learningnets

4-89 Forwarding vCenter Server Appliance Log Files to a Remote Host

vCenter Server and ESXi can stream their log information to a remote Syslog server:

- You can enable this feature in the vCenter Server Appliance Management Interface.
- With this feature, you can further analyze vCenter Server Appliance log files with log analysis products, such as vRealize Log Insight.

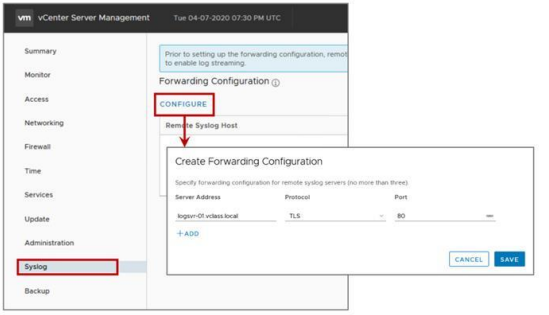

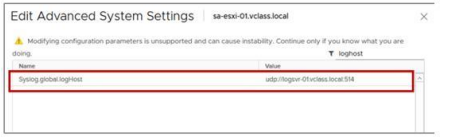

https://t.me/learningnets

4-90 vCenter Server Database Health

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

- By default, database health warnings trigger an alarm when the space used reaches 80 percent.
- The alarm changes from warning to error when the space used reaches 95 percent.
- vCenter Server services shut down so that you can configure more disk space or remove unwanted content.

You can also monitor database space utilization using the vCenter Server Appliance Management Interface.

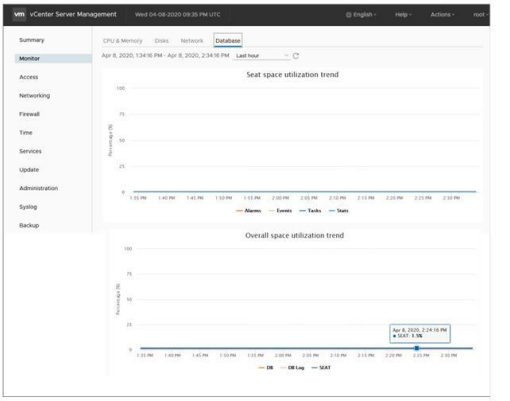

https://t.me/learningnets

4-91 Monitoring vCenter Server Appliance

The vCenter Server Appliance Management Interface has a built-in monitoring interface.

CPU and Memory View

Disks View

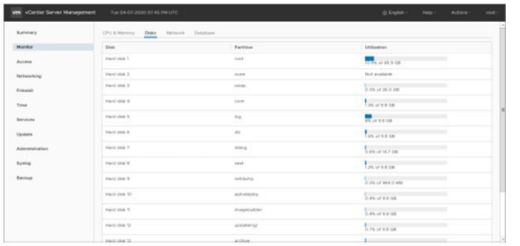

The CPU and Memory views provide a historical view of CPU and memory use.

Using the Disks view, you can monitor the available disk space.

https://t.me/learningnets

4-92 Monitoring vCenter Server Appliance Services

You can use the vCenter Server Appliance Management Interface to monitor the health and state of the vCenter Server Appliance services. You can restart, start, or stop services from this interface.

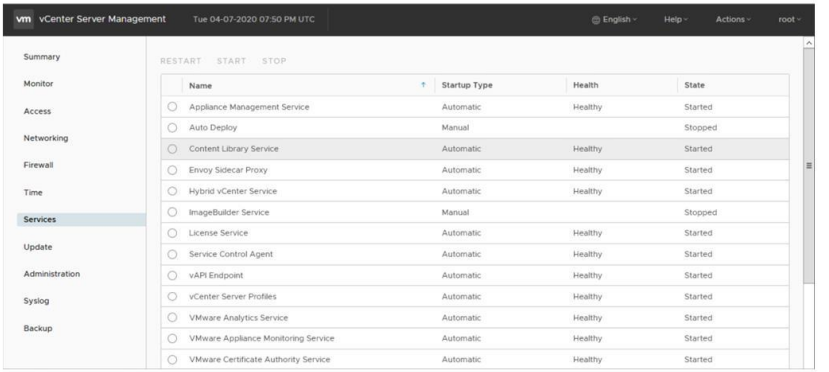

https://t.me/learningnets

4-93 Monthly Patch Updates for vCenter Server Appliance

VMware provides monthly security patches for vCenter Server Appliance:

- Critical vulnerability patches are delivered on a monthly release cycle.
- Important and low vulnerabilities are delivered with the next available vCenter Server patch or update.

You can configure the vCenter Server Appliance to perform automatic checks for available patches in the configured repository URL at a regular interval.

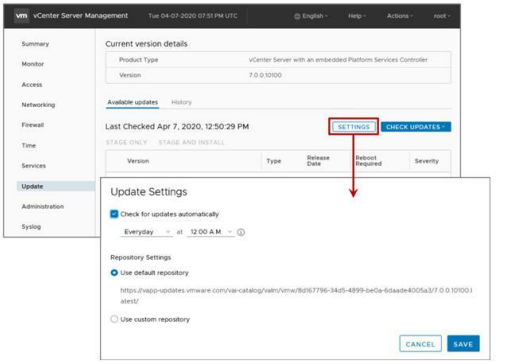

If a vCenter Server patch or update occurs in the same time period as the monthly security patch, the monthly security patch is rolled into the vCenter Server patch or update.

https://t.me/learningnets

4-94 Review of Learner Objectives

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

- View vCenter Server logs and events
- Manage vCenter Server services
- Monitor vCenter Server Appliance for service and disk space usage
- Use vSphere alarms for resource exhaustion and service failures

https://t.me/learningnets

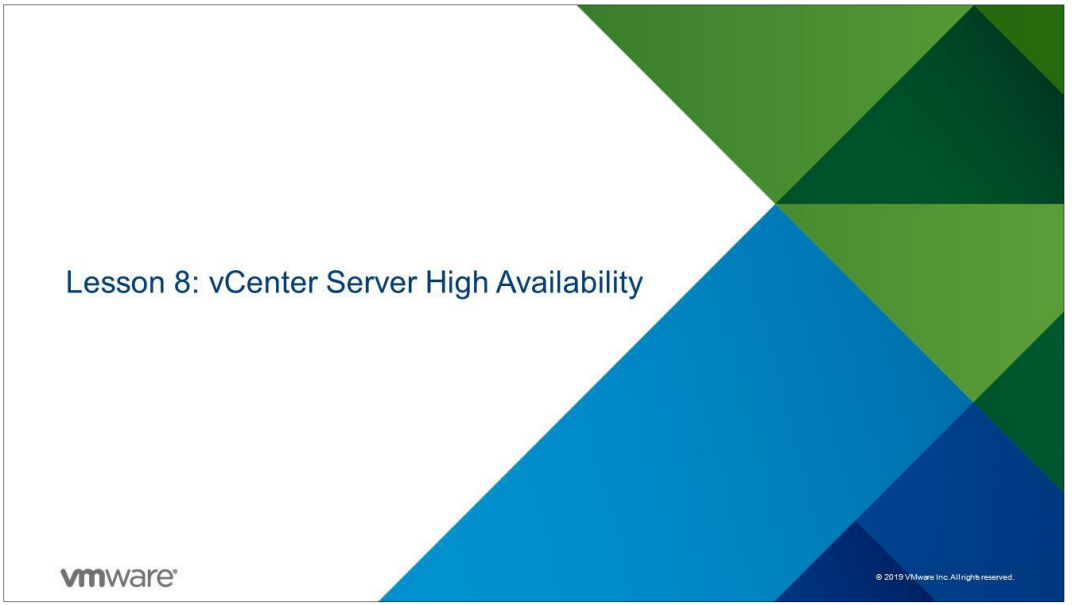

https://t.me/learningnets

4-96 Learner Objectives

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

- Recognize the importance of vCenter Server High Availability
- Explain how vCenter Server High Availability works
- Identify vCenter Server High Availability requirements

https://t.me/learningnets

4-97 Importance of Keeping vCenter Server Highly Available

High availability is an important characteristic for many VMware and third-party solutions that depend on vCenter Server as the primary management platform:

- vCenter Server is the foundation for the virtual desktop infrastructure.
- Backup and disaster recovery solutions rely on vCenter Server.

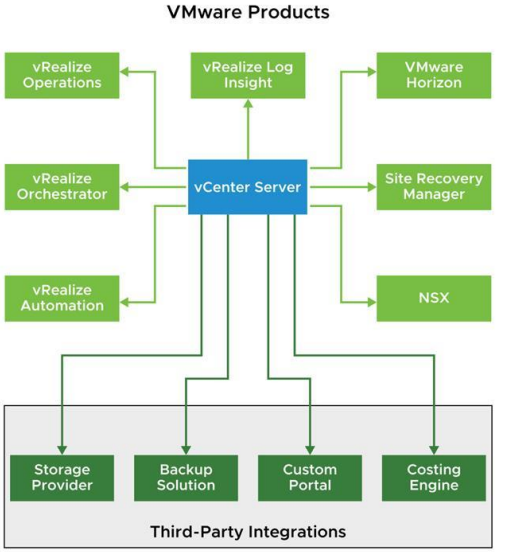

vSphere is a virtualization platform that forms the foundation for building and managing an organization's virtual, public, and private cloud infrastructures. vCenter Server Appliance sits at the heart of vSphere and provides services to manage various components of a virtual infrastructure, such as ESXi hosts, virtual machines, and storage and networking resources. As large virtual infrastructures are built using vSphere, vCenter Server becomes an important element in ensuring the business continuity of an organization. vCenter Server must protect itself from a set of hardware and software failures in an environment and must recover transparently from such failures.

https://t.me/learningnets

4-98 About vCenter Server High Availability

vCenter Server High Availability protects vCenter Server Appliance against both hardware and software failures.
vCenter Server High Availability forms a cluster of nodes:

- Active node: Runs the active vCenter Server Appliance instance
- Passive node: Automatically takes over the role of the Active node if a failure occurs
- Witness node: Provides a quorum to protect against a split-brain situation

vCenter Server High Availability is built in to vCenter Server Appliance and is included with the standard license.

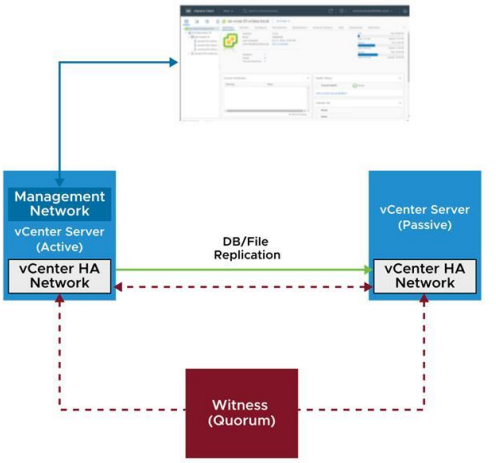

With vCenter Server High Availability, you can recover quickly from a vCenter Server failure. Using automated failover, vCenter Server failover occurs with minimal downtime.

https://t.me/learningnets

4-99 Scenario: Active Node Failure

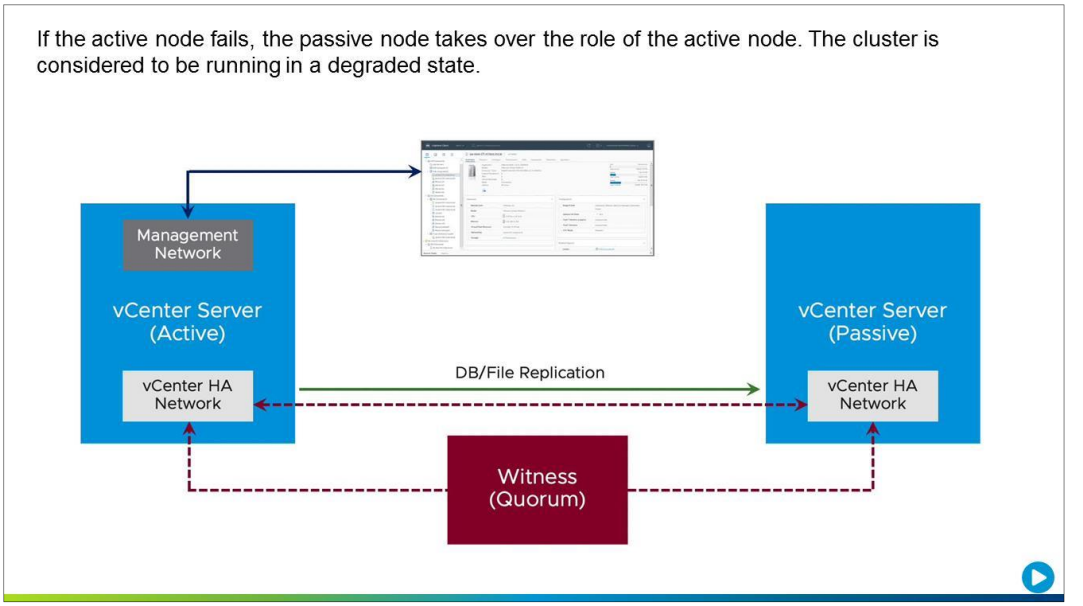

The animation demonstrates what happens if an active node fails. To play the animation, go to https://vmware.bravais.com/s/PlUBZn2zCO7HE5qN2fm4.

The active node runs the active instance of vCenter Server Appliance. The node uses an IP address on the Management network for the vSphere Client to connect to.

If the active node fails (because of a hardware, software, or network failure), the passive node takes over the role of the active node. The IP address to which the vSphere Client was connected is switched from the failed node to the new active node. The new active node starts serving client requests. Meanwhile, the user must log back in to the vSphere Client for continued access to vCenter Server.

Because only two nodes are up and running, the vCenter Server High Availability cluster is considered to be running in a degraded state and subsequent failover cannot occur. A subsequent failure in a degraded cluster means vCenter Server services are no longer available. A passive node is required to return the cluster to a healthy state.

https://t.me/learningnets

4-100 Scenario: Passive Node Failure

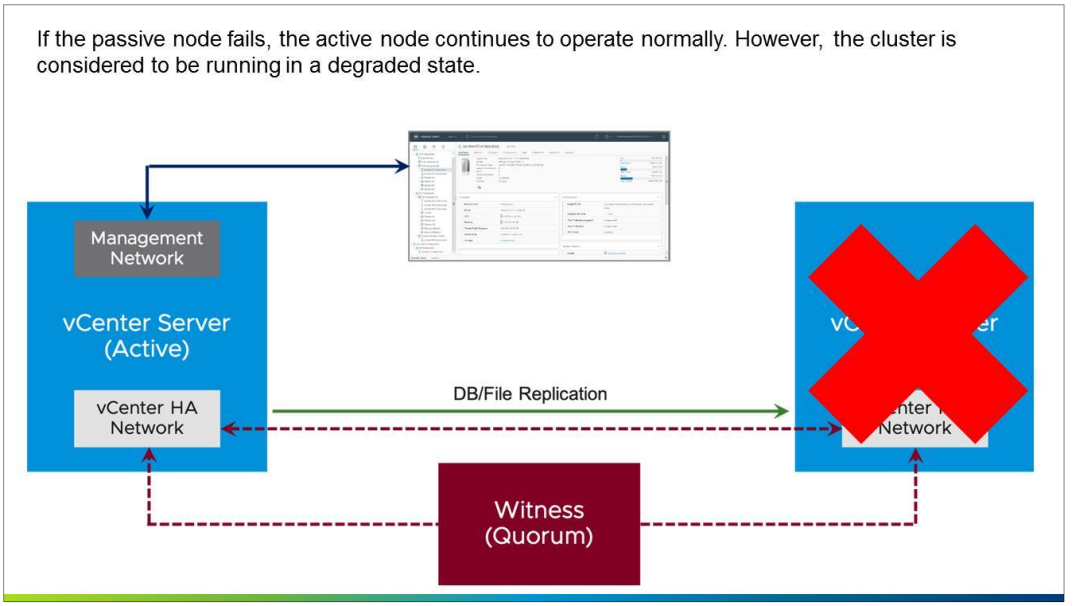

If the passive node fails, the active node continues to operate as normal. Because no disruption in service occurs, users can continue to access the active node using the vSphere Client.

Because the passive node is down, the active node is no longer protected. The cluster is considered to be running in a degraded state because only two nodes are up and running. A subsequent failure in a degraded cluster means vCenter Server services are no longer available. A passive node is required to return the cluster to a healthy state.

https://t.me/learningnets

4-101 Scenario: Witness Node Failure

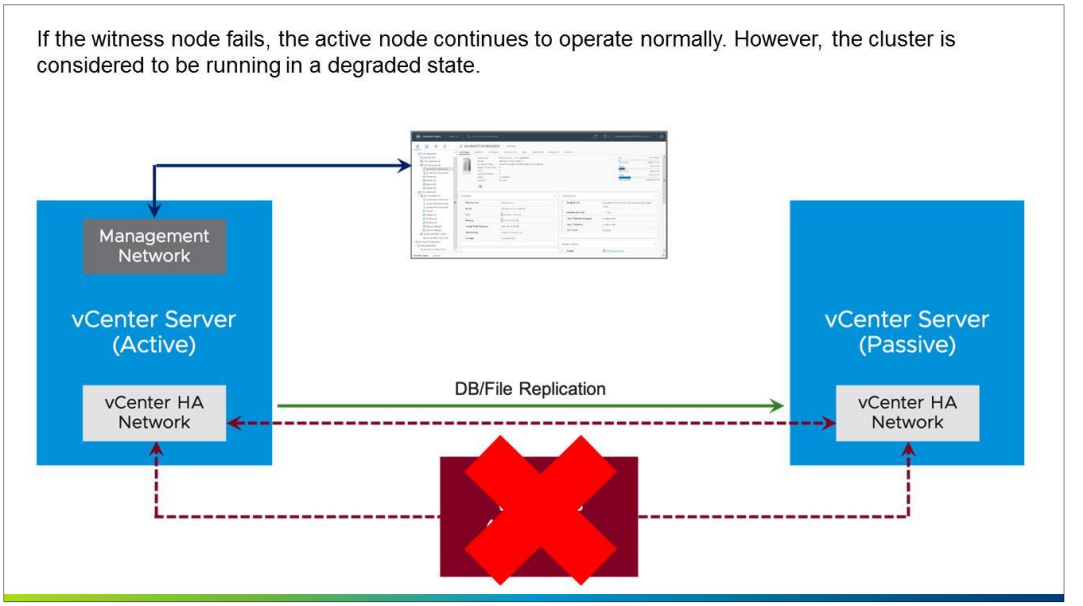

The witness node is used to maintain quorum.

If the witness node fails, the active node continues to operate without disruption in service. Because only two nodes are up and running, the cluster is considered to be running in a degraded state and failover cannot occur. A subsequent failure in a degraded cluster means vCenter Server services are no longer available. The witness node is required to return the cluster to a healthy state.

https://t.me/learningnets

4-102 Benefits of vCenter Server High Availability

vCenter Server High Availability provides many benefits:

- vCenter Server Appliance is made more resilient.
- Protection against hardware, host, and application failures is provided.
- Recovery occurs in minutes. End-to-end downtime is minimized.
- Active-passive architecture provides transparent failover:
  - Recovery point objective: No data loss.
  - Recovery time objective: Within minutes.
- One-click automated high availability is set up.
- Nodes can be geographically distant (less than 10 milliseconds latency).
- No shared storage is required.
- No third-party technologies are required.

https://t.me/learningnets

4-103 vCenter Server High Availability Requirements

Component	Requirements
ESXi	• Version 6.0 or later. • Minimum of three ESXi hosts is recommended.
vCenter Server Appliance	• Version 6.5 or later. • Deployment size small or larger is required to meet the RTO. • Enough disk space to collect and store support bundles for all three nodes on the active node.
Network connectivity	• Network latency between the three nodes must be less than 10 milliseconds. • The vCenter HA network must be on a different subnet than the management network.
Licensing	• A single vCenter Server Standard license.

For more information about the vCenter Server High Availability requirements, see *vSphere Availability* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-63F459B7-8884-4818-8872-C9753B2E0215.html.

https://t.me/learningnets

4-104 Demonstration: Configuring vCenter Server High Availability

Your instructor will run a demonstration.

https://t.me/learningnets

4-105 Review of Learner Objectives

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

- Recognize the importance of vCenter Server High Availability
- Explain how vCenter Server High Availability works
- Identify vCenter Server High Availability requirements

https://t.me/learningnets

4-106 Virtual Beans: vCenter Server Maintenance and Operations

As a Virtual Beans administrator, you plan to maintain vCenter Server and keep it up and running.

Virtual Beans Requirement	Plan
Back up vCenter Server data monthly.	• Use the VAMI to schedule monthly backups of vCenter Server. • Because the NFS protocol is supported, you can use one of the NFS filesystems in your data center to store the backups.
Make vCenter Server highly available.	• Configure vCenter Server High Availability to protect against vCenter Server failures. • Because recovery time is minimal, you are confident that 99.99 percent availability (downtime per year of 52.56 minutes) can be achieved.
Monitor vCenter Server regularly.	• Use the vSphere Client and VAMI daily to monitor vCenter Server health and performance.

https://t.me/learningnets

4-107 Key Points

- vCenter Server Appliance uses the Photon operating system and the PostgreSQL database.
- You use the vSphere Client to connect to vCenter Server instances and manage vCenter Server inventory objects.
- A permission, defined in vCenter Server, gives one user or group a role (set of privileges) for a selected object.
- You can use the vCenter Server Appliance Management Interface to monitor appliance resource use and perform a file-based backup of the appliance.
- vCenter Server High Availability is built in to vCenter Server Appliance and protects the appliance from both hardware and software failures.

Questions?

https://t.me/learningnets

Module 5

Configuring and Managing Virtual Networks

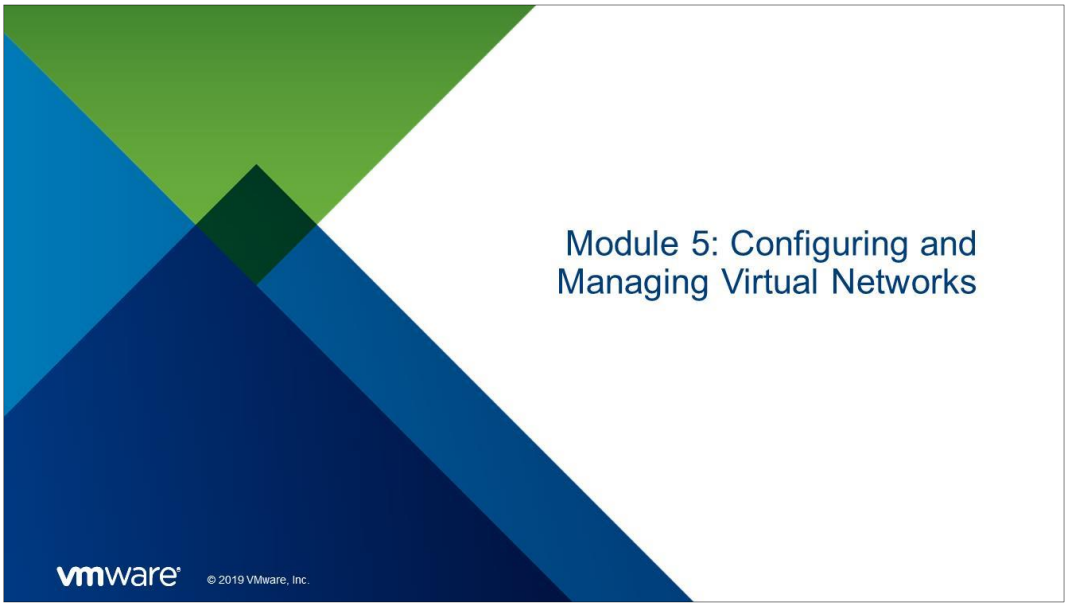

https://t.me/learningnets

5-2 Importance

When you configure ESXi networking properly, virtual machines can communicate with other virtual, and physical, machines. In this way, remote host management and IP-based storage operate effectively.

https://t.me/learningnets

5-3 Module Lessons

1. Introduction to vSphere Standard Switches
2. Configuring Standard Switch Policies

https://t.me/learningnets

5-4 Virtual Beans: Networking Requirements

Virtual Beans has the following requirements for its network infrastructure:

- Use the existing VLAN infrastructure and create VLANs as needed for the vSphere environment.
- Use the available bandwidth efficiently:
  - Infrastructure services must get enough bandwidth.
  - Infrastructure traffic should not interfere with the performance of business-critical and nonbusiness-critical application traffic.
- Avoid single points of failure.

As the Virtual Beans administrator, you must configure vSphere networking to meet these requirements.

https://t.me/learningnets

https://t.me/learningnets

5-6 Learner Objectives

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

- Identify virtual switch connection types
- Configure and view standard switch configurations
- Distinguish between the features of standard and distributed switches

https://t.me/learningnets

5-7 About Virtual Switches

Virtual switches connect VMs to the physical network.

They provide connectivity between VMs on the same ESXi host or on different ESXi hosts.

They also support VMkernel services, such as vSphere vMotion migration, iSCSI, NFS, and access to the management network.

https://t.me/learningnets

5-8 Types of Virtual Switch Connections

A virtual switch has specific connection types:

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

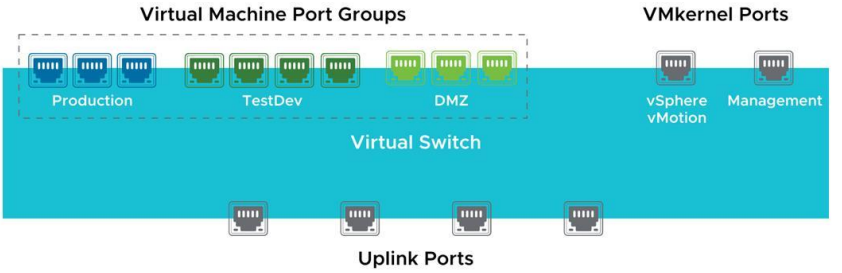

The ESXi management network port is a VMkernel port that connects to network or remote services, including vpxd on vCenter Server and VMware Host Client.

Each ESXi management network port and each VMkernel port must be configured with its own IP address, netmask, and gateway.

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

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

https://t.me/learningnets

5-9 Virtual Switch Connection Examples

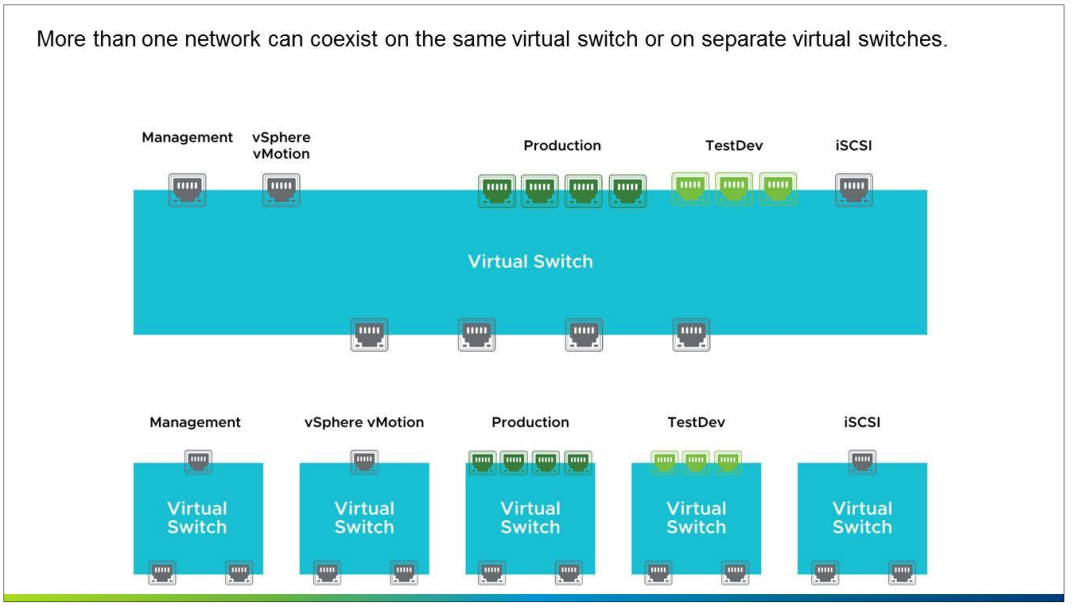

When you design your networking environment, you can team all your networks on a single virtual switch. Alternatively, you can opt for multiple virtual switches, each with a separate network. The decision partly depends on the layout of your physical networks.

For example, you might not have enough network adapters to create a separate virtual switch for each network. Instead, you might place your network adapters in a single virtual switch and isolate the networks by using VLANs.

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

https://t.me/learningnets

5-10 About VLANs

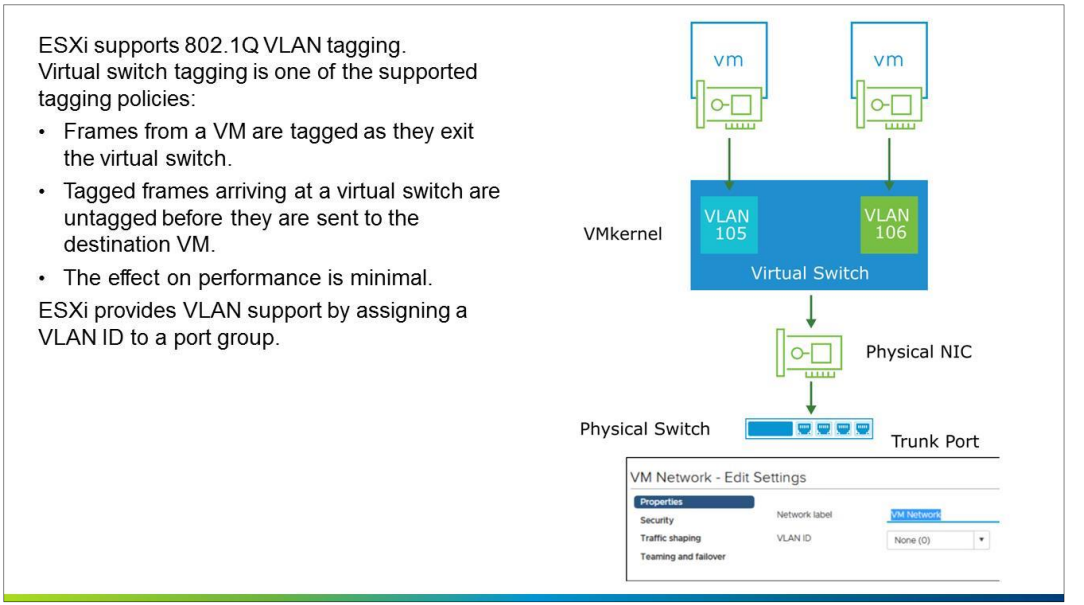

VLANs provide for logical groupings of switch ports. All virtual machines or ports in a VLAN communicate as if they are on the same physical LAN segment. A VLAN is a software-configured broadcast domain. Using a VLAN provides the following benefits:

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

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

The port on a physical switch to which an ESXi host is connected must be defined as a static trunk port. A trunk port is a port on a physical Ethernet switch that is configured to send and receive

https://t.me/learningnets

packets tagged with a VLAN ID. No VLAN configuration is required in the VM. In fact, the VM does not know that it is connected to a VLAN.

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

https://t.me/learningnets

5-11 Types of Virtual Switches

A virtual network supports standard and distributed switches. Both switch types are elastic: Ports are created and removed automatically.

- Standard switch:
  - Virtual switch that is configured for a single host.
- Distributed switch:
  - Virtual switch that is configured for an entire data center.
  - Up to 2,000 hosts can be attached to the same distributed switch.
  - The configuration is consistent across all attached hosts.
  - Hosts must either have an Enterprise Plus license or belong to a vSAN cluster.

https://t.me/learningnets

5-12 Adding ESXi Networking

You can add new standard switches to an ESXi host or configure existing ones using the vSphere Client or VMware Host Client.

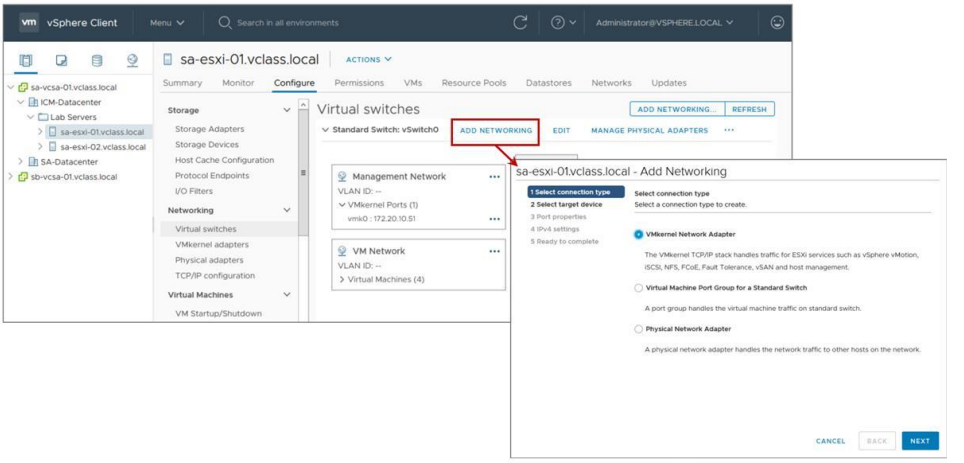

https://t.me/learningnets

5-13 Viewing the Configuration of Standard Switches

In the vSphere Client, you can view a host's standard switch configuration by selecting **Virtual Switches** on the **Configure** tab.

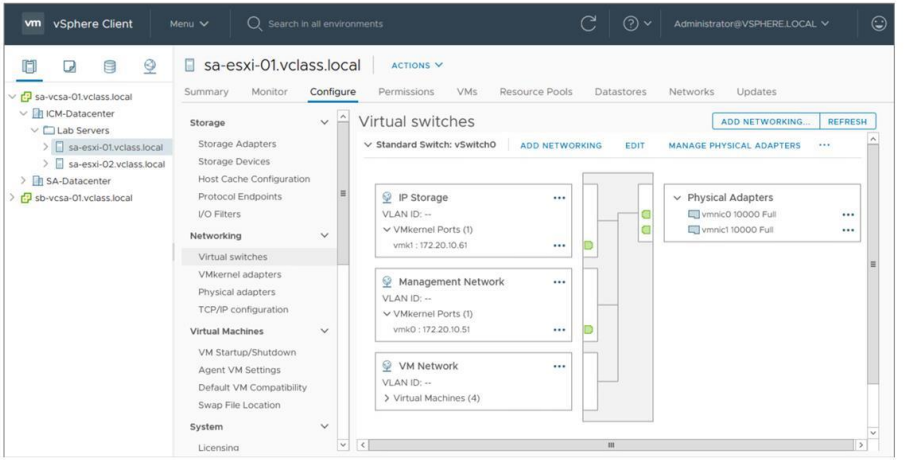

The slide shows the standard switch vSwitch0 on the sa-esxi-01.vclass.local ESXi host. By default, the ESXi installation creates a virtual machine port group named VM Network and a VMkernel port named Management Network. You can create additional port groups such as the Production port group, which you can use for the production virtual machine network.

For performance and security, you should remove the VM Network virtual machine port group and keep VM networks and management networks separated.

https://t.me/learningnets

5-14 Network Adapter Properties

The Physical adapters pane shows adapter details such as speed, duplex, and MAC address settings.

Although the speed and duplex settings are configurable, the best practice is to leave the settings at autonegotiate.

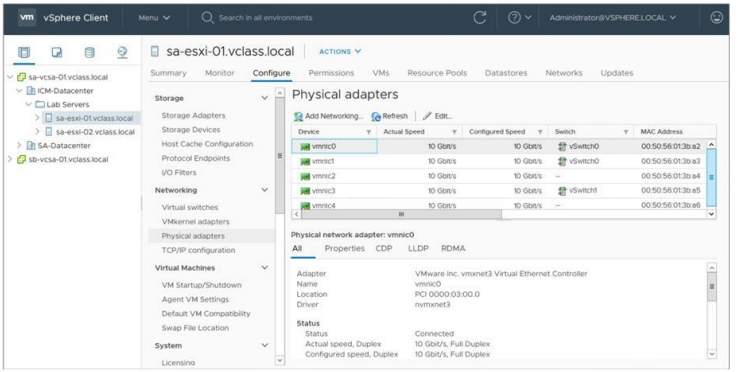

You can change the connection speed and duplex of a physical adapter to transfer data in compliance with the traffic rate.

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

https://t.me/learningnets

5-15 Distributed Switch Architecture

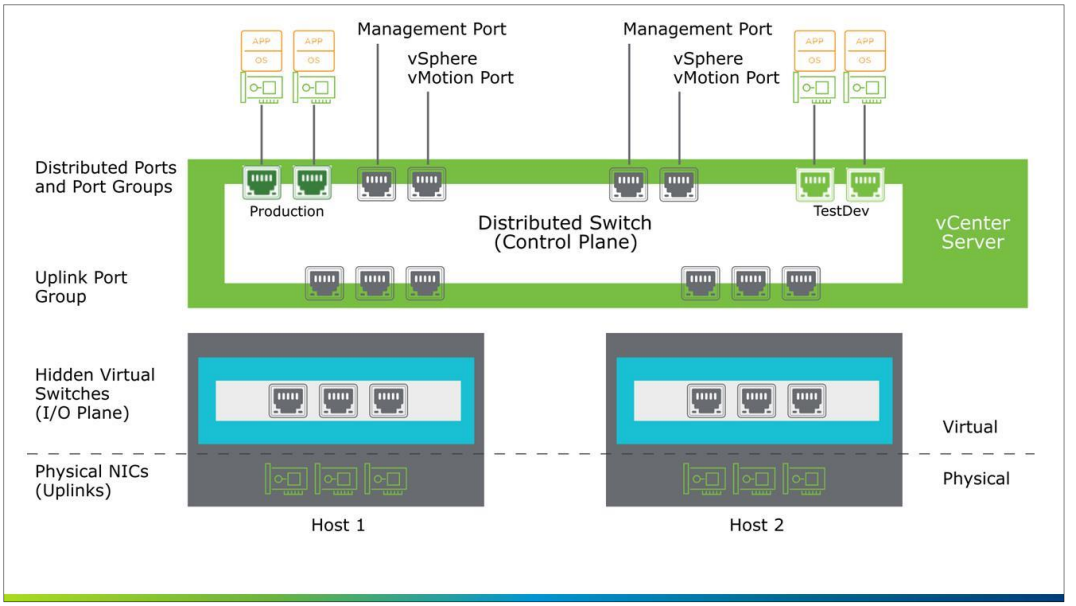

vCenter Server owns the configuration of the distributed switch. The configuration is consistent across all hosts that use the distributed switch.

https://t.me/learningnets

5-16 Standard and Distributed Switches: Shared Features

Standard and distributed switches have several features in common.

Feature	Standard Switch	Distributed Switch
Layer 2 switch	Yes	Yes
VLAN segmentation (802.1Q tagging)	Yes	Yes
IPv6 support	Yes	Yes
NIC teaming	Yes	Yes
Outbound traffic shaping	Yes	Yes
Cicso Discovery Protocol (CDP)	Yes	Yes

https://t.me/learningnets

5-17 Additional Features of Distributed Switches

Distributed switches include several features that are not part of standard switches.

Feature	Standard Switch	Distributed Switch
Inbound traffic shaping	No	Yes
VM network port block	No	Yes
Private VLANs	No	Yes
Load-based teaming	No	Yes
Data center level management	No	Yes
vSphere vMotion migration of virtual networking state	No	Yes
Per-port policy settings	No	Yes
Port state monitoring of network statistics	No	Yes
NetFlow	No	Yes
Port mirroring	No	Yes
Access to NSX-T port groups	No	Yes
Link Layer Discovery Protocol (LLDP)	No	Yes

During a vSphere vMotion migration, a distributed switch tracks the virtual networking state (for example, counters and port statistics) as the virtual machine moves from host to host. The tracking provides a consistent view of a virtual network interface, regardless of the virtual machine location or vSphere vMotion migration history. Tracking simplifies network monitoring and troubleshooting activities where vSphere vMotion is used to migrate virtual machines between hosts.

https://t.me/learningnets

5-18 Lab 11: Using Standard Switches

Create a standard switch and a port group for virtual machines:

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

https://t.me/learningnets

5-19 Review of Learner Objectives

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

- Identify virtual switch connection types
- Configure and view standard switch configurations
- Distinguish between the features of standard and distributed switches

https://t.me/learningnets

https://t.me/learningnets

5-21 Learner Objectives

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

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

https://t.me/learningnets

5-22 Network Switch and Port Policies

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

Available network policies:

- Security
- Traffic shaping
- NIC teaming and failover

Policy levels:

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

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

Traffic shaping is useful when you want to limit the amount of traffic to a VM or a group of VMs.

Use the teaming and failover policy to determine the following information:

- How the network traffic of VMs and VMkernel adapters that are connected to the switch is distributed between physical adapters
- How the traffic should be rerouted if an adapter fails.

https://t.me/learningnets

5-23 Configuring Security Policies

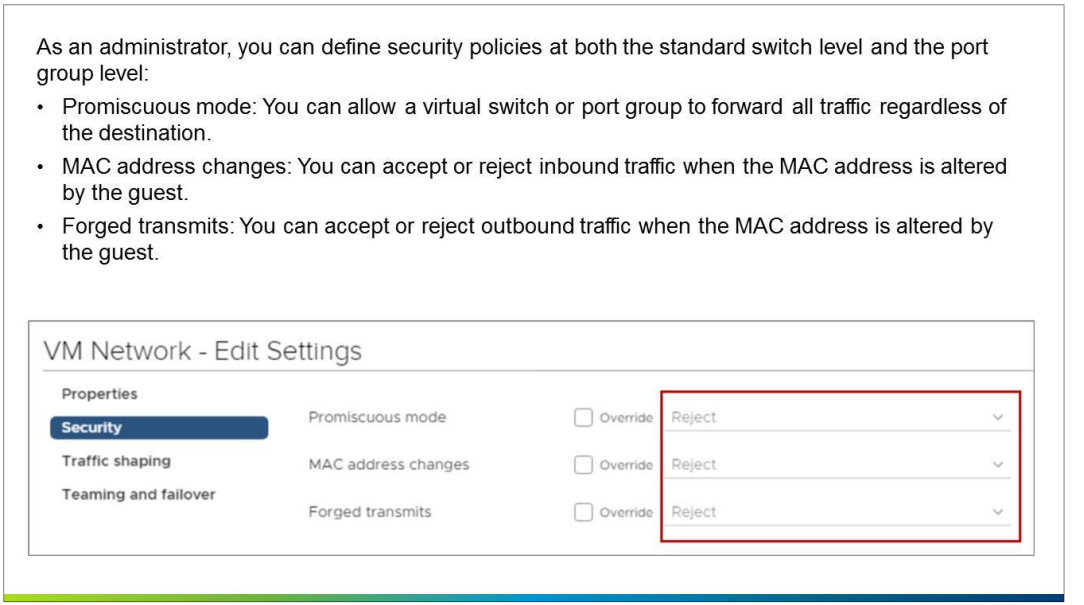

The network security policy contains the following exceptions:

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

In vSphere 7, these security settings are set to Reject by default.

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

https://t.me/learningnets

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

Keep the **MAC address changes** and **Forged transmits** set to **Reject** to help protect against attacks launched by a rogue guest operating system.

Set **MAC address changes** and **Forged transmits** to **Accept** if your applications change the mapped MAC address, as do some guest operating system-based firewalls.

https://t.me/learningnets

5-24 Traffic-Shaping Policies

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

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

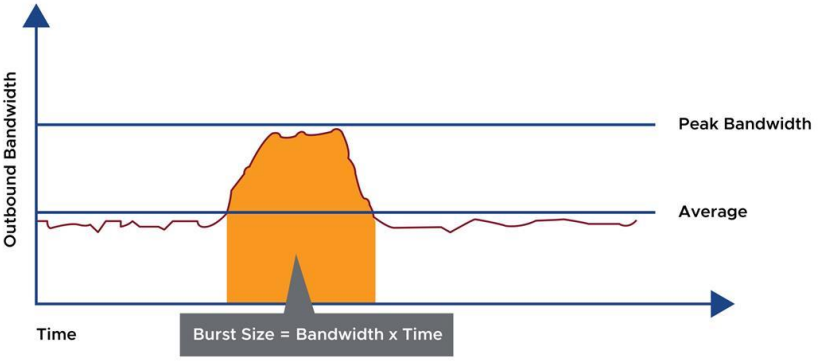

A virtual machine's network bandwidth can be controlled by enabling the network traffic shaper. The network traffic shaper, when used on a standard switch, shapes only outbound network traffic. To control inbound traffic, use a load-balancing system or turn on rate-limiting features on your physical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

https://t.me/learningnets

5-25 Configuring Traffic Shaping

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

- Traffic shaping is disabled by default.
- Parameters apply to each virtual NIC in the standard switch.
- On a standard switch, traffic shaping controls only outbound traffic, that is, traffic traveling from the VMs to the virtual switch and out onto the physical network.

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

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

https://t.me/learningnets

Network traffic shaping is off by default.

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

https://t.me/learningnets

5-26 NIC Teaming and Failover Policies

With NIC teaming, you can increase the network capacity of a virtual switch by including two or more physical NICs in a team.

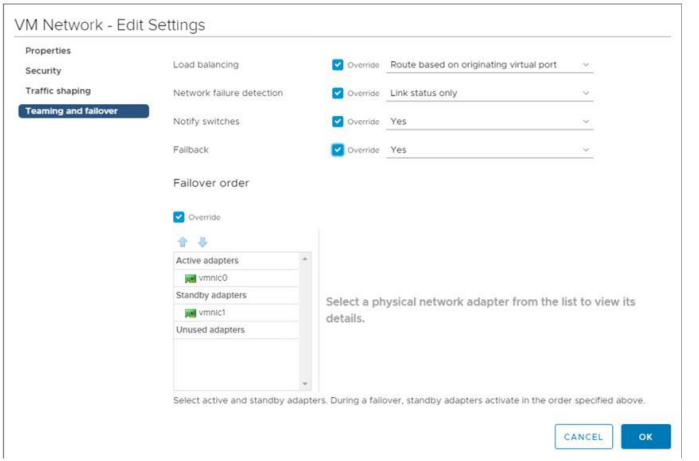

NIC teaming increases the network bandwidth of the switch and provides redundancy. To determine how the traffic is rerouted when an adapter fails, you include physical NICs in a failover order.

To determine how the virtual switch distributes the network traffic between the physical NICs in a team, you select load-balancing algorithms depending on the needs and capabilities of your environment:

- Load-balancing policy: This policy determines how network traffic is distributed between the network adapters in a NIC team. Virtual switches load balance only the outgoing traffic. Incoming traffic is controlled by the load-balancing policy on the physical switch.
- Failback policy: By default, a failback policy is enabled on a NIC team. If a failed physical NIC returns online, the virtual switch sets the NIC back to active by replacing the standby NIC that took over its slot.

https://t.me/learningnets

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

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

Default NIC teaming and failover policies are set for the entire standard switch. These default settings can be overridden at the port group level. The policies show what is inherited from the settings at the switch level.

https://t.me/learningnets

5-27 Load-Balancing Method: Originating Virtual Port ID

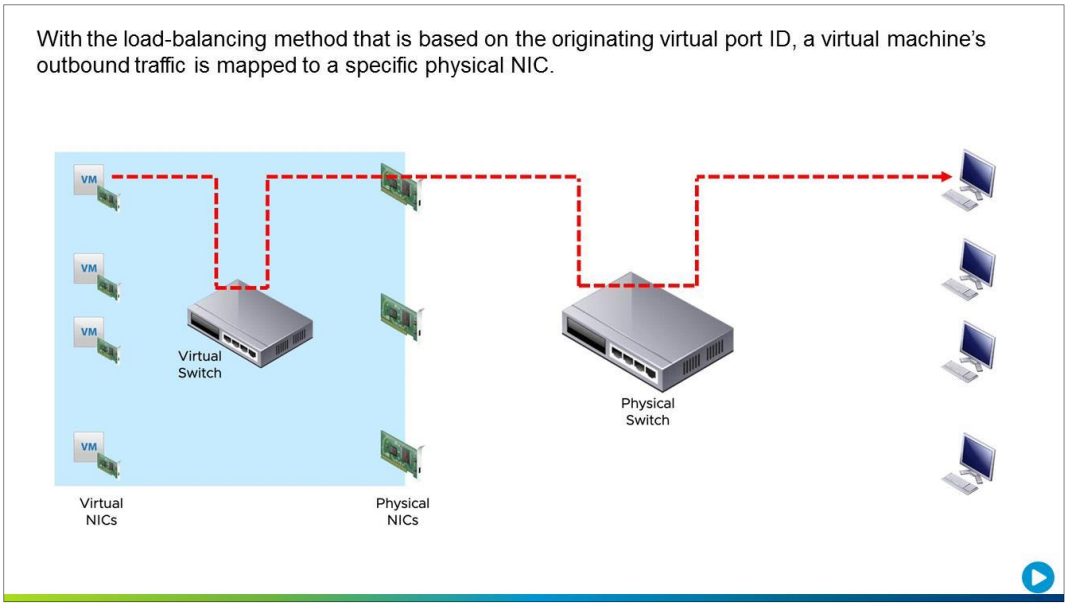

To play the animation, go to https://vmware.bravais.com/s/7jEkuYyY0f7OxeWbmmvZ.

The load-balancing method that uses the originating virtual port ID is simple and fast and does not require the VMkernel to examine the frame for the necessary information. The NIC is determined by the ID of the virtual port to which the VM is connected. With this method, no single-NIC VM gets more bandwidth than can be provided by a single physical adapter.

This method has advantages:

- Traffic is evenly distributed if the number of virtual NICs is greater than the number of physical NICs in the team.
- Resource consumption is low because, in most cases, the virtual switch calculates uplinks for the VM only once.
- No changes on the physical switch are required.

https://t.me/learningnets

This method also has disadvantages:

- The virtual switch is not aware of the traffic load on the uplinks, and it does not load balance the traffic to uplinks that are less used.
- The bandwidth that is available to a VM is limited to the speed of the uplink that is associated with the relevant port ID, unless the VM has more than one virtual NIC.

https://t.me/learningnets

5-28 Load-Balancing Method: Source MAC Hash

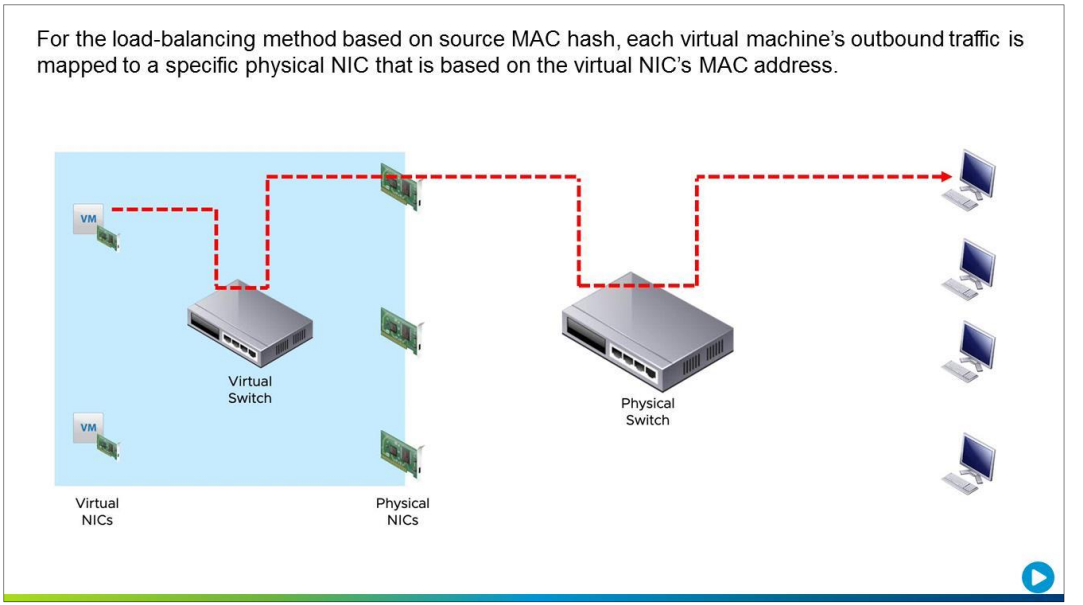

To play the animation, go to https://vmware.bravais.com/s/MmjsUVkaURaNJzMnsao2.

The load-balancing method based on source MAC hash has low overhead and is compatible with all switches, but it might not spread traffic evenly across all the physical NICs. In addition, no single-NIC virtual machine gets more bandwidth than a single physical adapter can provide.

This method has advantages:

- VMs use the same uplink because the MAC address is static. Powering a VM on or off does not change the uplink that the VM uses.
- No changes on the physical switch are required.

https://t.me/learningnets

This method has disadvantages:

- The bandwidth that is available to a VM is limited to the speed of the uplink that is associated with the relevant port ID, unless the VM uses multiple source MAC addresses.
- Resource consumption is higher than with a route based on the originating virtual port because the virtual switch calculates an uplink for every packet.
- The virtual switch is not aware of the load of the uplinks, so uplinks might become overloaded.

https://t.me/learningnets

5-29 Load-Balancing Method: Source and Destination IP Hash

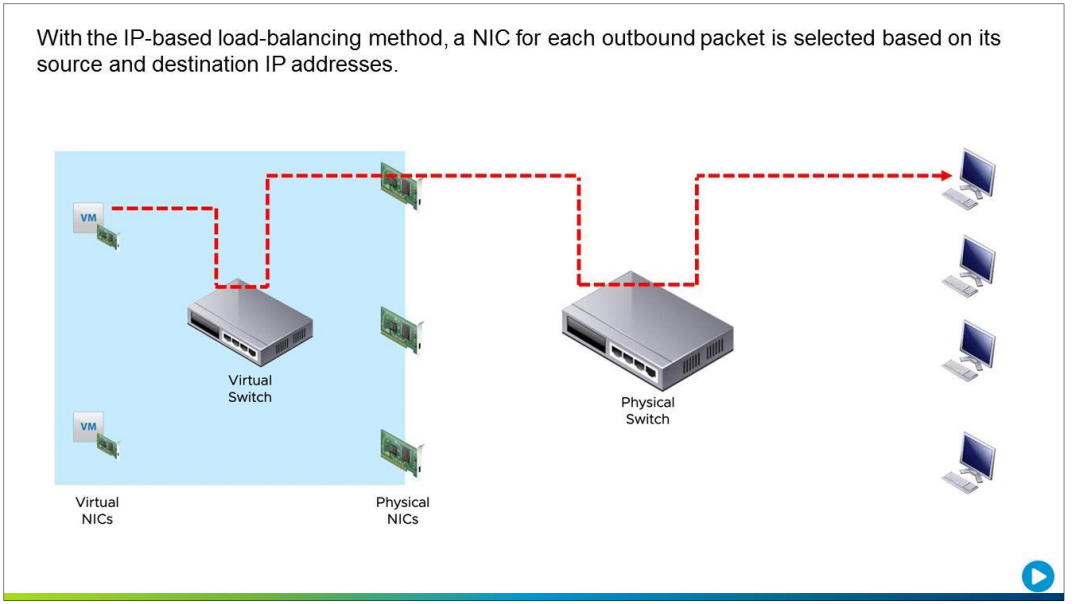

To play the animation, go to https://vmware.bravais.com/s/55sfU1JyzGzuBGWETPu9.

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

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

With this method, a single-NIC virtual machine might use the bandwidth of multiple physical adapters.

The IP-based load-balancing method only affects outbound traffic. For example, a VM might choose a particular NIC to communicate with a particular destination VM. The return traffic might not arrive on the same NIC as the outbound traffic. The return traffic might arrive on another NIC in the same NIC team.

https://t.me/learningnets

This method has advantages:

- The load is more evenly distributed compared to the route based on the originating virtual port and the route based on source MAC hash because the virtual switch calculates the uplink for every packet.
- VMs that communicate with multiple IP addresses have a potentially higher throughput.

This method has disadvantages:

- Resource consumption is the highest compared to the other load-balancing algorithms.
- The virtual switch is not aware of the actual load of the uplinks.
- Changes on the physical network are required.
- The method is complex to troubleshoot.

https://t.me/learningnets

5-30 Detecting and Handling Network Failure

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

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

The VMkernel notifies physical switches of changes in the physical location of a MAC address. Failover is implemented by the VMkernel based on configurable parameters:

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

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

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

A physical switch can be notified by the VMkernel whenever a virtual NIC is connected to a virtual switch. A physical switch can also be notified whenever a failover event causes a virtual NIC's traffic to be routed over a different physical NIC. The notification is sent over the network to update the lookup tables on physical switches. In most cases, this notification process is

https://t.me/learningnets

beneficial because, without it, VMs experience greater latency after failovers and vSphere vMotion operation.

Do not set this option when the VMs connected to the port group are running unicast-mode Microsoft Network Load Balancing (NLB). NLB in multicast mode is unaffected. For more information about the NLB issue, see VMware knowledge base article 1556 at http://kb.vmware.com/kb/1556.

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

The failback option determines how a physical adapter is returned to active duty after recovering from a failure:

- If **Failback** is set to **Yes**, the failed adapter is returned to active duty immediately on recovery, displacing the standby adapter that took its place at the time of failure.
- If **Failback** is set to **No**, a failed adapter is left inactive even after recovery, until another currently active adapter fails, requiring its replacement.

https://t.me/learningnets

5-31 Physical Network Considerations

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

The following issues are topics for discussion:

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

https://t.me/learningnets

5-32 Review of Learner Objectives

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

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

https://t.me/learningnets

5-33 Virtual Beans: Networking Requirements

As a Virtual Beans administrator, you have a few decisions to make about your network infrastructure.

As you plan your network, you consider these key takeaways about vSphere networking:

- You must create port groups for the VLANs that you want to use in your vSphere environment.
- You can use NIC teaming in the virtual switch to avoid a single point of failure.
- You can separate infrastructure service traffic from your application traffic by putting each traffic type on its own VLAN.

  Segmenting traffic can improve performance and enhance security by limiting network access to a specific traffic type.
- You should research the benefits of using distributed switches in your environment. Distributed switches have additional features over standard switches.

https://t.me/learningnets

5-34 Key Points

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

Questions?

https://t.me/learningnets

Module 6

Configuring and Managing Virtual Storage

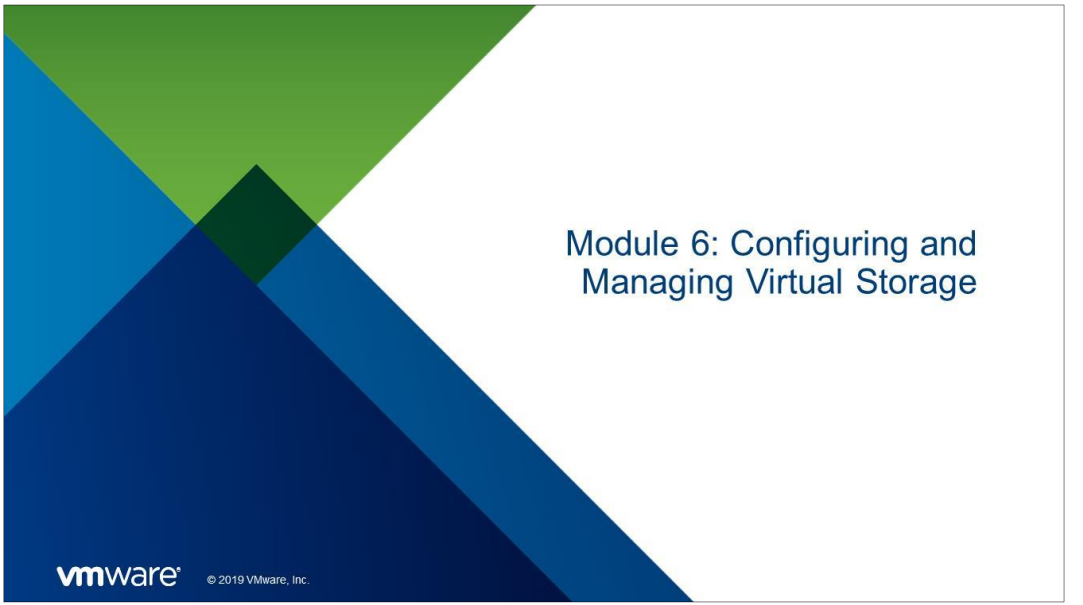

https://t.me/learningnets

6-2 Importance

Understanding the available storage options helps you set up your storage according to your cost, performance, and manageability requirements.

You can use shared storage for disaster recovery, high availability, and moving virtual machines between hosts.

https://t.me/learningnets

6-3 Module Lessons

1. Storage Concepts
2. Fibre Channel Storage
3. iSCSI Storage
4. VMFS Datastores
5. NFS Datastores
6. vSAN Datastores

https://t.me/learningnets

6-4 Virtual Beans: Storage

Virtual Beans current storage infrastructure consists of NAS storage and iSCSI storage arrays. Virtual Beans has the following goals for its storage infrastructure:

- Run the production and development workloads on shared storage systems.
- Use existing NAS and iSCSI storage arrays in the vSphere environment.
  These arrays are managed by the storage administrator at Virtual Beans.
- Evaluate other options in vSphere 7 for cost-effective, high-performance storage.

As a Virtual Beans vSphere administrator, you must configure storage for use in the vSphere environment and provide recommendations to management on other storage options in vSphere 7.

https://t.me/learningnets

6-5 Lesson 1: Storage Concepts

https://t.me/learningnets

6-6 Learner Objectives

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

- Recognize vSphere storage technologies
- Identify types of datastores

https://t.me/learningnets

6-7 About Datastores

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

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

vSphere supports the following types of datastores:

- VMFS
- NFS
- vSAN
- vSphere Virtual Volumes

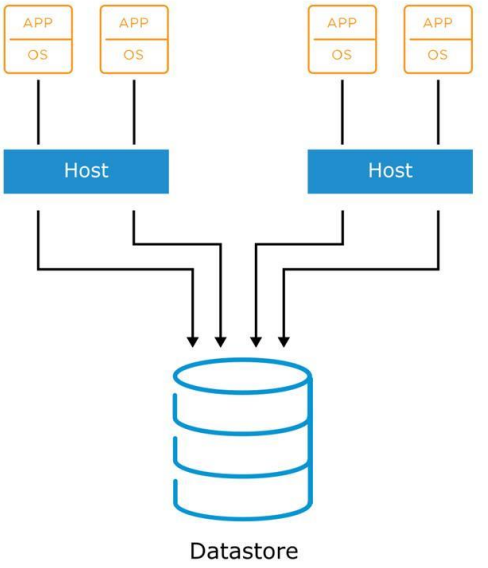

A datastore is a generic term for a container that holds files and objects. Datastores are logical containers, analogous to file systems, that hide the specifics of each storage device and provide a uniform model for storing virtual machine files. A VM is stored as a set of files in its own directory or as a group of objects in a datastore.

You can display all datastores that are available to your hosts and analyze their properties.

https://t.me/learningnets

6-8 Storage Overview

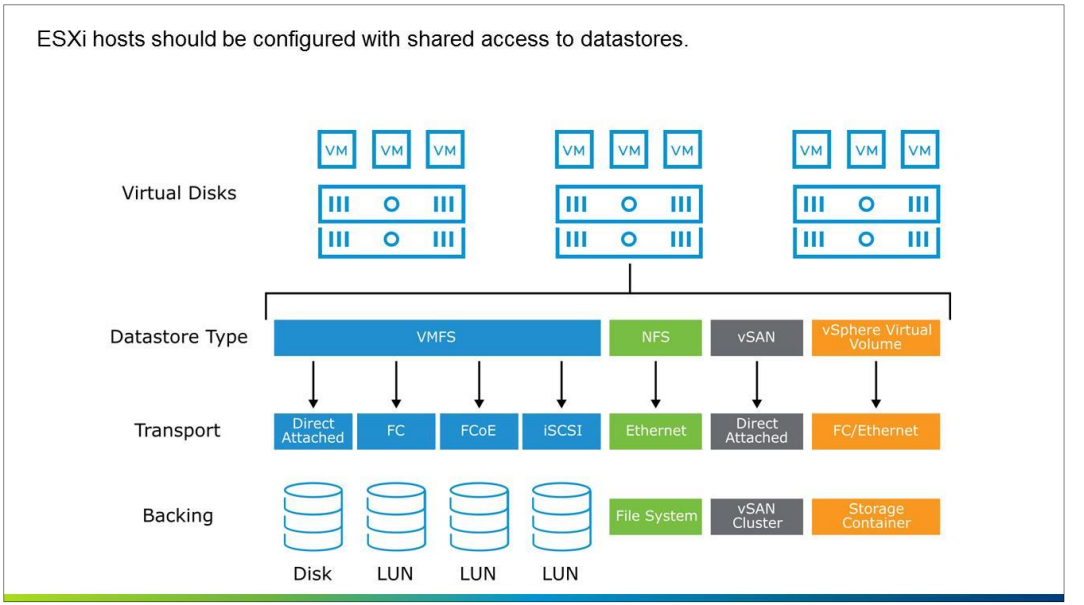

Depending on the type of storage that you use, datastores can be formatted with VMFS or NFS.

In the vSphere environment, ESXi hosts support several storage technologies:

- Direct-attached storage: Internal or external storage disks or arrays attached to the host through a direct connection instead of a network connection.
- Fibre Channel (FC): A high-speed transport protocol used for SANs. Fibre Channel encapsulates SCSI commands, which are transmitted between Fibre Channel nodes. In general, a Fibre Channel node is a server, a storage system, or a tape drive. A Fibre Channel switch interconnects multiple nodes, forming the fabric in a Fibre Channel network.
- FCoE: The Fibre Channel traffic is encapsulated into Fibre Channel over Ethernet (FCoE) frames. These FCoE frames are converged with other types of traffic on the Ethernet network.
- iSCSI: A SCSI transport protocol, providing access to storage devices and cabling over standard TCP/IP networks. iSCSI maps SCSI block-oriented storage over TCP/IP. Initiators,

https://t.me/learningnets

such as an iSCSI host bus adapter (HBA) in an ESXi host, send SCSI commands to targets, located in iSCSI storage systems.

- NAS: Storage shared over standard TCP/IP networks at the file system level. NAS storage is used to hold NFS datastores. The NFS protocol does not support SCSI commands.
- iSCSI, network-attached storage (NAS), and FCoE can run over high-speed networks providing increased storage performance levels and ensuring sufficient bandwidth. With sufficient bandwidth, multiple types of high-bandwidth protocol traffic can coexist on the same network. r For more information about physical NIC support and maximum ports supported, see VMware Configuration Maximums at https://configmax.vmware.com.

https://t.me/learningnets

6-9 Storage Protocol Overview

Each datastore uses a protocol with varying support features.

Datastore Type	Storage Protocol	Boot from SAN Support	vSphere vMotion Support	vSphere HA Support	vSphere DRS Support
VMFS	Fibre Channel	Yes	Yes	Yes	Yes
	FCoE	Yes	Yes	Yes	Yes
	iSCSI	Yes	Yes	Yes	Yes
	iSER/NVMe-oF (RDMA)	No	Yes	Yes	Yes
	DAS (SAS, SATA, NVMe)	N/A	Yes*	No	No
NFS	NFS	No	Yes	Yes	Yes
vSphere Virtual Volumes	FC/Ethernet (iSCSI, NFS)	No	Yes	Yes	Yes
vSAN Datastore	vSAN	No	Yes	Yes	Yes

* Direct-attached storage (DAS) supports vSphere vMotion when combined with vSphere Storage vMotion.

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

- CD/DVD ISO images
- Decommissioned VMs
- VM templates

https://t.me/learningnets

In comparison, storage LUNs must be pooled and shared so that all ESXi hosts can access them. Shared storage provides the following vSphere features:

- vSphere vMotion
- vSphere HA
- vSphere DRS

Using shared SAN storage also provides robust features in vSphere:

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

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

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

https://t.me/learningnets

6-10 About VMFS

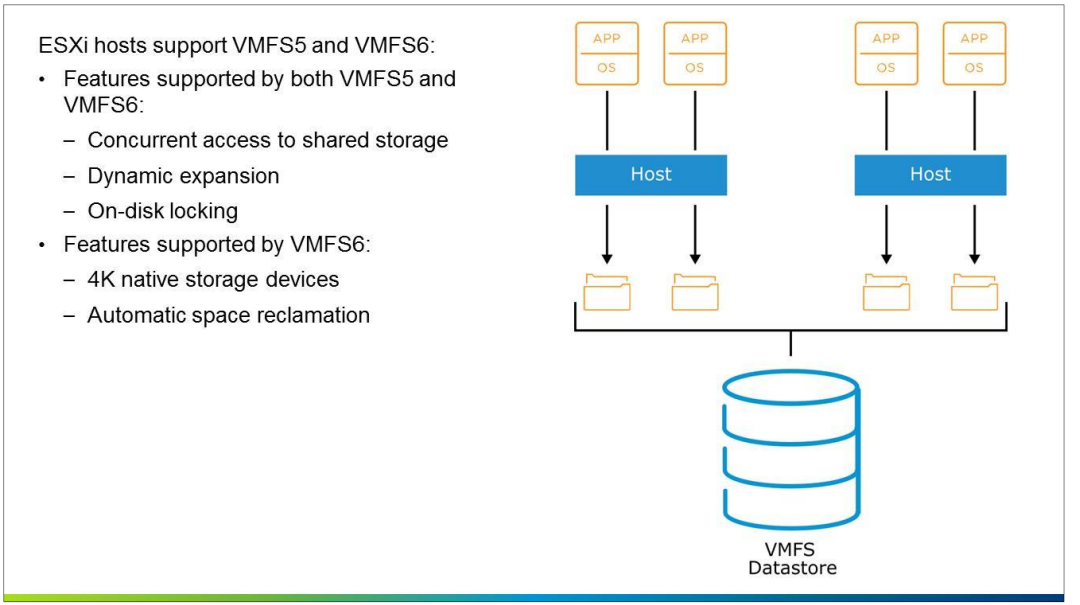

VMFS is a clustered file system where multiple ESXi hosts can read and write to the same storage device simultaneously. The clustered file system provides unique, virtualization-based services:

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

Using VMFS, IT organizations can simplify VM provisioning by efficiently storing the entire VM state in a central location. Multiple ESXi hosts can access shared VM storage concurrently.

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

https://t.me/learningnets

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

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

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

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

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

For the operating system in the VM, VMFS preserves the internal file system semantics. As a result, the operating system running in the VM sees a native file system, not VMFS. These semantics ensure correct behavior and data integrity for applications running on the VMs.

https://t.me/learningnets

6-11 About NFS

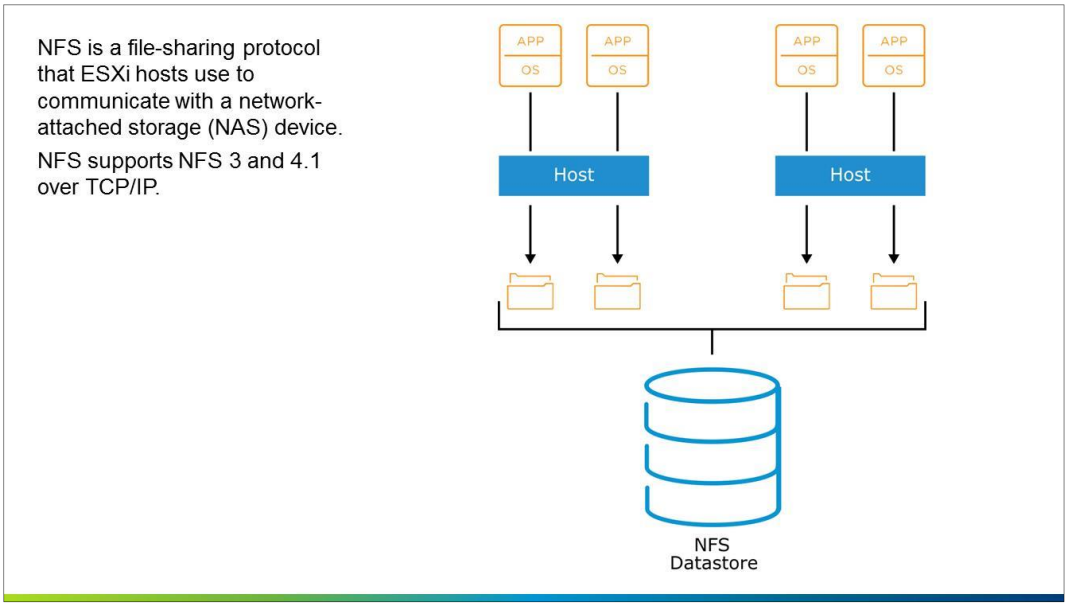

NAS is a specialized storage device that connects to a network and can provide file access services to ESXi hosts.

NFS datastores are treated like VMFS datastores because they can hold VM files, templates, and ISO images. In addition, like a VMFS datastore, an NFS volume allows the vSphere vMotion migration of VMs whose files reside on an NFS datastore. The NFS client built in to ESXi uses NFS protocol versions 3 and 4.1 to communicate with the NAS or NFS servers.

ESXi hosts do not use the Network Lock Manager protocol, which is a standard protocol that is used to support the file locking of NFS-mounted files. VMware has its own locking protocol. NFS 3 locks are implemented by creating lock files on the NFS server. NFS 4.1 uses server-side file locking.

Because NFS 3 and NFS 4.1 clients do not use the same locking protocol, you cannot use different NFS versions to mount the same datastore on multiple hosts. Accessing the same virtual disks from two incompatible clients might result in incorrect behavior and cause data corruption.

https://t.me/learningnets

6-12 About vSAN

vSAN is hypervisor-converged, software-defined storage for virtual environments that does not use traditional external storage.

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

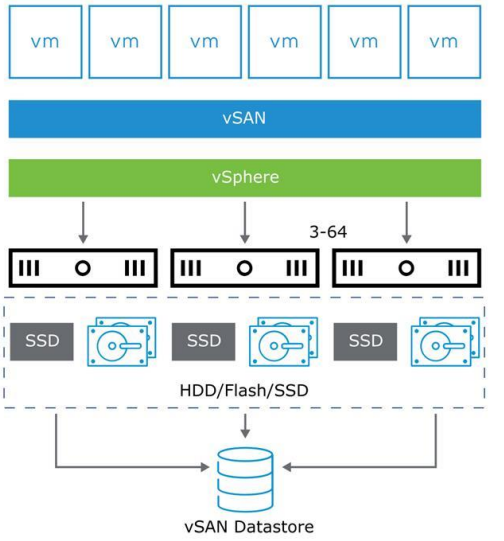

When vSAN is enabled on a cluster, a single vSAN datastore is created. This datastore uses the storage components of each host in the cluster.

vSAN can be configured as hybrid or all-flash storage.

In a hybrid storage architecture, vSAN pools server-attached HDDs and SSDs to create a distributed shared datastore. This datastore abstracts the storage hardware to provide a software-defined storage tier for VMs. Flash is used as a read cache/write buffer to accelerate performance, and magnetic disks provide capacity and persistent data storage.

Alternately, vSAN can be deployed as an all-flash storage architecture in which flash devices are used as a write cache. SSDs provide capacity, data persistence, and consistent, fast response times. In the all-flash architecture, the tiering of SSDs results in a cost-effective implementation: a write-intensive, enterprise-grade SSD cache tier and a read-intensive, lower-cost SSD capacity tier.

https://t.me/learningnets

6-13 About vSphere Virtual Volumes

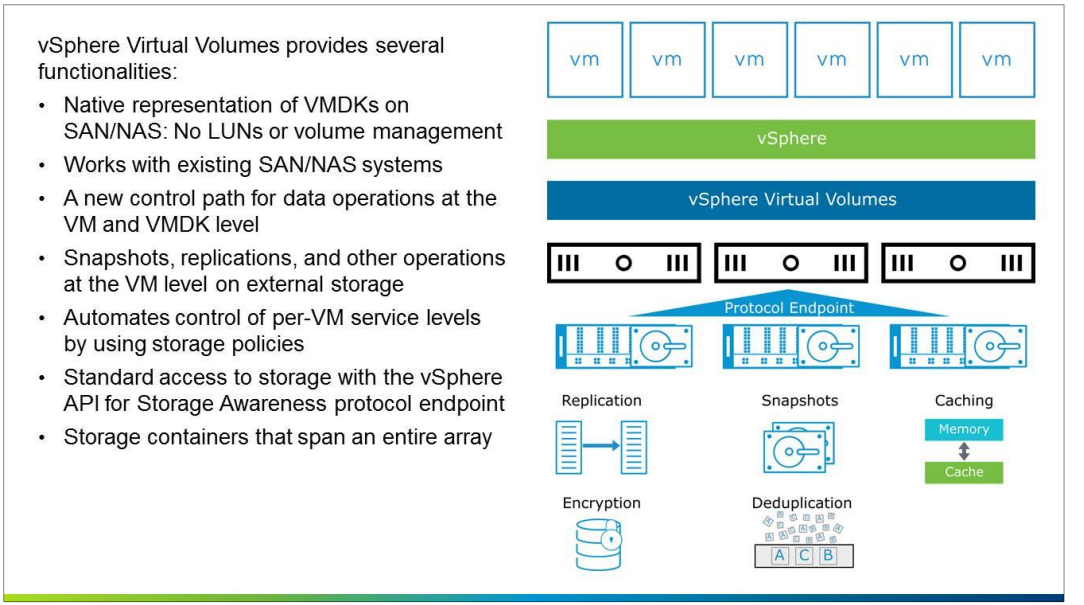

vSphere Virtual Volumes virtualizes SAN and NAS devices by abstracting physical hardware resources into logical pools of capacity.

vSphere Virtual Volumes provides the following benefits:

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

https://t.me/learningnets

6-14 About Raw Device Mapping

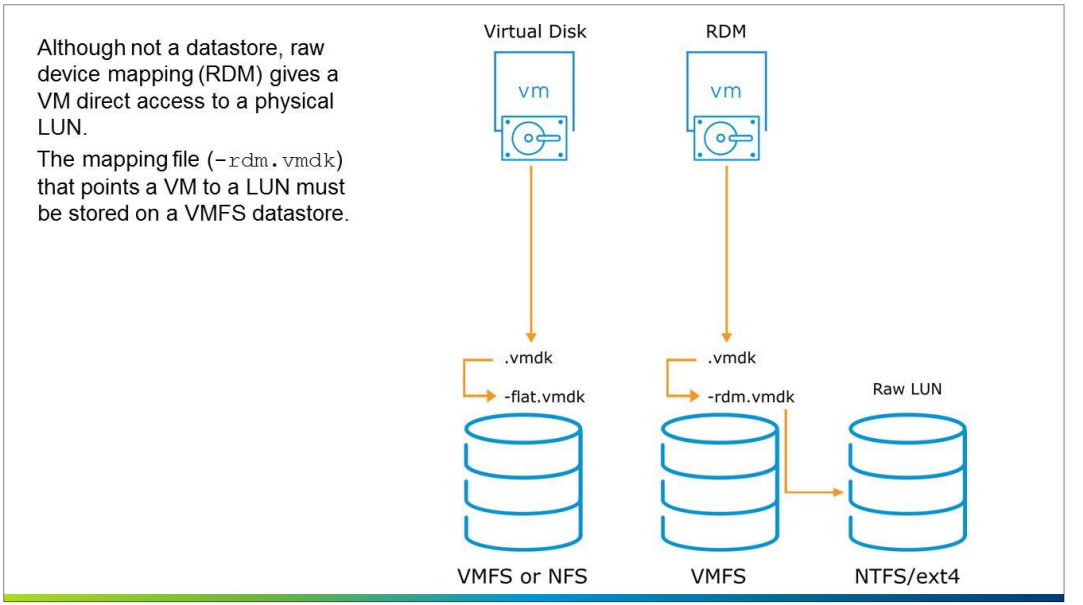

Raw device mapping (RDM) is a file stored in a VMFS volume that acts as a proxy for a raw physical device.

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

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

https://t.me/learningnets

6-15 Physical Storage Considerations

Before implementing your vSphere environment, discuss the storage needs with your storage administration team. Consider the following factors:

- LUN sizes
- I/O bandwidth required by your applications
- I/O requests per second that a LUN is capable of
- Disk cache parameters
- Zoning and masking
- Multipathing setting for your storage arrays (active-active or active-passive)
- Export properties for NFS datastores

For information to help you plan for your storage needs, 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.

https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.htmlAnother good source of information is the vSphere Storage page at https://storagehub.vmware.com/.

https://t.me/learningnets

6-16 Review of Learner Objectives

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

- Recognize vSphere storage technologies
- Identify types of datastores

https://t.me/learningnets

https://t.me/learningnets

6-18 Learner Objectives

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

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

https://t.me/learningnets

6-19 About Fibre Channel

Fibre Channel stores VM files remotely on a Fibre Channel SAN.

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

The network uses the Fibre Channel protocol to transport SCSI traffic from VMs to the Fibre Channel SAN devices.
ESXi supports:

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

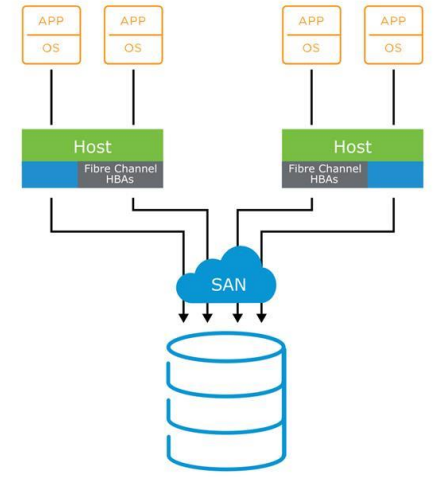

To connect to the Fibre Channel SAN, your host should be equipped with Fibre Channel host bus adapters (HBAs).

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

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

Alternatively, you can access a storage array that supports vSphere Virtual Volumes and create vSphere Virtual Volumes datastores on the array's storage containers.

https://t.me/learningnets

6-20 Fibre Channel SAN Components

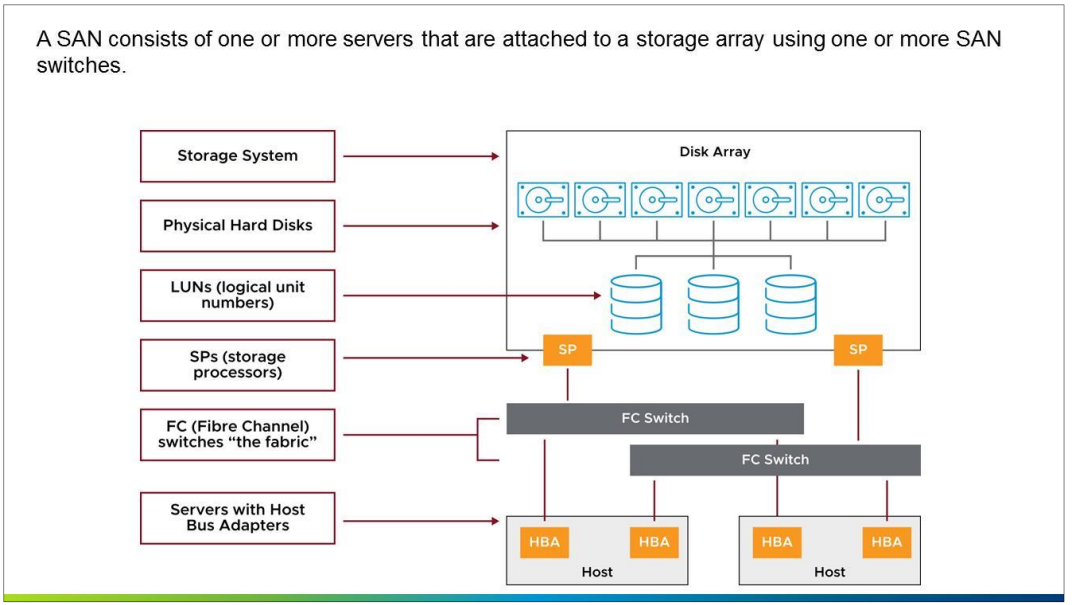

Each SAN server might host numerous applications that require dedicated storage for applications processing.

https://t.me/learningnets

The following components are involved:

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

https://t.me/learningnets

6-21 Fibre Channel Addressing and Access Control

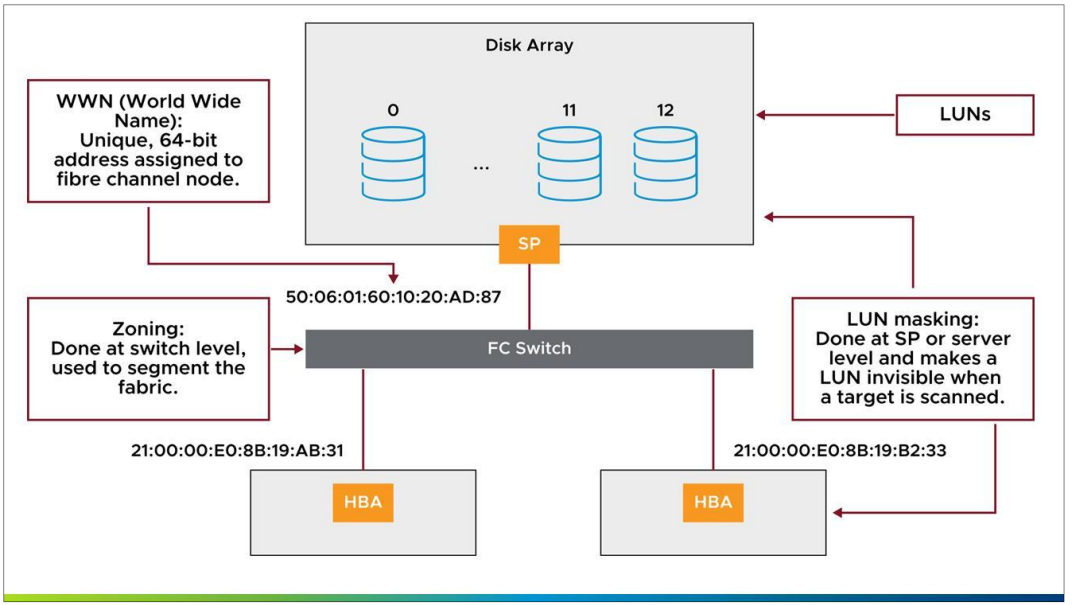

A port connects from a device into the SAN. Each node in the SAN includes each host, storage device, and fabric component (router or switch). Each node in the SAN has one or more ports that connect it to the SAN. Ports can be identified in the following ways:

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

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

You can protect access to storage in your vSphere environment by using zoning and LUN masking with your SAN resources. For example, you might manage zones defined for testing

https://t.me/learningnets

independently within the SAN so that they do not interfere with activity in the production zones. Similarly, you might set up different zones for different departments.

When you set up zones, consider host groups that are set up on the SAN device.

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

See your SAN vendor's documentation and *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-8AE88758-20C1-4873-99C7-181EF9ACFA70.html.

https://t.me/learningnets

6-22 Multipathing with Fibre Channel

Multipathing is having more than one path from a host to a LUN. Multipathing provides the following functions:

- Continued access to SAN LUNs if hardware fails
- Load balancing

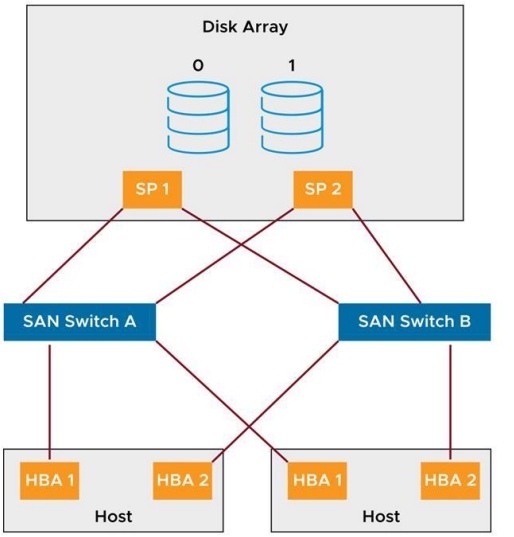

A Fibre Channel path describes a route:

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

By default, ESXi hosts use only one path from a host to a given LUN at any one time. If the path actively being used by the ESXi host fails, the server selects another available path.

The process of detecting a failed path and switching to another is called path failover. A path fails if any of the components along the path (HBA, cable, switch port, or storage processor) fail.

https://t.me/learningnets

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

- An active-active disk array allows access to the LUNs simultaneously through the available storage processors without significant performance degradation. All the paths are active at all times (unless a path fails).
- In an active-passive disk array, one storage processor is actively servicing a given LUN. The other storage processor acts as a backup for the LUN and might be actively servicing other LUN I/O.

  I/O can be sent only to an active processor. If the primary storage processor fails, one of the secondary storage processors becomes active, either automatically or through administrative intervention.

https://t.me/learningnets

6-23 FCoE Adapters

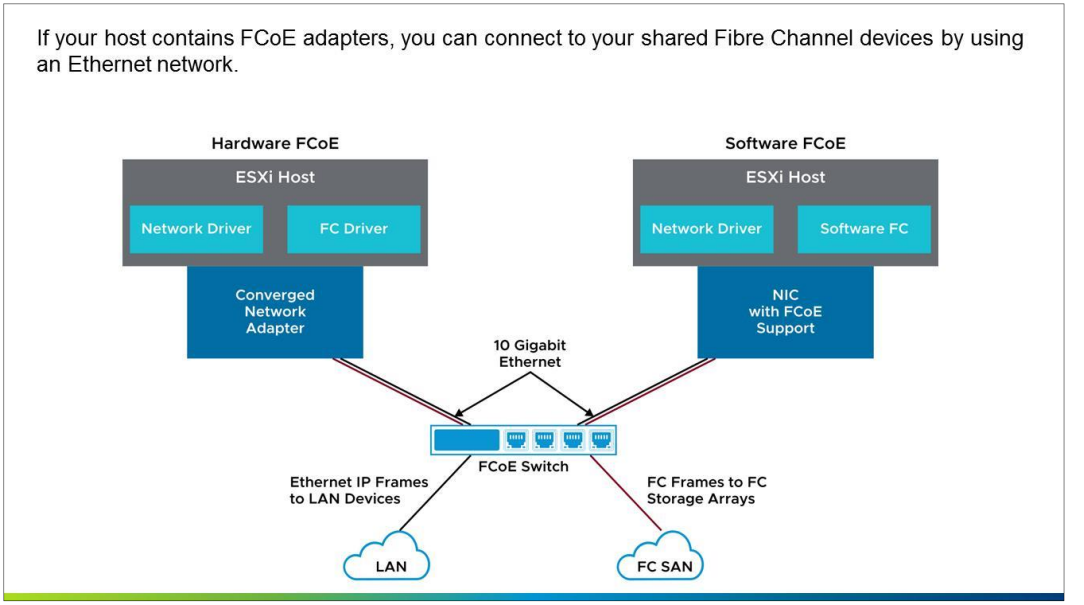

The Fibre Channel traffic is encapsulated into FCoE frames. These FCoE frames are converged with other types of traffic on the Ethernet network.

When both Ethernet and Fibre Channel traffic are carried on the same Ethernet link, use of the physical infrastructure increases. FCoE also reduces the total number of network ports and cabling.

https://t.me/learningnets

6-24 Configuring Software FCoE: Creating VMkernel Ports

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

- The VLAN ID and the priority class are discovered during FCoE initialization. The priority class is not configured in vSphere.
- ESXi supports a maximum of four network adapter ports for software FCoE.

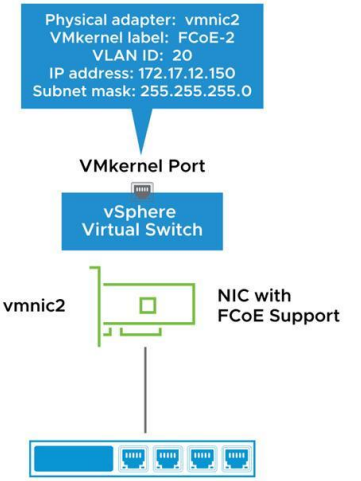

https://t.me/learningnets

6-25 Configuring Software FCoE: Activating Software FCoE Adapters

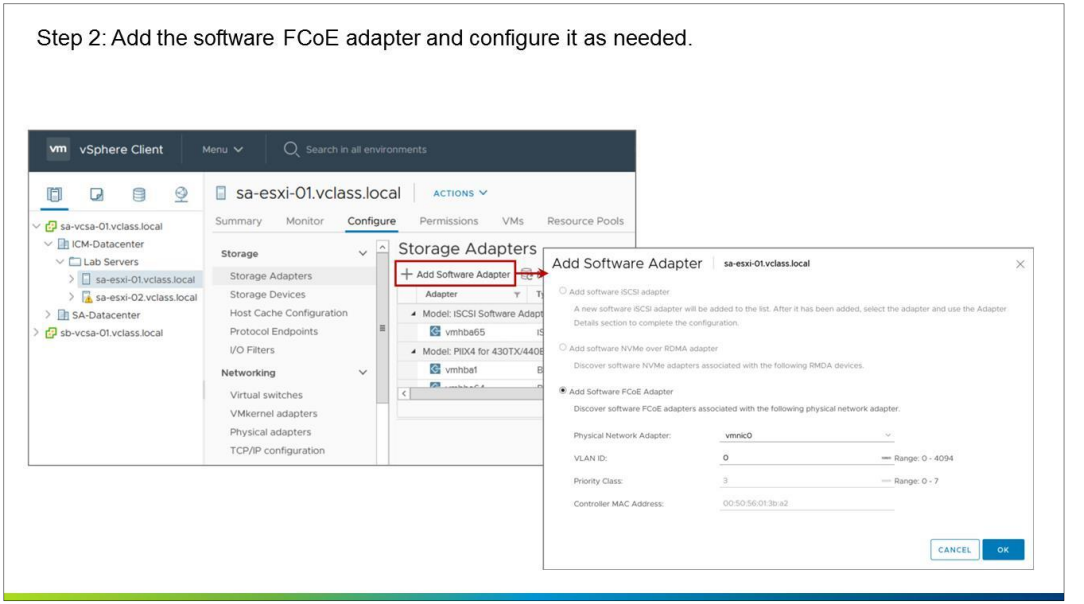

You add the software FCoE adapter by selecting the host, clicking the **Configure** tab, selecting **Storage Adapters**, and clicking **Add Software Adapter**.

https://t.me/learningnets

6-26 Review of Learner Objectives

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

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

https://t.me/learningnets

https://t.me/learningnets

6-28 Learner Objectives

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

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

https://t.me/learningnets

6-29 iSCSI Components

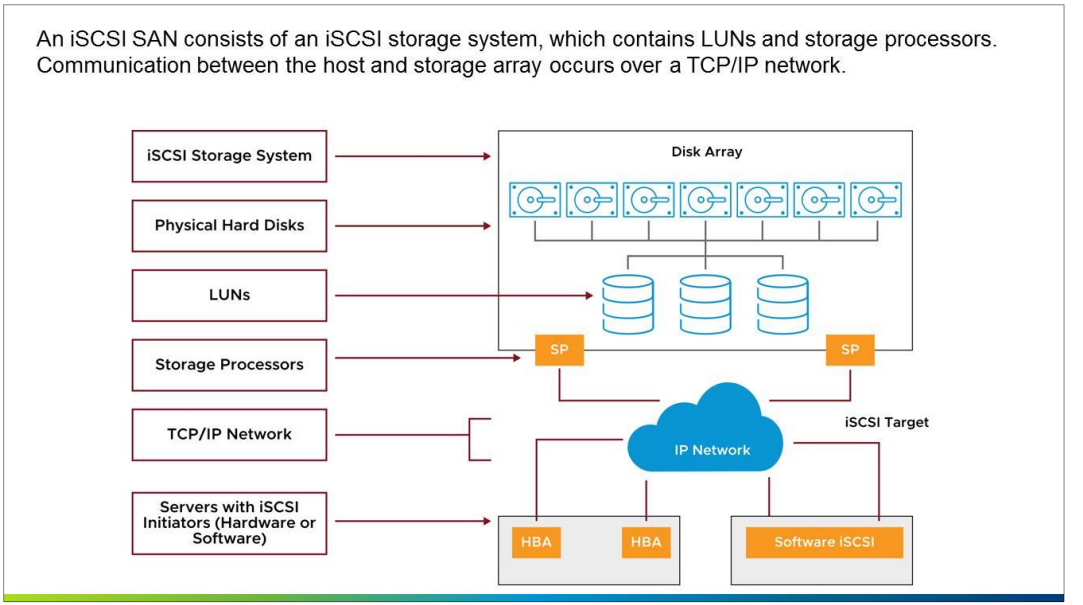

An iSCSI SAN consists of an iSCSI storage system, which contains one or more LUNs and one or more storage processors. Communication between the host and the storage array occurs over a TCP/IP network.

The ESXi host is configured with an iSCSI initiator. An initiator can be hardware-based, where the initiator is an iSCSI HBA. Or the initiator can be software-based, known as the iSCSI software initiator.

An initiator transmits SCSI commands over the IP network. A target receives SCSI commands from the IP network. Your iSCSI network can include multiple initiators and targets. iSCSI is SAN-oriented for the following reasons:

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

https://t.me/learningnets

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

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

https://t.me/learningnets

6-30 iSCSI Addressing

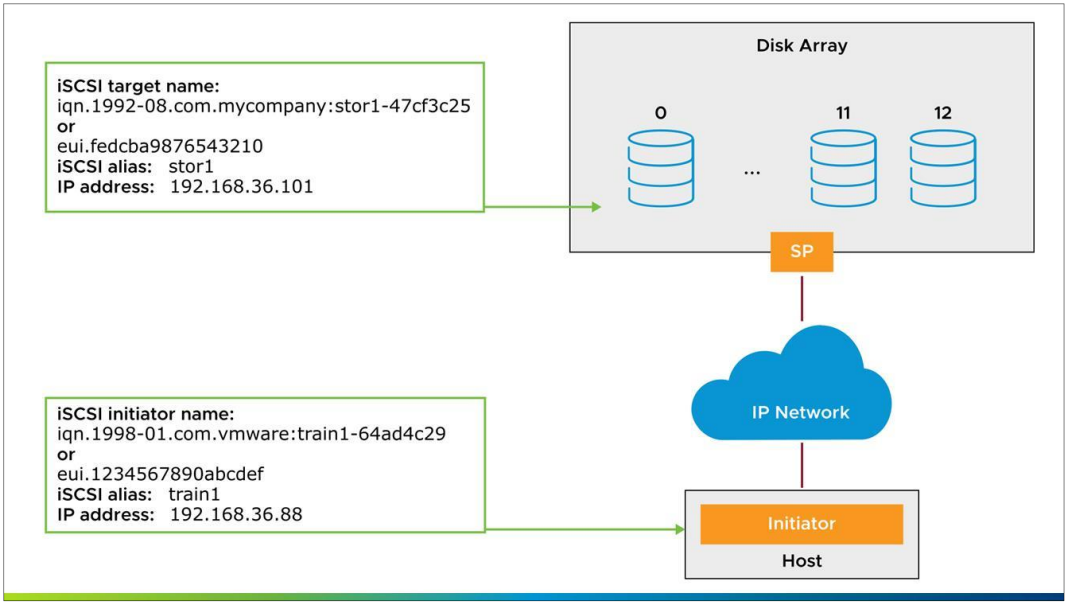

The main addressable, discoverable entity is an iSCSI node. An iSCSI node can be an initiator or a target. An iSCSI node requires a name so that storage can be managed regardless of address.

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

The IQN can be up to 255 characters long. Several naming conventions are used:

- Prefix `iqn`
- Date code specifying the year and month in which the organization registered the domain or subdomain name that is used as the naming authority string
- Organizational naming authority string, which consists of a valid, reversed domain or subdomain name
- (Optional) Colon (:), followed by a string of the assigning organization's choosing, which must make each assigned iSCSI name unique

https://t.me/learningnets

EUI naming conventions are as follows:

- Prefix is `eui`.
- A 16-character name follows the prefix.

  The name includes 24 bits for a company name that is assigned by the IEEE and 40 bits for a unique ID, such as a serial number.

https://t.me/learningnets

6-31 Storage Device Naming Conventions

Storage devices are identified in several ways:

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

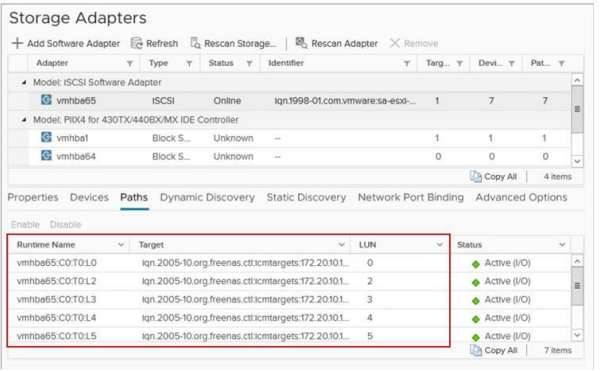

On ESXi hosts, SCSI storage devices use various identifiers. Each identifier serves a specific purpose. For example, the VMkernel requires an identifier, generated by the storage device, which is guaranteed to be unique to each LUN. If the storage device cannot provide a unique identifier, the VMkernel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

The following SCSI storage device identifiers are available:

- Runtime name: The name of the first path to the device. The runtime name is a user-friendly name that is created by the host after each reboot. It is not a reliable identifier for the disk device because it is not persistent. The runtime name might change if you add HBAs to the ESXi host. However, you can use this name when you use command-line utilities to interact with storage that an ESXi host recognizes.
- iSCSI name: A worldwide unique name for identifying the node. iSCSI uses the IQN and EUI. IQN uses the format `iqn.yyyy-mm.naming-authority:unique name`.

Storage device names appear in various panels in the vSphere Client.

https://t.me/learningnets

6-32 iSCSI Adapters

You must set up software or hardware iSCSI adapters before an ESXi host can work with iSCSI storage.

To access iSCSI targets, your host uses iSCSI initiators.

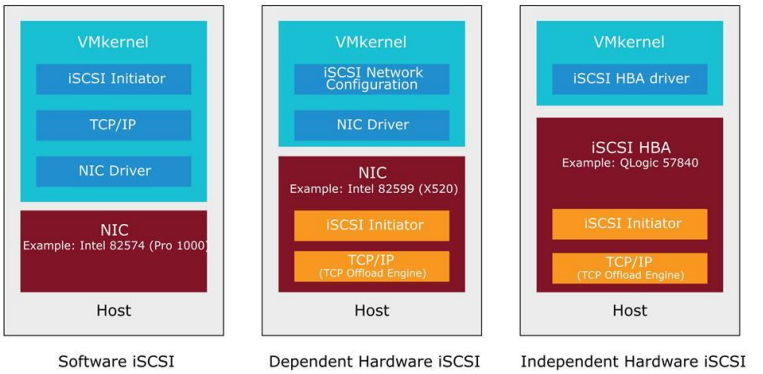

The iSCSI initiators transport SCSI requests and responses, encapsulated in the iSCSI protocol, between the host and the iSCSI target. Your host supports two types of initiators: software iSCSI and hardware iSCSI.

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

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

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

https://t.me/learningnets

An independent hardware iSCSI adapter handles all iSCSI and network processing and management for your ESXi host. In this case, a VMkernel iSCSI port is not required.

For configuration information, 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.

https://t.me/learningnets

6-33 ESXi Network Configuration for IP Storage

A VMkernel port must be created for ESXi to access software iSCSI. The same port can be used to access NAS and NFS storage.
To optimize your vSphere networking setup, separate iSCSI networks from NAS and NFS networks:

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

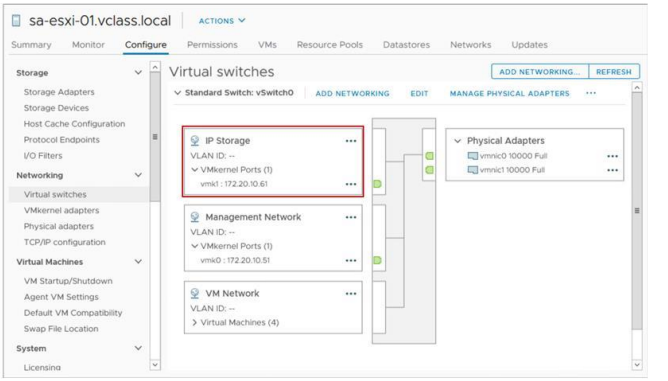

Networking configuration for software iSCSI involves creating a VMkernel port on a virtual switch to handle your iSCSI traffic.

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

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

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

https://t.me/learningnets

6-34 Activating the Software iSCSI Adapter

To add the software iSCSI adapter:

1. Select the host and click the **Configure** tab.
2. Select **Storage Adapters** and click **Add Software Adapter**.

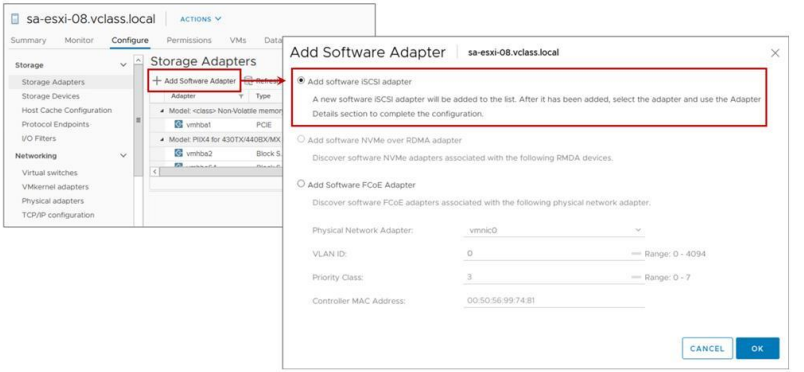

You must activate your software iSCSI adapter so that your host can use it to access iSCSI storage.

You can activate only one software iSCSI adapter.

NOTE

If you boot from iSCSI using the software iSCSI adapter, the adapter is enabled, and the network configuration is created at the first boot. If you disable the adapter, it is reenabled each time you boot the host.

https://t.me/learningnets

6-35 Discovering iSCSI Targets

The iSCSI adapter discovers storage resources on the network and determines which resources are available for access.
An ESXi host supports the following discovery methods:

- Static
- Dynamic or SendTargets

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

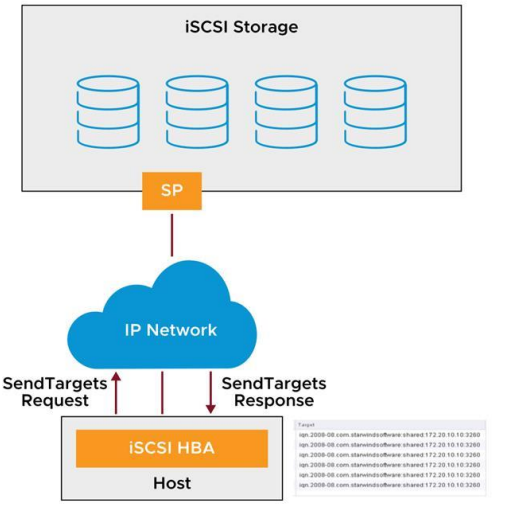

The ESXi host supports the following iSCSI target-discovery methods:

- Static discovery: The initiator does not have to perform discovery. The initiator knows in advance all the targets that it will contact. It uses their IP addresses and domain names to communicate with them.
- Dynamic discovery or SendTargets discovery: Each time the initiator contacts a specified iSCSI server, it sends the SendTargets request to the server. The server responds by supplying a list of available targets to the initiator.

  The names and IP addresses of these targets appear as static targets in the vSphere Client. You can remove a static target that is added by dynamic discovery. If you remove the target, the target might be returned to the list during the next rescan operation. The target might also be returned to the list if the HBA is reset or the host is rebooted.

https://t.me/learningnets

6-36 iSCSI Security: CHAP

iSCSI initiators use CHAP for authentication purposes.

By default, CHAP is not configured. ESXi supports two types of CHAP authentication:

- Unidirectional
- Bidirectional

ESXi also supports per-target CHAP authentication.

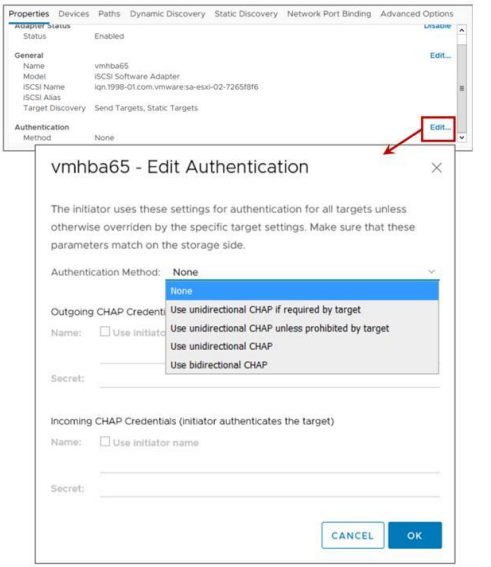

You can implement CHAP to provide authentication between iSCSI initiators and targets.

ESXi supports the following CHAP authentication methods:

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

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

https://t.me/learningnets

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

Before configuring CHAP, check whether CHAP is enabled at the iSCSI storage system and check the CHAP authentication method that the system supports. If CHAP is enabled, you must enable it for your initiators, verifying that the CHAP authentication credentials match the credentials on the iSCSI storage.

Using CHAP in your iSCSI SAN implementation is recommended, but consult with your storage vendor to ensure that best practices are followed.

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

https://t.me/learningnets

6-37 Multipathing with iSCSI Storage

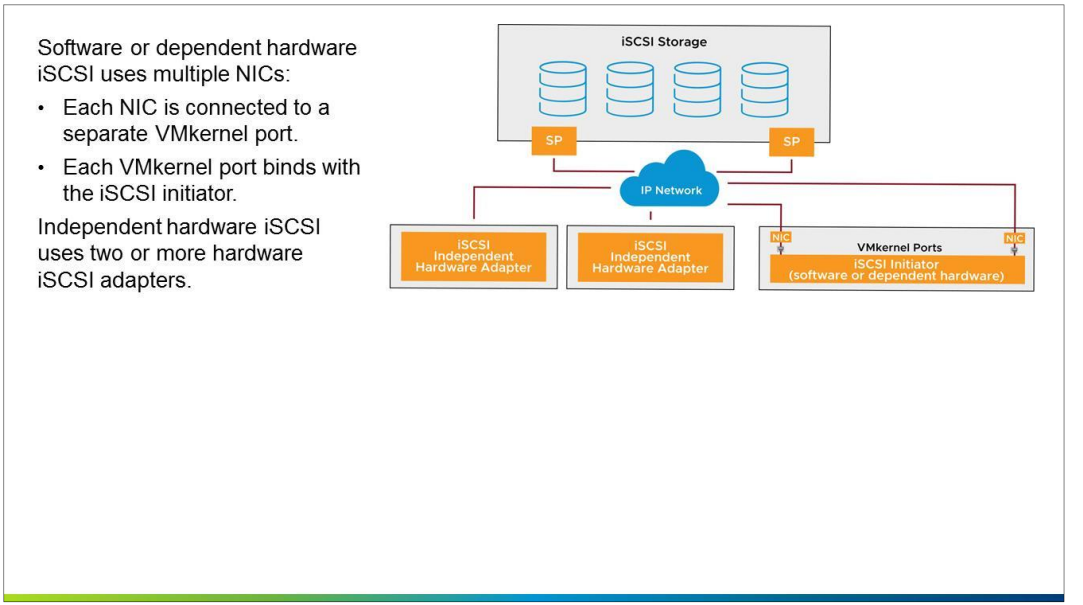

When setting up your ESXi host for multipathing and failover, you can use multiple hardware iSCSI adapters or multiple NICs. The choice depends on the type of iSCSI initiators on your host.

With software iSCSI and dependent hardware iSCSI, you can use multiple NICs that provide failover for iSCSI connections between your host and iSCSI storage systems.

With independent hardware iSCSI, the host typically has two or more available hardware iSCSI adapters, from which the storage system can be reached by using one or more switches.

Alternatively, the setup might include one adapter and two storage processors so that the adapter can use a different path to reach the storage system.

After iSCSI multipathing is set up, each port on the ESXi system has its own IP address, but the ports share the same iSCSI initiator IQN. When iSCSI multipathing is configured, the VMkernel routing table is not consulted for identifying the outbound NIC to use. Instead, iSCSI multipathing is managed using vSphere multipathing modules. Because of the latency that can be incurred, routing iSCSI traffic is not recommended.

https://t.me/learningnets

6-38 Binding VMkernel Ports with the iSCSI Initiator

With port binding, each VMkernel port that is connected to a separate NIC becomes a different path that the iSCSI storage stack can use.

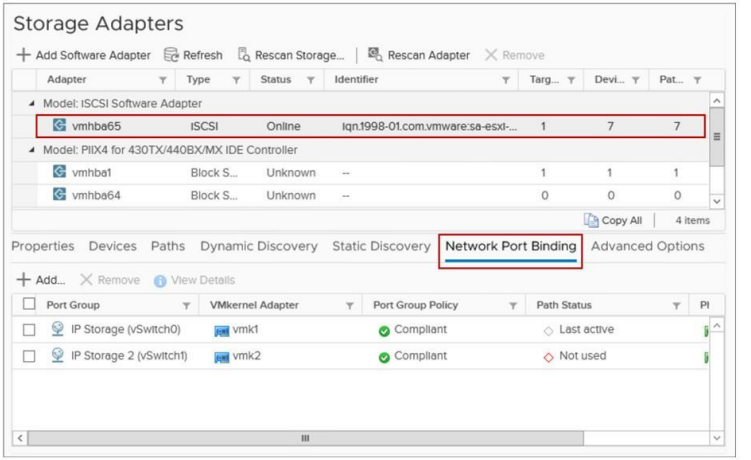

With software iSCSI and dependent hardware iSCSI, multipathing plug-ins do not have direct access to physical NICs on your host. For this reason, you must first connect each physical NIC to a separate VMkernel port. Then you use a port-binding technique to associate all VMkernel ports with the iSCSI initiator.

For dependent hardware iSCSI, you must correctly install the physical network card, which should appear on the host's **Configure** tab in the Virtual Switches view.

https://t.me/learningnets

6-39 Lab 12: Accessing iSCSI Storage

Configure access to an iSCSI datastore:

1. View an Existing ESXi Host iSCSI Configuration
2. Add a VMkernel Port Group to a Standard Switch
3. Add the iSCSI Software Adapter to an ESXi Host
4. Connect the iSCSI Software Adapters to Storage

https://t.me/learningnets

6-40 Review of Learner Objectives

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

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

https://t.me/learningnets

https://t.me/learningnets

6-42 Learner Objectives

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

- Create a VMFS datastore
- Increase the size of a VMFS datastore
- Delete a VMFS datastore

https://t.me/learningnets

6-43 Creating a VMFS Datastore

You can create VMFS datastores on any SCSI-based storage devices that the host discovers, including Fibre Channel, iSCSI, and local storage devices.

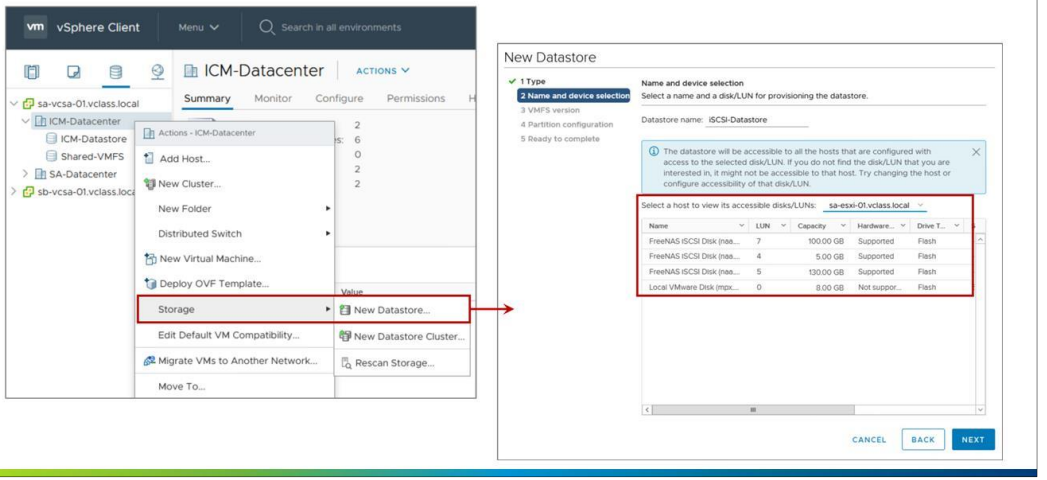

https://t.me/learningnets

6-44 Browsing Datastore Contents

You use the datastore file browser to manage the contents of your datastores.

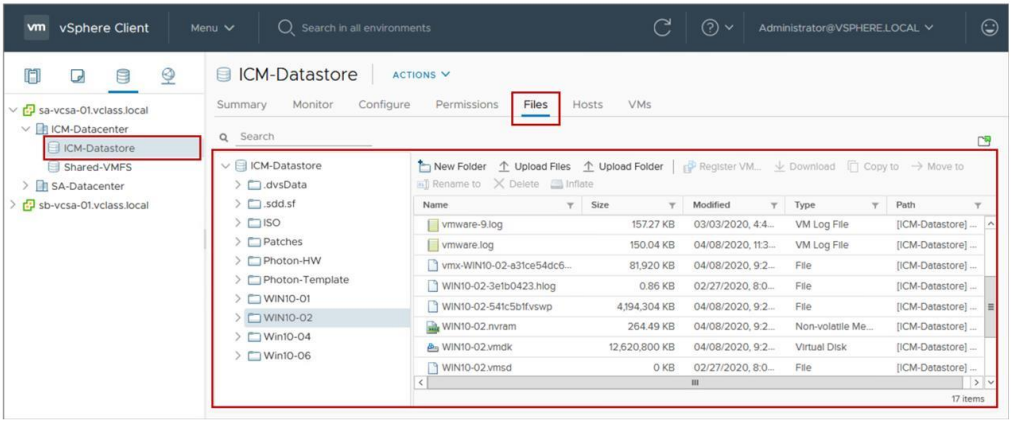

The Datastores pane lists all datastores currently configured for all managed ESXi hosts.

The example shows the contents of the VMFS datastore named Class-Datastore. The contents of the datastore are folders that contain the files for virtual machines or templates.

https://t.me/learningnets

6-45 About VMFS Datastores

A VMFS datastore primarily serves as a repository for VM files.

This type of datastore is optimized for storing and accessing large files, such as virtual disks and memory images of suspended VMs.

A VMFS datastore can have a maximum volume size of 64 TB.

https://t.me/learningnets

6-46 Managing Overcommitted Datastores

A datastore becomes overcommitted when the total provisioned space of thin-provisioned disks is greater than the size of the datastore.

To actively monitor datastore capacity:

- Set alarms to send notifications about:
  - Datastore disk overallocation
  - VM disk use
- Use reporting to view space usage.

To actively manage datastore capacity:

- Increase datastore capacity when necessary.
- Use vSphere Storage vMotion to mitigate space use problems on a particular datastore.

Using thin-provisioned virtual disks for your VMs is a way to make the most of your datastore capacity. But if your datastore is not sized properly, it can become overcommitted. A datastore becomes overcommitted when the full capacity of its thin-provisioned virtual disks is greater than the datastore's capacity.

When a datastore reaches capacity, the vSphere Client prompts you to provide more space on the underlying VMFS datastore and all VM I/O is paused.

Monitor your datastore capacity by setting alarms to alert you about how much a datastore's disks are fully allocated or how much disk space a VM is using.

Manage your datastore capacity by dynamically increasing the size of your datastore when necessary. You can also use vSphere Storage vMotion to mitigate space use issues.

For example, with vSphere Storage vMotion, you can migrate a VM off a datastore. The migration can be done by changing from virtual disks of thick format to thin format at the target datastore.

https://t.me/learningnets

6-47 Increasing the Size of VMFS Datastores

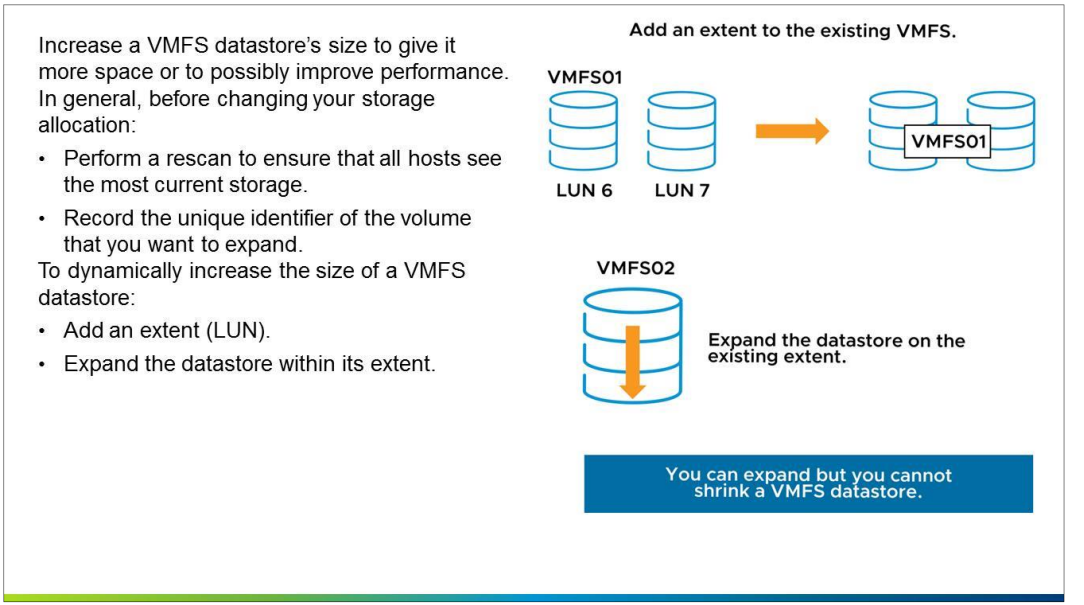

An example of the unique identifier of a volume is the NAA ID. You require this information to identify the VMFS datastore that must be increased.

You can dynamically increase the capacity of a VMFS datastore if the datastore has insufficient disk space. You discover whether insufficient disk space is an issue when you create a VM or you try to add more disk space to a VM.

Use one of the following methods:

- Add an extent to the VMFS datastore: An extent is a partition on a LUN. You can add an extent to any VMFS datastore. The datastore can stretch over multiple extents, up to 32.
- Expand the VMFS datastore: You expand the size of the VMFS datastore by expanding its underlying extent first.

https://t.me/learningnets

6-48 Datastore Maintenance Mode

Before taking a datastore out of service, place the datastore in maintenance mode.

Before placing a datastore in maintenance mode, you must migrate all VMs (powered on and powered off) and templates to a different datastore.

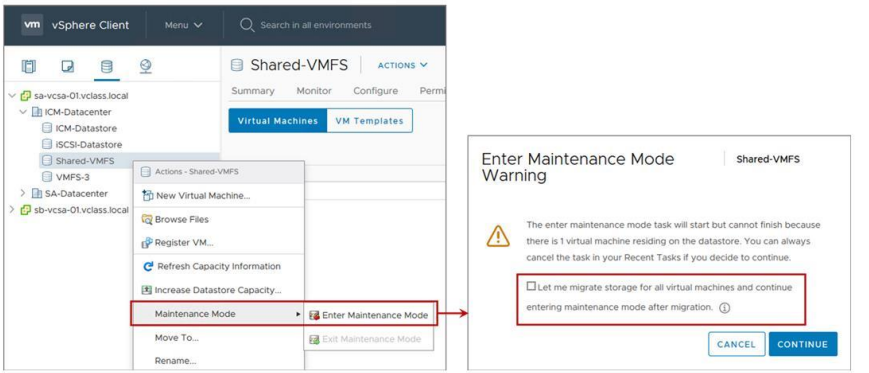

By selecting the **Let me migrate storage for all virtual machines and continue entering maintenance mode after migration** check box, all VMs and templates on the datastore are automatically migrated to the datastore of your choice. The datastore enters maintenance mode after all VMs and templates are moved off the datastore.

Datastore maintenance mode is a function of the vSphere Storage DRS feature, but you can use maintenance mode without enabling vSphere Storage DRS. For more information on vSphere Storage DRS, see *vSphere Resource Management* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.resmgmt.doc/GUID-98BD5A8A-260A-494F-BAAE-74781F5C4B87.html.

https://t.me/learningnets

6-49 Deleting or Unmounting a VMFS Datastore

An unmounted datastore remains intact but cannot be seen from the hosts that you specify.

It continues to appear on other hosts, where it remains mounted.

A deleted datastore is destroyed and disappears from all hosts that have access to it.

The deleted datastore permanently removes all files on the datastore.

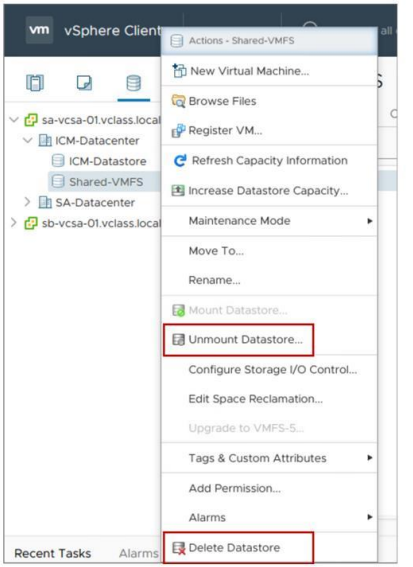

Unmounting a VMFS datastore preserves the files on the datastore but makes the datastore inaccessible to the ESXi host.

Do not perform any configuration operations that might result in I/O to the datastore while the unmounting is in progress.

You can delete any type of VMFS datastore, including copies that you mounted without resignaturing. Although you can delete the datastore without unmounting, you should unmount the datastore first. Deleting a VMFS datastore destroys the pointers to the files on the datastore, so the files disappear from all hosts that have access to the datastore.

https://t.me/learningnets

Before you delete or unmount a VMFS datastore, power off all VMs whose disks reside on the datastore. If you do not power off the VMs and you try to continue, an error message tells you that the resource is busy. Before you unmount a VMFS datastore, use the vSphere Client to verify the following conditions:

- No virtual machines reside on the datastore.
- The datastore is not part of a datastore cluster.
- The datastore is not managed by vSphere Storage DRS.
- vSphere Storage I/O Control is disabled.
- The datastore is not used for vSphere HA heartbeat.

To keep your data, back up the contents of your VMFS datastore before you delete the datastore.

https://t.me/learningnets

6-50 Multipathing Algorithms

Arrays provide active-active and active-passive storage processors. Multipathing algorithms interact with these storage arrays:

- vSphere offers native path selection, load-balancing, and failover mechanisms.
- Third-party vendors can create software for ESXi hosts to properly interact with the storage arrays.

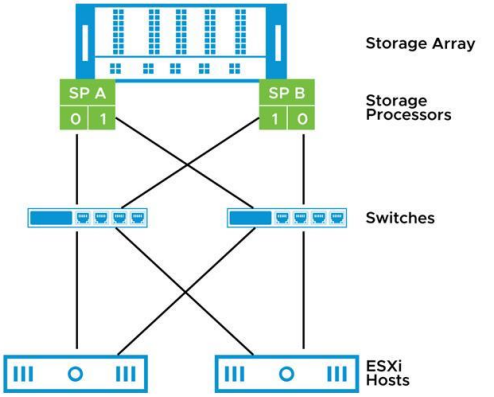

The Pluggable Storage Architecture is a VMkernel layer responsible for managing multiple storage paths and providing load balancing. An ESXi host can be attached to storage arrays with either active-active or active-passive storage processor configurations.

VMware offers native load-balancing and failover mechanisms. VMware path selection policies include the following examples:

- Round Robin
- Most Recently Used (MRU)
- Fixed

Third-party vendors can design their own load-balancing techniques and failover mechanisms for particular storage array types to add support for new arrays. Third-party vendors do not need to provide internal information or intellectual property about the array to VMware.

https://t.me/learningnets

6-51 Configuring Storage Load Balancing

Path selection policies provide:

- Scalability:
  - Round Robin
- Availability:
  - Most Recently Used
  - Fixed

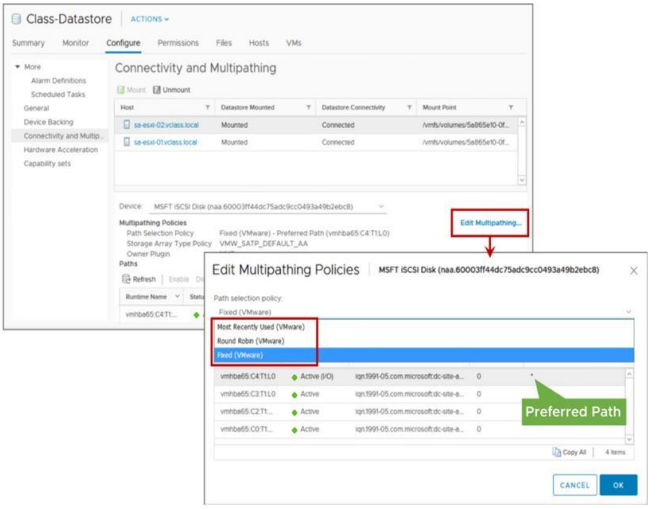

Multiple paths from an ESXi host to a datastore are possible.

For multipathing with Fibre Channel or iSCSI, the following path selection policies are supported:

- Fixed: The host always uses the preferred path to the disk when that path is available. If the host cannot access the disk through the preferred path, it tries the alternative paths. This policy is the default policy for active-active storage devices.
- Most Recently Used: The host selects the first working path discovered at system boot time. When the path becomes unavailable, the host selects an alternative path. The host does not revert to the original path when that path becomes available. The Most Recently Used policy does not use the preferred path setting. This policy is the default policy for active-passive storage devices and is required for those devices.
- Round Robin: The host uses a path selection algorithm that rotates through all available paths. In addition to path failover, the Round Robin multipathing policy supports load balancing

https://t.me/learningnets

across the paths. Before using this policy, check with storage vendors to find out whether a Round Robin configuration is supported on their storage.

https://t.me/learningnets

6-52 Lab 13: Managing VMFS Datastores

Create and manage VMFS datastores:

1. Create VMFS Datastores for the ESXi Host
2. Expand a VMFS Datastore to Consume Unused Space on a LUN
3. Remove a VMFS Datastore
4. Extend a VMFS Datastore
5. Create a Second VMFS Datastore

https://t.me/learningnets

6-53 Review of Learner Objectives

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

- Create a VMFS datastore
- Increase the size of a VMFS datastore
- Delete a VMFS datastore

https://t.me/learningnets

https://t.me/learningnets

6-55 Learner Objectives

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

- Identify NFS components
- Recognize the differences between NFS 3 and NFS 4.1
- Configure and manage NFS datastores

https://t.me/learningnets

6-56 NFS Components

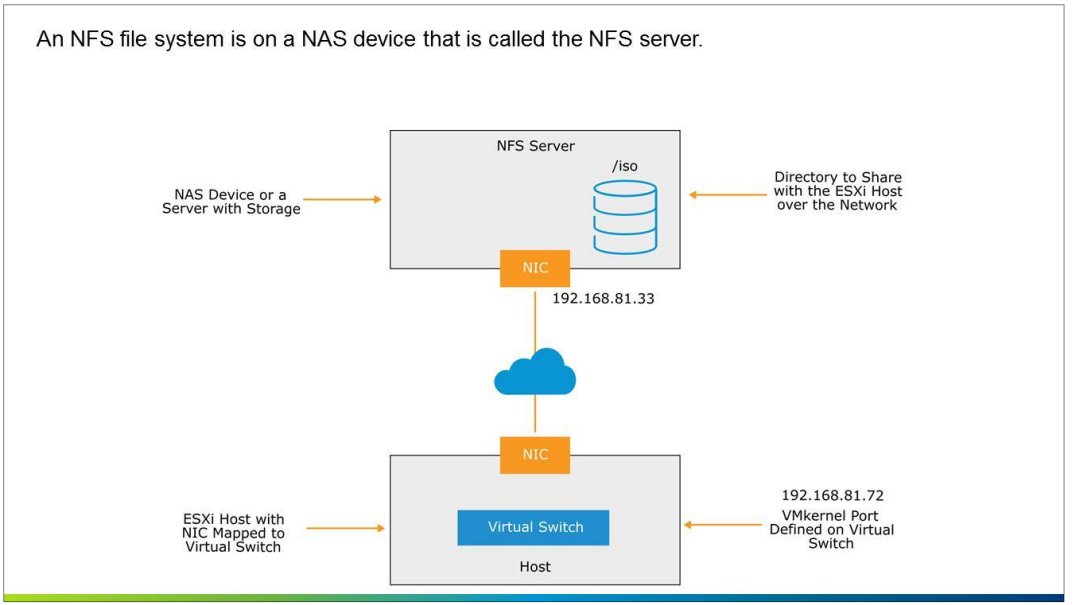

The NFS server contains one or more directories that are shared with the ESXi host over a TCP/IP network. An ESXi host accesses the NFS server through a VMkernel port that is defined on a virtual switch.

https://t.me/learningnets

6-57 NFS 3 and NFS 4.1

An NFS datastore can be created as either NFS 3 or NFS 4.1.

NFS 3	NFS 4.1
ESXi managed multipathing	Native multipathing and session trunking
AUTH_SYS (root) authentication	Optional Kerberos authentication
VMware proprietary client-side file locking	Server-side file locking
Client-side error tracking	Server-side error tracking

Compatibility issues between the two NFS versions prevent access to datastores using both protocols at the same time from different hosts. If a datastore is configured as NFS 4.1, all hosts that access that datastore must mount the share as NFS 4.1. Data corruption can occur if hosts access a datastore with the wrong NFS version.

https://t.me/learningnets

6-58 NFS Version Compatibility with Other vSphere Technologies

vSphere supports NFS 4.1 to overcome many limitations when using NFS 3. Both NFS 3 and NFS 4.1 shares can be used, but you must consider important constraints when designing a vSphere environment in which both versions are used.

vSphere Technology	NFS 3	NFS 4.1
vSphere vMotion and vSphere Storage vMotion	Yes	Yes
vSphere HA and vSphere Fault Tolerance	Yes	Yes
vSphere DRS and vSphere DPM	Yes	Yes
Stateless ESXi and Host Profiles	Yes	Yes
vSphere Storage DRS and Storage I/O Control	Yes	No
Site Recovery Manager	Yes	No
vSphere Virtual Volumes and vSphere Replication	Yes	Yes
vRealize Operations Manager	Yes	Yes
Host Profiles	Yes	Yes

NFS 4.1 provides the following enhancements:

- Native multipathing and session trunking: NFS 4.1 provides multipathing for servers that support session trunking. When trunking is available, you can use multiple IP addresses to access a single NFS volume. Client ID trunking is not supported.
- Kerberos authentication: NFS 4.1 introduces Kerberos authentication in addition to the traditional AUTH_SYS method used by NFS 3.
- Improved built-in file locking.
- Enhanced error recovery using server-side tracking of open files and delegations.
- Many general efficiency improvements including session leases and less protocol overhead.

https://t.me/learningnets

The NFS 4.1 client offers the following new features:

- Stateful locks with share reservation using a mandatory locking semantic
- Protocol integration, side-band (auxiliary) protocol no longer required to lock and mount
- Trunking (true NFS multipathing), where multiple paths (sessions) to the NAS array can be created and load-distributed across those sessions
- Enhanced error recovery to mitigate server failure and loss of connectivity

https://t.me/learningnets

6-59 Configuring NFS Datastores

To configure an NFS datastore:

1. Create a VMkernel port:
   - For better performance and security, separate your NFS network from the iSCSI network.
2. Create the NFS datastore by providing the following information:
   - NFS version: 3 or 4.1
   - Datastore name
   - NFS server names or IP addresses
   - Folder on the NFS server, for example, `/templates` or `/nfs_share`
   - Hosts that mount the datastore
   - Whether to mount the NFS file system as read only
   - Authentication parameters

For each ESXi host that accesses an NFS datastore over the network, a VMkernel port must be configured on a virtual switch. The name of this port can be anything that you want.

For performance and security reasons, isolate your NFS networks from the other networks, such as your iSCSI network and your virtual machine networks.

https://t.me/learningnets

6-60 Configuring ESXi Host Authentication and NFS Kerberos Credentials

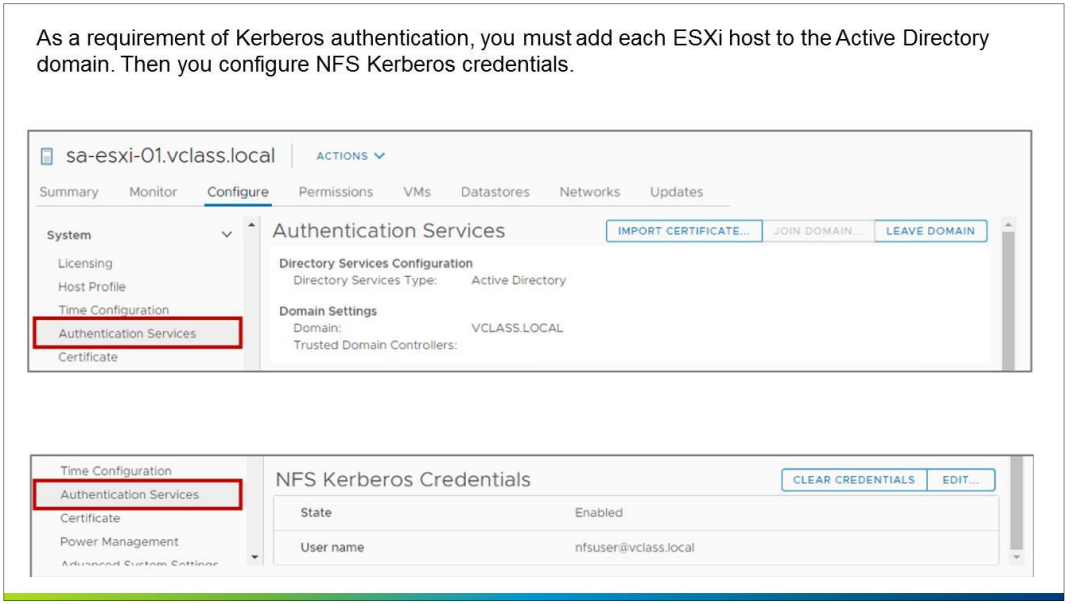

You must take several configuration steps to prepare each ESXi host to use Kerberos authentication.

Kerberos authentication requires that all nodes involved (the Active Directory server, the NFS servers, and the ESXi hosts) be synchronized so that little to no time drift exists. Kerberos authentication fails if any significant drift exists between the nodes.

To prepare your ESXi host to use Kerberos authentication, configure the NTP client settings to reference a common NTP server (or the domain controller, if applicable).

When planning to use NFS Kerberos, consider the following points:

- NFS 3 and 4.1 use different authentication credentials, resulting in incompatible UID and GID on files.
- Using different Active Directory users on different hosts that access the same NFS share can cause the vSphere vMotion migration to fail.

https://t.me/learningnets

- NFS Kerberos configuration can be automated by using host profiles to reduce configuration conflicts.
- Time must be synchronized between all participating components.

https://t.me/learningnets

6-61 Configuring the NFS Datastore to Use Kerberos

When creating each NFS datastore, you enable Kerberos authentication by selecting one of the security modes:

- Kerberos5 authentication
- Kerberos5i authentication and data integrity

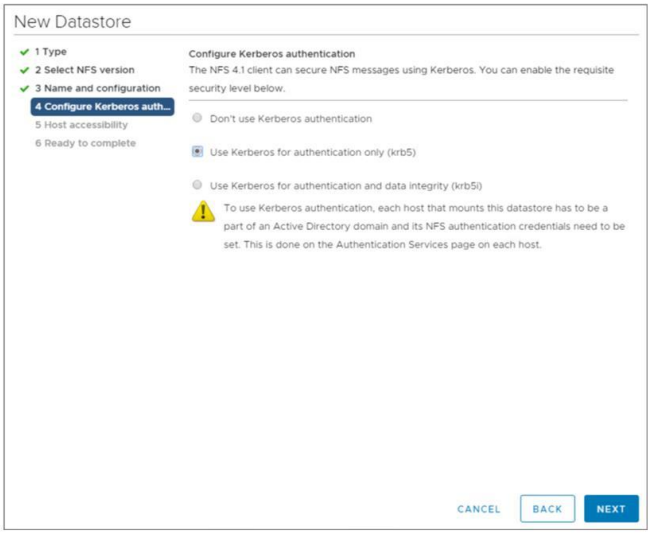

After performing the initial configuration steps, you can configure the datastore to use Kerberos authentication.

The screenshot shows a choice of Kerberos authentication only (krb5) or authentication with data integrity (krb5i). The difference is whether only the header or the header and the body of each NFS operation is signed using a secure checksum.

For more information about how to configure the ESXi hosts for Kerberos authentication, 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.

https://t.me/learningnets

6-62 Unmounting an NFS Datastore

Unmounting an NFS datastore causes the files on the datastore to become inaccessible to the ESXi host.

Before unmounting an NFS datastore, you must stop all VMs whose disks reside on the datastore.

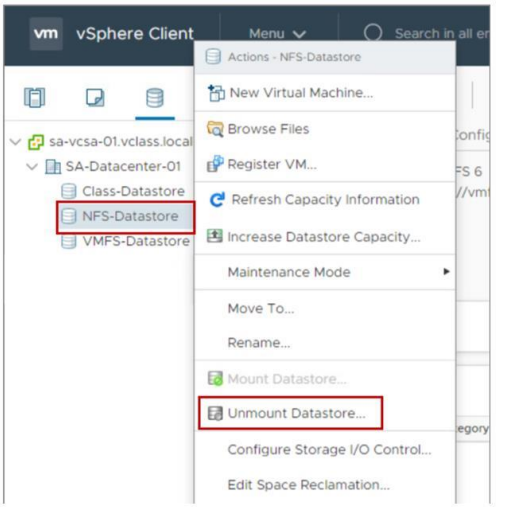

https://t.me/learningnets

6-63 Multipathing and NFS Storage

For a highly available NAS architecture, configure NFS multipathing to avoid single points of failure.

Example of a multipathing configuration:

- Configure one VMkernel port.
- Attach NICs to the same physical switch to configure NIC teaming.
- Configure the NFS server with multiple IP addresses (same subnet is OK).
- To better use multiple links, configure NIC teams with the IP hash load-balancing policy.

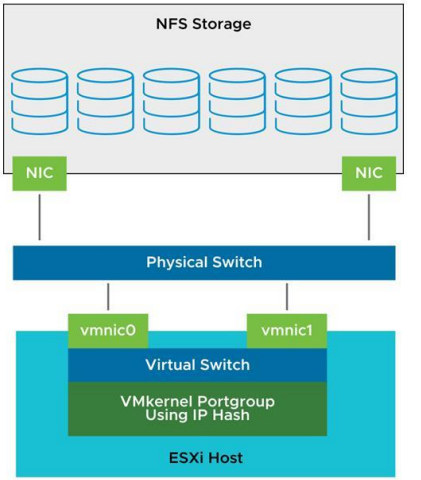

Examples of a single point of failure in the NAS architecture include the NIC card in an ESXi host, and the cable between the NIC card and the switch. To avoid single points of failure and to create a highly available NAS architecture, configure the ESXi host with redundant NIC cards and redundant physical switches.

The best approach is to install multiple NICs on an ESXi host and configure them in NIC teams. NIC teams should be configured on separate external switches, with each NIC pair configured as a team on the respective external switch.

In addition, you might apply a load-balancing algorithm, based on the link aggregation protocol type supported on the external switch, such as 802.3ad or EtherChannel.

An even higher level of performance and high availability can be achieved with cross-stack, EtherChannel-capable switches. With certain network switches, you can team ports across two or more separate physical switches that are managed as one logical switch.

https://t.me/learningnets

NIC teaming across virtual switches provides additional resilience and some performance optimization. Having more paths available to the ESXi host can improve performance by enabling distributed load sharing.

Only one active path is available for the connection between the ESXi host and a single storage target (LUN or mount point). Although alternative connections might be available for failover, the bandwidth for a single datastore and the underlying storage is limited to what a single connection can provide.

To use more available bandwidth, an ESXi host requires multiple connections from the ESXi host to the storage targets. You might need to configure multiple datastores, each using separate connections between the ESXi host and the storage.

The table shows the recommended configuration for NFS multipathing.

External Switches Support Cross-Stack EtherChannel	**External Switches Do Not Support Cross-Stack EtherChannel**
Configure one VMkernel port.	Configure two or more VMkernel ports on different virtual switches on different subnets.
Configure NIC teaming by using adapters attached to separate physical switches.	Configure NIC teaming with adapters attached to the same physical switch.
Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.	Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.
To use multiple links, configure NIC teams with the IP hash load-balancing policy.	To use multiple links, allow the VMkernel routing table to decide which link to send packets (requires multiple datastores).

https://t.me/learningnets

6-64 Enabling Multipathing for NFS 4.1

NFS 4.1 supports native multipathing and session trunking.

To enable multipathing, enter multiple server IP addresses when configuring the datastore.

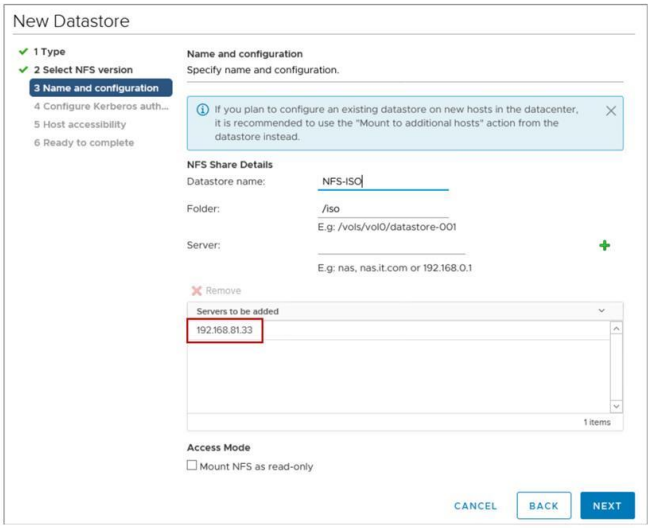

NFS 4.1 provides multipathing for servers that support the session trunking. When trunking is available, you can use multiple IP addresses to access a single NFS volume. Client ID trunking is not supported.

https://t.me/learningnets

6-65 Lab 14: Accessing NFS Storage

Create an NFS datastore and record its storage information:

1. Configure Access to an NFS Datastore
2. View NFS Storage Information

https://t.me/learningnets

6-66 Review of Learner Objectives

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

- Identify NFS components
- Recognize the differences between NFS 3 and NFS 4.1
- Configure and manage NFS datastores

https://t.me/learningnets

https://t.me/learningnets

6-68 Learner Objectives

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

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Explain the purpose of vSAN storage policies

https://t.me/learningnets

6-69 About vSAN Datastores

vSAN is a software-defined storage solution providing shared storage for vSphere clusters without using traditional external storage.
A vSAN cluster requires:

- A minimum of three hosts to be part of the vSphere cluster and enabled for vSAN
- A vSAN network
- Local disks on each host that are pooled to create a virtual shared vSAN datastore

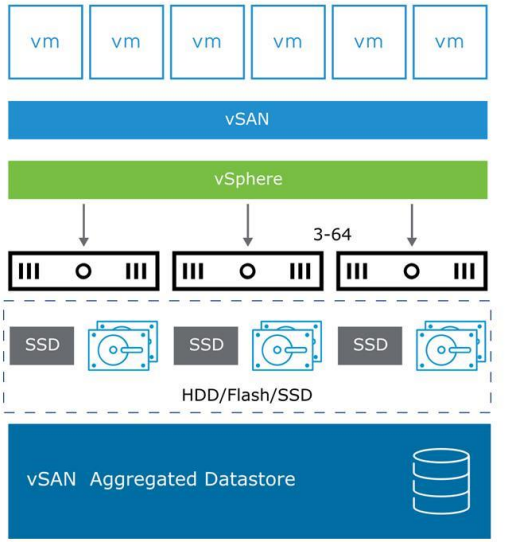

vSAN datastores help administrators use software-defined storage in the following ways:

- Storage policy per VM architecture: With multiple policies per datastore, each VM can have different storage.
- vSphere and vCenter Server integration: vSAN capability is built in and requires no appliance. You create a vSAN cluster, like vSphere HA or vSphere DRS.
- Scale-out storage: Up to 64 ESXi hosts can be in a cluster. Scale out by populating new nodes in the cluster.
- Built-in resiliency: The default vSAN storage policy establishes RAID 1 redundancy for all VMs.

https://t.me/learningnets

6-70 Disk Groups

Disk groups are vSAN management constructs on all ESXi hosts in a vSAN cluster. A host can include a maximum of five disk groups.

The disk groups are combined to create a single vSAN datastore. A disk group requires:

- One flash device for caching
- One to seven capacity devices for storage

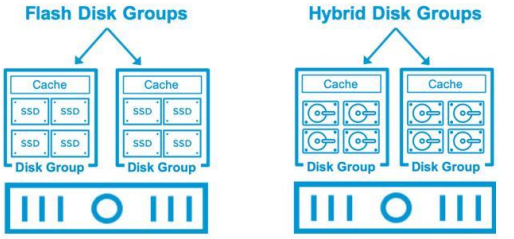

vSAN uses the concept of disk groups to pool together cache devices and capacity devices as single management constructs. A disk group is a pool of one cache device and one to seven capacity devices

https://t.me/learningnets

6-71 vSAN Hardware Requirements

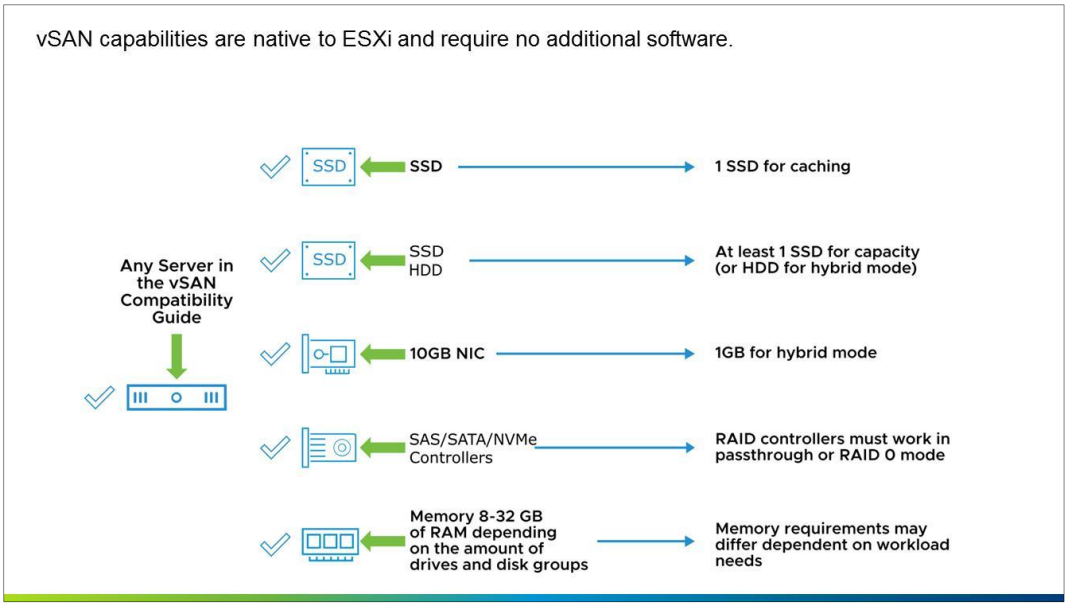

vSAN requires several hardware components that hosts do not normally have:

- One Serial Attached SCSI (SAS), SATA solid-state drive (SSD), or PCIe flash device and one to seven magnetic drives for each hybrid disk group.
- One SAS, SATA SSD, or PCIe flash device and one to seven flash disks with flash capacity enabled for all-flash disk groups.
- Dedicated 1 Gbps network (10 Gbps is recommended) for hybrid disk groups.
- Dedicated 10 Gbps network for all-flash disk groups.

  1 Gbps network speeds result in detrimental congestion for an all-flash architecture and are unsupported.
- The vSAN network must be configured for IPv4 or IPv6 and support unicast.

https://t.me/learningnets

In addition, each host should have a minimum of 32 GB of memory to accommodate a maximum number of five disk groups and a maximum number of seven capacity devices per disk group.

https://t.me/learningnets

6-72 Viewing the vSAN Datastore Summary

The **Summary** tab of the vSAN datastore shows the general vSAN configuration information.

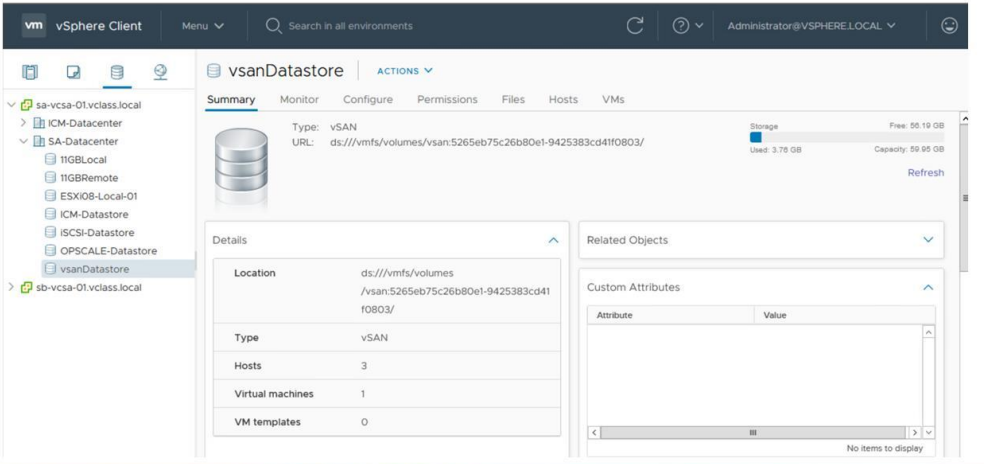

https://t.me/learningnets

6-73 Objects in vSAN Datastores

vSAN storage is object-based and policy-driven.

VMs created on a vSAN datastore include the following objects:

- One VM home namespace
- One or more VMDK objects
- Thin-provisioned VM swap object
- One or more VM memory objects
- vSAN performance data objects
- iSCSI LUN objects

vSAN Object	Traditional VM Files
VM home namespace	`.nvram`, `.vmsd`, `.vmx`, `vmx-*.vswp`, `.log`, `.hlog`
VMDK	`-flat.vmdk`
VM swap	`.vswp`
VM memory	`.vmem`
Snapshot delta	`-00000#-delta.vmdk`, `-00000#-sesparse.vmdk`

A vSAN cluster stores and manages data as flexible data containers called objects. When you provision a VM on a vSAN datastore, a set of objects is created:

- VM home namespace: Stores the virtual machine metadata (configuration files)
- VMDK: Virtual machine disk
- VM swap: Virtual machine swap file, which is created when the VM is powered on
- VM memory: Virtual machine's memory state when a VM is suspended or when a snapshot is taken of a VM and its memory state is preserved
- Snapshot delta: Created when a virtual machine snapshot is taken

https://t.me/learningnets

6-74 VM Storage Policies

Storage policies define how objects that are included in a VM are stored.

Storage policies have the following characteristics:

- Based on storage capabilities
- Defined for a VM or disk at the time of deployment
- Can be applied later
- Can be changed at any time
- Cannot be deleted if they are in use

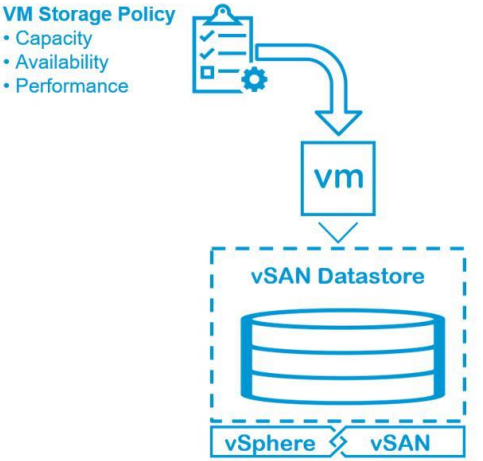

VM storage policies are a set of rules that you configure for VMs. Each storage policy reflects a set of capabilities that meet the availability, performance, and storage requirements of the application or service-level agreement for that VM.

You should create storage policies before deploying the VMs that require these storage policies. You can apply and update storage policies after deployment.

A vSphere administrator who is responsible for the deployment of VMs can select policies that are created based on storage capabilities.

Based on the policy that is selected for the object VM, these capabilities are pushed back to the vSAN datastore. The object is created across ESXi hosts and disk groups to satisfy these policies.

https://t.me/learningnets

6-75 Viewing VM Settings for vSAN Information

The consumption of vSAN storage is based on the VM's storage policy.
The VM's hard disk view provides the following information:

- A display of the VM storage policy
- The location of disk files on a vSAN datastore

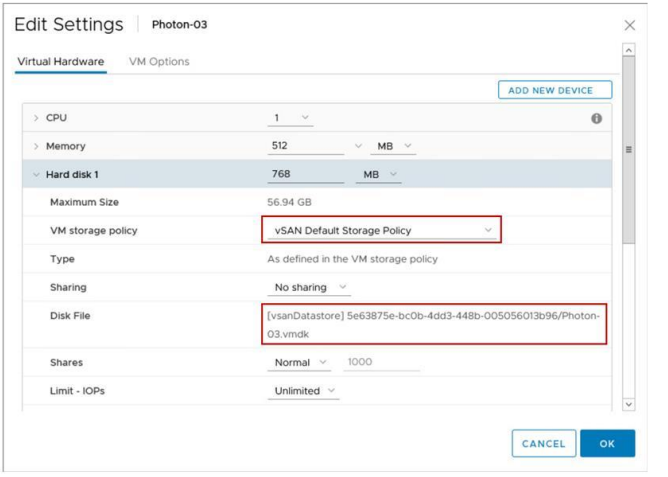

https://t.me/learningnets

6-76 Lab 15: Using a vSAN Datastore

View a vSAN datastore configuration and a virtual machine's components on the vSAN datastore:

1. View a vSAN Datastore Configuration
2. View the vSAN Default Storage Policy
3. View a Virtual Machine on the vSAN Datastore

https://t.me/learningnets

6-77 Review of Learner Objectives

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

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Explain the purpose of vSAN storage policies

https://t.me/learningnets

6-78 Virtual Beans: Storage

As a Virtual Beans administrator, you are planning how to use NAS and iSCSI storage with vSphere:

- For NAS storage, you can create one or more NFS datastores and share them across ESXi hosts:
  - Use the datastores to hold templates, VMs, and vCenter Server Appliance backups.
- For iSCSI storage, you can create one or more iSCSI datastores and share them across ESXi hosts:
  - Use the datastores to hold templates and VMs.

https://t.me/learningnets

6-79 Activity: Using vSAN Storage at Virtual Beans (1)

As a Virtual Beans administrator, you think that vSAN storage is the best option for the company's new storage requirements. What are the benefits to Virtual Beans of using vSAN storage?

https://t.me/learningnets

6-80 Activity: Using vSAN Storage at Virtual Beans (2)

What are the benefits to Virtual Beans of using vSAN storage?
Benefits include (but are not limited to):

- You can use the vSphere Client to manage the vSAN configuration. No separate user interface is necessary.
- vSphere administrators do not need special storage hardware training.
- You can use vSAN storage policies to define specific levels of service for a VM.
- You can expand the vSAN capacity by adding one or more hosts to the vSAN cluster (also known as scale out).

https://t.me/learningnets

6-81 Key Points

- ESXi hosts support various storage technologies: Direct-attached storage, Fibre Channel, FCoE, iSCSI, and NAS.
- You use VMFS and NFS datastores to hold VM files.
- Shared storage is integral to vSphere features such as vSphere vMotion, vSphere HA, and vSphere DRS.
- vSAN clusters direct-attached server disks to create shared storage designed for VMs.

Questions?

https://t.me/learningnets

https://t.me/learningnets

Module 7

Virtual Machine Management

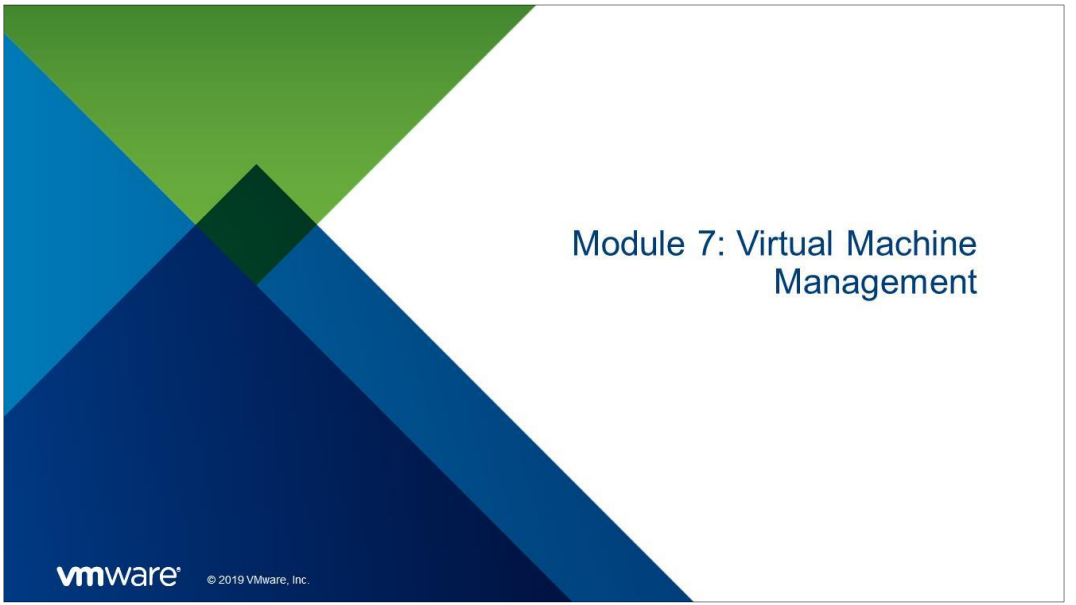

https://t.me/learningnets

7-2 Importance

Virtual machines are the foundation of your virtual infrastructure. Managing VMs effectively requires skills in creating templates and clones, modifying VMs, migrating VMs, taking snapshots, and protecting the VMs through replication and backups.

https://t.me/learningnets

7-3 Module Lessons

1. Creating Templates and Clones
2. Working with Content Libraries
3. Modifying Virtual Machines
4. Migrating VMs with vSphere vMotion
5. Enhanced vMotion Compatibility
6. Migrating VMs with vSphere Storage vMotion
7. Creating Virtual Machine Snapshots
8. vSphere Replication and Backup

https://t.me/learningnets

7-4 Virtual Beans: VM Management

Virtual Beans wants to automate its processes. It requires the following processes for the vSphere infrastructure:

- Provisioning and deploying VMs: Performing fast, error-free, consistent deployments
- Maintaining VMs: Patching and upgrading operating systems and applications
- Backing up VMs: Making daily backups of application data
- Disaster recovery and business continuity: Moving VMs between the primary and secondary data center

As a Virtual Beans administrator, you must recognize the options available for these processes. Then, you can create effective processes for managing VMs in your data center.

https://t.me/learningnets

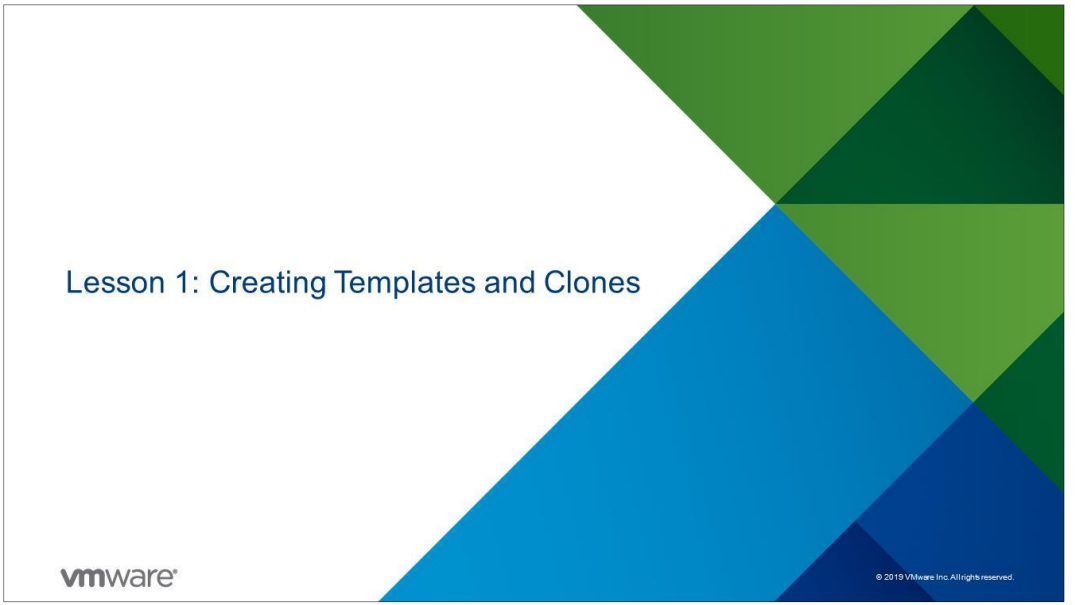

https://t.me/learningnets

7-6 Learner Objectives

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

- Create a template of a virtual machine
- Deploy a virtual machine from a template
- Clone a virtual machine
- Create customization specifications for guest operating systems

https://t.me/learningnets

7-7 About Templates

A template is a master copy of a virtual machine. You use templates to create and provision new VMs.

A template typically includes:

- A guest operating system
- One or more applications
- A specific VM configuration
- VMware Tools

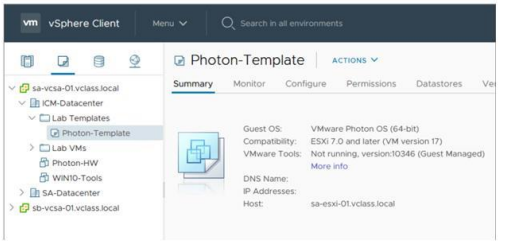

Creating templates makes the provisioning of virtual machines much faster and less error-prone than provisioning physical machines and creating a VM by using the New Virtual Machine wizard.

Templates coexist with VMs in the inventory. You can organize collections of VMs and templates into arbitrary folders and apply permissions to VMs and templates. You can change VMs into templates without having to make a full copy of the VM files and create an object.

You can deploy a VM from a template. The deployed VM is added to the folder that you selected when creating the template.

https://t.me/learningnets

7-8 Creating a Template: Clone VM to Template

You can create templates using different methods. One method is to clone the VM to a template. The VM can be powered on or off.

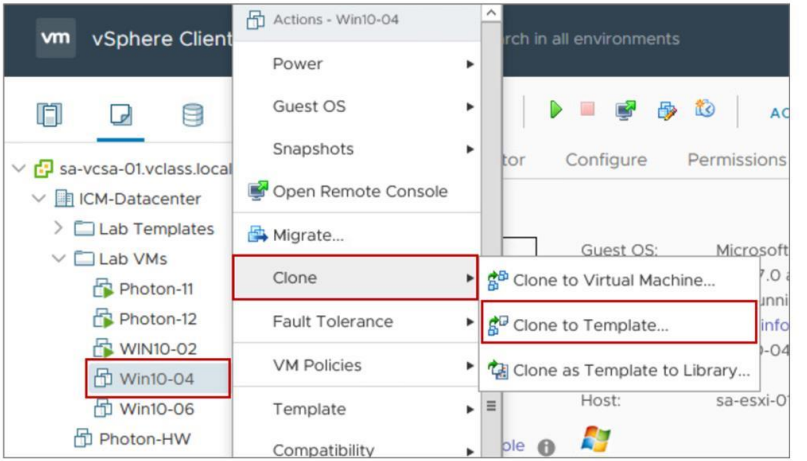

The **Clone to Template** option offers you a choice of format for storing the VM's virtual disks:

- Same format as source
- Thin-provisioned format
- Thick-provisioned lazy-zeroed format
- Thick-provisioned eager-zeroed format

https://t.me/learningnets

7-9 Creating a Template: Convert VM to Template

You can create a template by converting a VM to a template. In this case, the VM must be powered off.

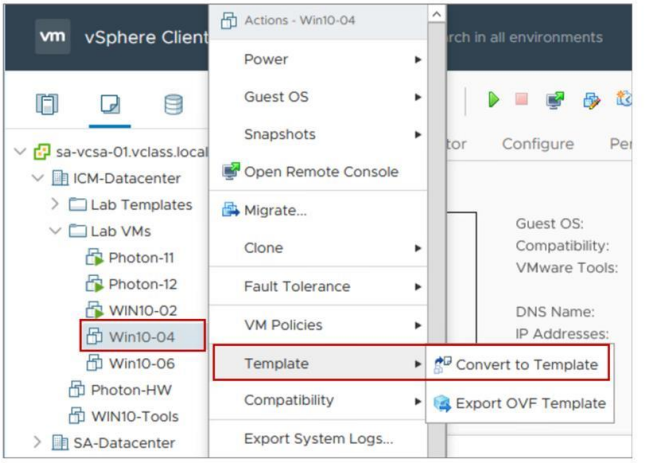

The **Convert to Template** option does not offer a choice of format and leaves the VM's disk file intact.

https://t.me/learningnets

7-10 Creating a Template: Clone a Template

You can create a template from an existing template, or clone a template.

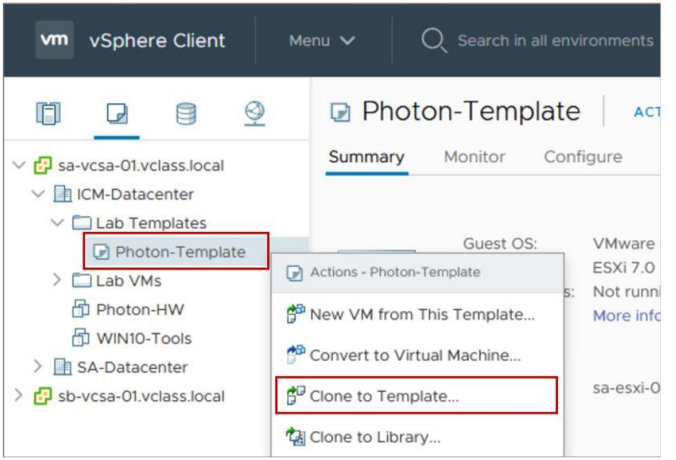

https://t.me/learningnets

7-11 Updating Templates

You update a template to include new patches, make system changes, and install new applications.

To update a template:

1. Convert the template to a VM.
2. Place the VM on an isolated network to prevent user access.
3. Make appropriate changes to the VM.
4. Convert the VM to a template.

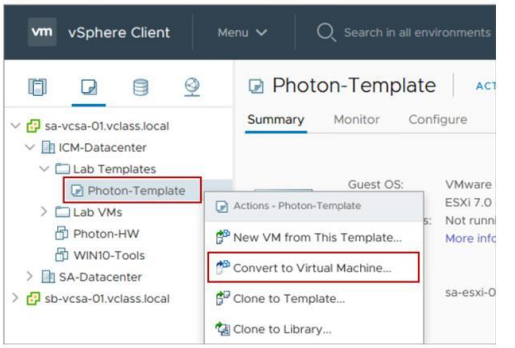

To update your template to include new patches or software, you do not need to create a template. Instead, you convert the template to a VM. You can then power on the VM.

For added security, you might want to prevent users from accessing the VM while you update it. To prevent access, either disconnect the VM from the network or place it on an isolated network.

Log in to the VM's guest operating system and apply the patch or install the software. When you finish, power off the VM and convert it to a template again.

https://t.me/learningnets

7-12 Deploying VMs from a Template

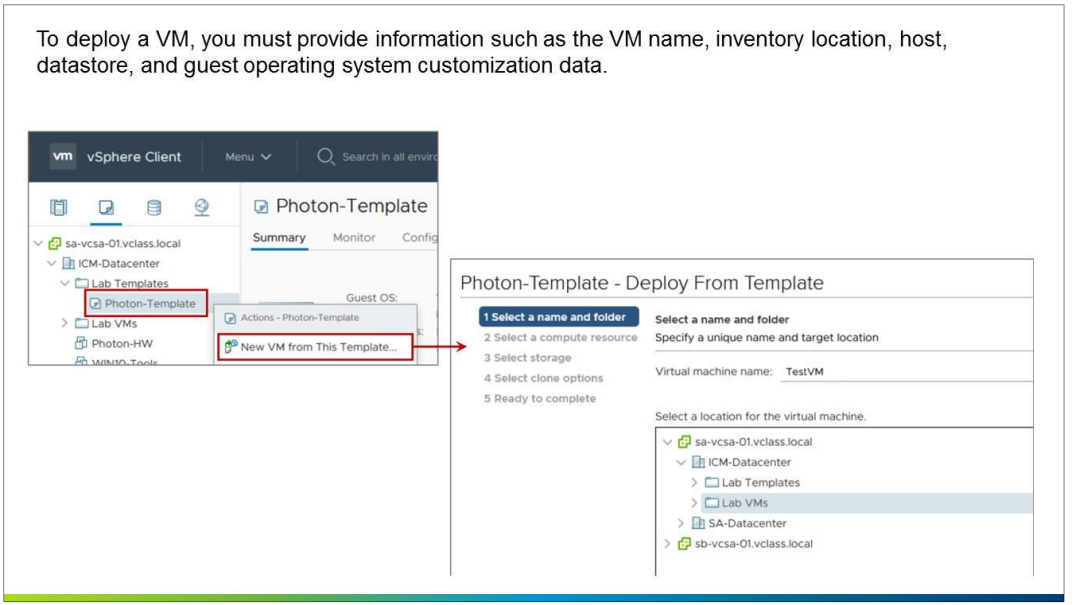

When you place ISO files in a content library, the ISO files are available only to VMs that are registered on an ESXi host that can access the datastore where the content library is located. These ISO files are not available to VMs on hosts that cannot see the datastore on which the content library is located.

https://t.me/learningnets

7-13 Cloning Virtual Machines

Cloning a VM creates a VM that is an exact copy of the original:

- Cloning is an alternative to deploying a VM.
- During cloning, the VM can be powered on or off.

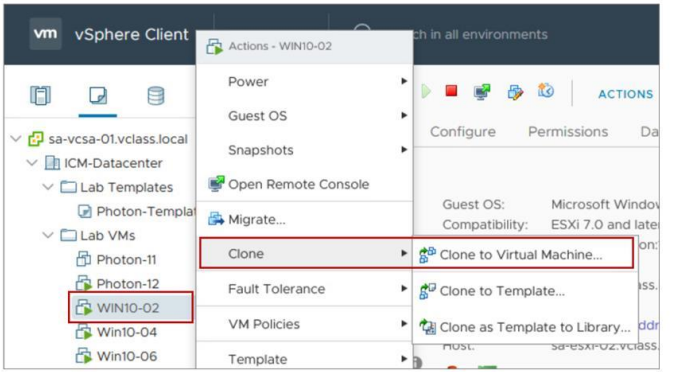

To clone a VM, you must be connected to vCenter Server. You cannot clone VMs if you use VMware Host Client to manage a host directly.

When you clone a VM that is powered on, services and applications are not automatically quiesced when the VM is cloned.

When deciding whether to clone a VM or deploy a VM from a template, consider the following points:

- VM templates use storage space, so you must plan your storage space requirements accordingly.
- Deploying a VM from a template is quicker than cloning a running VM, especially when you must deploy many VMs at a time.
- When you deploy many VMs from a template, all the VMs start with the same base image. Cloning many VMs from a running VM might not create identical VMs, depending on the activity happening within the VM when the VM is cloned.

https://t.me/learningnets

7-14 Guest Operating System Customization

You customize the guest operating system to make VMs, created from the same template or clone, unique.
By customizing a guest operating system, you can change information, including the following details:

- Computer name
- Network settings
- License settings
- Windows Security Identifier

Customizing the guest operating system prevents conflicts that might occur when you deploy a VM and a clone with identical guest OS settings simultaneously.

https://t.me/learningnets

7-15 About Customization Specifications

You can create a customization specification to prepare the guest operating system:

- Specifications are stored in the vCenter Server database.
- Windows and Linux guests are supported.

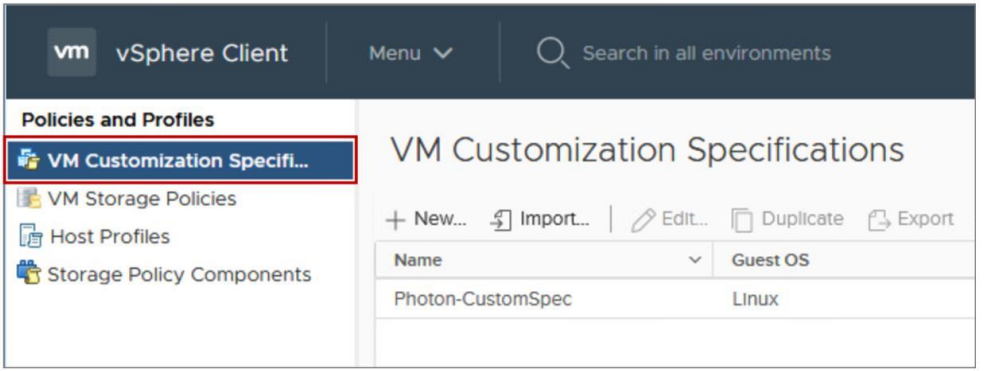

To manage customization specifications, select **Policies and Profiles** from the **Menu** drop-down menu.

On the VM Customization Specifications pane, you can create specifications or manage existing ones.

https://t.me/learningnets

7-16 Customizing the Guest Operating System

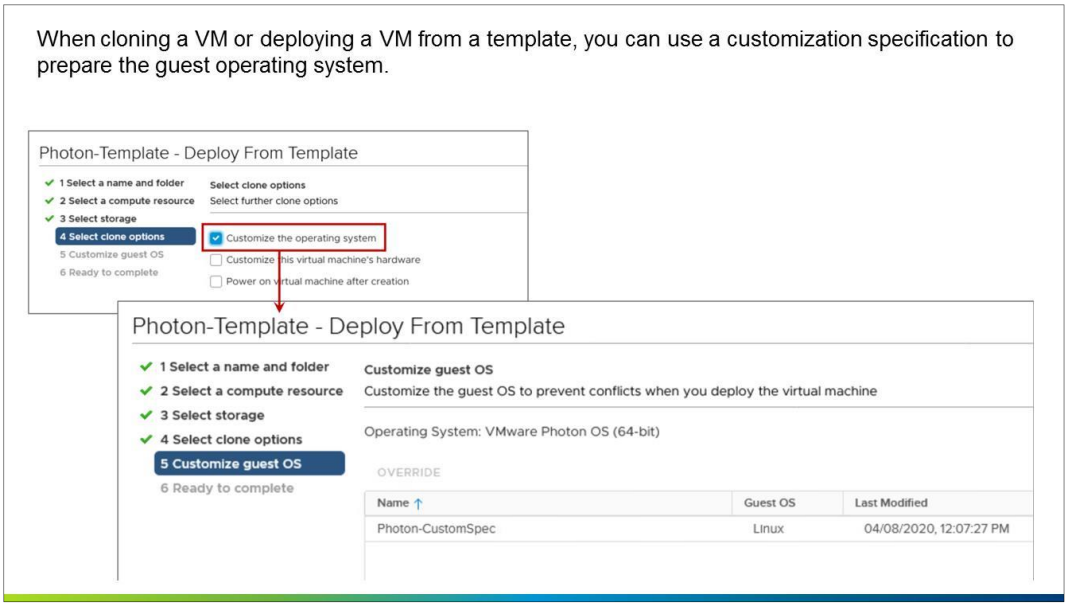

You can define the customization settings by using an existing customization specification during cloning or deployment. You create the specification ahead of time. During cloning or deployment, you can select the customization specification to apply to the new VM.

VMware Tools must be installed on the guest operating system that you want to customize.

The guest operating system must be installed on a disk attached to SCSI node 0:0 in the VM configuration.

For more about guest operating system customization, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html.

https://t.me/learningnets

7-17 About Instant Clones

You can use Instant Clone Technology to create a powered-on VM from the running state of another powered-on VM:

- The processor state, virtual device state, memory state, and disk state of the destination (child) VM are identical to the states of the source (parent) VM.
- Snapshot-based disk sharing is used to provide storage efficiency and to improve the speed of the cloning process.

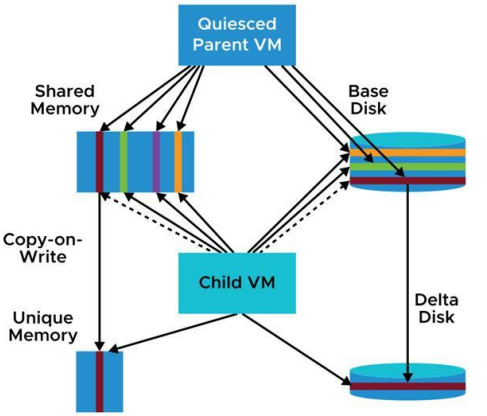

Through instant cloning, the source (parent) VM does not lose its state because of the cloning process. You can move to just-in-time provisioning, given the speed and state-persisting nature of this operation.

During an instant clone operation, the source VM is stunned for a short time, less than 1 second.

While the source VM is stunned, a new writable delta disk is generated for each virtual disk, and a checkpoint is taken and transferred to the destination VM.

The destination VM powers on by using the source's checkpoint.

After the destination VM is fully powered on, the source VM resumes running.

Instant clone VMs are fully independent vCenter Server inventory objects. You can manage instant clone VMs like regular VMs, without any restrictions.

https://t.me/learningnets

7-18 Use Cases for Instant Clones

Instant clone VMs have various uses:

- Virtual desktop infrastructure
- Rapid scale-out: Container hosts, big data, and Hadoop worker nodes
- DevTest: Quickly and efficiently replicate VMs and test beds with the same running state
- DevOps: Replicate VMs from staging to production, and the converse, with the identical running state

In vSphere 6.7 and later, you can create instant clones of VMs only through API calls.

In vSphere 7, instant clone guest customization is supported (Linux only).

Instant cloning is convenient for large-scale application deployments because it ensures memory efficiency, and you can create many VMs on a single host.

To avoid network conflicts, you can customize the virtual hardware of the destination VM during the instant cloning operation. For example, you can customize the MAC addresses of the virtual NICs or the serial and parallel port configurations of the destination VM.

Starting with vSphere 7, you can customize the guest operating system for Linux VMs only. You can customize networking settings such as IP address, DNS server, and the gateway. You can change these settings without having to power off or restart the VM.

https://t.me/learningnets

7-19 Lab 16: Using VM Templates: Creating Templates and Deploying VMs

Create a VM template, create a customization specification, and deploy VMs from a template:

1. Create a Virtual Machine Template
2. Create Customization Specifications
3. Deploy Virtual Machines from a Template

https://t.me/learningnets

7-20 Review of Learner Objectives

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

- Create a template of a virtual machine
- Deploy a virtual machine from a template
- Clone a virtual machine
- Create customization specifications for guest operating systems

https://t.me/learningnets

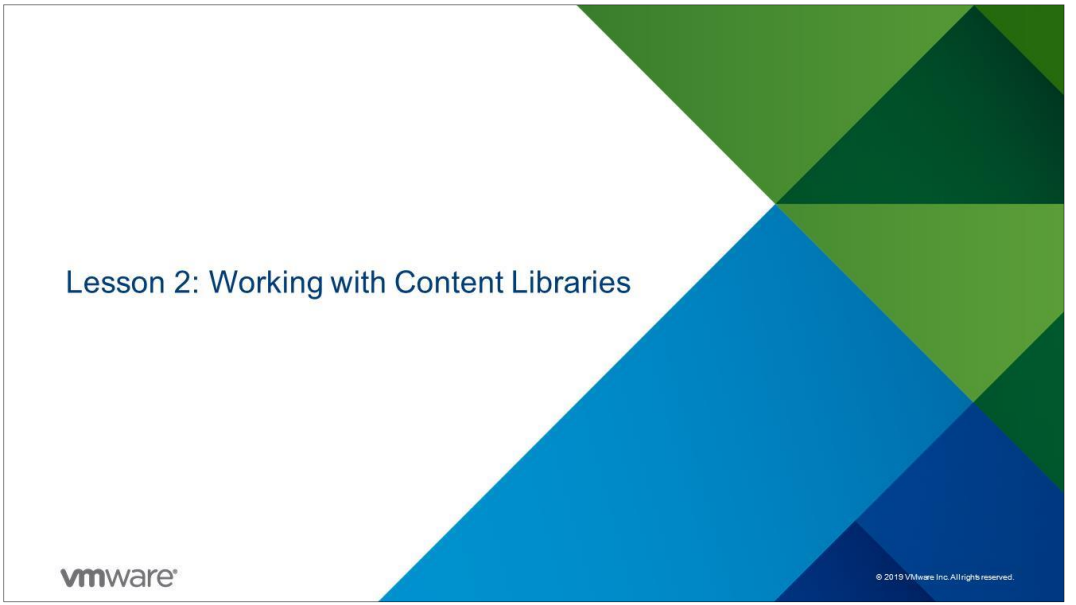

https://t.me/learningnets

7-22 Learner Objectives

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

- Identify the benefits of a content library
- Recognize types of content libraries
- Deploy a virtual machine from a content library

https://t.me/learningnets

7-23 About Content Libraries

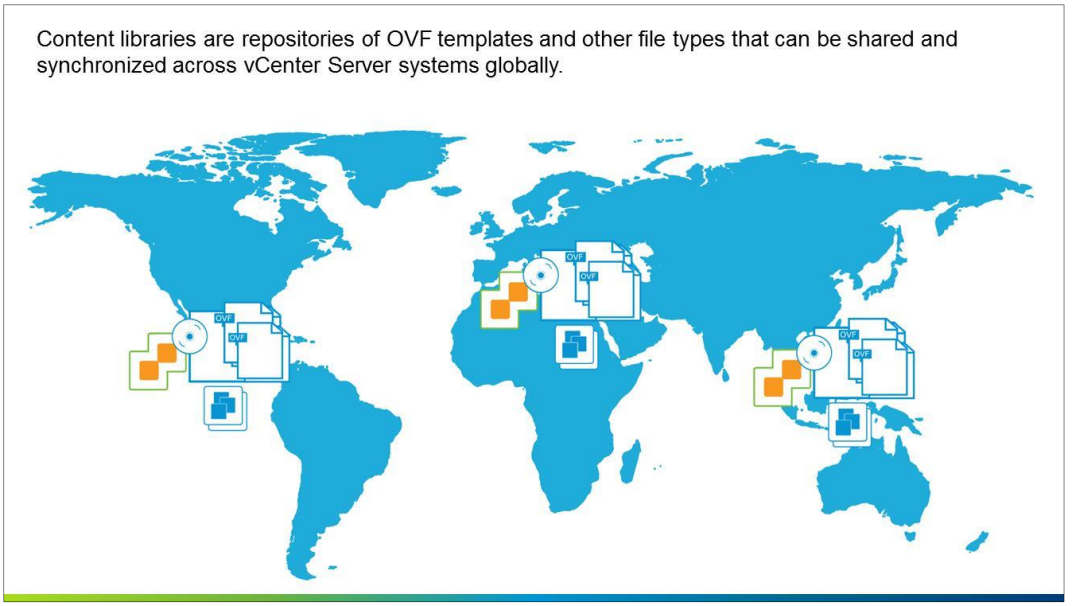

Organizations might have multiple vCenter Server instances in data centers around the globe. On these vCenter Server instances, organizations might have a collection of templates, ISO images, and so on. The challenge is that all these items are independent of one another, with different versions of these files and templates on various vCenter Server instances.

The content library is the solution to this challenge. IT can store OVF templates, ISO images, or any other file types in a central location. The templates, images, and files can be published, and other content libraries can subscribe to and download content. The content library keeps content up to date by periodically synchronizing with the publisher, ensuring that the latest version is available.

https://t.me/learningnets

7-24 Benefits of Content Libraries

Storage and consistency are key reasons to install and use a content library.
Using content libraries, administrators can perform the following functions:

- Store, version, and share content.
- Perform distributed file management.
- Synchronize content libraries across sites and vCenter Server instances.
- Mount an ISO file directly from a content library.
- Perform live updates of VM templates

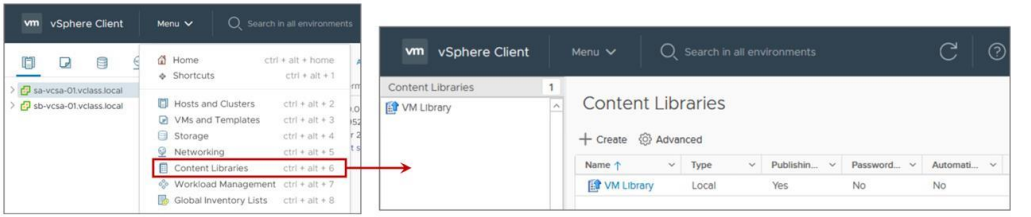

Sharing content and ensuring that the content is kept up to date are major tasks.

For example, for a main vCenter Server instance, you create a central content library to store the master copies of OVF templates, ISO images, and other file types. When you publish this content library, other libraries, which might be located anywhere in the world, can subscribe and download an exact copy of the data.

When an OVF template is added, modified, or deleted from the published catalog, the subscriber synchronizes with the publisher, and the libraries are updated with the latest content.

Starting with vSphere 7, you can update a template while simultaneously deploying VMs from the template. In addition, the content library keeps two copies of the VM template, the previous and current versions. You can roll back the template to reverse changes made to the template.

https://t.me/learningnets

7-25 Types of Content Libraries

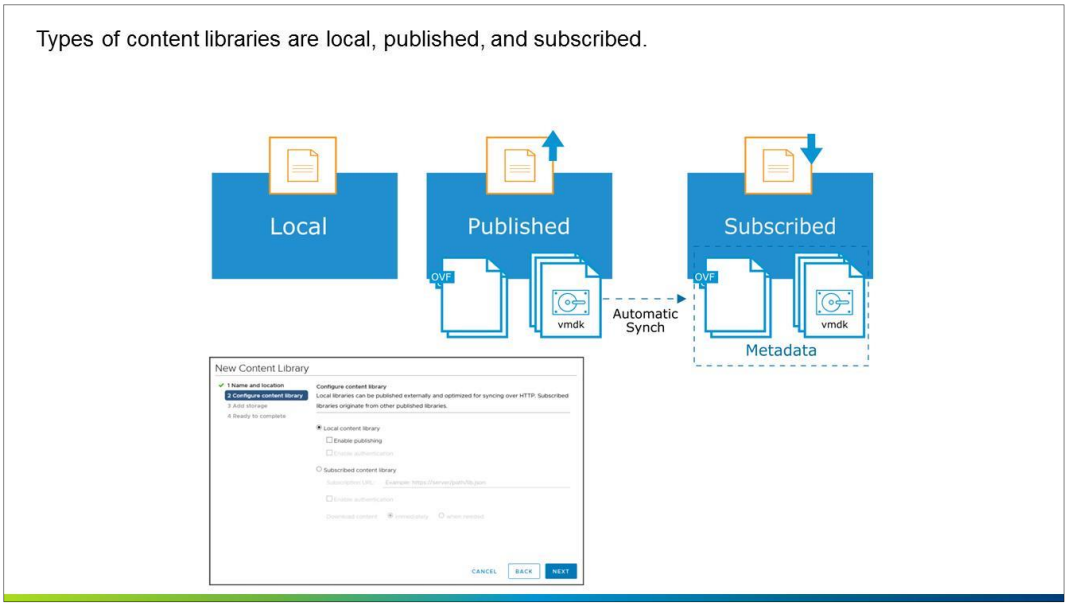

You can create a local library as the source for content that you want to save or share. You create the local library on a single vCenter Server instance. You can then add or remove items to and from the local library.

You can publish a local library, and this content library service endpoint can be accessed by other vCenter Server instances in your virtual environment. When you publish a library, you can configure the authentication method, which a subscribed library must use to authenticate to it.

You can create a subscribed library and populate its content by synchronizing it to a published library. A subscribed library contains copies of the published library files or only the metadata of the library items.

The published library can be on the same vCenter Server instance as the subscribed library, or the subscribed library can reference a published library on a different vCenter Server instance.

You cannot add library items to a subscribed library. You can add items only to a local or published library.

https://t.me/learningnets

After synchronization, both libraries contain the same items, or the subscribed library contains the metadata for the items.

https://t.me/learningnets

7-26 Adding VM Templates to a Content Library

Library items include VM templates, vApp templates, or other VMware objects that can be contained in a content library.

Clone a template to the content library.

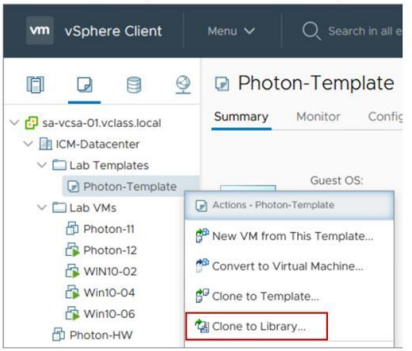

Clone a VM to a template in the content library.

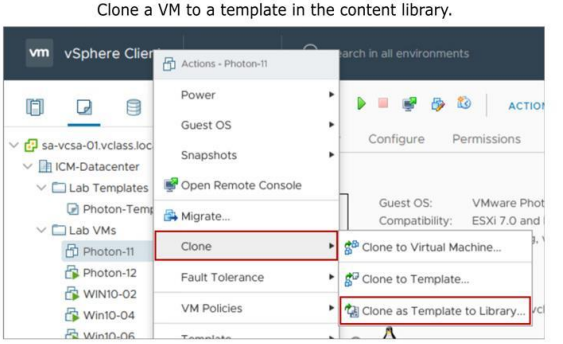

VMs and vApps have several files, such as log files, disk files, memory files, and snapshot files that are part of a single library item. You can create library items in a specific local library or remove items from a local library. You can also upload files to an item in a local library so that the libraries subscribed to it can download the files to their NFS or SMB server, or datastore.

https://t.me/learningnets

7-27 Deploying VMs from Templates in a Content Library

The templates in the content library can be used to deploy VMs and vApps.

Each VM template, vApp template, or other type of file in a library is a library item.

You can also mount an ISO file directly from a content library.

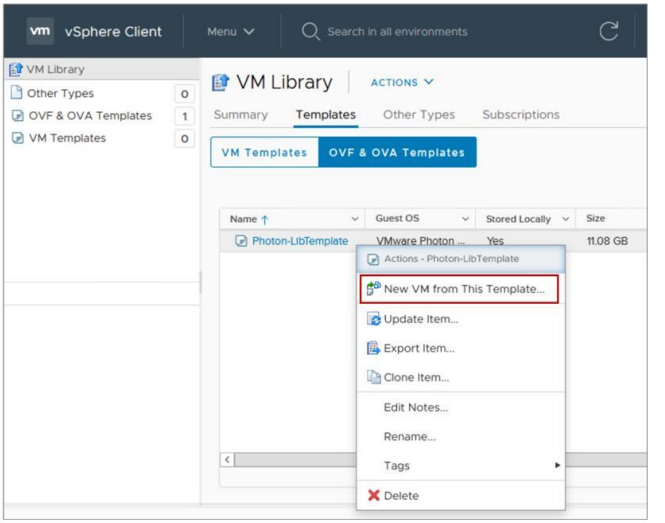

https://t.me/learningnets

7-28 Lab 17: Using Content Libraries

Create a content library to clone and deploy virtual machines:

1. Create a Content Library
2. Clone a VM Template to a Template in a Content Library
3. Deploy a VM from a VM Template in the Content Library

https://t.me/learningnets

7-29 Review of Learner Objectives

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

- Identify the benefits of a content library
- Recognize types of content libraries
- Deploy a virtual machine from a content library

https://t.me/learningnets

https://t.me/learningnets

7-31 Learner Objectives

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

- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk

https://t.me/learningnets

7-32 Modifying Virtual Machine Settings

You can modify a VM's configuration by editing the VM's settings:

- Add virtual hardware:
  - You can add some hardware while the VM is powered on.
- Remove virtual hardware:
  - You can remove some hardware only when the VM is powered off.
- Set VM options.
- Control a VM's CPU and memory resources.

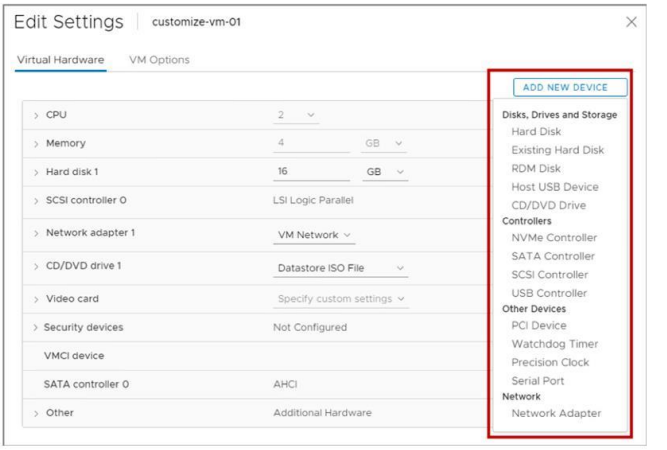

You might have to modify a VM's configuration, for example, to add a network adapter or a virtual disk. You can make all VM changes while the VM is powered off. Some VM hardware changes can be made while the VM is powered on.

vSphere 7.0 makes the following virtual devices available:

- Watchdog timer: Virtual device used to detect and recover from operating system problems. If a failure occurs, the watchdog timer 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 watchdog timer is useful with high availability solutions such as Red Hat High Availability and the MS SQL failover cluster. This device is also useful on VMware Cloud and in hosted environments for implementing custom failover logic to reset or power off VMs.

https://t.me/learningnets

- Precision Clock: Virtual device that presents the ESXi host's system time to the guest OS. Precision Clock helps the guest operating system achieve clock accuracy in the 1 millisecond range. The guest operating system uses Precision Clock time as reference time. Precision Clock is not directly involved in guest OS time synchronization.

  Precision Clock is useful when precise timekeeping is a requirement for the application, such as for financial services applications. Precision Clock is also useful when precise time stamps are required on events that track financial transactions.

- Virtual SGX: Virtual device that exposes Intel's SGX technology to VMs. Intel's SGX technology prevents unauthorized programs or processes from accessing certain regions in memory. Intel SGX meets the needs of the Trusted Computing Industry.

  Virtual SGX is useful for applications that must conceal proprietary algorithms and encryption keys from unauthorized users. For example, cloud service providers cannot inspect a client's code and data in a virtual SGX-protected environment.

https://t.me/learningnets

7-33 Hot-Pluggable Devices

With the hot plug option, you can add resources to a running VM. Examples of hot-pluggable devices:

- USB controllers
- Ethernet adapters
- Hard disk devices

With supported guest operating systems, you can also add CPU and memory while the VM is powered on.

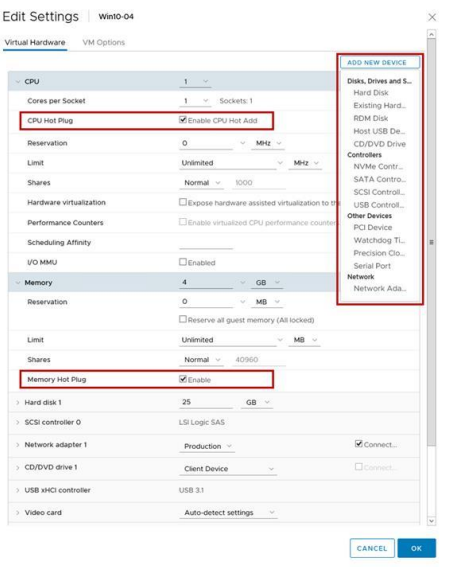

Adding devices to a physical server or removing devices from a physical server requires that you physically interact with the server in the data center. When you use VMs, resources can be added dynamically without a disruption in service. You must shut down a VM to remove hardware, but you can reconfigure the VM without entering the data center.

You can add CPU and memory while the VM is powered on. These features are called the CPU Hot Add and Memory Hot Plug, which are supported only on guest operating systems that support hot-pluggable functionality. These features are disabled by default. To use these hot-plug features, the following requirements must be satisfied:

- You must install VMware Tools.
- The VM must use hardware version 11 or later.
- The guest operating system in the VM must support CPU and memory hot-plug features.
- The hot-plug features must be enabled in the **CPU** or **Memory** settings on the **Virtual Hardware** tab.

https://t.me/learningnets

If virtual NUMA is configured with virtual CPU hot-plug settings, the VM is started without virtual NUMA. Instead, the VM uses UMA (Uniform Memory Access).

https://t.me/learningnets

7-34 Dynamically Increasing Virtual Disk Size

You can increase the size of a virtual disk that belongs to a powered-on VM.

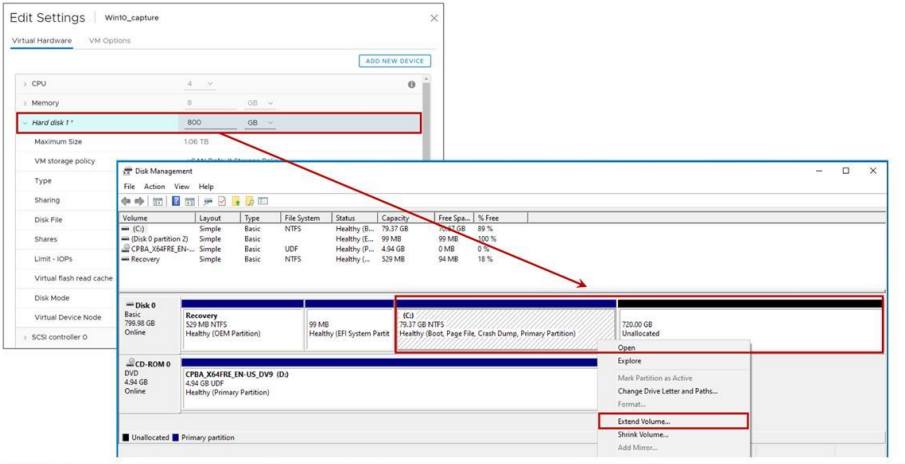

When you increase the size of a virtual disk, the VM must not have snapshots attached.

After you increase the size of a virtual disk, you might need to increase the size of the file system on this disk. Use the appropriate tool in the guest OS to enable the file system to use the newly allocated disk space.

https://t.me/learningnets

7-35 Inflating Thin-Provisioned Disks

Thin-provisioned virtual disks can be converted to a thick, eager-zeroed format.
To inflate a thin-provisioned disk:

- The VM must be powered off.
- Right-click the VM's file with the .vmdk extension and select **Inflate**.

Or you can use vSphere Storage vMotion and select a thick-provisioned disk as the destination.

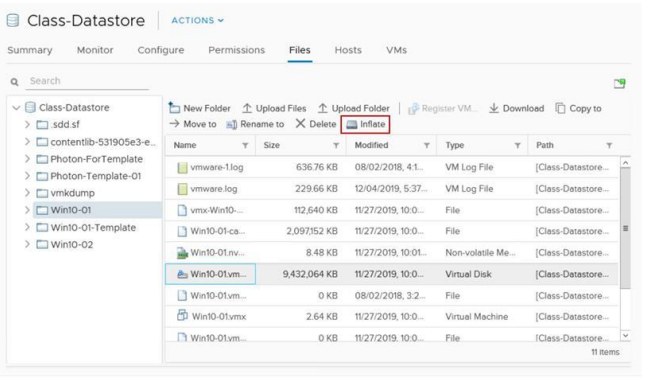

When you inflate a thin-provisioned disk, the inflated virtual disk occupies the entire datastore space originally provisioned to it. Inflating a thin-provisioned disk converts a thin disk to a virtual disk in thick-provisioned format.

https://t.me/learningnets

7-36 VM Options: General Settings

You can use the **VM Options** tab to modify properties such as the display name for the VM and the type of guest operating system that is installed.

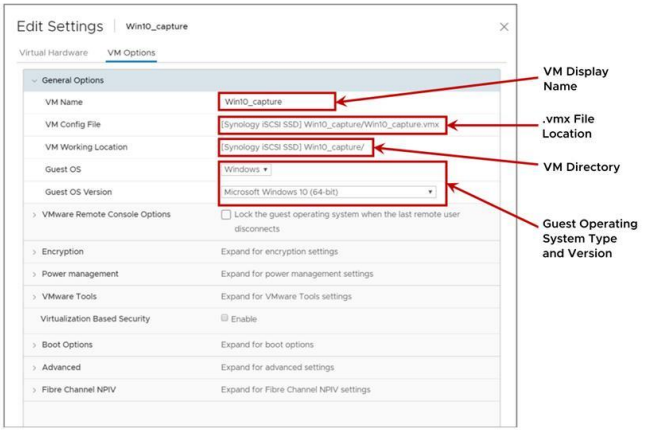

Under General Options, you can view the location and name of the configuration file (with the .vmx extension) and the location of the VM's directory.

You can select the text for the configuration file and the working location to copy and paste them into a document. However, only the display name and the guest operating system type can be modified.

Changing the display name does not change the names of all the VM files or the directory that the VM is stored in. When a VM is created, the filenames and the directory name associated with the VM are based on its display name. But changing the display name later does not modify the filename and the directory name.

https://t.me/learningnets

7-37 VM Options: VMware Tools Settings

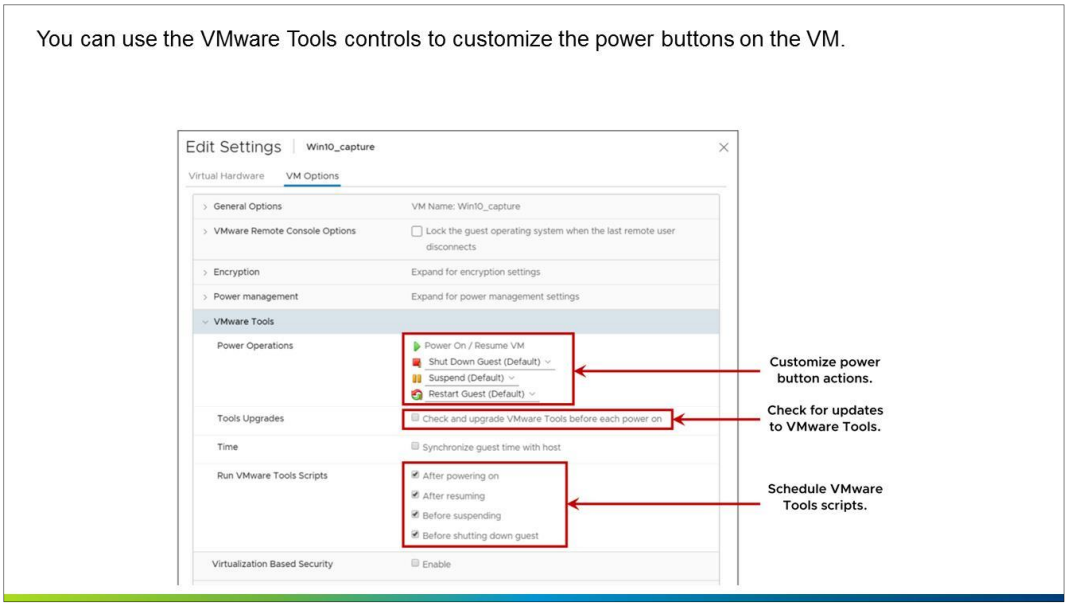

When you use the VMware Tools controls to customize the power buttons on the VM, the VM must be powered off.

You can select the **Check and upgrade VMware Tools before each power on** check box to check for a newer version of VMware Tools. If a newer version is found, VMware Tools is upgraded when the VM is power cycled.

When you select the **Synchronize guest time with host** check box, the guest operating system's clock synchronizes with the host.

For information about time keeping best practices for the guest operating systems that you use, see VMware knowledge base articles 1318 at http://kb.vmware.com/kb/1318 and 1006427 at http://kb.vmware.com/kb/1006427.

https://t.me/learningnets

7-38 VM Options: VM Boot Settings

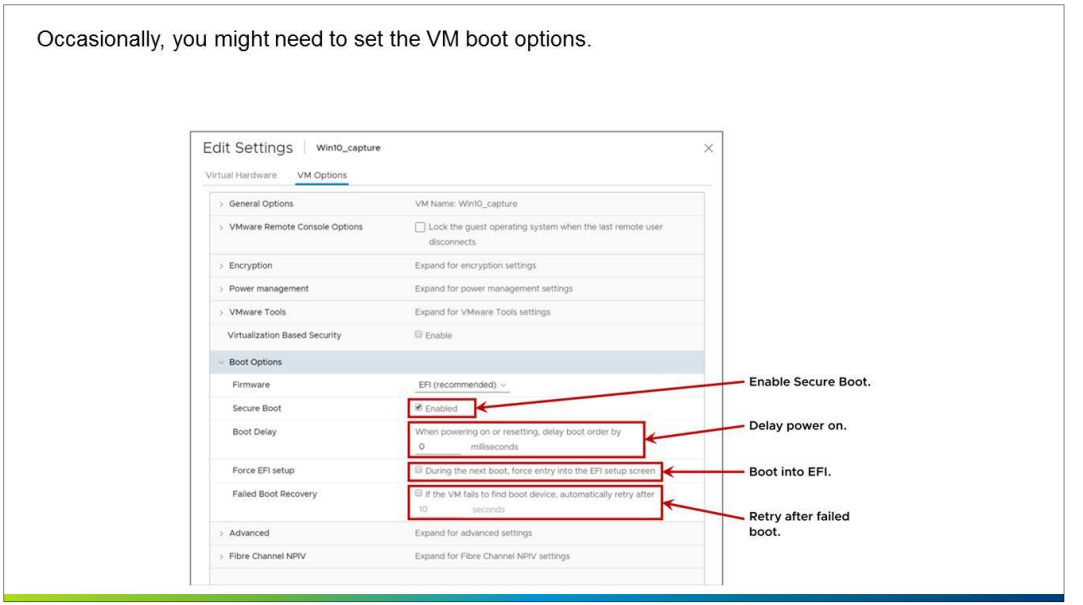

When you build a VM and select a guest operating system, **BIOS** or **EFI** is selected automatically, depending on the firmware supported by the operating system. Mac OS X Server guest operating systems support only Extensible Firmware Interface (EFI). If the operating system supports BIOS and EFI, you can change the boot option as needed. However, you must change the option before installing the guest OS.

UEFI Secure Boot is a security standard that helps ensure that your PC boots use only software that is trusted by the PC manufacturer. In an OS that supports UEFI Secure Boot, each piece of boot software is signed, including the bootloader, the operating system kernel, and operating system drivers. If you enable Secure Boot for a VM, you can load only signed drivers into that VM.

With the **Boot Delay** value, you can set a delay between the time when a VM is turned on and the guest OS starts to boot. A delayed boot can help stagger VM start ups when several VMs are powered on.

https://t.me/learningnets

You can change the BIOS or EFI settings. For example, you might want to force a VM to start from a CD/DVD. The next time the VM powers on, it goes straight into the BIOS. A forced entry into the firmware setup is much easier than powering on the VM, opening a console, and quickly trying to press the F2 key.

With the Failed Boot Recovery setting, you can configure the VM to retry booting after 10 seconds (the default) if the VM fails to find a boot device.

https://t.me/learningnets

7-39 Removing VMs

You can remove a VM in the following ways:

- Remove from the inventory:
  - The VM is unregistered from the ESXi host and vCenter Server.
  - The VM's files remain on the disk.
  - The VM can later be registered (added) to the inventory.
- Delete from disk:
  - All VM files are permanently deleted from the datastore.
  - The VM is unregistered from the ESXi host and vCenter Server.

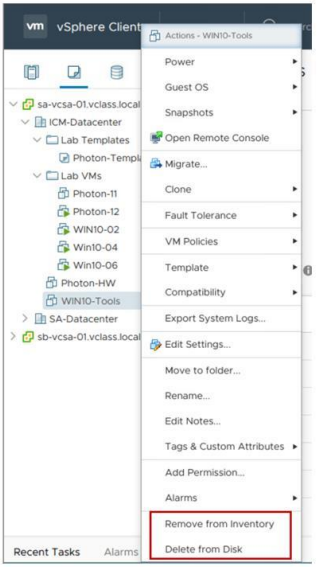

When a VM is removed from the inventory, its files remain at the same storage location, and the VM can be re-registered in the datastore browser.

https://t.me/learningnets

7-40 Lab 18: Modifying Virtual Machines

Modify a virtual machine's hardware and rename a virtual machine:

1. Adjust Memory Allocation on a Virtual Machine
2. Increase the Size of a VMDK File
3. Rename a Virtual Machine in the vCenter Server Inventory

https://t.me/learningnets

7-41 Review of Learner Objectives

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

- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk

https://t.me/learningnets

https://t.me/learningnets

7-43 Learner Objectives

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

- Recognize the types of VM migrations that you can perform within a vCenter Server instance and across vCenter Server instances
- Explain how vSphere vMotion works
- Verify vSphere vMotion requirements
- Migrate virtual machines using vSphere vMotion

https://t.me/learningnets

7-44 About VM Migration

Migration means moving a VM from one host, datastore, or vCenter Server instance to another host, datastore, or vCenter Server instance.
Depending on the power state of the VM that you migrate, migration can be cold or hot:

- A cold migration involves moving a powered-off or suspended VM to a new host.
- A hot migration involves moving a powered-on VM to a new host.

Depending on the VM resource type, you can perform different types of migrations.

Migration Type	Description
Compute resource only	Move VM, but not its storage, to another host. For a hot migration, vSphere vMotion is used to move the VM.
Storage only	Move a VM's storage, but not its host, to a new datastore. For a hot migration, vSphere Storage vMotion is used to move the VM.
Both compute resource and storage	Move a VM to another host and datastore. For a hot migration, vSphere vMotion and vSphere Storage vMotion are used to move the VM.

A deciding factor for using a particular migration technique is the purpose of performing the migration. For example, you might need to stop a host for maintenance but keep the VMs running. You use vSphere vMotion to migrate the VMs instead of performing a cold or suspended VM migration. If you must move a VM's files to another datastore to better balance the disk load or transition to another storage array, you use vSphere Storage vMotion.

Some migration techniques, such as vSphere vMotion migration, have special hardware requirements that must be met to function properly. Other techniques, such as a cold migration, do not have special hardware requirements to function properly.

You can perform the different types of migration on either powered-off (cold) or powered-on (hot) VMs.

https://t.me/learningnets

7-45 About vSphere vMotion

A vSphere vMotion migration moves a powered-on VM from one host to another. vSphere vMotion changes the compute resource only.
vSphere vMotion provides the following capabilities:

- Improvement in overall hardware use
- Continuous VM operation while accommodating scheduled hardware downtime
- vSphere DRS to balance VMs across hosts

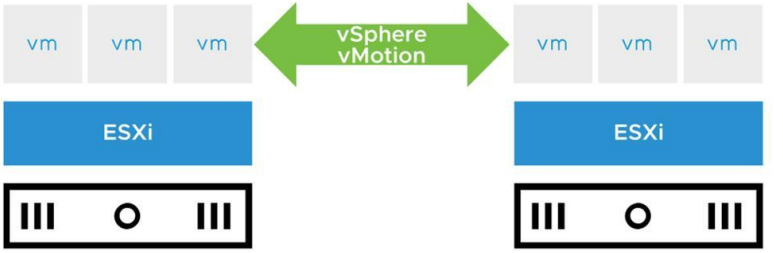

Using vSphere vMotion, you can migrate running VMs from one ESXi host to another ESXi host with no disruption or downtime. With vSphere vMotion, vSphere DRS can migrate running VMs from one host to another to ensure that the VMs have the resources that they require.

With vSphere vMotion, the entire state of the VM is moved from one host to another, but the data storage remains in the same datastore.

The state information includes the current memory content and all the information that defines and identifies the VM. The memory content includes transaction data and whatever bits of the operating system and applications are in memory. The definition and identification information stored in the state includes all the data that maps to the VM hardware elements, such as the BIOS, devices, CPU, and MAC addresses for the Ethernet cards.

https://t.me/learningnets

7-46 Enabling vSphere vMotion

To enable vSphere vMotion, you must configure a VMkernel port with the vSphere vMotion service enabled on the source and destination host.

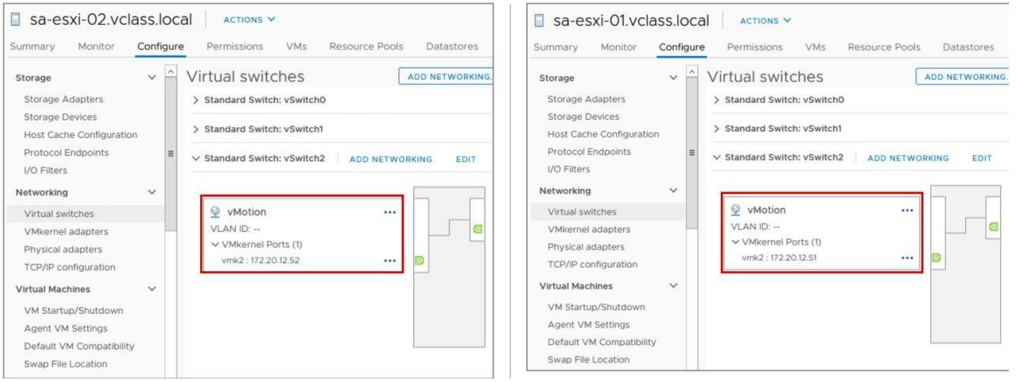

https://t.me/learningnets

7-47 vSphere vMotion Migration Workflow

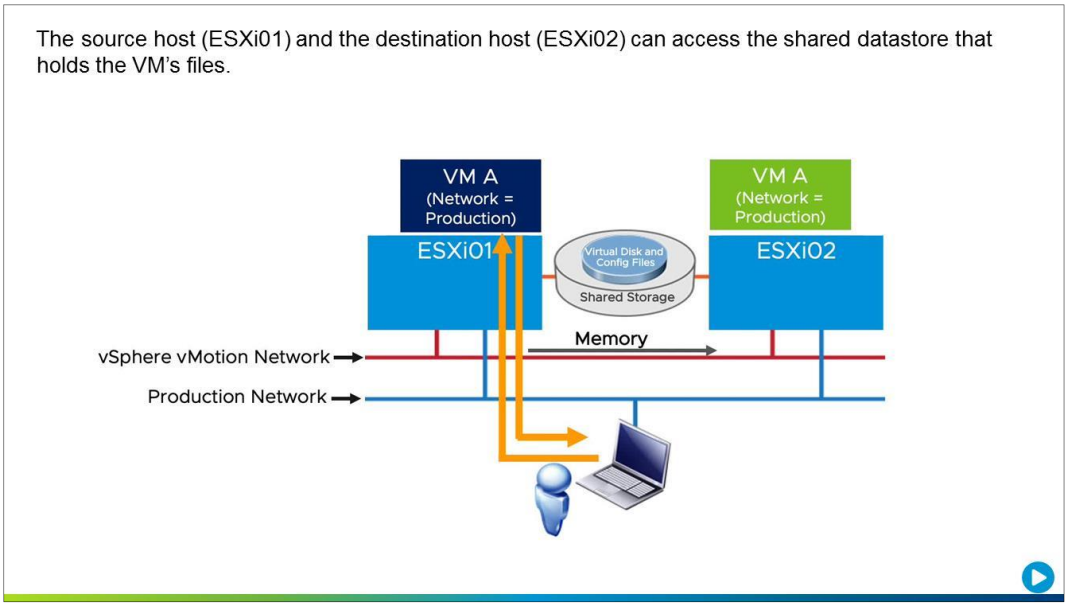

To play the animation, go to https://vmware.bravais.com/s/FbzaDb6owpSMKyKc940F.

A vSphere vMotion migration consists of the following steps:

1. A shadow VM is created on the destination host.
2. The VM's memory state is copied over the vSphere vMotion network from the source host to the target host through the vSphere vMotion network. Users continue to access the VM and, potentially, update pages in memory. A list of modified pages in memory is kept in a memory bitmap on the source host.
3. After the first pass of memory state copy completes, another pass of memory copy is performed to copy any pages that changed during the last iteration. This iterative memory copying continues until no changed pages remain.
4. After most of the VM's memory is copied from the source host to the target host, the VM is quiesced. No additional activity occurs on the VM. In the quiesce period, vSphere vMotion transfers the VM device state and memory bitmap to the destination host.

https://t.me/learningnets

5. Immediately after the VM is quiesced on the source host, the VM is initialized and starts running on the target host. A Gratuitous Address Resolution Protocol (GARP) request notifies the subnet that VM A's MAC address is now on a new switch port.
6. Users access the VM on the target host instead of the source host.
7. The memory pages that the VM was using on the source host are marked as free.

https://t.me/learningnets

7-48 VM Requirements for vSphere vMotion Migration

For migration with vSphere vMotion, a VM must meet these requirements:

- If it uses an RDM disk, the RDM file and the LUN to which it maps must be accessible by the destination host.
- It must not have a connection to a virtual device, such as a CD/DVD or floppy drive, with a host-local image mounted.

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

Remote devices include physical devices or disk images on the client machine running the remote console.

For the complete list of vSphere vMotion migration requirements, see *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vcenterhost.doc/GUID-3B5AF2B1-C534-4426-B97A-D14019A8010F.html.

https://t.me/learningnets

7-49 Host Requirements for vSphere vMotion Migration (1)

Source and destination hosts must have the following characteristics:

- Accessibility to all the VM's storage:
  - 128 concurrent migrations are possible per VMFS or NFS datastore.
  - If the swap file location on the destination host differs from the swap file location on the source host, the swap file is copied to the new location.
- VMkernel port with vSphere vMotion enabled
- Matching management network IP address families (IPv4 or IPv6) between the source and destination hosts

You cannot migrate a VM from a host that is registered to vCenter Server with an IPv4 address to a host that is registered with an IPv6 address.

Copying a swap file to a new location can result in slower migrations. If the destination host cannot access the specified swap file location, it stores the swap file with the VM configuration file.

https://t.me/learningnets

7-50 Host Requirements for vSphere vMotion Migration (2)

- At least a 1 Gigabit Ethernet (1 GigE) network:
  - Each active vSphere vMotion process requires a minimum throughput of 250 Mbit/second on the vSphere vMotion network.
  - Concurrent migrations are limited to four on a 1 Gbps network.
  - Concurrent migrations are limited to eight on a 10 Gbps (or faster) network.
  - For better performance, dedicate at least two port groups to the vSphere vMotion traffic.
- Compatible CPUs:
  - The CPU feature sets of both the source host and the destination host must be compatible.
  - Some features can be hidden by using Enhanced vMotion Compatibility or compatibility masks.

Using 1 GbE network adapters for the vSphere vMotion network might result in migration failure, if you migrate VMs with large vGPU profiles.

https://t.me/learningnets

7-51 Checking vSphere vMotion Errors

When you select the host and cluster, a validation check is performed to verify that most vSphere vMotion requirements are met.

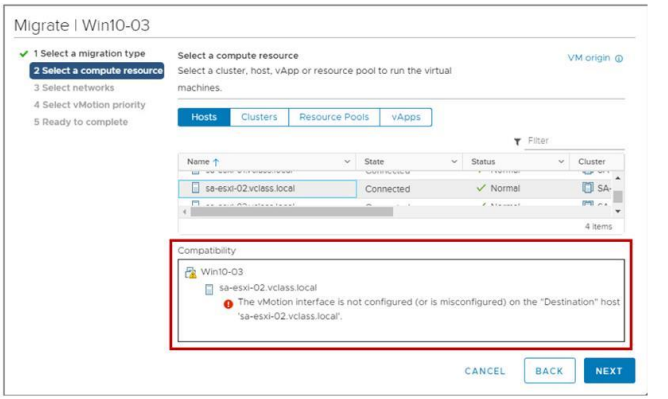

If validation succeeds, you can continue in the wizard. If validation does not succeed, a list of vSphere vMotion errors and warnings displays in the Compatibility pane.

With warnings, you can still perform a vSphere vMotion migration. But with errors, you cannot continue. You must exit the wizard and fix all errors before retrying the migration.

If a failure occurs during the vSphere vMotion migration, the VM is not migrated and continues to run on the source host.

https://t.me/learningnets

7-52 Encrypted vSphere vMotion

When migrating encrypted VMs, you always use encrypted vSphere vMotion.

For VMs that are not encrypted, select one of the following encrypted vSphere vMotion menu items:

- **Disabled**.
- **Opportunistic** (default): Encrypted vSphere vMotion is used if the source and destination hosts support it.
- **Required**: If the source or destination host does not support encrypted vSphere vMotion, the migration fails.

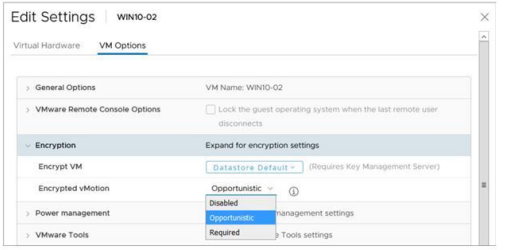

Encrypted vSphere vMotion secures confidentiality, integrity, and authenticity of data that is transferred with vSphere vMotion. Encrypted vSphere vMotion supports all variants of vSphere vMotion, including migration across vCenter Server systems. Encrypted vSphere Storage vMotion is not supported.

You cannot turn off encrypted vSphere vMotion for encrypted VMs.

https://t.me/learningnets

7-53 Cross vCenter Migrations

With vSphere vMotion, you can migrate VMs between linked vCenter Server systems.
Migration of VMs across vCenter Server instances is helpful in the following cases:

- Balancing workloads across clusters and vCenter Server instances that are in the same site or in another geographical area.
- Moving VMs between environments that have different purposes, for example, from a development environment to production environment.
- Moving VMs to meet different Service Level Agreements (SLAs) for storage space, performance, and so on.

https://t.me/learningnets

7-54 Cross vCenter Migration Requirements

Cross vCenter migrations have the following requirements:

- ESXi hosts and vCenter Server systems must be at vSphere 6.0 or later.
- vCenter Server instances must be in Enhanced Linked Mode.
- Hosts must be time-synchronized.

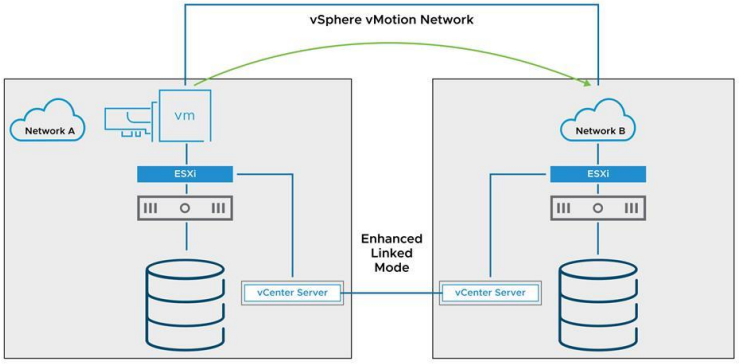

You can perform cross vCenter migrations between vCenter Server instances of different versions. For information on the supported versions, see VMware knowledge base article 2106952 at http://kb.vmware.com/kb/2106952.

https://t.me/learningnets

7-55 Network Checks for Cross vCenter Migrations

vCenter Server performs several network compatibility checks to prevent the following configuration problems:

- MAC address incompatibility on the destination host
- vSphere vMotion migration from a distributed switch to a standard switch
- vSphere vMotion migration between distributed switches of different versions

https://t.me/learningnets

7-56 VMkernel Networking Layer and TCP/IP Stacks

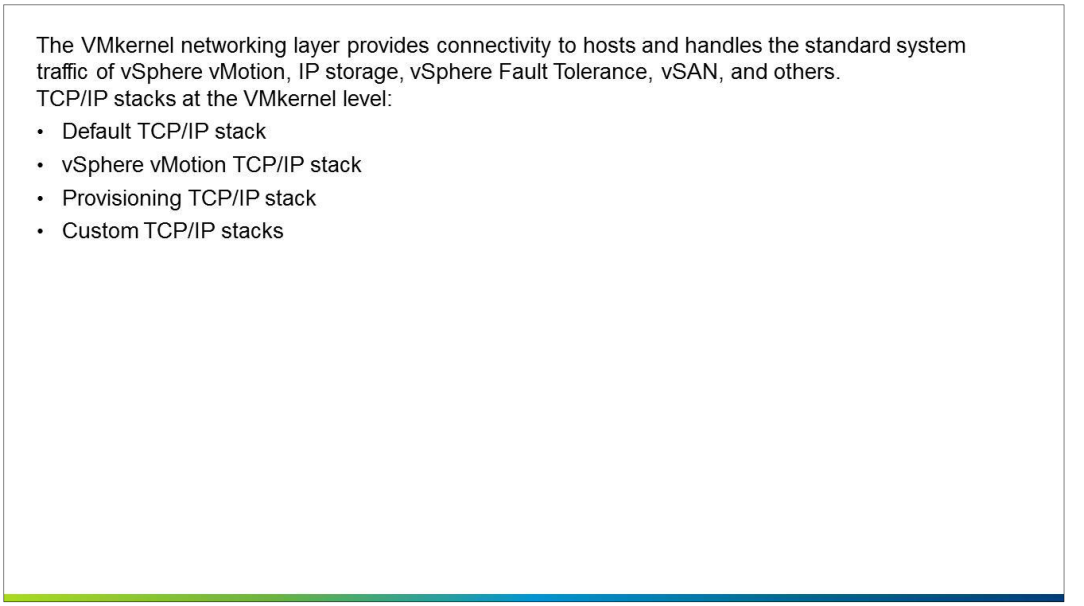

Consider the following key points about TCP/IP stacks at the VMkernel level:

- Default TCP/IP stack: Provides networking support for the management traffic between vCenter Server and ESXi hosts and for system traffic such as vSphere vMotion, IP storage, and vSphere Fault Tolerance.
- vSphere vMotion TCP/IP stack: Supports the traffic for hot migrations of VMs.
- Provisioning TCP/IP stack: Supports the traffic for VM cold migration, cloning, and snapshot creation. You can use the provisioning TPC/IP stack to handle NFC traffic during long-distance vSphere vMotion migration. VMkernel adapters configured with the provisioning TCP/IP stack handle the traffic from cloning the virtual disks of the migrated VMs in long-distance vSphere vMotion.

  By using the provisioning TCP/IP stack, you can isolate the traffic from the cloning operations on a separate gateway. After you configure a VMkernel adapter with the

https://t.me/learningnets

provisioning TCP/IP stack, all adapters on the default TCP/IP stack are disabled for the provisioning traffic.

- Custom TCP/IP stacks: You can create a custom TCP/IP stack on a host to forward networking traffic through a custom application. Open an SSH connection to the host and run the vSphere CLI command:

```
esxcli network ip netstack add -N="stack_name"
```

Take appropriate security measures to prevent unauthorized access to the management and system traffic in your vSphere environment. For example, isolate the vSphere vMotion traffic in a separate network that includes only the ESXi hosts that participate in the migration. Isolate the management traffic in a network that only network and security administrators can access.

https://t.me/learningnets

7-57 vSphere vMotion TCP/IP Stacks

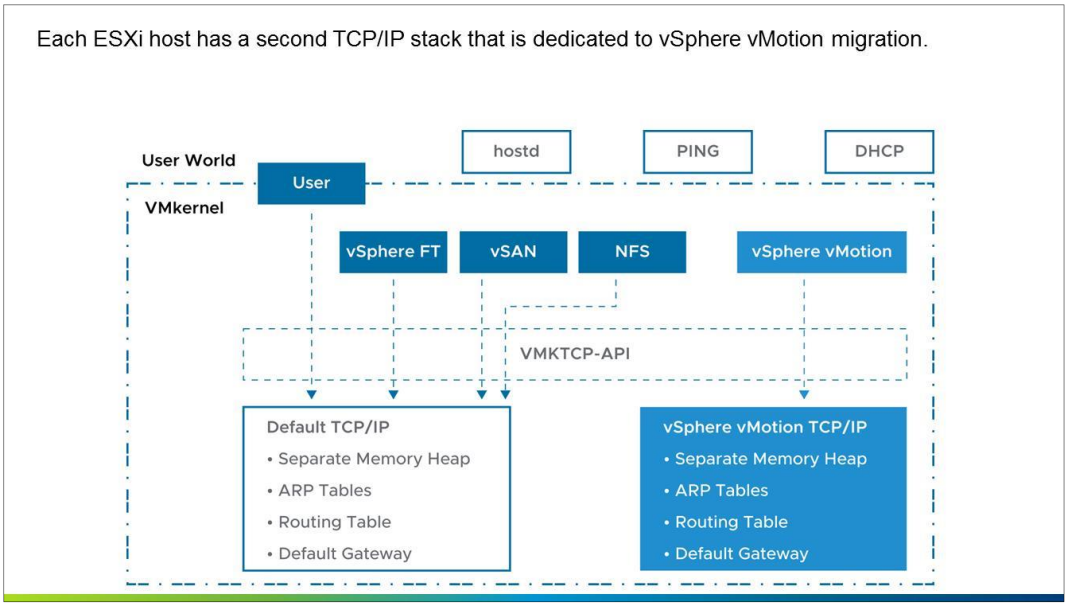

vSphere vMotion TCP/IP stacks support the traffic for hot migrations of VMs. Use the vSphere vMotion TCP/IP stack to provide better isolation for the vSphere vMotion traffic. After you create a VMkernel adapter on the vSphere vMotion TCP/IP stack, you can use only this stack for vSphere vMotion migration on this host.

The VMkernel adapters on the default TCP/IP stack are disabled for the vSphere vMotion service after you create a VMkernel adapter on the vSphere vMotion TCP/IP stack. If a hot migration uses the default TCP/IP stack while you configure VMkernel adapters with the vMotion TCP/IP stack, the migration completes successfully. However, these VMkernel adapters on the default TCP/IP stack are disabled for future vSphere vMotion sessions.

https://t.me/learningnets

7-58 Long-Distance vSphere vMotion Migration

Long-distance vSphere vMotion migration is an extension of cross vCenter migration.

vCenter Server instances are spread across large geographic distances and where the latency across sites is high.

Use cases for long-distance vSphere vMotion migration:

- Permanent migrations
- Disaster avoidance
- Site Recovery Manager and disaster avoidance testing
- Multisite load balancing
- Follow-the-sun scenario support

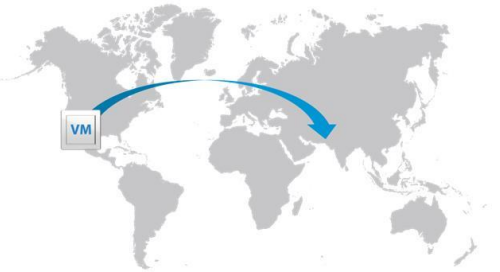

In the follow-the-sun scenario, a global support team might support a certain set of VMs. As one support team ends their workday, another support team in a different timezone takes over support duty. The VMs being supported can be moved from one geographical location to another so that the support team on duty can access those VMs locally instead of long distance.

https://t.me/learningnets

7-59 Networking Prerequisites for Long-Distance vSphere vMotion

Long-distance vSphere vMotion migrations must connect over layer 3 connections:

- Virtual machine network:
  - L2 connection.
  - Same VM IP address is available at the destination.
- vSphere vMotion network:
  - L3 connection.
  - Secure (if you are not using vSphere 6.5 or later encrypted vSphere vMotion).
  - 250 Mbps per vSphere vMotion operation.
  - Round-trip time between hosts can take up to 150 milliseconds.

https://t.me/learningnets

7-60 Lab 19: vSphere vMotion Migrations

Configure vSphere vMotion networking and migrate virtual machines using vSphere vMotion:

1. Configure vSphere vMotion Networking on sa-esxi-01.vclass.local
2. Configure vSphere vMotion Networking on sa-esxi-02.vclass.local
3. Prepare Virtual Machines for vSphere vMotion Migration
4. Migrate Virtual Machines Using vSphere vMotion

https://t.me/learningnets

7-61 Review of Learner Objectives

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

- Recognize the types of VM migrations that you can perform within a vCenter Server instance and across vCenter Server instances
- Explain how vSphere vMotion works
- Verify vSphere vMotion requirements
- Migrate virtual machines using vSphere vMotion

https://t.me/learningnets

https://t.me/learningnets

7-63 Learner Objectives

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

- Describe Enhanced vMotion Compatibility
- Configure EVC mode on a vSphere cluster
- Explain how per-VM EVC mode works with vSphere vMotion

https://t.me/learningnets

7-64 CPU Constraints on vSphere vMotion Migration

CPU compatibility between source and target hosts is a vSphere vMotion requirement that must be met.

CPU Characteristics	Exact Match Required By Source Host and Target Host	Reason
Clock speeds, cache sizes, hyperthreading, and number of cores	N/A	The VMkernel virtualizes these characteristics.
Manufacturer (Intel or AMD) family and generation (Opteron4, Intel Westmere)	Applicable	Instruction sets contain many small differences.
Presence or absence of SSE3, SSSE3, or SSE4.1 instructions	Applicable	Multimedia instructions are usable directly by applications.
Virtualization hardware assist	For 32-bit VMs: N/A	The VMkernel virtualizes this characteristic.
	For 64-bit VMs on Intel: Applicable	Intel 64-bit with VMware implementation uses Intel VT.

Depending on the CPU characteristic, an exact match between the source and target host might or might not be required.

For example, if hyperthreading is enabled on the source host and disabled on the destination host, the vSphere vMotion migration continues because the VMkernel handles this difference in characteristics.

But, if the source host processor supports SSE4.1 instructions and the destination host processor does not support them, the hosts are considered incompatible and the vSphere vMotion migration fails.

SSE4.1 instructions are application-level instructions that bypass the virtualization layer and might cause application instability if mismatched after a migration with vSphere vMotion.

https://t.me/learningnets

7-65 About Enhanced vMotion Compatibility

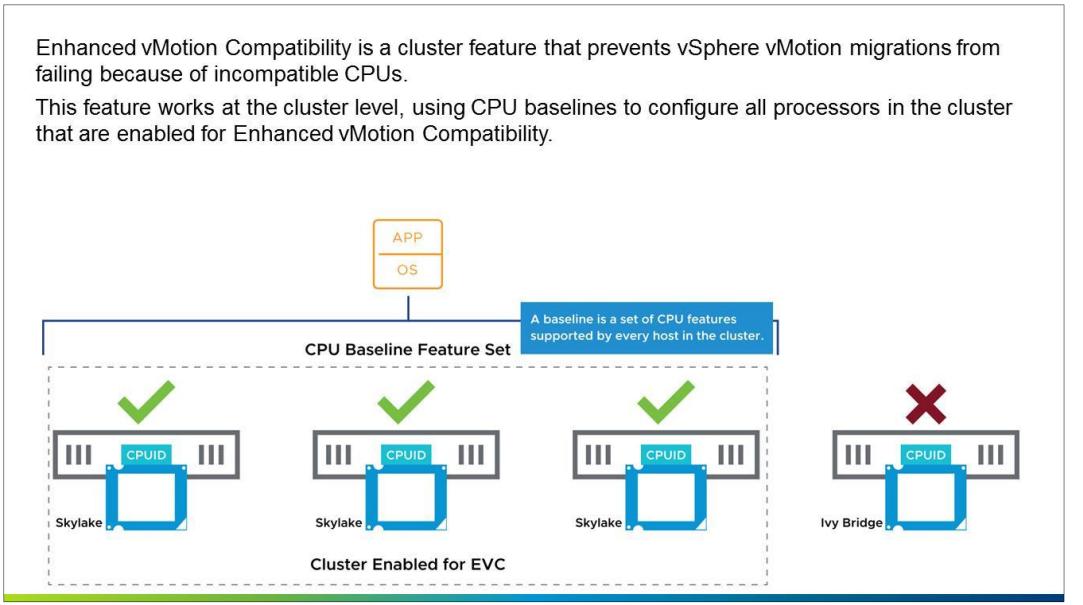

Enhanced vMotion Compatibility ensures that all hosts in a cluster present the same CPU feature set to VMs, even if the CPUs on the hosts differ.

Enhanced vMotion Compatibility facilitates safe vSphere vMotion migration across a range of CPU generations. With Enhanced vMotion Compatibility, you can use vSphere vMotion to migrate VMs among CPUs that otherwise are considered incompatible.

Enhanced vMotion Compatibility allows vCenter Server to enforce vSphere vMotion compatibility among all hosts in a cluster by forcing hosts to expose a common set of CPU features (baseline) to VMs. A baseline is a set of CPU features that are supported by every host in the cluster. When you configure Enhanced vMotion Compatibility, you set all host processors in the cluster to present the features of a baseline processor. After the features are enabled for a cluster, hosts that are added to the cluster are automatically configured to the CPU baseline.

Hosts that cannot be configured to the baseline are not permitted to join the cluster. VMs in the cluster always see an identical CPU feature set, no matter which host they happen to run on.

https://t.me/learningnets

Because this process is automatic, Enhanced vMotion Compatibility is easy to use and requires no specialized knowledge of CPU features and masks.

https://t.me/learningnets

7-66 Enhanced vMotion Compatibility Cluster Requirements

All hosts in the cluster must meet several requirements:

- Use CPUs from a single vendor, either Intel or AMD:
  - Use Intel CPUs with Merom microarchitecture and later.
  - Use AMD first-generation Opteron CPUs and later.
- Be enabled for hardware virtualization: AMD-V or Intel VT.
- Be enabled for execution-disable technology: AMD No eXecute (NX) or Intel eXecute Disable (XD).
- Be configured for vSphere vMotion migration.

Applications in VMs must be CPU ID compatible.

Before you create an Enhanced vMotion Compatibility cluster, ensure that the hosts that you intend to add to the cluster meet the requirements.

Enhanced vMotion Compatibility automatically configures hosts whose CPUs have Intel FlexMigration and AMD-V Extended Migration technologies to be compatible with vSphere vMotion with hosts that use older CPUs.

For Enhanced vMotion Compatibility to function properly, the applications on the VMs must be written to use the CPU ID machine instruction for discovering CPU features as recommended by the CPU vendors. vSphere cannot support Enhanced vMotion Compatibility with applications that do not follow the CPU vendor recommendations for discovering CPU features.

To determine which EVC modes are compatible with your CPU, search the VMware Compatibility Guide at http://www.vmware.com/resources/compatibility. Search for the server model or CPU family, and click the entry in the CPU Series column to display the compatible EVC modes.

https://t.me/learningnets

7-67 Enabling EVC Mode on an Existing Cluster

You enable EVC mode on an existing cluster to ensure vSphere vMotion CPU compatibility between the hosts in the cluster.

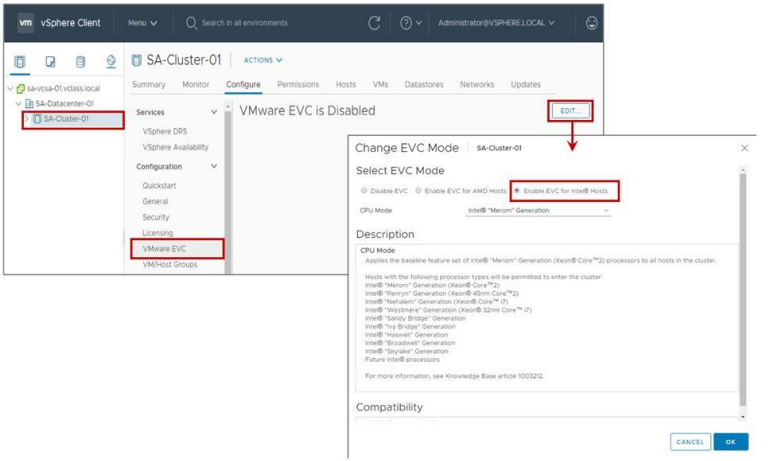

You can use one of the following methods to create an Enhanced vMotion Compatibility cluster:

- Create an empty cluster with EVC mode enabled and move hosts into the cluster.
- Enable EVC mode on an existing cluster.

For information about Enhanced vMotion Compatibility processor support, see VMware knowledge base article 1003212 at http://kb.vmware.com/kb/1003212.

https://t.me/learningnets

7-68 Changing the EVC Mode for a Cluster

Several EVC mode approaches are available to ensure CPU compatibility:

- If all the hosts in a cluster are compatible with a newer EVC mode, you can change the EVC mode of an existing Enhanced vMotion Compatibility cluster.
- You can enable EVC mode for a cluster that does not have EVC mode enabled.

You can raise or lower the EVC mode, but the VMs must be in the correct power state to do so.

EVC Mode	VM Power Action
Raise the EVC mode to a CPU baseline with more features.	• Running VMs can remain powered on. • New EVC mode features are not available to the VMs until they are powered off and powered back on again (Suspending and resuming the VM is not sufficient.)
Lower the EVC mode to a CPU baseline with fewer features.	• Power off VMs if they are powered on and running at a higher EVC mode than the one you intend to enable.

https://t.me/learningnets

7-69 Virtual Machine EVC Mode

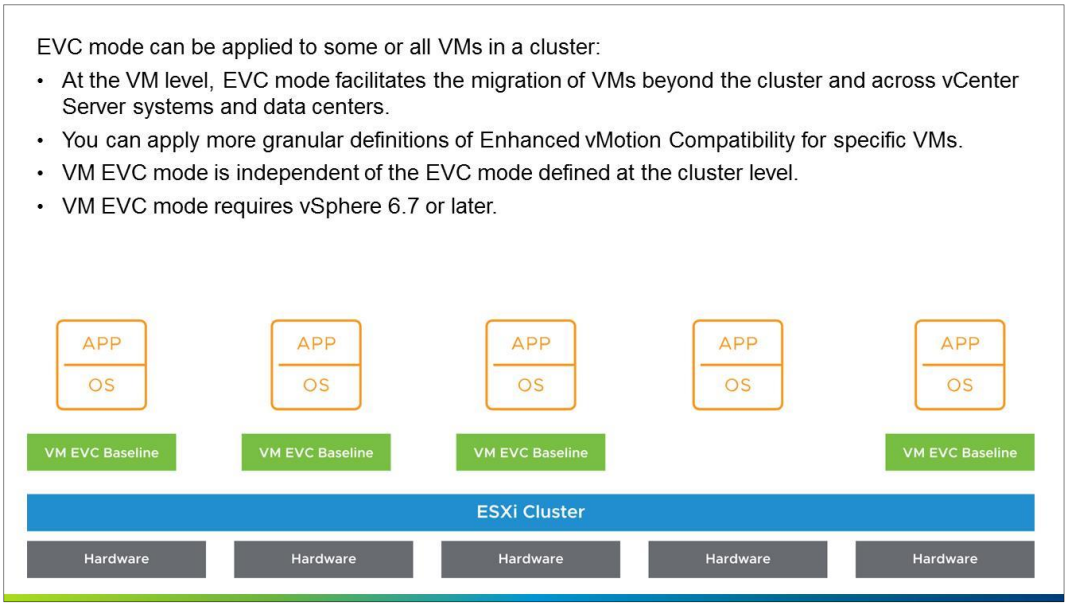

With per-VM EVC mode, the EVC mode becomes an attribute of the VM rather than the specific processor generation it happens to be booted on in the cluster. This feature supports seamless migration between two data centers that have different processors. Further, the feature is persisted per VM and does not lose the EVC mode during migrations across clusters or during power cycles.

In this diagram, EVC mode is not enabled on the cluster. The cluster consists of differing CPU models with different feature sets. The VMs with per-VM EVC mode can run on any ESXi host that can satisfy the defined EVC mode.

https://t.me/learningnets

7-70 Review of Learner Objectives

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

- Describe Enhanced vMotion Compatibility
- Configure EVC mode on a vSphere cluster
- Explain how per-VM EVC mode works with vSphere vMotion

https://t.me/learningnets

https://t.me/learningnets

7-72 Learner Objectives

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

- Explain how vSphere Storage vMotion works
- Recognize guidelines for using vSphere Storage vMotion
- Migrate virtual machines using vSphere Storage vMotion
- Migrate both the compute resource and storage of a virtual machine

https://t.me/learningnets

7-73 About vSphere Storage vMotion

With vSphere Storage vMotion, you can migrate a powered-on VM from one datastore to another.

Using vSphere Storage vMotion, you can perform the following tasks:

- Move VMs off arrays for maintenance or to upgrade.
- Change the disk provisioning type.
- Change VM files on the destination datastore to match the inventory name of the VM.
- Migrate between datastores to balance traffic across storage paths and reduce latencies.
- Redistribute VMs or virtual disks to different storage volumes to balance capacity or improve performance.

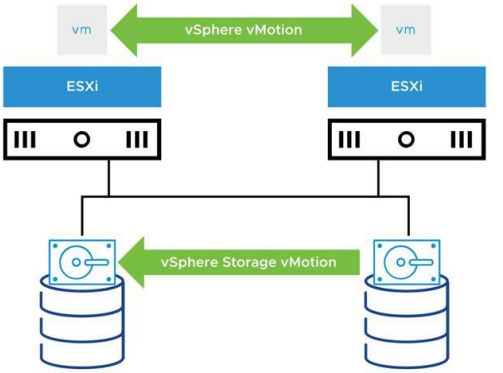

vSphere Storage vMotion provides flexibility to optimize disks for performance or transform disk types, which you can use to reclaim space.

You can place the VM and all its disks in a single location, or you can select separate locations for the VM configuration file and each virtual disk. During a migration with vSphere Storage vMotion, the VM does not change the host that it runs on.

With vSphere Storage vMotion, you can rename a VM's files on the destination datastore. The migration renames all virtual disk, configuration, snapshot, and `.nvram` files.

https://t.me/learningnets

7-74 vSphere Storage vMotion In Action

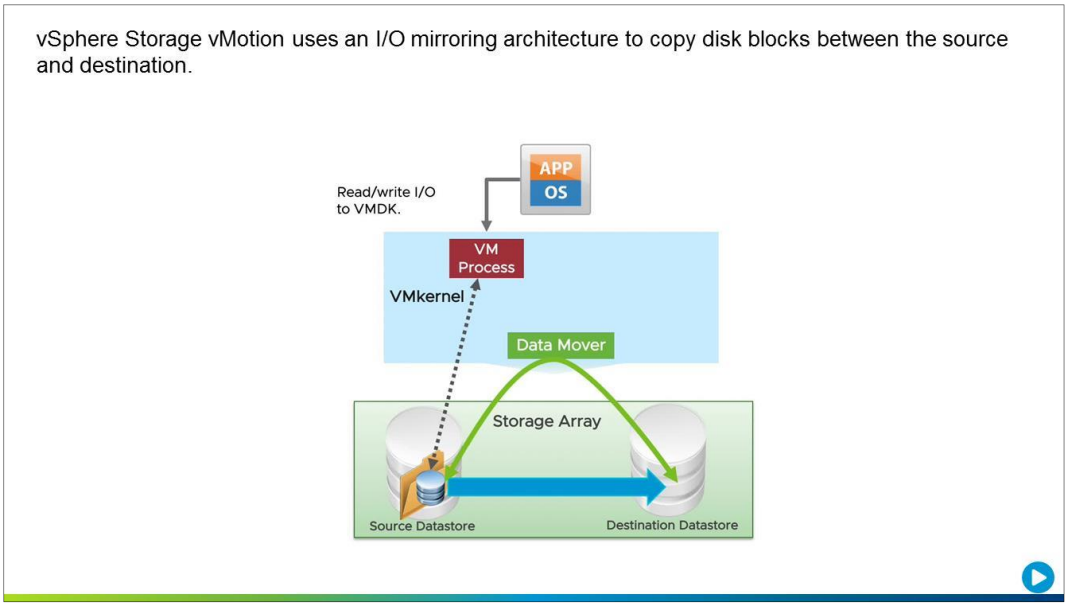

To play the animation, go to https://vmware.bravais.com/s/FnHZwq043PJ8dV3ZRV7p.

The vSphere Storage vMotion migration process includes the following steps:

1. Initiate storage migration.
2. Use the VMkernel data mover or vSphere Storage APIs - Array Integration to copy data.
3. Start a new VM process.
4. Mirror I/O calls to file blocks that are already copied to the virtual disk on the destination datastore.
5. Transition to the destination VM process to begin accessing the virtual disk copy.

The storage migration process does a single pass of the disk, copying all the blocks to the destination disk. If blocks are changed after they are copied, the blocks are synchronized from the source to the destination through the mirror driver, with no need for recursive passes.

https://t.me/learningnets

This approach guarantees complete transactional integrity and is fast enough to be unnoticeable to the end user. The mirror driver uses the VMkernel data mover to copy blocks of data from the source disk to the destination disk. The mirror driver synchronously mirrors writes to both disks during the vSphere Storage vMotion operation.

Finally, vSphere Storage vMotion operations are performed either internally on a single ESXi host or offloaded to the storage array. Operations performed internally on the ESXi host use a data mover built into the VMkernel. Operations are offloaded to the storage array if the array supports vSphere Storage APIs - Array Integration, also called hardware acceleration.

https://t.me/learningnets

7-75 Identifying Storage Arrays That Support vSphere Storage APIs - Array Integration

vSphere Storage vMotion offloads its operations to the storage array if the array supports VMware vSphere Storage APIs - Array Integration, also called hardware acceleration.

Use the vSphere Client to determine whether your storage array supports hardware acceleration.

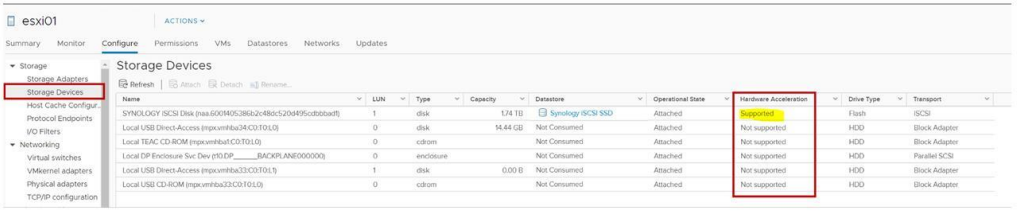

https://t.me/learningnets

7-76 vSphere Storage vMotion Guidelines and Limitations

Guidelines:

- Plan the migration and coordinate with administrators.
- Perform migrations during off-peak hours.

Limitation:

- Independent virtual machine disks must be in persistent mode.

A VM and its host must meet certain resource and configuration requirements for the virtual machine disks (VMDKs) to be migrated with vSphere Storage vMotion. One of the requirements is that the host on which the VM runs must have access both to the source datastore and to the target datastore.

During a migration with vSphere Storage vMotion, you can change the disk provisioning type. Migration with vSphere Storage vMotion changes VM files on the destination datastore to match the inventory name of the VM. The migration renames all virtual disk, configuration, snapshot, and .nvram-extension files. If the new names exceed the maximum filename length, the migration does not succeed.

https://t.me/learningnets

7-77 Changing Both Compute Resource and Storage During Migration (1)

When you change both compute resource and storage during migration, a VM changes its host, datastores, networks, and vCenter Server instances simultaneously:

- This technique combines vSphere vMotion and vSphere Storage vMotion into a single operation.
- You can migrate VMs across clusters, data centers, and vCenter Server instances.

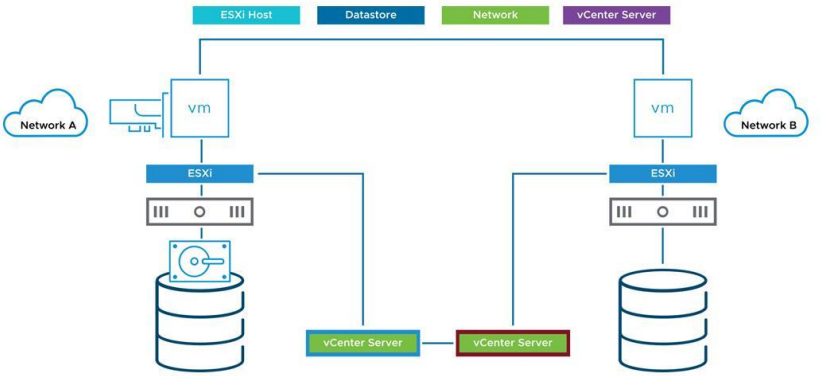

You can migrate VMs beyond storage accessibility boundaries and between hosts, within and across clusters, data centers, and vCenter Server instances.

This type of migration is useful for performing cross-cluster migrations, when the target cluster VMs might not have access to the source cluster's storage. Processes on the VM continue to run during the migration with vSphere vMotion.

https://t.me/learningnets

7-78 Changing Both Compute Resource and Storage During Migration (2)

Compute resource and storage migration is useful for virtual infrastructure administration tasks.

Task	Benefit of Using vSphere Storage vMotion
Host maintenance	You can move VMs from a host when you want to perform host maintenance.
Storage maintenance and reconfiguration	You can move VMs from a storage device so that you can perform maintenance or reconfigure the storage device without VM downtime.
Storage load redistribution	You can manually redistribute VMs or virtual disks to different storage volumes to balance capacity or to improve performance.

https://t.me/learningnets

7-79 Lab 20: vSphere Storage vMotion Migrations

Use vSphere Storage vMotion to migrate virtual machines:

1. Migrate Virtual Machine Files from Local Storage to Shared Storage
2. Migrate Both the Compute Resource and Storage of a Virtual Machine

https://t.me/learningnets

7-80 Review of Learner Objectives

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

- Explain how vSphere Storage vMotion works
- Recognize guidelines for using vSphere Storage vMotion
- Migrate virtual machines using vSphere Storage vMotion
- Migrate both the compute resource and storage of a virtual machine

https://t.me/learningnets

https://t.me/learningnets

7-82 Learner Objectives

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

- Take a snapshot of a virtual machine
- Manage multiple snapshots
- Delete virtual machine snapshots
- Consolidate snapshots

https://t.me/learningnets

7-83 VM Snapshots

With snapshots, you can preserve the state of the VM so that you can repeatedly return to the same state.

For example, if problems occur during the patching or upgrading process, you can stop the process and revert to the previous state.

VM snapshots are not recommended as a VM backup strategy.

Snapshots are useful when you want to revert repeatedly to the same state but do not want to create multiple VMs. Examples include patching or upgrading the guest operating system in a VM.

The relationship between snapshots is like the relationship between a parent and a child. Snapshots are organized in a snapshot tree. In a snapshot tree, each snapshot has one parent and one or more children, except for the last snapshot, which has no children.

https://t.me/learningnets

7-84 Taking Snapshots

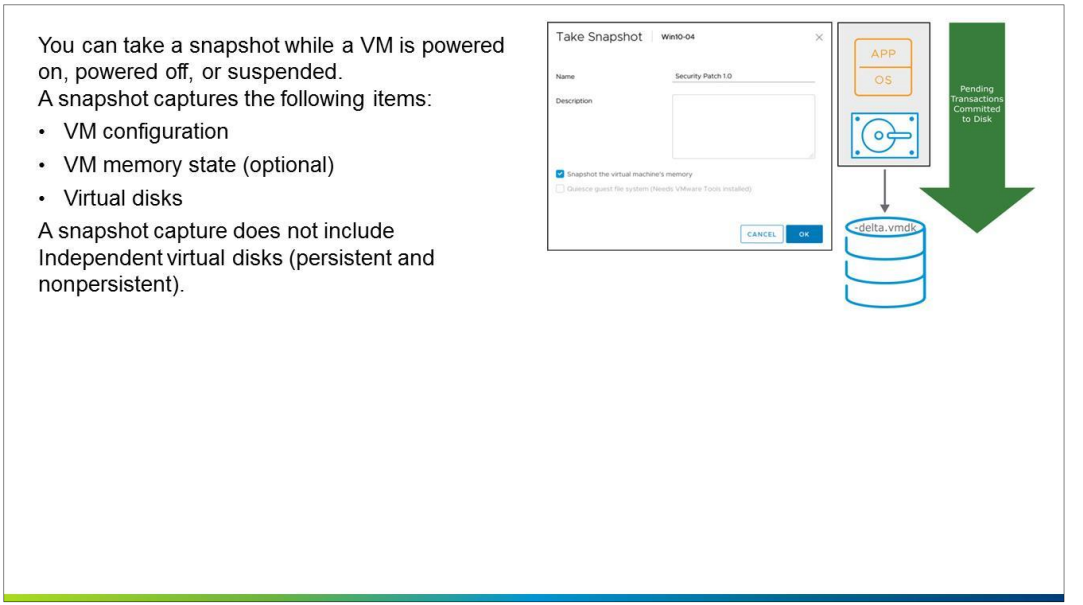

A snapshot captures the entire state of the VM at the time that you take the snapshot, including the following states:

- Memory state: The contents of the VM's memory. The memory state is captured only if the VM is powered on and if you select the **Snapshot the virtual machine's memory** check box (selected by default).
- Settings state: The VM settings.
- Disk state: The state of all the VM's virtual disks.

At the time that you take the snapshot, you can also quiesce the guest operating system. This action quiesces the file system of the guest operating system. This option is available only if you do not capture the memory state as part of the snapshot.

https://t.me/learningnets

7-85 Types of Snapshots

A delta or child disk is created when you create a snapshot:

- On the VMFS datastore, the delta disk is a sparse disk.
- Delta disks use different sparse formats depending on the type of datastore.

Snapshot Type	Notes	Filename	Block Size
VMFSsparse	VMFS5 with virtual disks smaller than 2 TB	`#-delta.vmdk`	512 bytes
SEsparse	• VMFS6 • VMFS5 with virtual disks larger than 2 TB • Space efficient (thin provisioned) • Supports disk reclamation (unmap)	`#-sesparse.vmdk`	4 KB
vsanSparse	vSAN	`Delta object`	4 MB

Delta disks use different sparse formats depending on the type of datastore.

- VMFSsparse: VMFS5 uses the VMFSsparse format for virtual disks smaller than 2 TB. VMFSsparse is implemented on top of VMFS. The VMFSsparse layer processes I/O operations issued to a snapshot VM. Technically, VMFSsparse is a redo log that starts empty, immediately after a VM snapshot is taken. The redo log expands to the size of its base VMDK, when the entire VMDK is rewritten with new data after the VM snapshot. This redo log is a file in the VMFS datastore. On snapshot creation, the base VMDK attached to the VM is changed to the newly created sparse VMDK.
- SEsparse: SEsparse is a default format for all delta disks on the VMFS6 datastores. On VMFS5, SEsparse is used for virtual disks of the size 2 TB and larger. SEsparse is a format that is like VMFSsparse with some enhancements. This format is space efficient and supports the space-reclamation technique. With space reclamation, blocks that the guest OS deletes are marked. The system sends commands to the SEsparse layer in the hypervisor to unmap those

https://t.me/learningnets

blocks. The unmapping helps to reclaim space allocated by SEsparse after the guest operating system deletes the data.

https://t.me/learningnets

7-86 VM Snapshot Files

A snapshot consists of a set of files:

- `-Snapshot#.vmsn`: Configuration state
- `-Snapshot#.vmem`: Memory state (optional)
- `-00000#.vmdk`: Disk descriptor
- `-00000#-delta.vmdk`: VMFS5 delta
- `-00000#-sesparse.vmdk`: VMFS6 delta
- `.vmsd`: Stores names, descriptions, and relationships for all the VM's snapshots

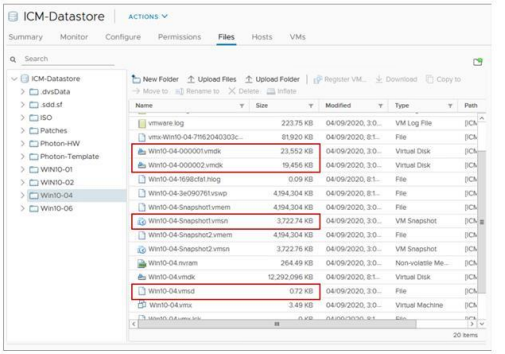

A VM can have one or more snapshots. For each snapshot, the following files are created:

- Snapshot delta file: This file contains the changes to the virtual disk's data since the snapshot was taken. When you take a snapshot of a VM, the state of each virtual disk is preserved. The VM stops writing to its `-flat.vmdk` file. Writes are redirected to `>-######-delta.vmdk` (or `-######-sesparse.vmdk`) instead (for which `######` is the next number in the sequence). You can exclude one or more virtual disks from a snapshot by designating them as independent disks. Configuring a virtual disk as independent is typically done when the virtual disk is created, but this option can be changed whenever the VM is powered off.
- Disk descriptor file: `-00000#.vmdk`. This file is a small text file that contains information about the snapshot.
- Configuration state file: `-.vmsn`. `#` is the next number in the sequence, starting with 1. This file holds the active memory state of the VM at the point that the snapshot was taken, including virtual hardware, power state, and hardware version.

https://t.me/learningnets

- Memory state file: `-.vmem`. This file is created if the option to include memory state was selected during the creation of the snapshot. It contains the entire contents of the VMs at the time that the snapshot of the VM was taken.
- Snapshot active memory file: `-.vmem`. This file contains the contents of the VM memory if the option to include memory is selected during the creation of the snapshot.
- The `.vmsd` file is the snapshot list file and is created at the time that the VM is created. It maintains snapshot information for a VM so that it can create a snapshot list in the vSphere Client. This information includes the name of the snapshot `.vmsn` file and the name of the virtual disk file.
- The snapshot state file has a `.vmsn` extension and is used to store the state of a VM when a snapshot is taken. A new `.vmsn` file is created for every snapshot that is created on a VM and is deleted when the snapshot is deleted. The size of this file varies, based on the options selected when the snapshot is created. For example, including the memory state of the VM in the snapshot increases the size of the `.vmsn` file.

You can exclude one or more of the VMDKs from a snapshot by designating a virtual disk in the VM as an independent disk. Placing a virtual disk in independent mode is typically done when the virtual disk is created. If the virtual disk was created without enabling independent mode, you must power off the VM to enable it.

Other files might also exist, depending on the VM hardware version. For example, each snapshot of a VM that is powered on has an associated `_.vmem` file, which contains the guest operating system main memory, saved as part of the snapshot.

https://t.me/learningnets

7-87 VM Snapshot Files Example (1)

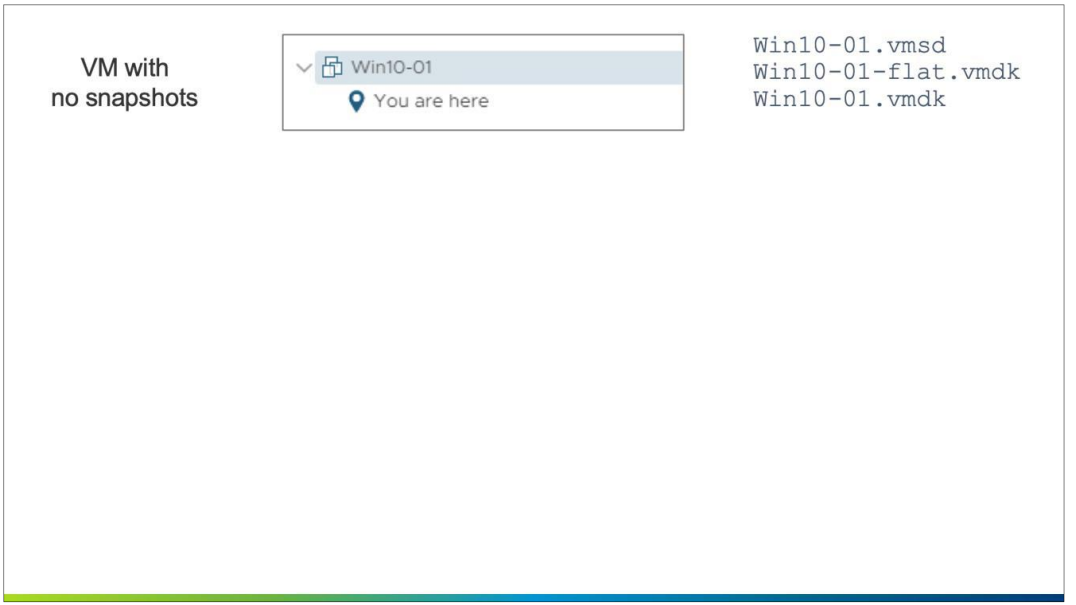

This example shows the snapshot and virtual disk files that are created when a VM has no snapshots, one snapshot, and two snapshots.

https://t.me/learningnets

7-88 VM Snapshot Files Example (2)

https://t.me/learningnets

7-89 VM Snapshot Files Example (3)

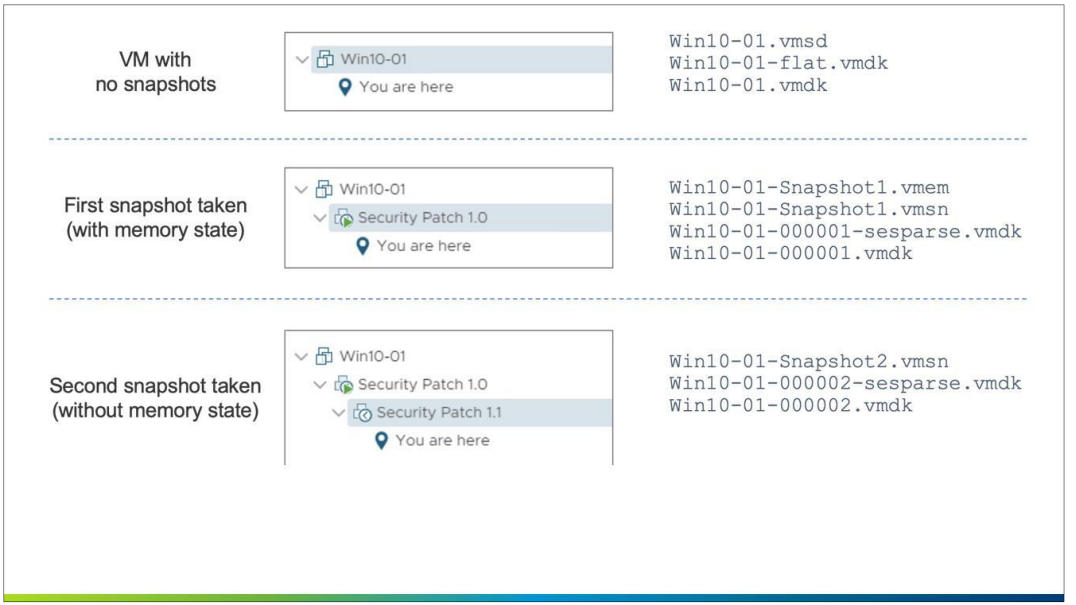

https://t.me/learningnets

7-90 Managing Snapshots

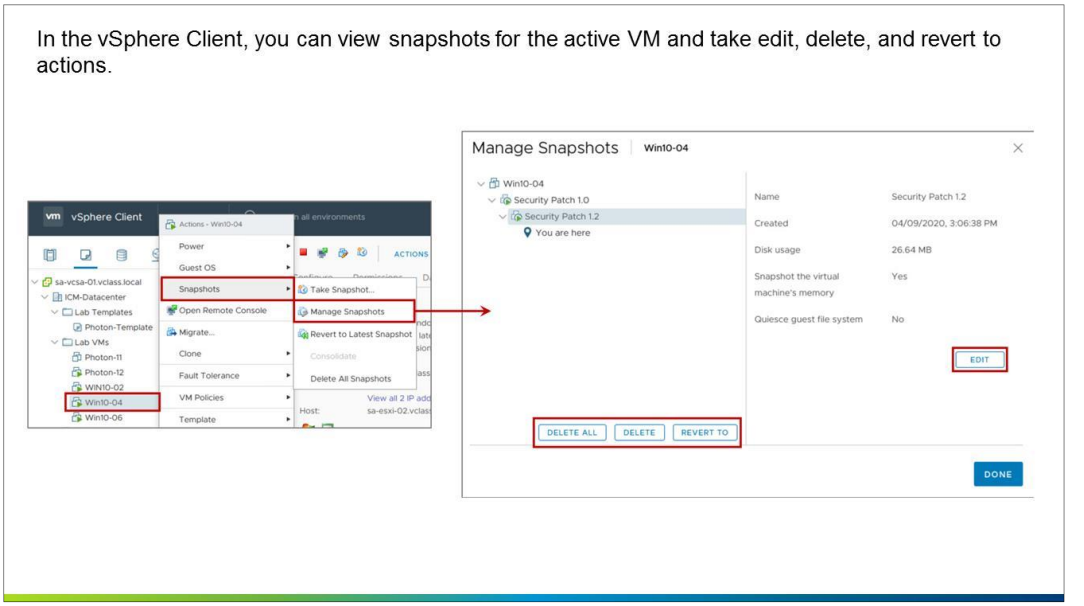

You can perform the following actions from the Manage Snapshots window:

- Edit the snapshot: Edit the snapshot name and description.
- Delete the snapshot: Remove the snapshot from the Snapshot Manager, consolidate the snapshot files to the parent snapshot disk, and merge with the VM base disk.
- Delete all snapshots: Commit all the intermediate snapshots before the current-state icon (You are here) to the VM and remove all snapshots for that VM.
- Revert to a snapshot: Restore, or revert to, a particular snapshot. The snapshot that you restore becomes the current snapshot.

When you revert to a snapshot, you return all these items to the state that they were in at the time that you took the snapshot. If you want the VM to be suspended, powered on, or powered off when you start it, ensure that the VM is in the correct state when you take the snapshot.

https://t.me/learningnets

Deleting a snapshot (**DELETE** or **DELETE ALL**) consolidates the changes between snapshots and previous disk states. Deleting a snapshot also writes to the parent disk all data from the delta disk that contains the information about the deleted snapshot. When you delete the base parent snapshot, all changes merge with the base VMDK.

https://t.me/learningnets

7-91 Deleting VM Snapshots (1)

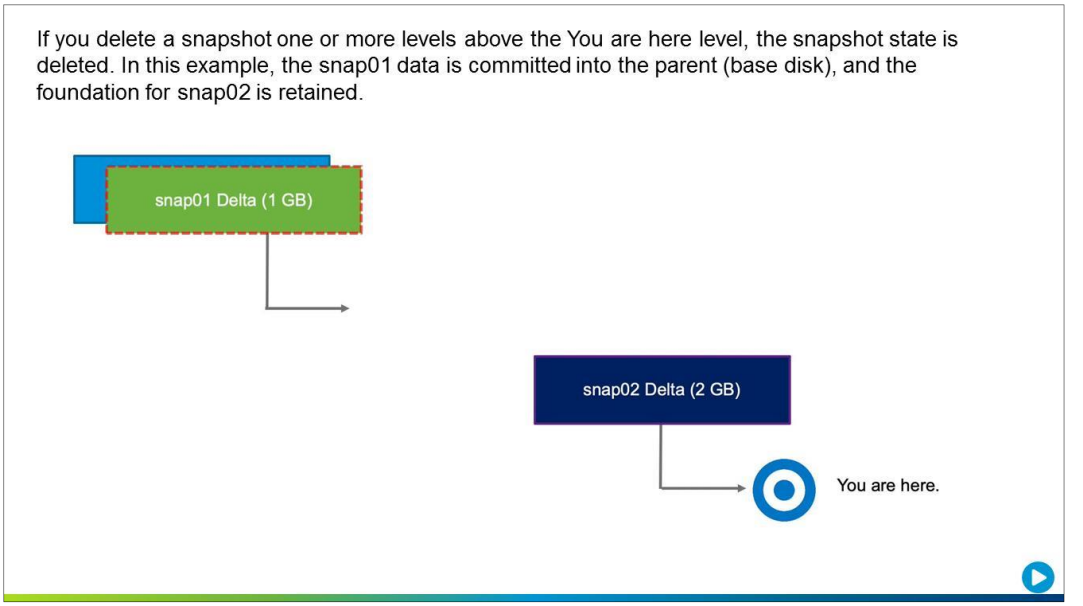

To play the animation, go to https://vmware.bravais.com/s/WhbcXR4sSwk2Vl7MeaXD.

https://t.me/learningnets

7-92 Deleting VM Snapshots (2)

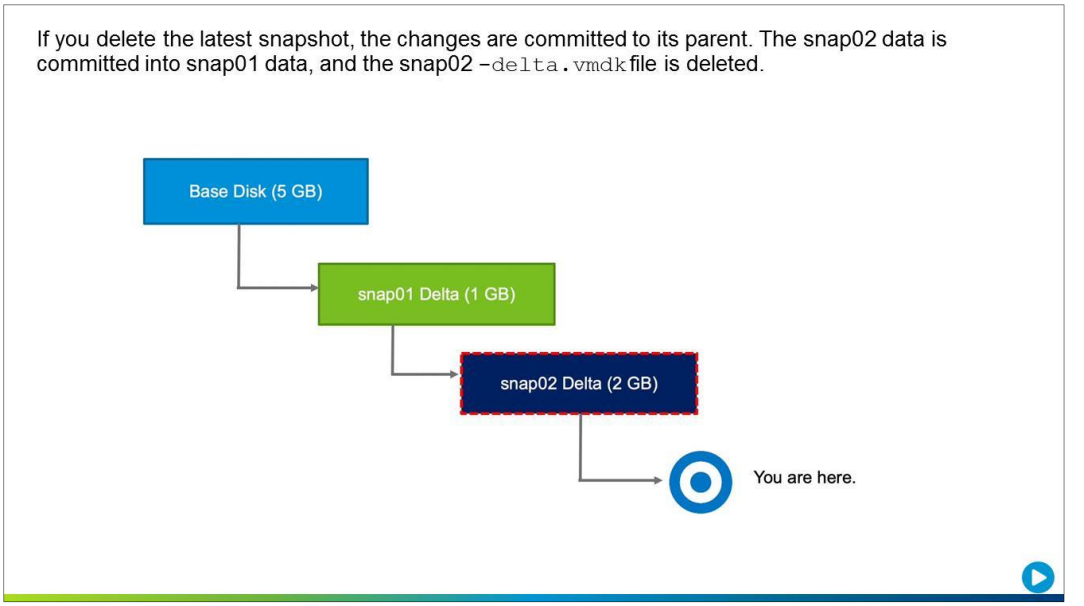

To play the animation, go to https://vmware.bravais.com/s/l0JYYQzMTv7pvxBqNcQp.

https://t.me/learningnets

7-93 Deleting VM Snapshots (3)

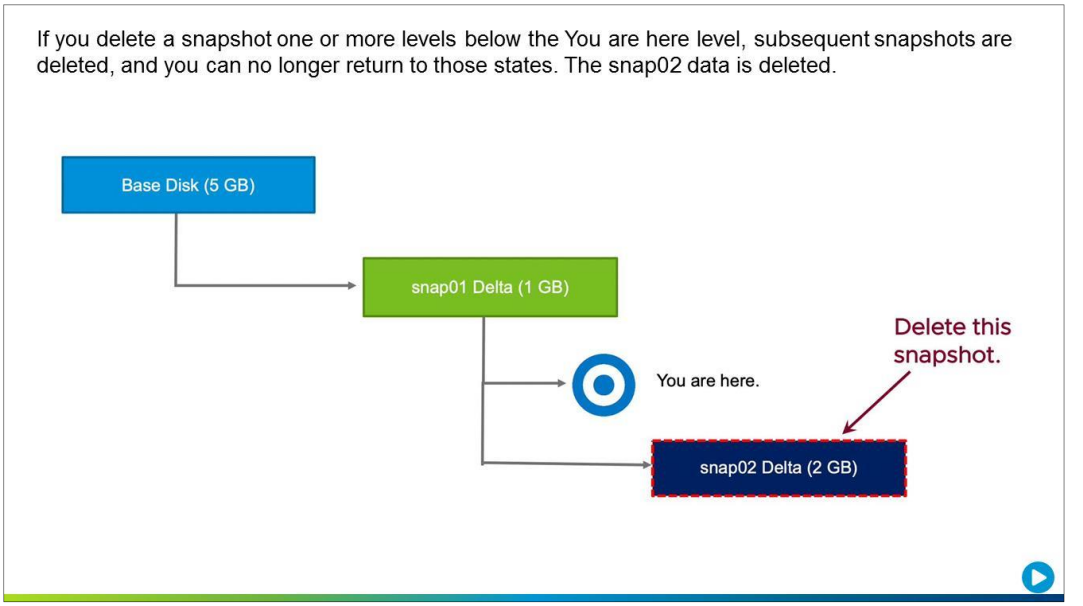

To play the animation, go to https://vmware.bravais.com/s/NiQxPT3iycemQ8WYXKom.

https://t.me/learningnets

7-94 Deleting All VM Snapshots

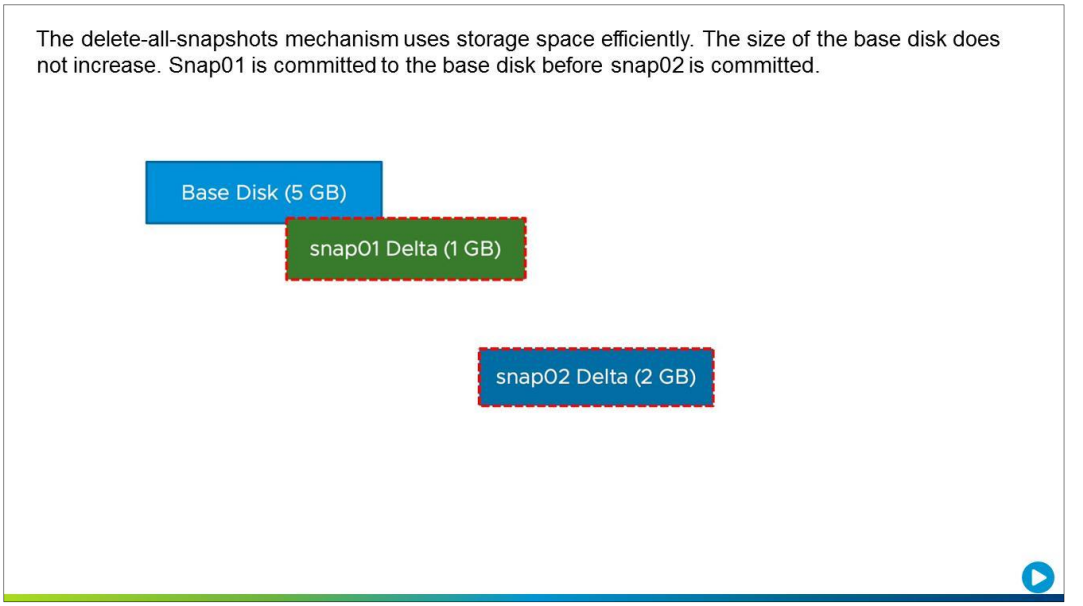

To play the animation, go to https://vmware.bravais.com/s/L3ilQHlrywEhIgr5p7RP.

All snapshots before the You are here point are committed all the way up to the base disk. All snapshots after You are here are discarded.

Like a single snapshot deletion, changed blocks in the snapshot overwrite their counterparts in the base disk.

https://t.me/learningnets

7-95 About Snapshot Consolidation

Snapshot consolidation is a method for committing a chain of delta disks to the base disks when the Snapshot Manager shows that no snapshots exist but the delta disk files remain on the datastore.

Snapshot consolidation resolves problems that might occur with snapshots:

- The snapshot descriptor file is committed correctly, and the Snapshot window shows that all the snapshots are deleted.
- The snapshot files (`-delta.vmdk`) are still part of the VM.
- Delta disk files continue to expand until the datastore on which the VM is located runs out of space.

Snapshot consolidation is a way to clean unneeded delta disk files from a datastore. If no snapshots are registered for a VM, but delta disk files exist, snapshot consolidation commits the chain of the delta disk files and removes them.

If consolidation is not performed, the delta disk files might expand to the point of consuming all the remaining space on the VM's datastore or the delta disk file reaches its configured size. The delta disk cannot be larger than the size configured for the base disk.

https://t.me/learningnets

7-96 Discovering When to Consolidate Snapshots

On the **Monitor** tab under **All Issues** for the VM, a warning notifies you that a consolidation is required.

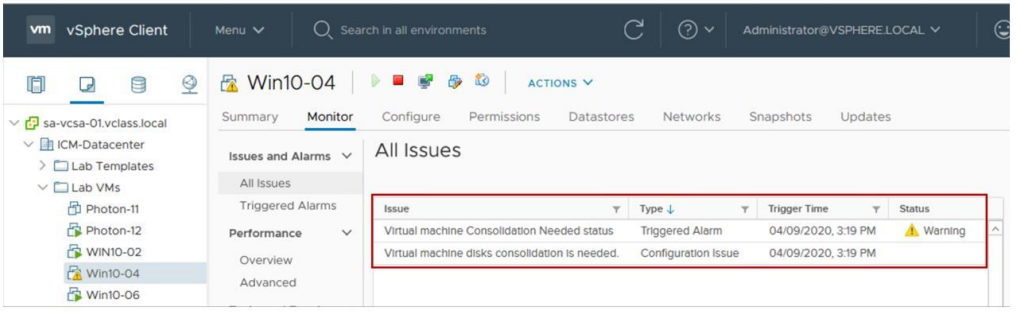

With snapshot consolidation, vCenter Server displays a warning when the descriptor and the snapshot files do not match. After the warning displays, you can use the vSphere Client to commit the snapshots.

https://t.me/learningnets

7-97 Consolidating Snapshots

After the snapshot consolidation warning appears, you can use the vSphere Client to consolidate the snapshots.

All snapshot delta disks are committed to the base disks.

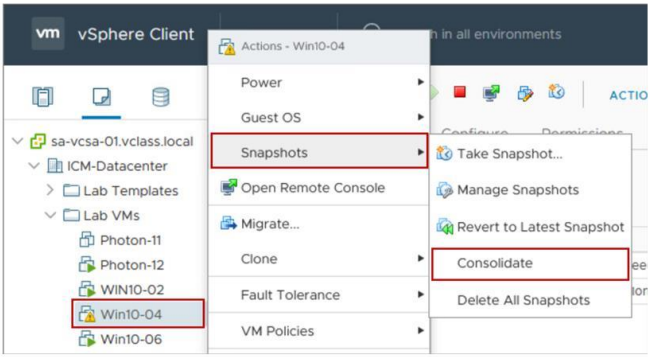

For a list of best practices for using snapshots in a vSphere environment, see VMware knowledge base article 1025279 at http://kb.vmware.com/kb/1025279.

https://t.me/learningnets

7-98 Lab 21: Working with Snapshots

Take VM snapshots, revert a VM to a different snapshot, and delete snapshots

1. Take Snapshots of a Virtual Machine
2. Add Files and Take Another Snapshot of a Virtual Machine
3. Revert the Virtual Machine to a Snapshot
4. Delete an Individual Snapshot
5. Delete All Snapshots

https://t.me/learningnets

7-99 Review of Learner Objectives

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

- Take a snapshot of a virtual machine
- Manage multiple snapshots
- Delete virtual machine snapshots
- Consolidate snapshots

https://t.me/learningnets

https://t.me/learningnets

7-101 Learner Objectives

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

- Identify the components in the vSphere Replication architecture
- Deploy and configure vSphere Replication
- Recover replicated virtual machines
- Explain the backup and restore solution for VMs
- Describe the benefits of vSphere Storage APIs - Data Protection

https://t.me/learningnets

7-102 About vSphere Replication

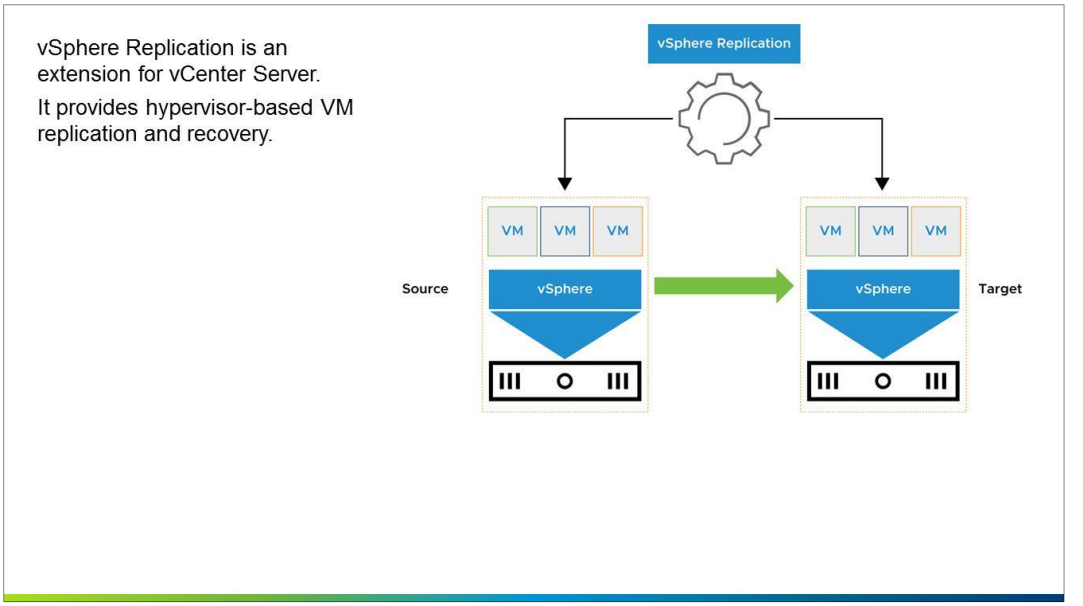

vSphere Replication is an alternative to array-based replication. vSphere Replication protects VMs from partial or complete site failures by replicating the VMs between the following sites:

- From a source site to a target site
- Within a single site from one cluster to another
- From multiple source sites to a shared remote target site

vSphere Replication provides several benefits as compared to array-based replication:

- Data protection at lower cost per VM
- A replication solution that supports flexibility in storage vendor selection at the source and target sites
- Overall lower cost per replication

https://t.me/learningnets

7-103 About the vSphere Replication Appliance

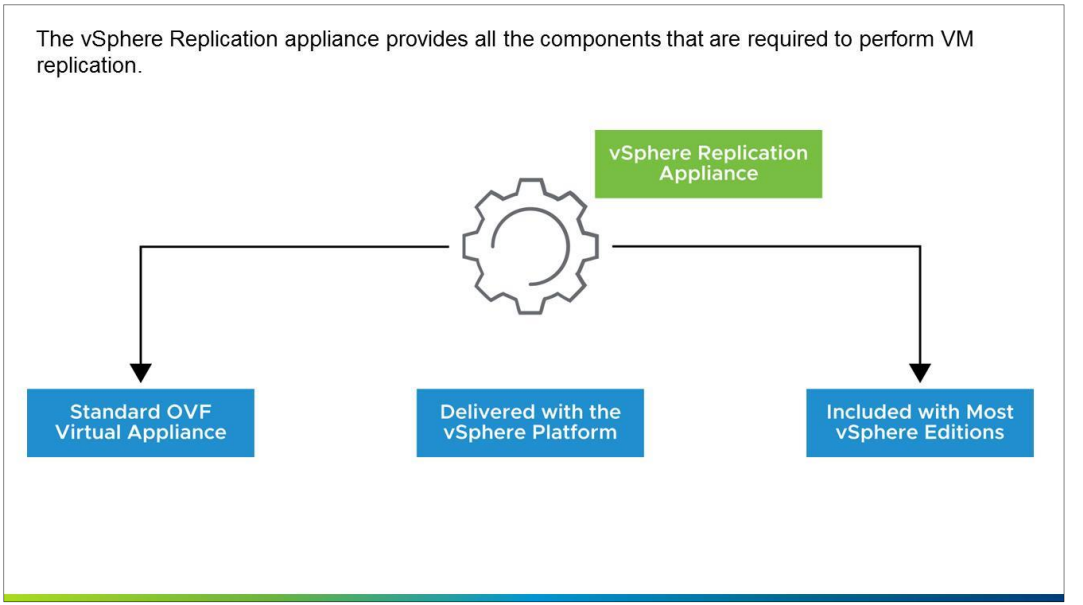

The vSphere Replication appliance includes the following components:

- A vSphere Replication server that provides the core of the vSphere Replication infrastructure
- An embedded database that stores replication configuration and management information
- A vSphere Replication management server that performs the following functions:
  - Configures the vSphere Replication server
  - Enables, manages, and monitors replications
  - Authenticates users and checks their permissions to perform vSphere Replication operations
- A plug-in to the vSphere Client that provides a user interface for vSphere Replication

You can use vSphere Replication immediately after you deploy the appliance. The vSphere Replication appliance provides the virtual appliance management interface (VAMI) that is used to

https://t.me/learningnets

reconfigure the appliance after deployment. For example, you can use the VAMI to change the appliance security settings, change the network settings, or configure an external database. You can deploy additional vSphere Replication servers by using a separate OVF package.

https://t.me/learningnets

7-104 Replication Functions

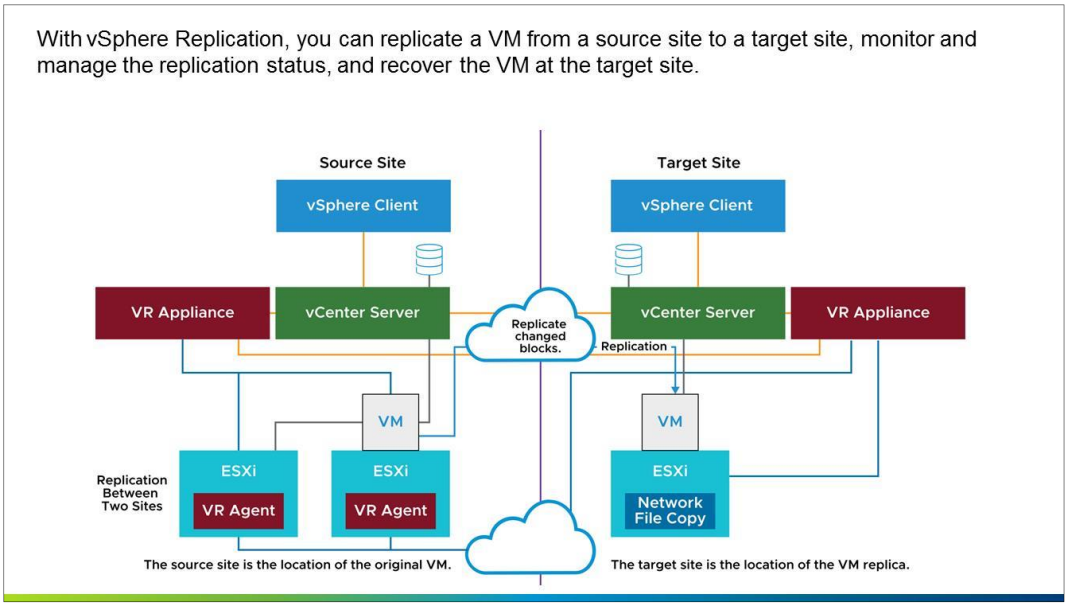

You can replicate a VM between two sites. vSphere Replication is installed on both source and target sites. Only one vSphere Replication appliance is deployed on each vCenter Server. The vSphere Replication (VR) appliance contains an embedded vSphere Replication server that manages the replication process. To meet the load-balancing needs of your environment, you might need to deploy additional vSphere Replication servers at each site.

When you configure a VM for replication, the vSphere Replication agent sends changed blocks in the VM disks from the source site to the target site. The changed blocks are applied to the copy of the VM. This process occurs independently of the storage layer. vSphere Replication performs an initial full synchronization of the source VM and its replica copy. You can use replication seeds to reduce the network traffic that is generated by data transfer during the initial full synchronization.

https://t.me/learningnets

7-105 Deploying the vSphere Replication Appliance

You use the vSphere Client to deploy the vSphere Replication appliance on an ESXi host:

1. Download the installation package to a local directory or obtain its URL.
2. Use the standard vSphere OVF deployment wizard to deploy the appliance.
   a. Select **Menu** > **Hosts and Clusters**.
   b. Right-click an ESXi host and select **Deploy OVF template.**

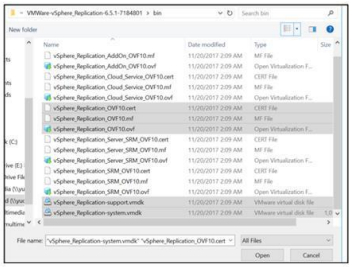

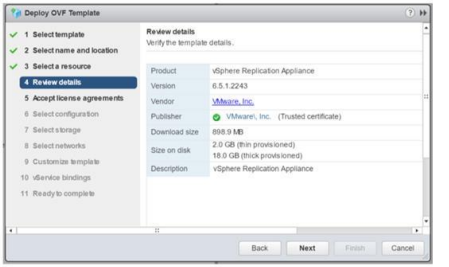

You can deploy vSphere Replication with either an IPv4 or IPv6 address. Mixing IP addresses, for example having a single appliance with an IPv4 and an IPv6 address, is not supported.

After you deploy the vSphere Replication appliance, you use the VAMI to register the endpoint and the certificate of the vSphere Replication management server with the vCenter Lookup Service. You also use the VAMI to register the vSphere Replication solution user with the vCenter Single Sign-On administration server.

For more details on deploying the vSphere Replication appliance, see *VMware vSphere Replication Documentation* at https://docs.vmware.com/en/vSphere-Replication/index.html.

https://t.me/learningnets

7-106 Configuring vSphere Replication for a Single VM

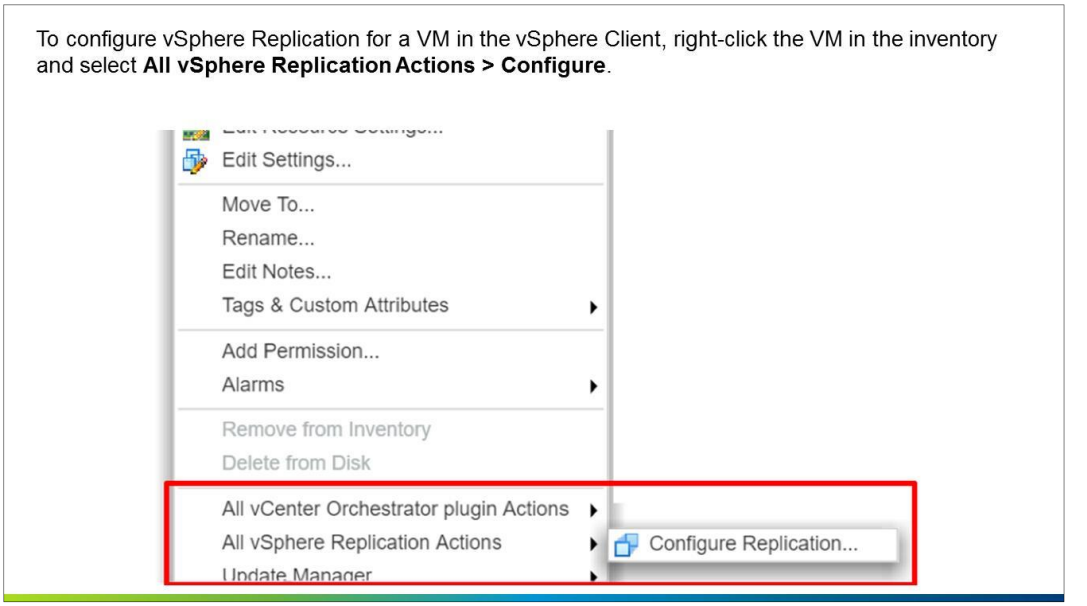

vSphere Replication can protect individual VMs and their virtual disks by replicating them to another location.

https://t.me/learningnets

7-107 Configuring Recovery Point Objective and Point in Time Instances

During replication configuration, you can set an RPO and enable retention of instances from multiple points in time.

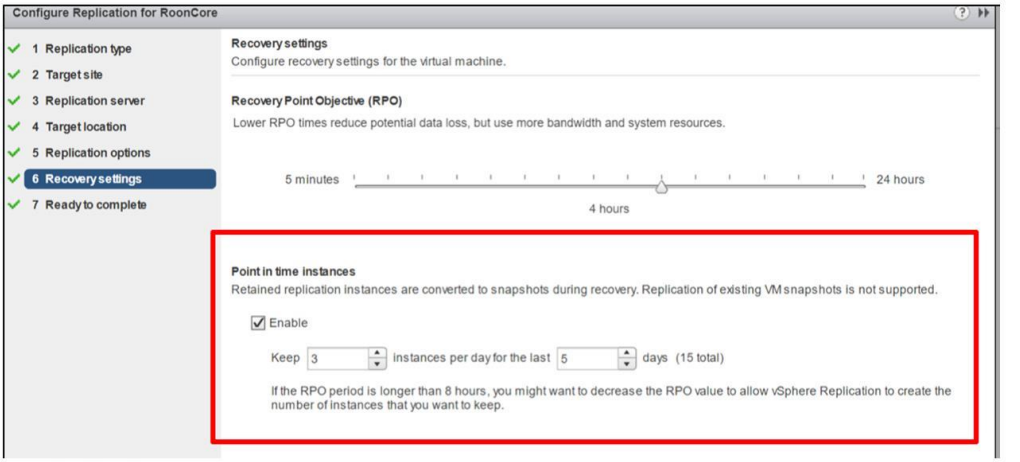

The value that you set for the recovery point objective (RPO) affects replication scheduling.

When you configure replication, you set an RPO to determine the time between replications. For example, an RPO of 1 hour aims to ensure that a VM loses no more than 1 hour of data during the recovery. For smaller RPOs, less data is lost in a recovery, but more network bandwidth is consumed to keep the replica up to date.

For a discussion about how the RPO affects replication scheduling, see *vSphere Replication Administration* at https://docs.vmware.com/en/vSphere-Replication/8.3/com.vmware.vsphere.replication-admin.doc/GUID-35C0A355-C57B-430B-876E-9D2E6BE4DDBA.html.

https://t.me/learningnets

7-108 Recovering Replicated VMs

With vSphere Replication, you can recover VMs that were successfully replicated at the target site. You can recover one VM at a time on the **Incoming Replications** tab.

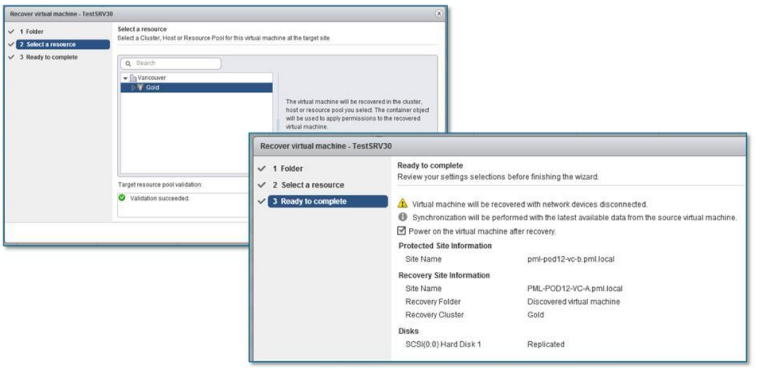

To perform the recovery, you use the Recover virtual machine wizard in the vSphere Client at the target site.

You are asked to select either to recover the VM with all the latest data or to recover the VM with the most recent data available on the target site:

- If you select **Recover with recent changes to avoid data loss**, vSphere Replication performs a full synchronization of the VM from the source site to the target site before recovering the VM. This option requires that the data of the source VM be accessible. You can select this option only if the VM is powered off.
- If you select **Recover with latest available data**, vSphere Replication recovers the VM by using the data from the most recent replication on the target site, without performing synchronization. Selecting this option results in the loss of any data that changed since the most recent replication. Select this option if the source VM is inaccessible or if its disks are corrupted.

https://t.me/learningnets

vSphere Replication validates the input that you provide and recovers the VM. If successful, the VM status changes to Recovered. The VM appears in the inventory of the target site.

https://t.me/learningnets

7-109 Backup and Restore Solution for VMs

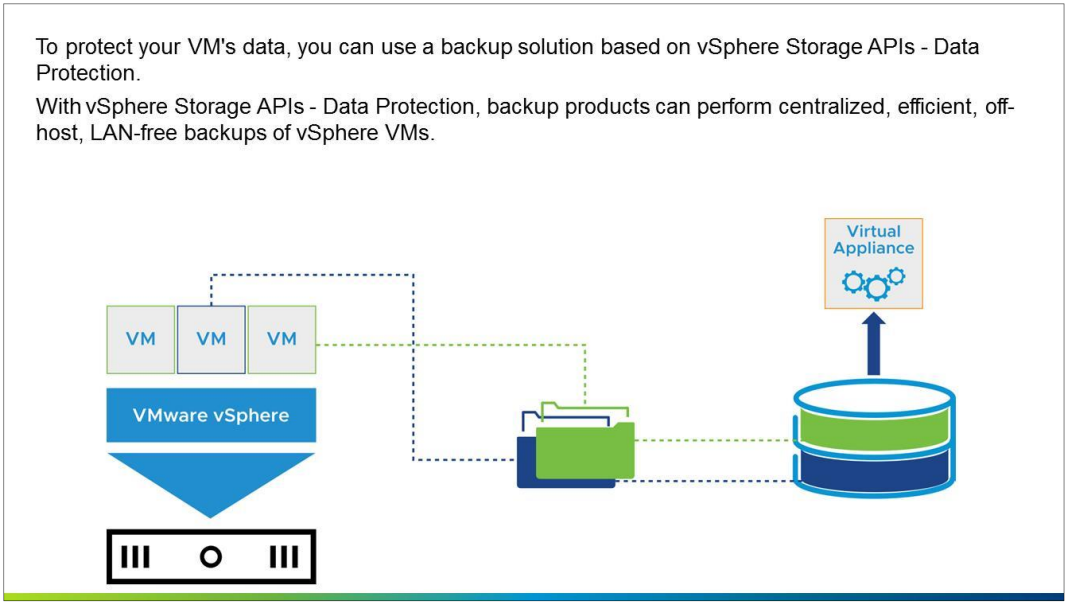

vSphere Storage APIs – Data Protection is VMware's data protection framework, which was introduced in vSphere 4.0. A backup product that uses this API can back up VMs from a central backup system (physical or virtual system). The backup does not require backup agents or any backup processing to be done inside the guest operating system.

Backup processing is offloaded from the ESXi host. In addition, vSphere snapshot capabilities are used to support backups across the SAN without requiring downtime for VMs. As a result, backups can be performed nondisruptively at any time of the day without requiring extended backup windows.

For frequently asked questions about vSphere Storage APIs - Data Protection, see VMware knowledge base article 1021175 at https://kb.vmware.com/s/article/1021175.

https://t.me/learningnets

7-110 vSphere Storage APIs - Data Protection: Offloaded Backup Processing

Configure the storage environment so that the backup server can access the storage volumes that are managed by the ESXi hosts.

Backup processing is offloaded from the ESXi host to the backup server, which prevents local ESXi resources from becoming overloaded.

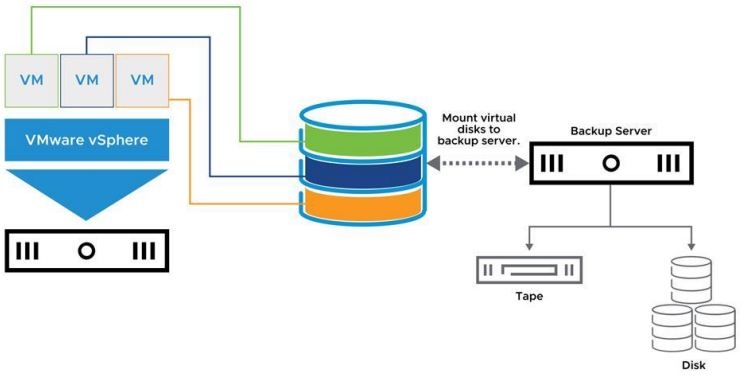

One of the biggest bottlenecks that limits backup performance is the backup server that is handling all the backup coordination tasks. One of these backup tasks is copying data from point A to point B. Other backup tasks do much CPU processing. For example, tasks are performed to determine what data to back up and what not to back up. Other tasks are performed to deduplicate data and compress data that is written to the target.

A server with insufficient CPU resources can greatly reduce backup performance. Provide enough resources for your backup server. A physical server or VM with an ample amount of memory and CPU capacity is necessary for the best backup performance possible.

The motivation to use LAN-free backups is to reduce the stress on the physical resources of the ESXi host when VMs are backed up. LAN-free backups reduce the stress by offloading backup processing from the ESXi host to a backup proxy server.

https://t.me/learningnets

You can configure your environment for LAN-free backups to the backup server, also called the backup proxy server. For LAN-free backups, the backup server must be able to access the storage managed by the ESXi hosts on which the VMs for backup are running.

If you use NAS or direct-attached storage, ensure that the backup proxy server accesses the volumes with a network-based transport. If you run a direct SAN backup, zone the SAN and configure the disk subsystem host mappings. The host mappings must be configured so that all ESXi hosts and the backup proxy server access the same disk volumes.

https://t.me/learningnets

7-111 vSphere Storage APIs - Data Protection: Changed-Block Tracking

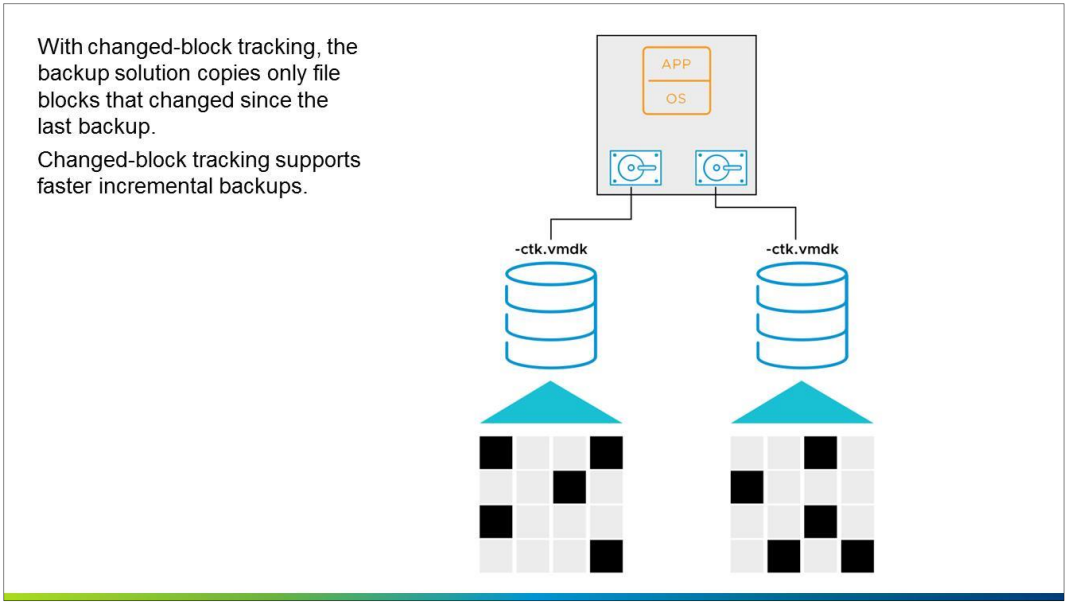

Changed-block tracking (CBT) is a VMkernel feature that tracks the storage blocks of VMs as they change over time. The VMkernel tracks block changes on VMs, enhancing the backup process for applications that are developed to exploit vSphere Storage APIs - Data Protection.

By using CBT during restores, vSphere Data Protection offers fast and efficient recoveries of VMs to their original location. During a restore process, the backup solution uses CBT to determine which blocks changed since the last backup. The use of CBT reduces data transfer within the vSphere environment during a recovery operation and, more important, reduces the recovery time.

https://t.me/learningnets

7-112 Review of Learner Objectives

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

- Identify the components in the vSphere Replication architecture
- Deploy and configure vSphere Replication
- Recover replicated virtual machines
- Explain the backup and restore solution for VMs
- Describe the benefits of vSphere Storage APIs - Data Protection

https://t.me/learningnets

7-113 Activity: Virtual Beans VM Management (1)

As a Virtual Beans administrator, you work with your team to consider which vSphere features to use for key VM management processes. Provide one or more suggestions for each process.

Virtual Beans Process	vSphere Suggestions
Provisioning and deploying VMs	
Maintaining VMs (patching and upgrading operating systems and applications)	
Backing up VMs	
Disaster recovery and business continuity	

https://t.me/learningnets

7-114 Activity: Virtual Beans VM Management (2)

As a Virtual Beans administrator, you work with your team to consider which vSphere features to use for key VM management processes. Provide one or more suggestions for each process.

Virtual Beans Process	vSphere Suggestions
Provisioning and deploying VMs	• Use VM templates. Consider creating a template and a customization specification for each guest operating system type. • Manage all templates with the content library. Have the primary data center publish its library to the secondary data center, when it comes online.
Maintaining VMs (patching and upgrading operating systems and applications)	• Take a snapshot of the VM before applying any patches or updates. • Manage all templates with the content library. Using the content library, you can update templates while VMs are deployed from the template.

https://t.me/learningnets

7-115 Activity: Virtual Beans VM Management (3)

As a Virtual Beans administrator, you work with your team to consider which vSphere features to use for key VM management processes. Provide one or more suggestions for each process.

Virtual Beans Process	vSphere Suggestions
Backing up VMs	• Use a vSphere Storage APIs - Data Protection solution. Check knowledge base article 1021175 for backup solutions that are currently available.
Disaster recovery and business continuity	• Use vSphere Replication, which protects VMs from partial or complete site failure. • For planned downtime, use the various types of vSphere vMotion migrations to move VMs between hosts, between vCenter Server instances, and even between data centers.

https://t.me/learningnets

7-116 Key Points

- vCenter Server provides features for provisioning virtual machines, such as templates, cloning, and content libraries.
- By deploying VMs from a template, you can create many VMs easily and quickly.
- You can dynamically manage a VM's configuration by adding hot-pluggable devices and increasing the size of a VM's virtual disk.
- Hot migrations use vSphere vMotion, vSphere Storage vMotion, or both.
- You can use VM snapshots to preserve the state of the VM so that you can return repeatedly to the same state.
- You can use vSphere Replication to protect VMs as part of a disaster recovery strategy.
- Backup products that use vSphere Storage APIs - Data Protection can be used to back up VM data.

Questions?

https://t.me/learningnets

Module 8

Resource Management and Monitoring

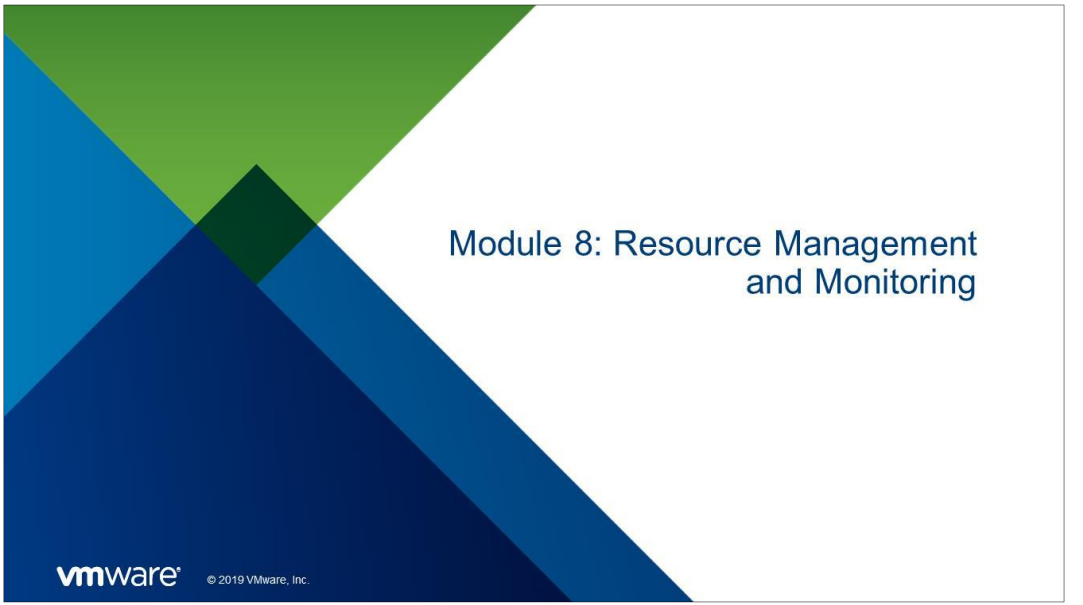

https://t.me/learningnets

8-2 Importance

Although the VMkernel works proactively to avoid resource contention, maximizing performance requires both analysis and ongoing monitoring. Developing skills in resource management, you can dynamically reallocate resources so that you can use available capacity more efficiently.

https://t.me/learningnets

8-3 Module Lessons

1. Virtual CPU and Memory Concepts
2. Resource Controls
3. Resource Monitoring Tools
4. Monitoring Resource Use
5. Using Alarms

https://t.me/learningnets

8-4 Virtual Beans: Resource Management and Monitoring

Virtual Beans wants to proactively manage and monitor its vSphere environment.
Virtual Beans administrators must be able to:

- Increase allocation of CPU and memory resources for business-critical workloads, particularly during peak months.
- Monitor VM performance to troubleshoot user problems.
- Monitor ESXi host performance to avoid potential problems in the infrastructure.
- Create monthly reports, for management, that contain graphs of VM resource usage.
- Set notifications for when ESXi hosts experience high resource use.

As a Virtual Beans administrator, you must use the available tools in vSphere for managing and monitoring the vSphere environment.

https://t.me/learningnets

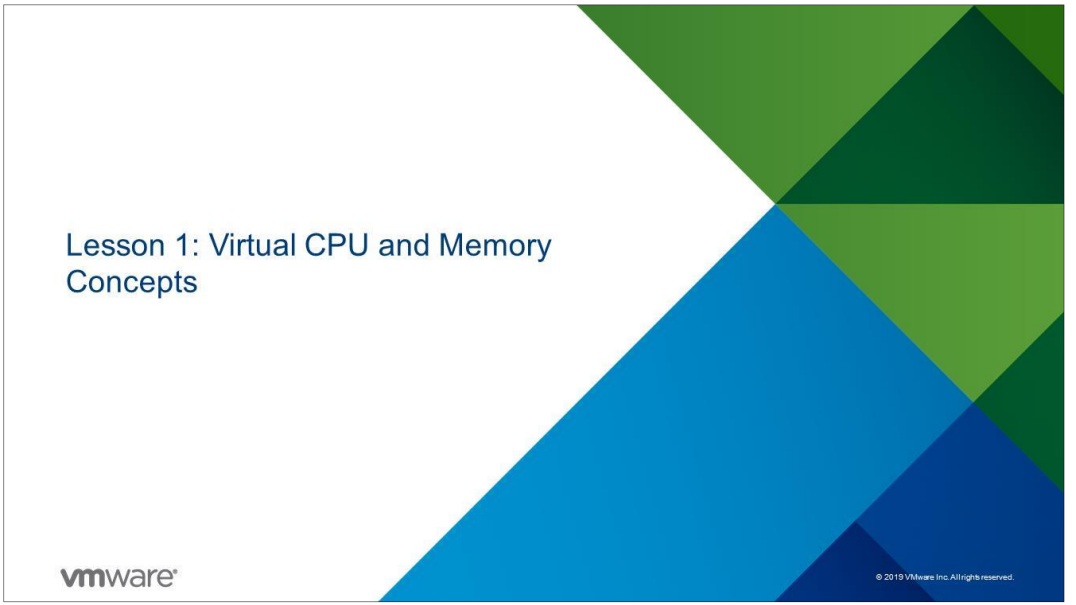

https://t.me/learningnets

8-6 Learner Objectives

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

- Describe CPU and memory concepts in relation to a virtualized environment
- Recognize techniques for addressing memory resource overcommitment
- Identify additional technologies that improve memory usage
- Describe how VMware Virtual SMP works
- Explain how the VMkernel uses hyperthreading

https://t.me/learningnets

8-7 Memory Virtualization Basics

vSphere has the following layers of memory:

- Guest OS virtual memory is presented to applications by the operating system.
- Guest OS physical memory is presented to the virtual machine by the VMkernel.
- Host machine memory that is managed by the VMkernel provides a contiguous, addressable memory space that is used by the VM.

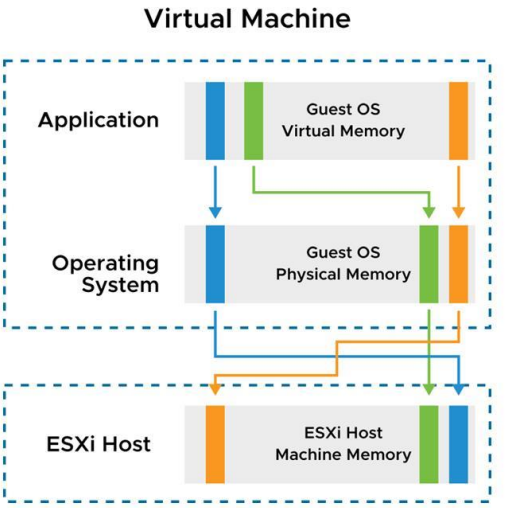

When running a virtual machine, the VMkernel creates a contiguous addressable memory space for the VM. This memory space has the same properties as the virtual memory address space presented to applications by the guest operating system. This memory space allows the VMkernel to run multiple VMs simultaneously while protecting the memory of each VM from being accessed by others. From the perspective of an application running in the VM, the VMkernel adds an extra level of address translation that maps the guest physical address to the host physical address.

https://t.me/learningnets

8-8 VM Memory Overcommitment

Memory is overcommitted when the combined configured memory footprint of all powered-on VMs exceeds that of the host memory sizes. When memory is overcommitted:

- VMs do not always use their full allocated memory.
- To improve memory usage, an ESXi host transfers memory from idle VMs to VMs that need more memory.
- Overcommitted memory is stored in the `.vswp` file.
- Memory overhead is stored in the `vmx-*.vswp` file.

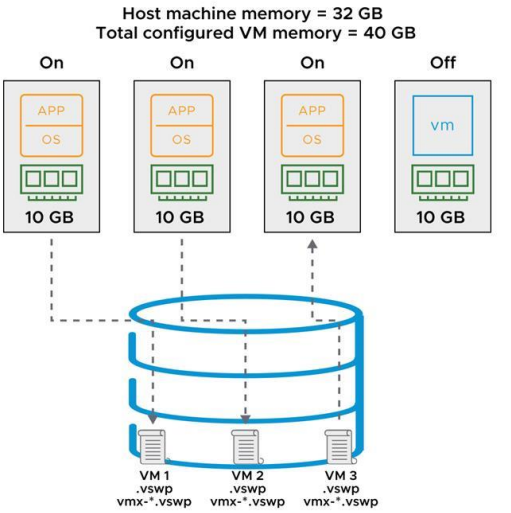

The total configured memory sizes of all VMs might exceed the amount of available physical memory on the host. However, this condition does not necessarily mean that memory is overcommitted. Memory is overcommitted when the working memory size of all VMs exceeds that of the ESXi host's physical memory size.

Because of the memory management techniques used by the ESXi host, your VMs can use more virtual RAM than the available physical RAM on the host. For example, you can have a host with 32 GB of memory and run four VMs with 10 GB of memory each. In that case, the memory is overcommitted. If all four VMs are idle, the combined consumed memory is below 32 GB. However, if all VMs are actively consuming memory, then their memory footprint might exceed 32 GB and the ESXi host becomes overcommitted. An ESXi host can run out of memory if VMs consume all reservable memory in an overcommitted-memory environment. Although the powered-on VMs are not affected, a new VM might fail to power on because of lack of memory. Overcommitment makes sense because, typically, some VMs are lightly loaded whereas others are more heavily loaded, and relative activity levels vary over time.

https://t.me/learningnets

Extra memory from a VM is gathered into a swap file with the `.vswp` extension. The memory overcommitment process on the host uses the `vmx-*.vswp` swap file to gather and track memory overhead. Memory from this file is swapped out to disk when host machine memory is overcommitted.

https://t.me/learningnets

8-9 Memory Overcommit Techniques

An ESXi host uses memory overcommit techniques to allow the overcommitment of memory while possibly avoiding the need to page memory out to disk.

Methods Used by the ESXi Host	Details
Transparent page sharing	This method economizes the use of physical memory pages. In this method, pages with identical contents are stored only once.
Ballooning	This method uses the VMware Tools balloon driver to deallocate memory from one VM to another. The ballooning mechanism becomes active when memory is scarce, forcing VMs to use their own paging areas.
Memory compression	This method tries to reclaim some memory performance when memory contention is high.
Host-level SSD swapping	Use of a solid-state drive on the ESXi host for a host cache swap file might increase performance.
VM memory paging to disk	Using VMkernel swap space is the last resort because of poor performance.

The VMkernel uses various techniques to dynamically reduce the amount of physical RAM that is required for each VM. Each technique is described in the order that the VMkernel uses it:

- Page sharing: ESXi can use a proprietary technique to transparently share memory pages between VMs, eliminating redundant copies of memory pages. Although pages are shared by default within VMs, as of vSphere 6.0, pages are no longer shared by default among VMs.
- Ballooning: If the host memory begins to get low and the VM's memory use approaches its memory target, ESXi uses ballooning to reduce that VM's memory demands. Using the VMware-supplied `vmmemctl` module installed in the guest operating system as part of VMware Tools, ESXi can cause the guest operating system to relinquish the memory pages it considers least valuable. Ballooning provides performance closely matching that of a native system under similar memory constraints. To use ballooning, the guest operating system must be configured with sufficient swap space.

https://t.me/learningnets

- Memory compression: If the VM's memory use approaches the level at which host-level swapping is required, ESXi uses memory compression to reduce the number of memory pages that it must swap out. Because the decompression latency is much smaller than the swap-in latency, compressing memory pages has significantly less impact on performance than swapping out those pages.
- Swap to host cache: Host swap cache is an optional memory reclamation technique that uses local flash storage to cache a virtual machine's memory pages. By using local flash storage, the virtual machine avoids the latency associated with a storage network that might be used if it swapped memory pages to the virtual swap (`.vswp`) file.
- Regular host-level swapping: When memory pressure is severe and the hypervisor must swap memory pages to disk, the hypervisor swaps to a host swap cache rather than to a `.vswp` file. When a host runs out of space on the host cache, a virtual machine's cached memory is migrated to a virtual machine's regular `.vswp` file. Each host must have its own host swap cache configured.

https://t.me/learningnets

8-10 Configuring Multicore VMs

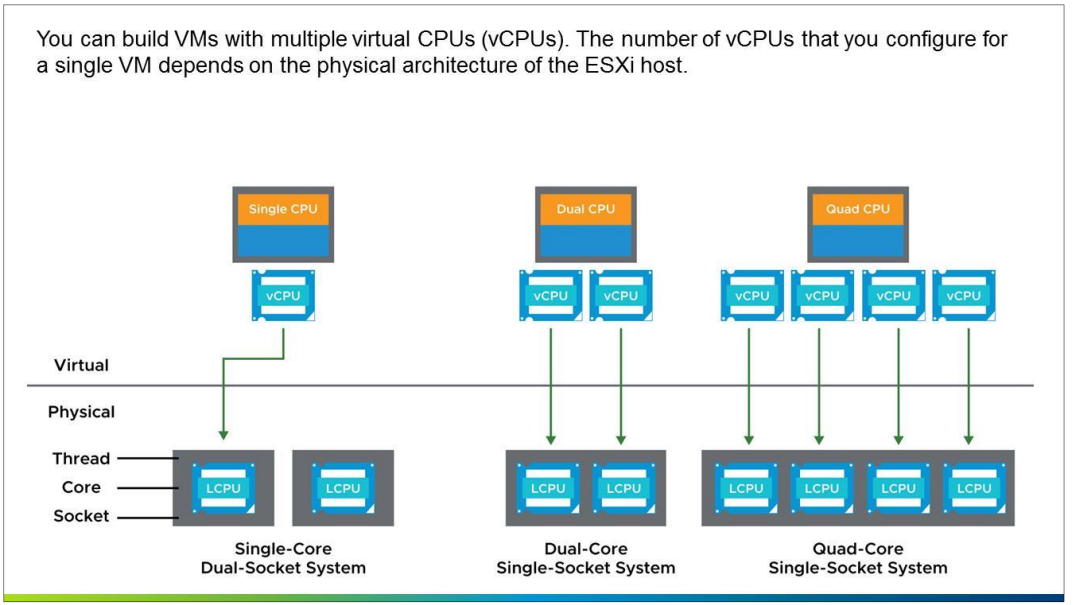

You can configure a VM with up to 256 virtual CPUs (vCPUs). The VMkernel includes a CPU scheduler that dynamically schedules vCPUs on the physical CPUs of the host system.

The VMkernel scheduler considers socket-core-thread topology when making scheduling decisions. Intel and AMD processors combine multiple processor cores into a single integrated circuit, called a socket in this discussion.

A socket is a single package with one or more physical CPUs. Each core has one or more logical CPUs (LCPU in the diagram) or threads. With logical CPUs, the core can schedule one thread of execution.

On the slide, the first system is a single-core, dual-socket system with two cores and, therefore, two logical CPUs.

When a vCPU of a single-vCPU or multi-vCPU VM must be scheduled, the VMkernel maps the vCPU to an available logical processor.

https://t.me/learningnets

In addition to the physical host configuration, the number of vCPUs configured for a VM also depends on the guest operating system, the applications, and the specific use case for the VM itself.

https://t.me/learningnets

8-11 About Hyperthreading

With hyperthreading, a core can execute two threads or sets of instructions at the same time.

- Hyperthreading provides more scheduler throughput.
- Hyperthreading is enabled by default.

To enable hyperthreading:

- Verify that the host system supports hyperthreading.
- Enable hyperthreading in the system BIOS.
- Ensure that hyperthreading for the ESXi host is turned on.

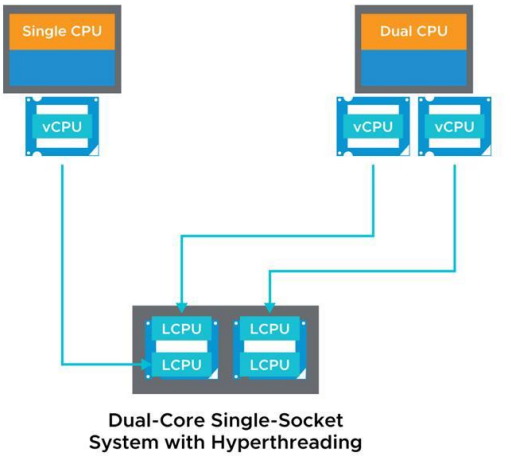

Dual-Core Single-Socket System with Hyperthreading

If hyperthreading is enabled, ESXi can schedule two threads at the same time on each processor core (physical CPU). Hyperthreading provides more scheduler throughput. That is, hyperthreading provides more logical CPUs on which vCPUs can be scheduled.

The drawback of hyperthreading is that it does not double the power of a core. So, if both threads of execution need the same on-chip resources at the same time, one thread has to wait. Still, on systems that use hyperthreading technology, performance is improved.

An ESXi host that is enabled for hyperthreading should behave almost exactly like a standard system. Logical processors on the same core have adjacent CPU numbers. Logical processors 0 and 1 are on the first core, logical processors 2 and 3 are on the second core, and so on.

Consult the host system hardware documentation to verify whether the BIOS includes support for hyperthreading. Then, enable hyperthreading in the system BIOS. Some manufacturers call this option Logical Processor and others call it Enable Hyperthreading.

Use the vSphere Client to ensure that hyperthreading for your host is turned on. To access the hyperthreading option, go to the host's **Summary** tab and select **CPUs** under Hardware.

https://t.me/learningnets

8-12 CPU Load Balancing

The VMkernel balances processor time to guarantee that the load is spread smoothly across processor cores in the system.

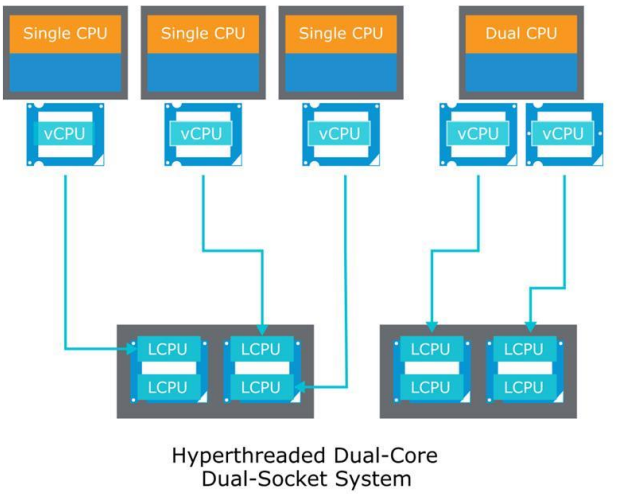

The CPU scheduler can use each logical processor independently to execute VMs, providing capabilities that are similar to traditional symmetric multiprocessing (SMP) systems. The VMkernel intelligently manages processor time to guarantee that the load is spread smoothly across processor cores in the system. Every 2 milliseconds to 40 milliseconds (depending on the socket-core-thread topology), the VMkernel seeks to migrate vCPUs from one logical processor to another to keep the load balanced.

The VMkernel does its best to schedule VMs with multiple vCPUs on two different cores rather than on two logical processors on the same core. But, if necessary, the VMkernel can map two vCPUs from the same VM to threads on the same core.

If a logical processor has no work, it is put into a halted state. This action frees its execution resources, and the VM running on the other logical processor on the same core can use the full execution resources of the core. Because the VMkernel scheduler accounts for this halt time, a VM running with the full resources of a core is charged more than a VM running on a half core. This approach to processor management ensures that the server does not violate the ESXi resource allocation rules.

https://t.me/learningnets

8-13 Review of Learner Objectives

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

- Describe CPU and memory concepts in relation to a virtualized environment
- Recognize techniques for addressing memory resource overcommitment
- Identify additional technologies that improve memory usage
- Describe how VMware Virtual SMP works
- Explain how the VMkernel uses hyperthreading

https://t.me/learningnets

https://t.me/learningnets

8-15 Learner Objectives

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

- Assign share values for CPU and memory resources
- Describe how virtual machines compete for resources
- Define CPU and memory reservations and limits

https://t.me/learningnets

8-16 Reservations, Limits, and Shares

Beyond the CPU and memory configured for a VM, you can apply resource allocation settings to a VM to control the amount of resources granted:

- A reservation specifies the guaranteed minimum allocation for a VM.
- A limit specifies an upper bound for CPU or memory that can be allocated to a VM.
- A share is a value that specifies the relative priority or importance of a VM's access to a given resource.

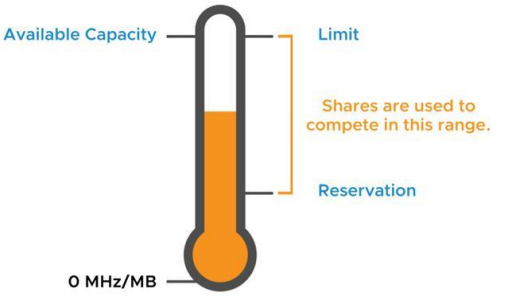

Because VMs simultaneously use the resources of an ESXi host, resource contention can occur.

To manage resources efficiently, vSphere provides mechanisms to allow less, more, or an equal amount of access to a defined resource. vSphere also prevents a VM from consuming large amounts of a resource. vSphere grants a guaranteed amount of a resource to a VM whose performance is not adequate or that requires a certain amount of a resource to run properly.

When host memory or CPU is overcommitted, a VM's allocation target is somewhere between its specified reservation and specified limit, depending on the VM's shares and the system load. vSphere uses a share-based allocation algorithm to achieve efficient resource use for all VMs and to guarantee a given resource to the VMs that need it most.

https://t.me/learningnets

8-17 Resource Allocation Reservations: RAM

RAM reservations:

- Memory reserved to a VM is guaranteed never to swap or balloon.
- If an ESXi host does not have enough unreserved RAM to support a VM with a reservation, the VM does not power on.
- Reservations are measured in MB, GB, or TB.
- The default is 0 MB.
- Adding a vSphere DirectPath I/O device to a VM sets memory reservation to the memory size of the VM.

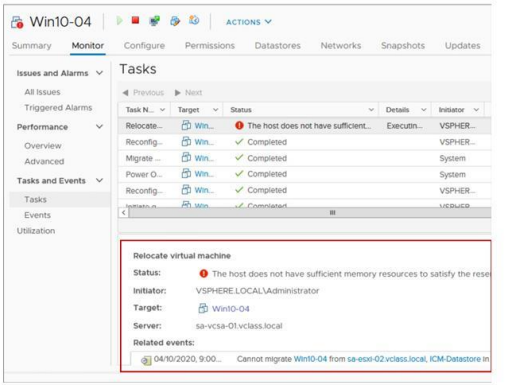

When configuring a memory reservation for a VM, you can specify the VM's configured amount of memory to reserve all of the VM's memory. For example, if a VM is configured with 4 GB of memory, you can set a memory reservation of 4 GB for the VM. You might configure such a memory reservation for a critical VM that must maintain a high level of performance.

Alternatively, you can select the **Reserve All Guest Memory (All locked)** check box. Selecting this check box ensures that all of the VM's memory gets reserved even if you change the total amount of memory for the VM. The memory reservation is immediately readjusted when the VM's memory configuration changes.

https://t.me/learningnets

8-18 Resource Allocation Reservations: CPU

CPU reservations:

- CPU that is reserved for a VM is guaranteed to be immediately scheduled on physical cores. The VM is never placed in a CPU ready state.
- If an ESXi host does not have enough unreserved CPU to support a VM with a reservation, the VM does not power on.
- Reservations are measured in MHz or GHz.
- The default is 0 MHz.

https://t.me/learningnets

8-19 Resource Allocation Limits

RAM limits:

- VMs never consume more physical RAM than is specified by the memory allocation limit.
- VMs might use the VM swap mechanism (`.vswp`) if the guest OS attempts to consume more RAM than is specified by the limit.

CPU limits:

- VMs never consume more physical CPU than is specified by the CPU allocation limit.
- CPU threads are placed in a ready state if the guest OS attempts to schedule threads faster than the limit allows.

Usually, specifying a limit is not necessary.

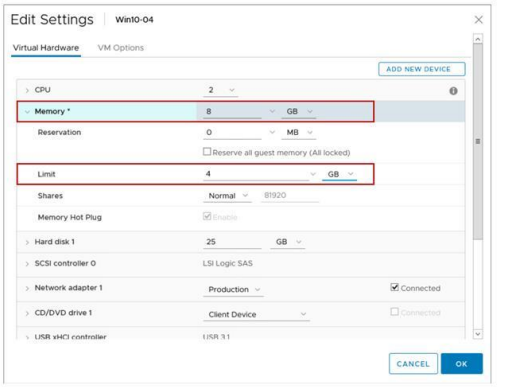

Specifying limits has the following benefits and drawbacks:

- Benefits: Assigning a limit is useful if you start with a few VMs and want to manage user expectations. The performance deteriorates as you add more VMs. You can simulate having fewer resources available by specifying a limit.
- Drawbacks: You might waste idle resources if you specify a limit. The system does not allow VMs to use more resources than the limit, even when the system is underused and idle resources are available. Specify the limit only if you have good reasons for doing so.

https://t.me/learningnets

8-20 Resource Allocation Shares

Shares define the relative importance of a VM:

- If a VM has twice as many shares of a resource as another VM, the VM is entitled to consume twice as much of that resource when these two VMs compete for resources.
- Share values apply only if an ESXi host experiences contention for a resource.

You can set shares to high, normal, or low. You can also select the custom setting to assign a specific number of shares to each VM.

Setting	CPU Share Values	Memory Share Values
High	2,000 shares per vCPU	20 shares per MB of configured VM memory
Normal	1,000 shares per vCPU	10 shares per MB of configured VM memory
Low	500 shares per vCPU	5 shares per MB of configured VM memory

High, normal, and low settings represent share values with a 4:2:1 ratio, respectively. A custom value of shares assigns a specific number of shares (which expresses a proportional weight) to each VM.

https://t.me/learningnets

8-21 Resource Shares Example (1)

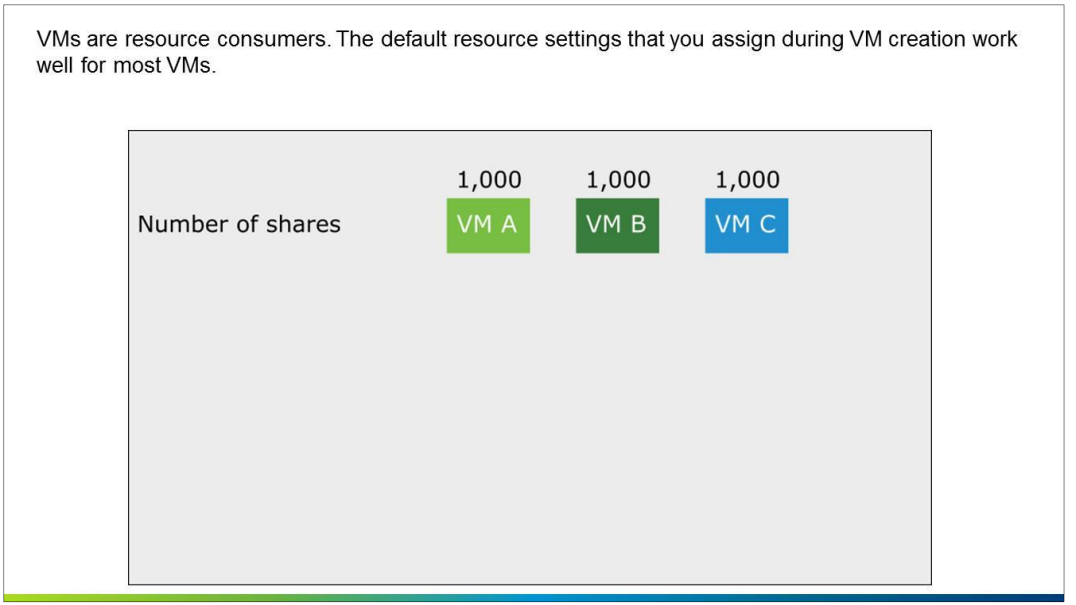

The proportional share mechanism applies to CPU, memory, storage I/O, and network I/O allocation. The mechanism operates only when VMs contend for the same resource.

https://t.me/learningnets

8-22 Resource Shares Example (2)

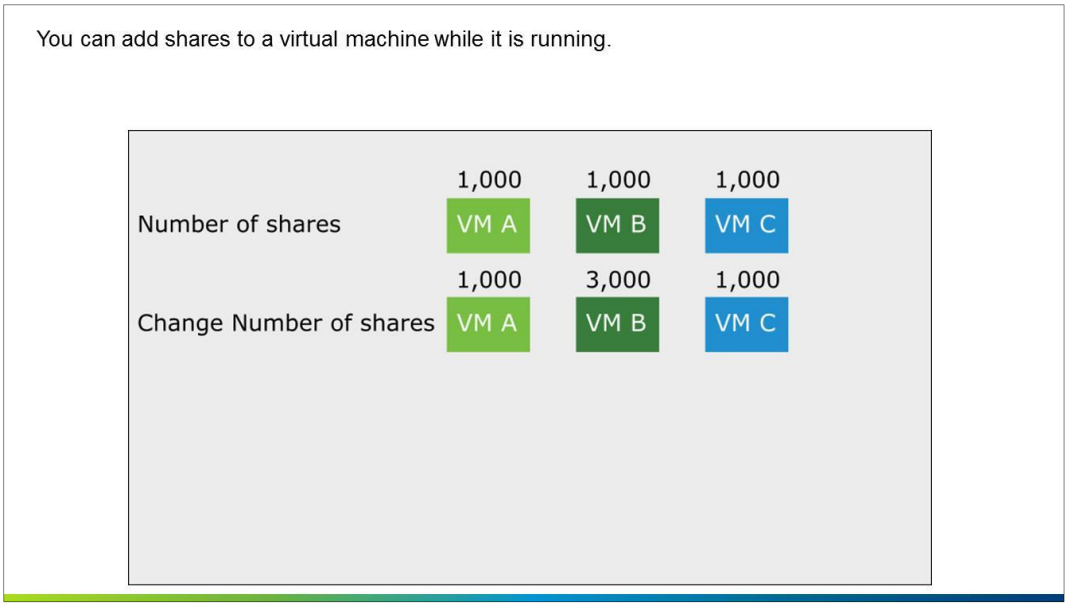

You can add shares to a VM while it is running, and the VM gets more access to that resource (assuming competition for the resource). When you add a VM, it gets shares too. The VM's share amount factors into the total number of shares, but existing VMs are guaranteed not to be starved for the resource.

https://t.me/learningnets

8-23 Resource Shares Example (3)

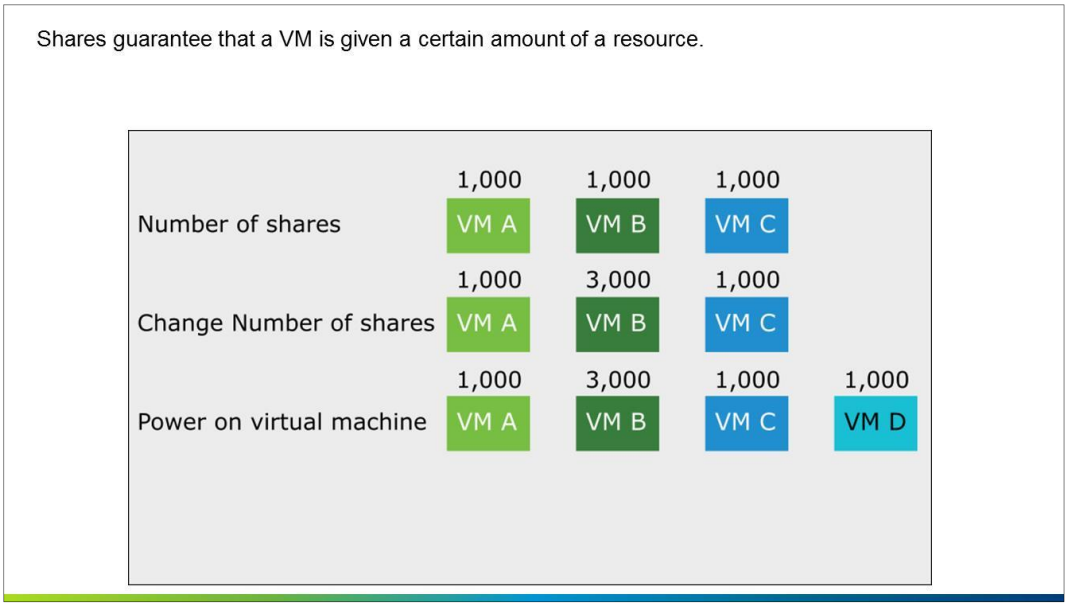

Shares guarantee that a VM is given a certain amount of a resource (CPU, RAM, storage I/O, or network I/O).

For example, consider the third row of VMs on the slide:

- VM D is powered on with 1,000 shares.
- Before VM D was powered on, a total of 5,000 shares were available, but VM D's addition increases the total shares to 6,000.
- The result is that the other VMs' shares decline in value. But each VM's share value still represents a minimum guarantee. VM A is still guaranteed one-sixth of the resource because it owns one-sixth of the shares.

https://t.me/learningnets

8-24 Resource Shares Example (4)

When you delete or power off a VM, fewer total shares remain, so the surviving VMs get more access.

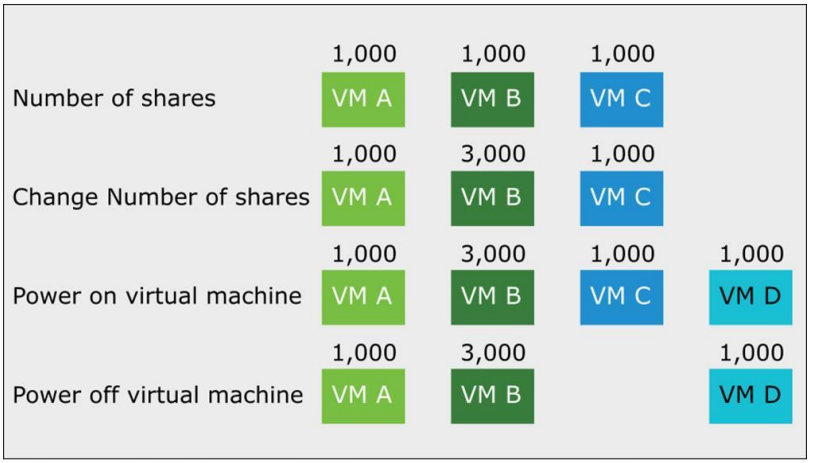

https://t.me/learningnets

8-25 Defining Resource Allocation Settings for a VM

You can edit a VM's settings to configure CPU and memory resource allocations.

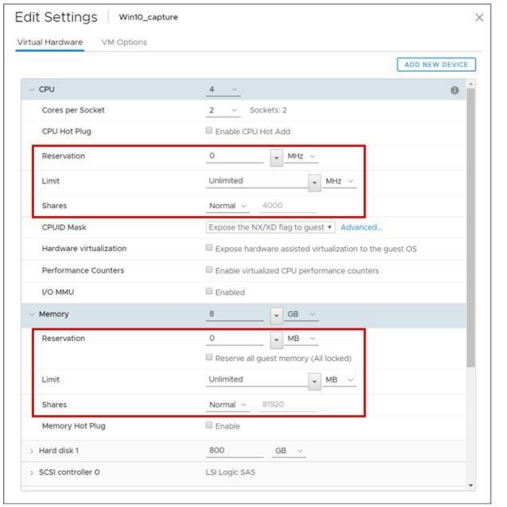

https://t.me/learningnets

8-26 Viewing VM Resource Allocation Settings

You can view reservations, limits, and shares settings for all VMs in a cluster.

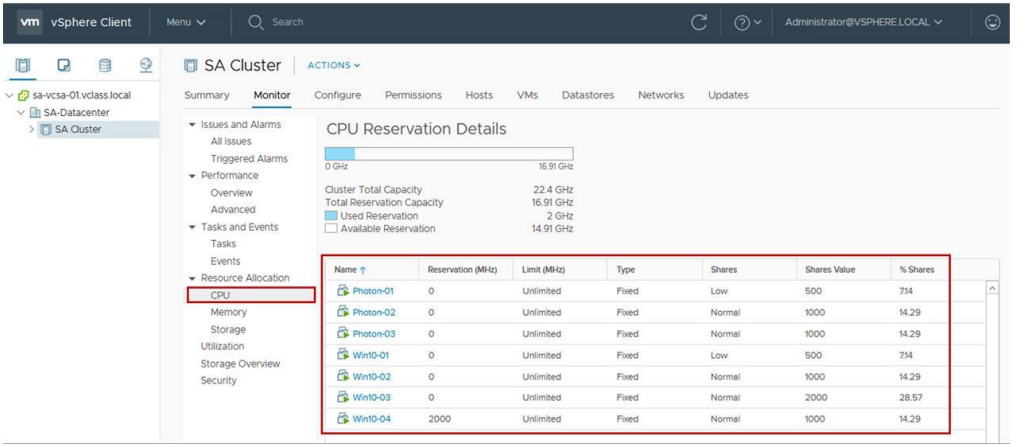

https://t.me/learningnets

8-27 Lab 22: Controlling VM Resources

Control VM CPU resources using shares:

1. Create CPU Contention
2. Verify CPU Share Functionality

https://t.me/learningnets

8-28 Review of Learner Objectives

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

- Assign share values for CPU and memory resources
- Describe how virtual machines compete for resources
- Define CPU and memory reservations and limits

https://t.me/learningnets

https://t.me/learningnets

8-30 Learner Objectives

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

- Describe the performance-tuning methodology
- Identify resource-monitoring tools
- Use vCenter Server performance charts to view performance

https://t.me/learningnets

8-31 Performance-Tuning Methodology

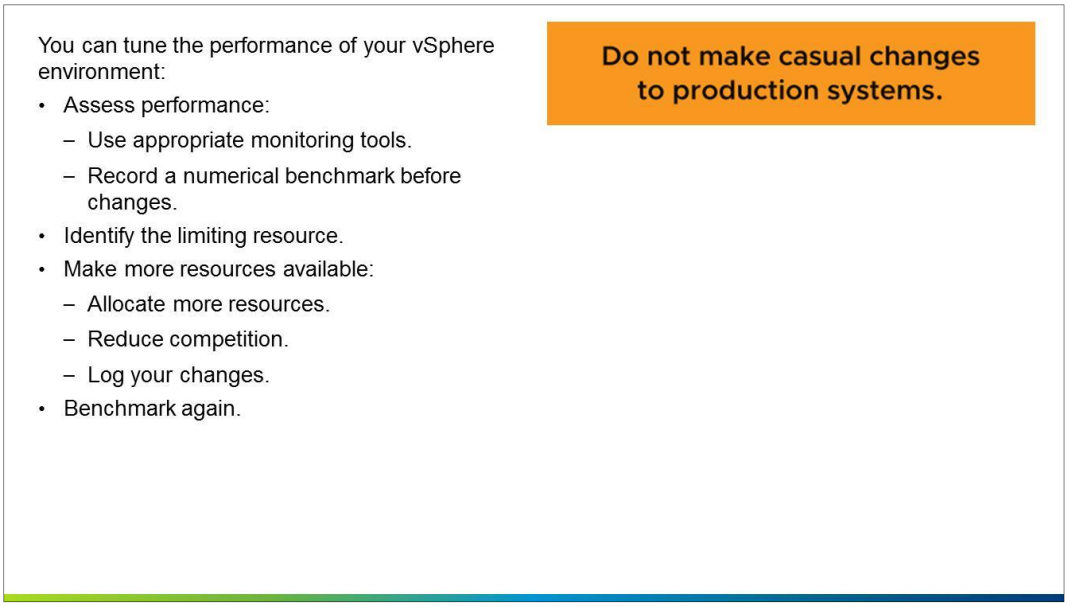

The best practice for performance tuning is to take a logical step-by-step approach:

- For a complete view of the performance situation of a VM, use monitoring tools in the guest operating system and in vCenter Server.
- Identify the resource that the VM relies on the most. This resource is most likely to affect the VM's performance if the VM is constrained by it.
- Give a VM more resources or decrease the resources of other VMs.
- After making more of the limiting resource available to the VM, take another benchmark and record changes.

Be cautious when making changes to production systems because a change might negatively affect the performance of the VMs.

https://t.me/learningnets

8-32 Resource-Monitoring Tools

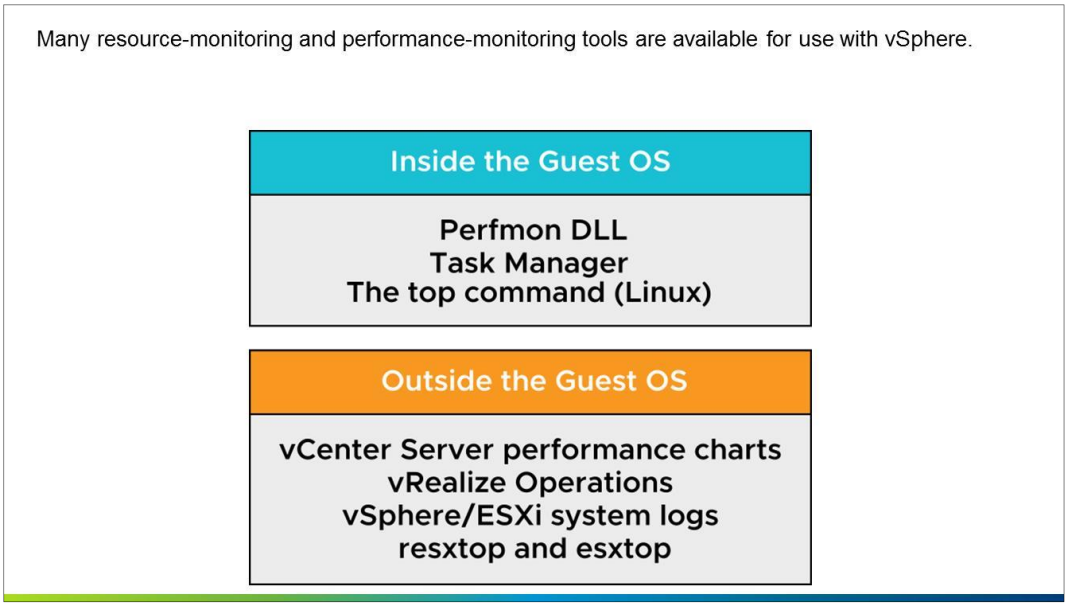

Tools in the guest operating system are available from sources external to VMware and are used in various VMware applications. Many tools used outside of the guest OS are made available by VMware for use with vSphere and other applications.

A partial list of these resource-monitoring tools is shown on the slide.

https://t.me/learningnets

8-33 Guest Operating System Monitoring Tools

To monitor performance in the guest operating system, use tools that you are familiar with, such as Windows Task Manager.

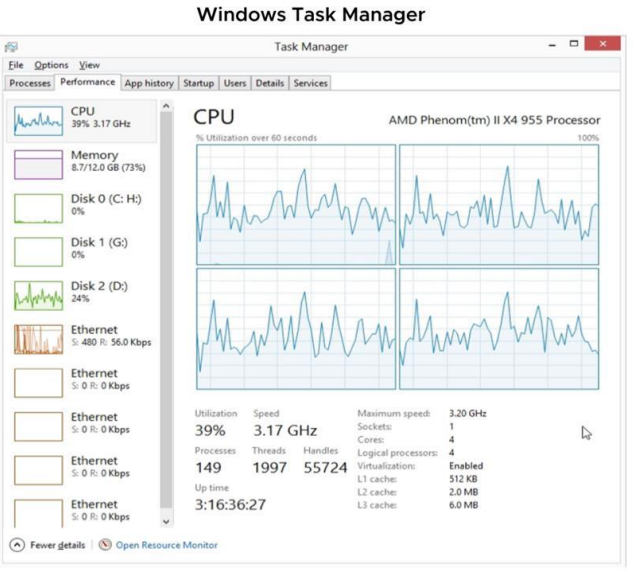

Windows Task Manager helps you measure CPU and memory use in the guest operating system.

The measurements that you take with tools in the guest operating system reflect resource usage of the guest operating system, not necessarily of the VM itself.

https://t.me/learningnets

8-34 Using Perfmon to Monitor VM Resources

The Perfmon DLL in VMware Tools provides VM processor and memory objects for accessing host statistics in a VM.

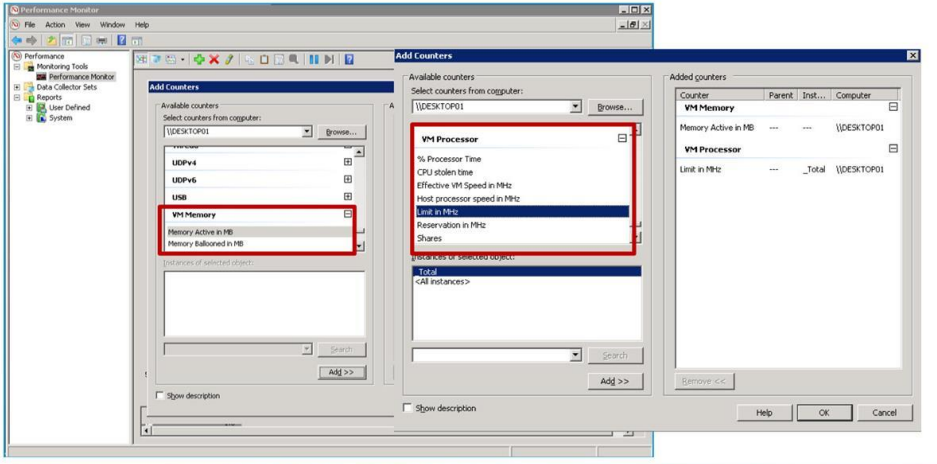

VMware Tools includes a library of functions called the Perfmon DLL. With Perfmon, you can access key host statistics in a guest VM. Using the Perfmon performance objects (VM Processor and VM Memory), you can view actual CPU and memory usage and observed CPU and memory usage of the guest operating system.

For example, you can use the VM Processor object to view the % Processor Time counter, which monitors the VM's current virtual processor load. Likewise, you can use the Processor object and view the % Processor Time counter (not shown), which monitors the total use of the processor by all running processes.

https://t.me/learningnets

8-35 Using esxtop to Monitor VM Resources

The `esxtop` utility is the primary real-time performance monitoring tool for vSphere:

- Can be run from the host's local vSphere ESXi Shell as `esxtop`
- Can be run remotely from vSphere CLI as `resxtop`
- Works like the `top` performance utility in Linux operating systems

In this example, you enter lowercase **c** and uppercase **V** to view CPU metrics for VMs.

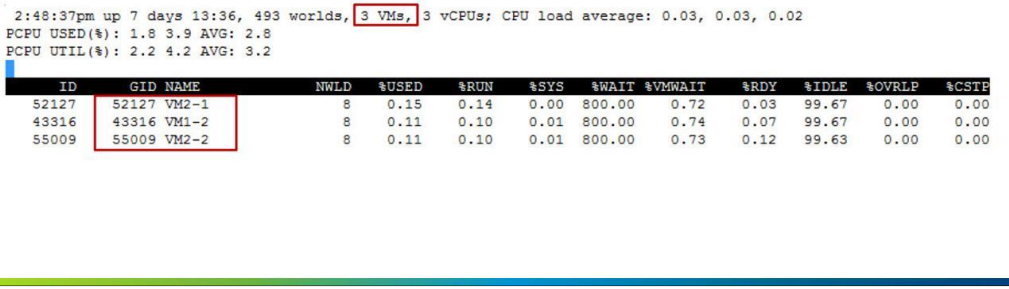

You can run the `esxtop` utility by using vSphere ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

https://t.me/learningnets

8-36 Monitoring Inventory Objects with Performance Charts

The vSphere statistics subsystem collects data on the resource usage of inventory objects, which include:

- Clusters
- Hosts
- Datastores
- Networks
- Virtual machines

Data on a wide range of metrics is collected at frequent intervals, processed, and archived in the vCenter Server database. You can access statistical information through command-line monitoring utilities or by viewing performance charts in the vSphere Client.

https://t.me/learningnets

8-37 Working with Overview Performance Charts

The overview performance charts display the most common metrics for an object in the inventory.

You can access overview and advanced performance charts in the vSphere Client.

Overview performance charts show the performance statistics that VMware considers most useful for monitoring performance and diagnosing problems.

Depending on the object that you select in the inventory, the performance charts provide a quick visual representation of how your host or VM is performing.

https://t.me/learningnets

8-38 Working with Advanced Performance Charts

Advanced charts support data counters that are not supported in other performance charts.

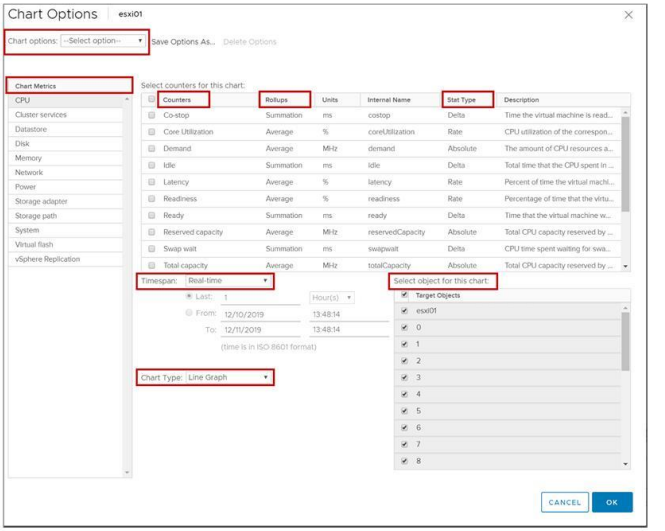

In the vSphere Client, you can customize the appearance of advanced performance charts.

Advanced charts have the following features:

- More information than overview charts: Point to a data point in a chart to display details about that specific data point.
- Customizable charts: Change chart settings. Save custom settings to create your own charts.
- Save data to an image file or a spreadsheet.

To customize advanced performance charts, select **Advanced** under Performance. Click the **Chart Options** link in the Advanced Performance pane.

https://t.me/learningnets

8-39 Chart Options: Real-Time and Historical

vCenter Server stores statistics at different specificities.

Time Interval	Data Frequency	Number of Samples
Real-time (past hour)	20 seconds	180
Past day	5 minutes	288
Past week	30 minutes	336
Past month	2 hours	360
Past year	1 day	365

Real-time information is information that is generated for the past hour at 20-second intervals. Historical information is generated for the past day, week, month, or year, at varying specificities.

By default, vCenter Server has four archiving intervals: day, week, month, and year. Each interval specifies a length of time that statistics are archived in the vCenter Server database.

You can configure which intervals are used and for what period of time. You can also configure the number of data counters that are used during a collection interval by setting the collection level.

Together, the collection interval and the collection level determine how much statistical data is collected and stored in your vCenter Server database.

For example, using the table, past-day statistics show one data point every 5 minutes, for a total of 288 samples. Past-year statistics show 1 data point per day, or 365 samples.

Real-time statistics are not stored in the database. They are stored in a flat file on ESXi hosts and in memory on vCenter Server instances. ESXi hosts collect real-time statistics only for the host or

https://t.me/learningnets

the VMs that are available on the host. Real-time statistics are collected directly on an ESXi host every 20 seconds.

If you query for real-time statistics, vCenter Server queries each host directly for the data. vCenter Server does not process the data at this point. vCenter Server only passes the data to the vSphere Client.

On ESXi hosts, the statistics are kept for 30 minutes, after which 90 data points are collected. The data points are aggregated, processed, and returned to vCenter Server. vCenter Server then archives the data in the database as a data point for the day collection interval.

To ensure that performance is not impaired when collecting and writing the data to the database, cyclical queries are used to collect data counter statistics. The queries occur for a specified collection interval. At the end of each interval, the data calculation occurs.

https://t.me/learningnets

8-40 Chart Types: Bar and Pie

Depending on the metric type and object, performance metrics are displayed in different types of charts, such as bar charts and pie charts.

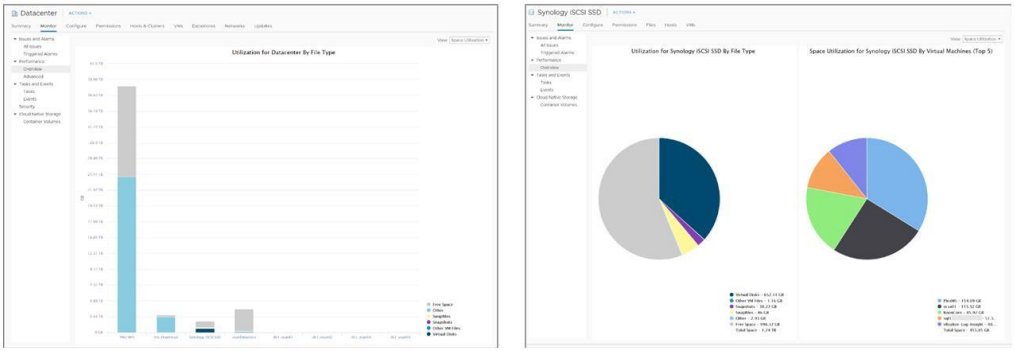

Bar charts display storage metrics for datastores in a selected data center. Each datastore is represented as a bar in the chart. Each bar displays metrics based on the file type: virtual disks, other VM files, snapshots, swap files, and other files.

Pie charts display storage metrics for a single object, based on the file types or VMs. For example, a pie chart for a datastore can display the amount of storage space occupied by the VMs that take up the largest space.

https://t.me/learningnets

8-41 Chart Types: Line

A line chart displays metrics for a single inventory object, for example, metrics for each CPU on an ESXi host.

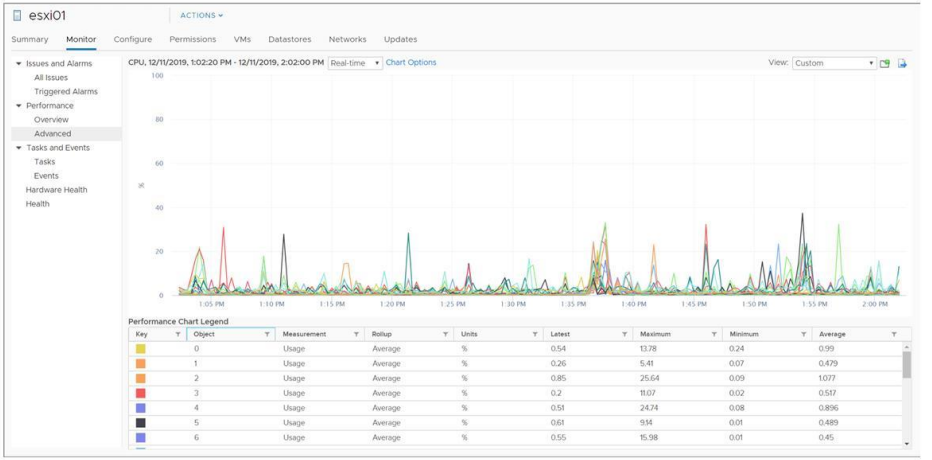

In a line chart, the data for each performance counter is plotted on a separate line in the chart. For example, a CPU chart for a host can contain a line for each of the host's CPUs. Each line plots the CPU's usage over time.

https://t.me/learningnets

8-42 Chart Types: Stacked

Stacked charts are useful for comparing resource allocation and usage across multiple hosts or VMs.

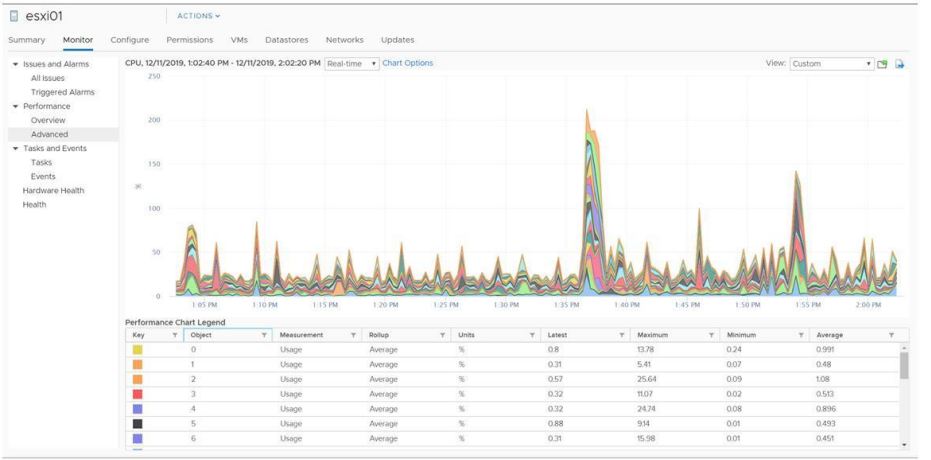

Stacked charts display metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term Other. For example, a host's stacked CPU usage chart displays CPU usage metrics for the five VMs on the host that are consuming the most CPU resources. The Other amount contains the total CPU usage of the remaining VMs. The metrics for the host itself are displayed in separate line charts. By default, the 10 child objects with the highest data counter values appear.

https://t.me/learningnets

8-43 Chart Types: Stacked Per VM

Per-VM stacked graphs are available only for hosts.

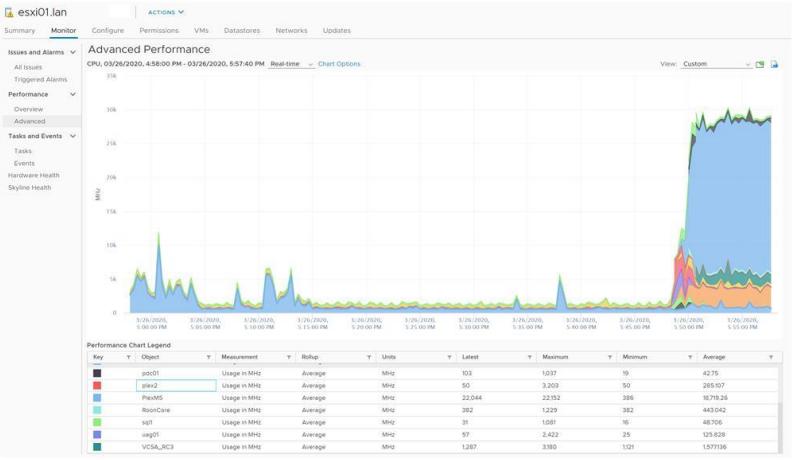

Stacked charts display metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term Other. For example, a host's stacked CPU usage chart displays CPU usage metrics for the five VMs on the host that are consuming the most CPU resources. The Other amount contains the total CPU usage of the remaining VMs. The metrics for the host itself are displayed in separate line charts. By default, the 10 child objects with the highest data counter values appear.

https://t.me/learningnets

8-44 Saving Charts

You click the **Save Chart** icon above the graph to save performance chart information.
You can save information in PNG, JPEG, SVG, and CSV formats.

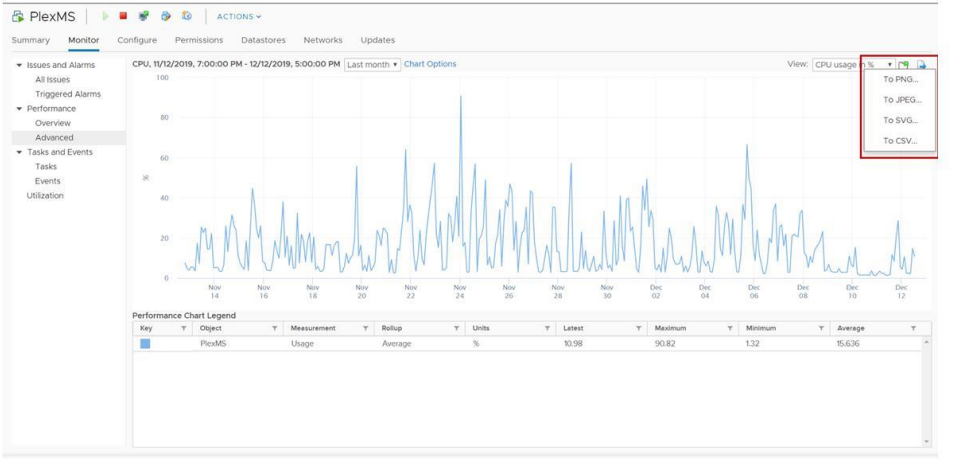

In the vSphere Client, you can save data from the advanced performance charts to a file in various graphics formats or in Microsoft Excel format. When you save a chart, you select the file type and save the chart to the location of your choice.

https://t.me/learningnets

8-45 About Objects and Counters

Performance charts graphically display CPU, memory, disk, network, and storage metrics for devices and entities managed by vCenter Server.

Objects are instances or aggregations of devices:

- Examples:
  - vCPU0
  - vCPU1
  - vmhba1:1:2
  - Aggregation over all NICs

Counters identify which statistics to collect:

- Examples:
  - CPU: Used time, ready time, usage (%)
  - NIC: Network packets received
  - Memory: Memory swapped

In vCenter Server, you can determine how much or how little information about a specific device type is displayed. You can control the amount of information a chart displays by selecting one or more objects and counters.

An object refers to an instance for which a statistic is collected. For example, you might collect statistics for an individual CPU, all CPUs, a host, or a specific network device.

A counter represents the actual statistic that you are collecting. An example is the amount of CPU used or the number of network packets per second for a given device.

https://t.me/learningnets

8-46 About Statistics Types

The statistics type is the unit of measurement that is used during the statistics interval.

Statistics Type	Description	Example
Rate	Value over the current interval	CPU use (MHz)
Delta	Change from previous interval	CPU ready time
Absolute	Absolute value, independent of interval	Memory active

The statistics type refers to the measurement that is used during the statistics interval and is related to the unit of measurement.

The statistics type is one of the following:

- Rate: Value over the current statistics interval
- Delta: Change from the previous statistics interval
- Absolute: Absolute value (independent of the statistics interval)

For example, CPU usage is a rate, CPU ready time is a delta, and memory active is an absolute value.

https://t.me/learningnets

8-47 About Rollup

Rollup is the conversion function between statistics intervals:

- 5 minutes of past-hour statistics are converted to one past-day value:
  - Fifteen 20-second statistics are rolled up into a single value.
- 30 minutes of past-day statistics are converted to 1 past-week value:
  - Six 5-minute statistics are rolled up into a single value.

Rollup Type	Conversion Function	Sample Statistic
Average	Average of data points	CPU use (average)
Summation	Sum of data points	CPU ready time (milliseconds)
Latest	Last data point	Uptime (days)
Minimum	Current or average minimum data point	Available in all counters when vCenter Server statistics collection is set to level 4
Maximum	Current or average maximum data point	

Data is displayed at different specificities according to the historical interval. Past-hour statistics are shown at a 20-second specificity, and past-day statistics are shown at a 5-minute specificity. The averaging that is done to convert from one time interval to another is called rollup.

Different rollup types are available. The rollup type determines the type of statistical values returned for the counter:

- Average: The data collected during the interval is aggregated and averaged.
- Minimum: The minimum value is rolled up.
- Maximum: The maximum value is rolled up.

The minimum and maximum values are collected and displayed only in collection level 4. Minimum and maximum rollup types are used to capture peaks in data during the interval. For real-time data, the value is the current minimum or current maximum. For historical data, the value is the average minimum or average maximum.

https://t.me/learningnets

For example, the following information for the CPU usage chart shows that the average is collected at collection level 1 and that the minimum and maximum values are collected at collection level 4:

- Counter: Usage
- Unit: Percentage (%)
- Rollup Type: Average (Minimum/Maximum)
- Collection Level: 1 (4)

Statistics levels include summation and latest:

- Summation: The collected data is summed. The measurement displayed in the performance chart represents the sum of data collected during the interval.
- Latest: The data that is collected during the interval is a set value. The value displayed in the performance chart represents the current value.

For example, if you look at the CPU Used counter in a CPU performance chart, the rollup type is summation. So, for a given 5-minute interval, the sum of all the 20-second samples in that interval is represented.

https://t.me/learningnets

8-48 Review of Learner Objectives

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

- Describe the performance-tuning methodology
- Identify resource-monitoring tools
- Use vCenter Server performance charts to view performance

https://t.me/learningnets

https://t.me/learningnets

8-50 Learner Objectives

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

- Monitor the key factors that can affect a virtual machine's performance
- Use performance charts to view and improve performance

https://t.me/learningnets

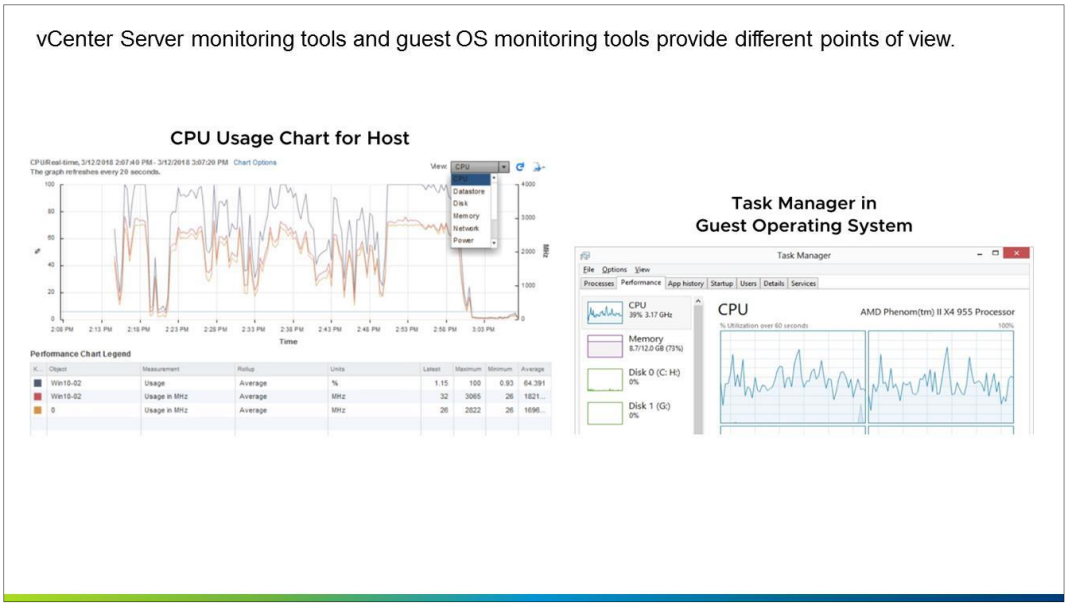

The key to interpreting performance data is to observe the range of data from the perspective of the guest operating system, the VM, and the host.

The CPU usage statistics in Task Manager, for example, do not give you the complete picture. View CPU usage for the VM and the host on which the VM is located.

Use the performance charts in the vSphere Client to view this data.

https://t.me/learningnets

8-52 CPU-Constrained VMs (1)

If CPU use is continuously high, the VM is constrained by the CPU. However, the host might have enough CPU for other VMs to run.

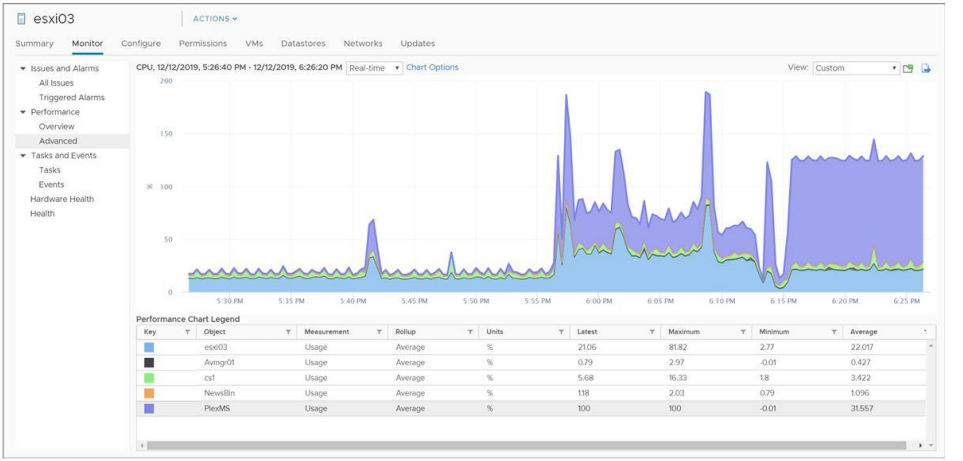

If CPU use is high, check the VM's CPU usage statistics. Use either the overview charts or the advanced charts to view CPU usage. The slide displays an advanced chart tracking a VM’s CPU usage.

If a VM’s CPU use remains high over a period of time, the VM is constrained by CPU. Other VMs on the host might have enough CPU resources to satisfy their needs.

If more than one VM is constrained by CPU, the key indicator is CPU ready time. Ready time refers to the interval when a VM is ready to execute instructions but cannot because it cannot get scheduled onto a CPU. Several factors affect the amount of ready time:

- Overall CPU use: You are more likely to see ready time when use is high because the CPU is more likely to be busy when another VM becomes ready to run.
- Number of resource consumers (in this case, guest operating systems): When a host is running a larger number of VMs, the scheduler is more likely to queue a VM behind VMs that are already running or queued.

https://t.me/learningnets

A good ready time value varies from workload to workload. To find a good ready time value for your workload, collect ready time data over time for each VM. When you have this ready time data for each VM, estimate how much of the observed response time is ready time. If the shortfalls in meeting response-time targets for the applications appear largely because of the ready time, take steps to address the excessive ready time.

https://t.me/learningnets

8-53 CPU-Constrained VMs (2)

Multiple VMs are constrained by the CPU if the following conditions are present:

- High CPU usage in the guest operating system
- Relatively high CPU readiness values for the VMs

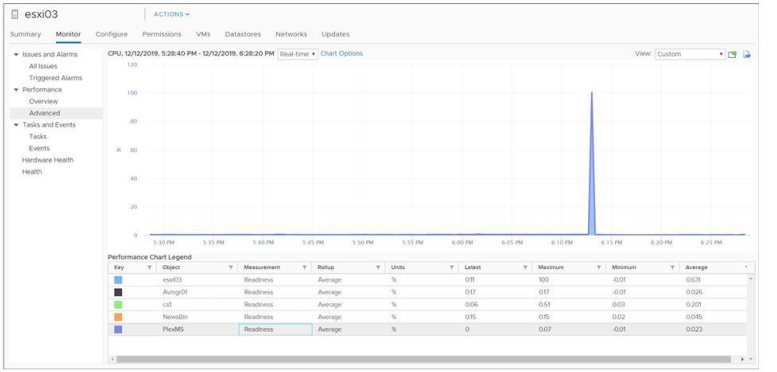

To determine whether a VM is being constrained by CPU resources, view CPU usage in the guest operating system using, for example, Task Manager.

If more than one VM is constrained by CPU, the key indicator is CPU readiness. CPU readiness is the percent of time that the VM cannot run because it is contending for access to the physical CPUs.

You are more likely to see readiness values when use is high because the CPU is more likely to be busy when another VM becomes ready to run. You are also more likely to see readiness values when a host is running many VMs. In this case, the scheduler is more likely to queue a VM behind VMs that are already running or queued.

A good readiness value varies from workload to workload.

https://t.me/learningnets

8-54 Memory-Constrained VMs (1)

Compare a VM's memory consumed and granted values to determine whether the VM is memory-constrained.

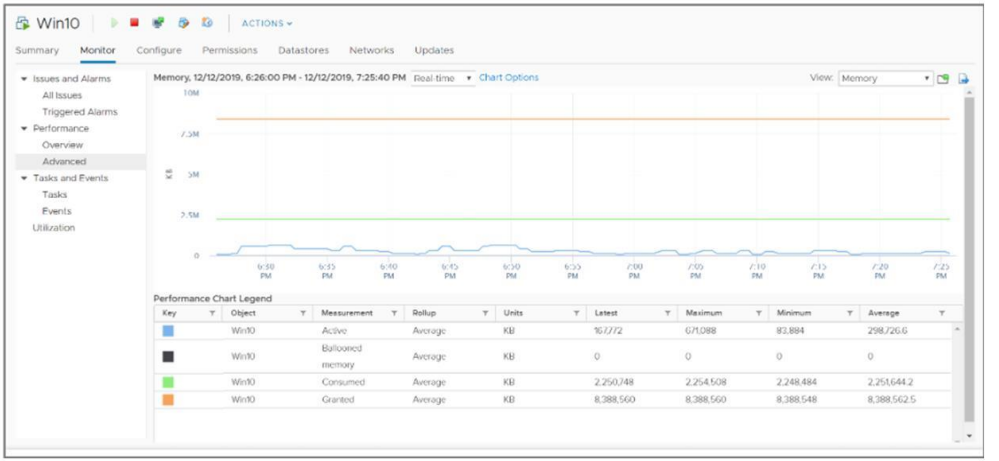

https://t.me/learningnets

8-55 Memory-Constrained VMs (2)

If a VM consumes its entire memory allocation, the VM might be memory-constrained, and you should consider increasing the VM's memory size.

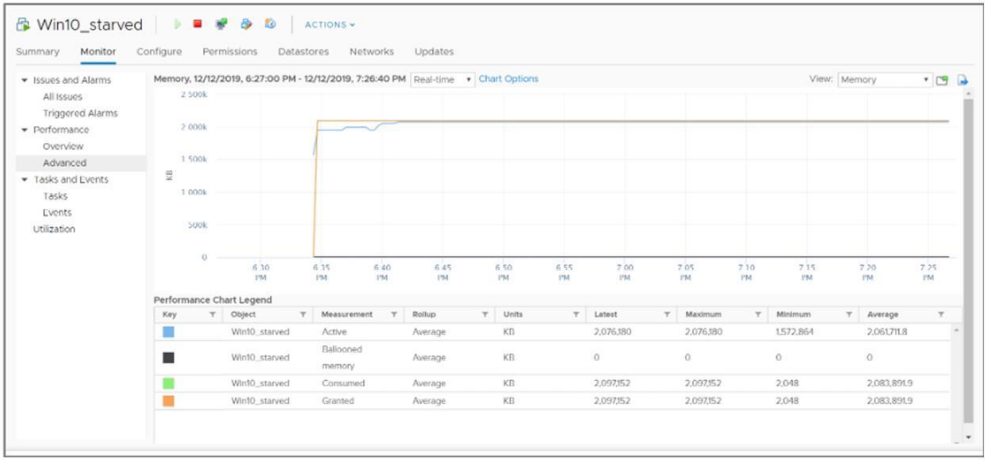

https://t.me/learningnets

8-56 Memory-Constrained Hosts

Any evidence of ballooning or swapping is a sign that your host might be memory-constrained.

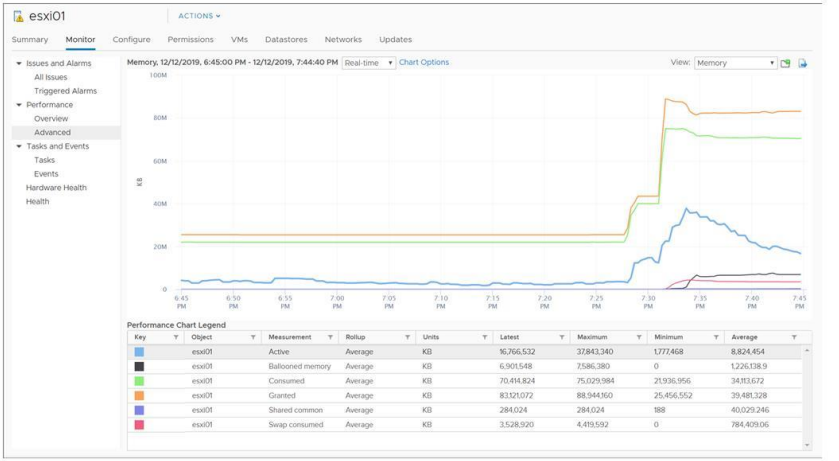

You might see VMs with high ballooning activity and VMs being swapped in and out by the VMkernel. This serious situation indicates that the host memory is overcommitted and must be increased.

https://t.me/learningnets

8-57 Disk-Constrained VMs

Disk-intensive applications can saturate the storage or the path.
If you suspect that a VM is constrained by disk access, take these actions:

- Measure the throughput and latency between the VM and storage.
- Use the advanced performance charts to monitor throughput and latency:
  - Read rate and write rate
  - Read latency and write latency

Disk performance problems are commonly caused by saturating the underlying physical storage hardware. You can use the vCenter Server advanced performance charts to measure storage performance at different levels. These charts provide insight about a VM performance. You can monitor everything from the VM's datastore to a specific storage path.

If you select a host object, you can view throughput and latency for a datastore, a storage adapter, or a storage path. The storage adapter charts are available only for Fibre Channel storage. The storage path charts are available for Fibre Channel and iSCSI storage, not for NFS.

If you select a VM object, you can view throughput and latency for the VM’s datastore or specific virtual disk.

To monitor throughput, view the Read rate and Write rate counters. To monitor latency, view the Read latency and Write latency counters.

https://t.me/learningnets

8-58 Monitoring Disk Latency

To determine disk performance problems, monitor two disk latency data counters:

- Kernel command latency:
  - This counter is the average time that is spent in the VMkernel per SCSI command.
  - High numbers (greater than 2 milliseconds or 3 milliseconds) represent either an overworked array or an overworked host.
- Physical device command latency:
  - This counter is the average time that the physical device takes to complete a SCSI command.
  - High numbers represent a slow or overworked array, for example:
    - For spinning disks (HDDs), greater than 15 milliseconds or 20 milliseconds
    - For SSDs, greater than 3 milliseconds or 4 milliseconds

To determine whether your vSphere environment is experiencing disk problems, monitor the disk latency data counters. Use the advanced performance charts to view these statistics. In particular, monitor the following counters:

- Kernel command latency: This data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0 through 1 millisecond. If the value is greater than 4 milliseconds, the VMs on the ESXi host are trying to send more throughput to the storage system than the configuration supports.
- Physical device command latency: This data counter measures the average amount of time, in milliseconds, for the physical device to complete a SCSI command.

https://t.me/learningnets

8-59 Network-Constrained VMs

Network-intensive applications often bottleneck on path segments outside the ESXi host:

- Example: WAN links between server and client

If you suspect that a VM is constrained by the network, take these actions:

- Verify that VMware Tools is installed and that VMXNET3 is the virtual network adapter.
- Measure the effective bandwidth between the VM and its peer system.
- Check for dropped receive packets and dropped transmit packets.

Like disk performance problems, network performance problems are commonly caused by saturating a network link between client and server. Use a tool such as Iometer, or a large file transfer, to measure the effective bandwidth.

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use the advanced performance charts to examine the droppedTx and droppedRx network counter values of a VM.

In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU that is required to process the data. In some instances, large packets can result in high network latency. When network packets are small, more packets are transferred, but the network speed is slower because more CPU is required to process the data.

https://t.me/learningnets

Lab 23: Monitoring Virtual Machine Performance

Use the system monitoring tools to review the CPU workload:

1. Create a CPU Workload
2. Use Performance Charts to Monitor CPU Use
3. Remove CPU Affinity and Change CPU Shares to Normal

https://t.me/learningnets

8-61 Review of Learner Objectives

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

- Monitor the key factors that can affect a virtual machine's performance
- Use performance charts to view and improve performance

https://t.me/learningnets

https://t.me/learningnets

8-63 Learner Objectives

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

- Use predefined alarms in vCenter Server
- View and acknowledge alarms
- Create custom alarms

https://t.me/learningnets

8-64 About Alarms

An alarm is a notification that is sent in response to an event or condition that occurs with an object in the inventory.

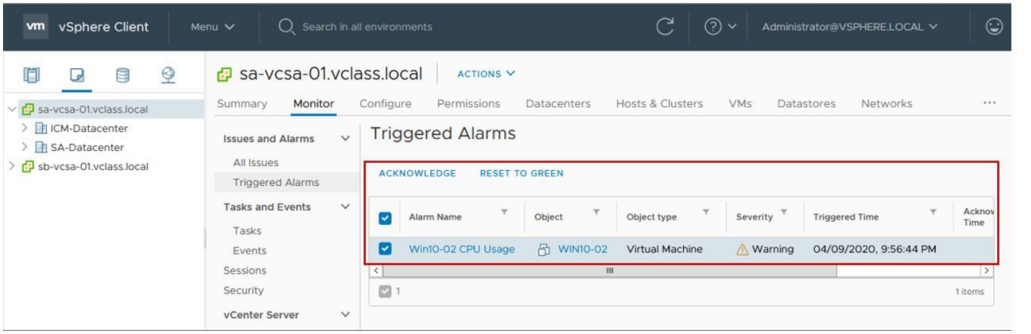

You can acknowledge an alarm to let other users know that you take ownership of the issue. For example, a VM has an alarm set to monitor CPU use. The alarm is configured to send an email to an administrator when the alarm is triggered. The VM CPU use spikes, triggering the alarm, which sends an email to the administrator. The administrator acknowledges the triggered alarm to let other administrators know the problem is being addressed

After you acknowledge an alarm, the alarm actions are discontinued, but the alarm does not get cleared or reset when acknowledged. You reset the alarm manually in the vSphere Client to return the alarm to a normal state.

https://t.me/learningnets

8-65 Predefined Alarms (1)

You can access many predefined alarms for various inventory objects, such as hosts, virtual machines, datastores, networks, and so on.

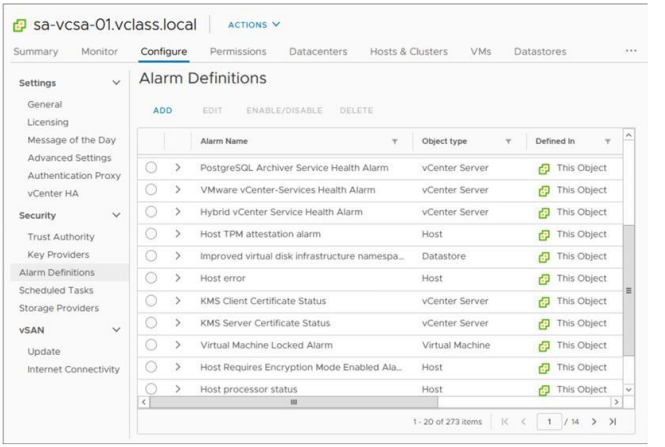

https://t.me/learningnets

8-66 Predefined Alarms (2)

You can edit predefined alarms, or you can make a copy of an existing alarm and modify the settings as needed.

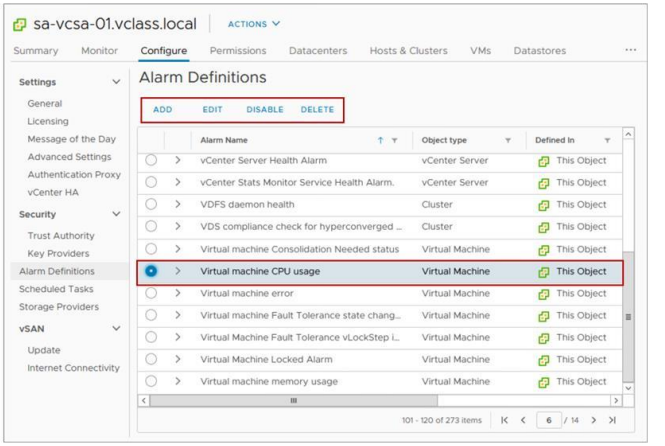

To make a copy of an alarm, select the alarm and click **ADD**.

https://t.me/learningnets

8-67 Creating a Custom Alarm

In addition to using predefined alarms, you can create custom alarms in the vSphere Client.

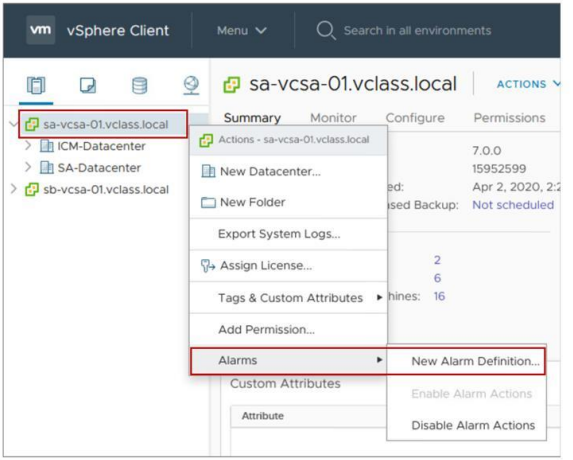

If the predefined alarms do not address the event, state, or condition that you want to monitor, define custom alarm definitions instead of modifying predefined alarms.

https://t.me/learningnets

8-68 Defining the Alarm Target Type

On the Name and Targets page, you name the alarm, give it a description, and select the type of inventory object that this alarm monitors.

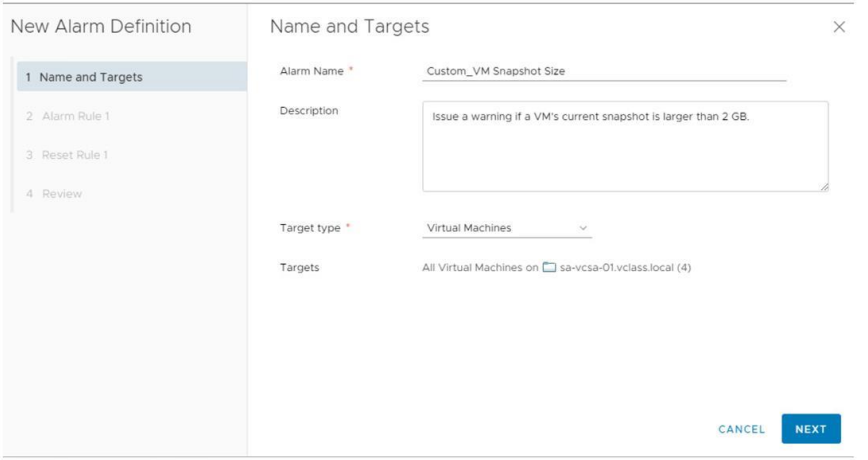

You can create custom alarms for the following target types:

- Virtual machines
- Hosts, clusters, and data centers
- Datastores and datastore clusters
- Distributed switches and distributed port groups
- vCenter Server

https://t.me/learningnets

8-69 Defining the Alarm Rule: Trigger (1)

An alarm rule must contain at least one trigger.

A trigger can monitor the current condition or state of an object, for example:

- A VM's current snapshot is more than 2 GB.
- A host is using 90 percent of its total memory.
- A datastore is disconnected from all hosts.

A trigger can monitor events that occur in response to operations occurring on a managed object, for example:

- The health of a host's hardware changes.
- A license expires in the data center.
- A host leaves the distributed switch.

You configure the alarm trigger to show as a warning or critical event when the specified criteria are met:

- You can monitor the current condition or state of virtual machines, hosts, and datastores. Conditions or states include power states, connection states, and performance metrics such as CPU and disk use.
- You can monitor events that occur in response to operations occurring with a managed object in the inventory or vCenter Server itself. For example, an event is recorded each time a VM (which is a managed object) is cloned, created, deleted, deployed, and migrated.

https://t.me/learningnets

8-70 Defining the Alarm Rule: Trigger (2)

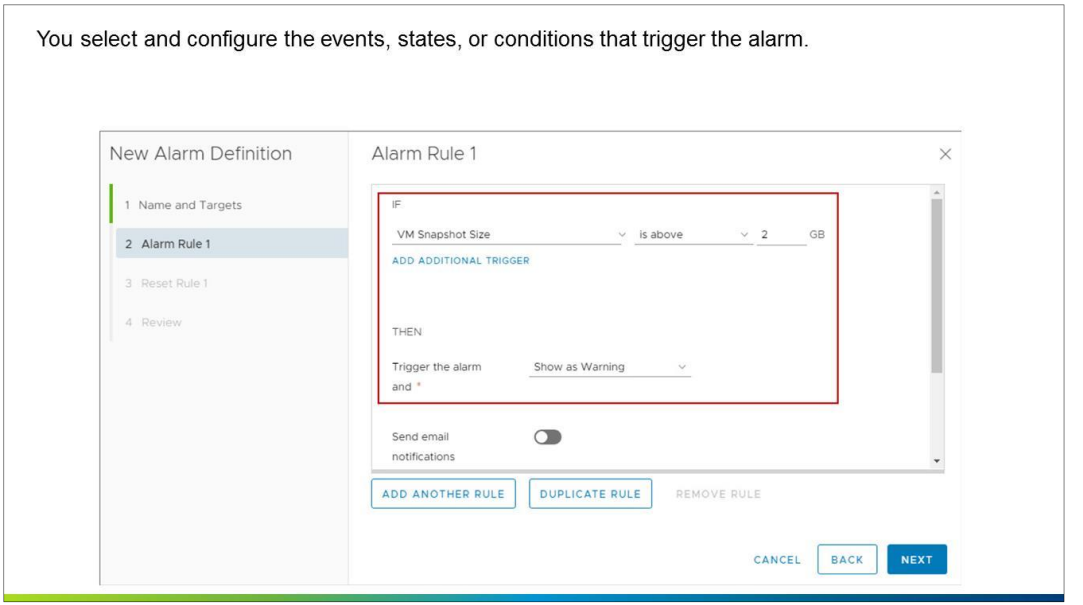

You must create a separate alarm definition for each trigger. The OR operator is not supported in the vSphere Client. However, you can combine more than one condition trigger with the AND operator.

https://t.me/learningnets

8-71 Defining the Alarm Rule: Setting the Notification

You configure the notification method to use when the alarm is triggered. The methods are sending an email, sending an SNMP trap, or running a script.

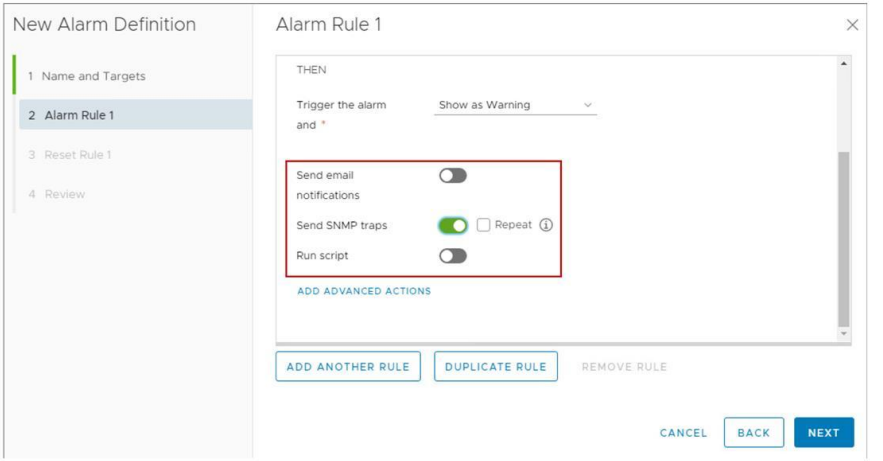

https://t.me/learningnets

8-72 Defining the Alarm Reset Rules

You can select and configure the events, states, or conditions to reset the alarm to normal. Sometimes, as in this example, you can access only one option to reset the alarm.

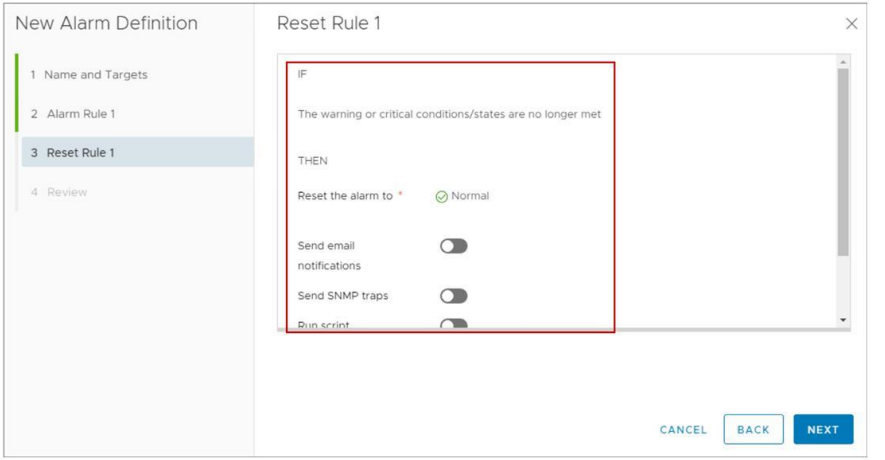

https://t.me/learningnets

8-73 Enabling the Alarm

On the Review page, the new alarm definition is enabled by default.

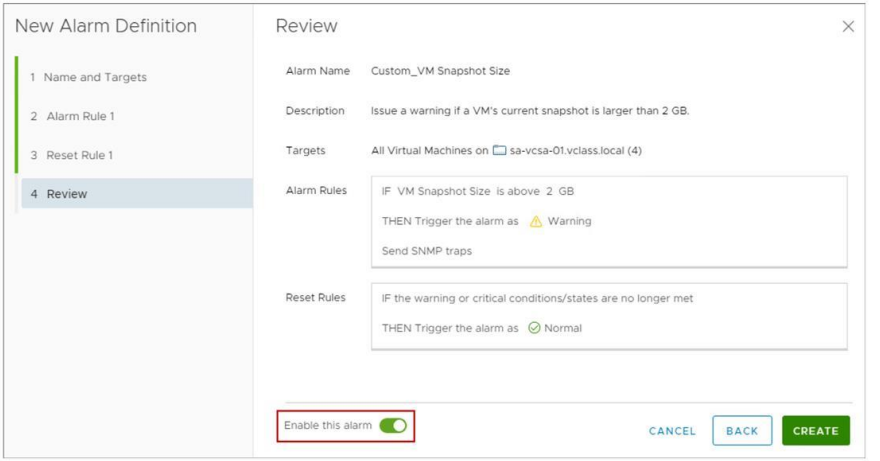

https://t.me/learningnets

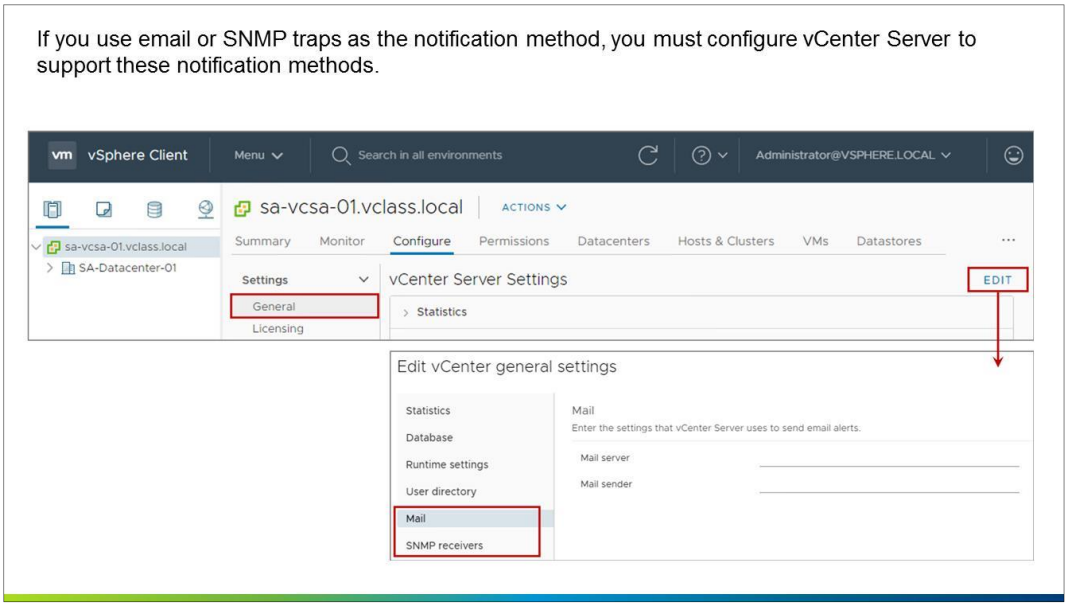

To configure email, specify the mail server FQDN or IP address and the email address of the sender account.

You can configure up to four receivers of SNMP traps. They must be configured in numerical order. Each SNMP trap requires a corresponding host name, port, and community.

https://t.me/learningnets

8-75 Lab 24: Using Alarms

Create alarms to monitor virtual machine events and conditions:

1. Create a Virtual Machine Alarm to Monitor a Condition
2. Trigger the Virtual Machine Alarm
3. Create a Virtual Machine Alarm to Monitor an Event
4. Trigger the Virtual Machine Alarm
5. Disable Virtual Machine Alarms

https://t.me/learningnets

8-76 Review of Learner Objectives

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

- Use predefined alarms in vCenter Server
- View and acknowledge alarms
- Create custom alarms

https://t.me/learningnets

8-77 Activity: Virtual Beans Resource Monitoring (1)

Which tools can Virtual Beans use to meet its goals for managing and monitoring the vSphere environment? Match each Virtual Beans requirement with the appropriate vSphere feature.

Virtual Beans Requirements

- Increase compute resources for business-critical workloads, particularly during peak months.
- Provide proactive recommendations to help avoid problems before they occur.
- Create monthly reports, for management, that contain graphs of VM resource usage.
- Be notified when ESXi hosts experience high CPU and memory usage.

vSphere Features

- Alarms
- Shares, limits, reservations
- VMware Skyline
- vCenter Server performance charts

https://t.me/learningnets

8-78 Activity: Virtual Beans Resource Management and Monitoring (2)

Which tools can Virtual Beans use to meet its goals for managing and monitoring the vSphere environment? Match each Virtual Beans requirement with the appropriate vSphere feature.

Virtual Beans Requirement	vSphere Feature
Increase compute resources for business-critical workloads, particularly during peak months.	Shares, limits, and reservations
Provide proactive recommendations to help avoid problems before they occur.	VMware Skyline
Create monthly reports, for management, that contain graphs of VM resource usage.	vCenter Server performance charts
Be notified when ESXi hosts experience high CPU and memory usage.	Alarms

https://t.me/learningnets

8-79 Key Points

- An ESXi host uses memory overcommit techniques to allow the overcommitment of memory while possibly avoiding the need to page memory out to disk.
- The VMkernel balances processor time to guarantee that the load is spread smoothly across processor cores in the system.
- You can apply reservations, limits, and shares against a VM to control the amount of CPU and memory resources granted.
- The key to interpreting performance data is to observe the range of data from the perspective of the guest operating system, the virtual machine, and the host.
- You use alarms to monitor the vCenter Server inventory objects and send notifications when selected events or conditions occur.

Questions?

https://t.me/learningnets

https://t.me/learningnets

Module 9

vSphere Clusters

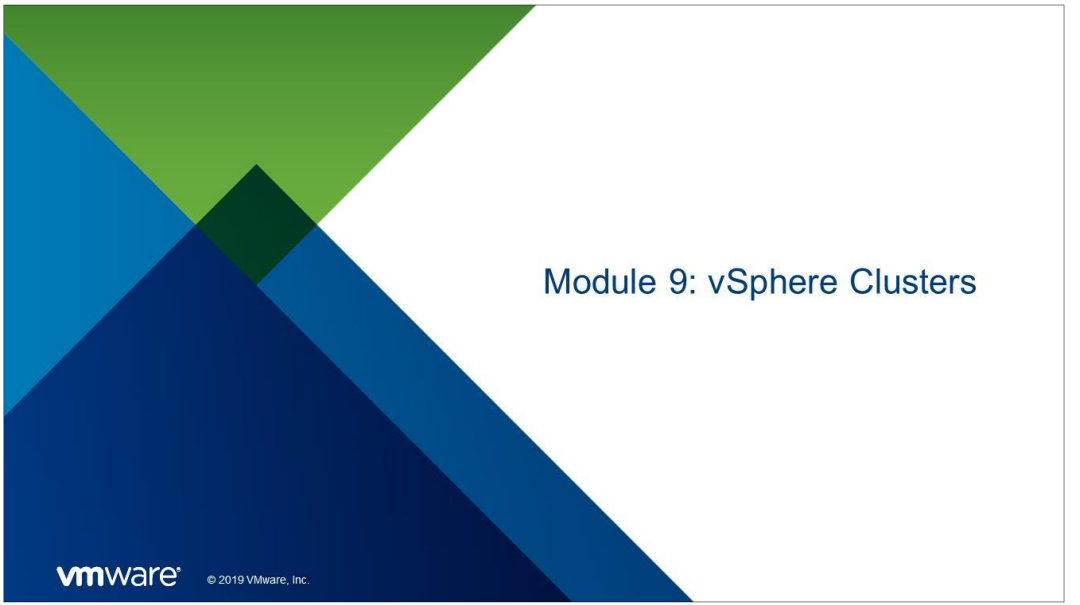

https://t.me/learningnets

9-2 Importance

Most organizations rely on computer-based services like email, databases, and web-based applications. The failure of any of these services can mean lost productivity and revenue.

By understanding and using vSphere HA, you can configure highly available, computer-based services, which are important for an organization to remain competitive in contemporary business environments. And by developing skills in using vSphere DRS, you can improve service levels by guaranteeing appropriate resources to virtual machines.

https://t.me/learningnets

9-3 Module Lessons

1. vSphere Clusters Overview
2. vSphere DRS
3. Introduction to vSphere HA
4. vSphere HA Architecture
5. Configuring vSphere HA
6. Introduction to vSphere Fault Tolerance

https://t.me/learningnets

9-4 Virtual Beans: vSphere Clusters

Virtual Beans has the following requirements for their data center:

- Infrastructure must be highly available:
  - Business-critical applications: 99.99 percent available (downtime per year of 52.56 minutes)
  - Nonbusiness-critical applications: 99 percent available (downtime per year of 3.65 days)
- Infrastructure must be scalable:
  - Virtual Beans expects huge growth over the next three years, so the virtual infrastructure must be easy to scale.
- Applications must perform well:
  - Applications must have enough resources to meet performance levels as defined in the service-level agreement.

As a Virtual Beans administrator, you create a vSphere cluster architecture for the data center that is highly available, scalable, and high-performing.

https://t.me/learningnets

https://t.me/learningnets

9-6 Learner Objectives

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

- Describe the benefits of vSphere clusters
- Create a vSphere cluster
- View information about a vSphere cluster

https://t.me/learningnets

9-7 About vSphere Clusters

A cluster is used in vSphere to share physical resources between a group of ESXi hosts. vCenter Server manages cluster resources as a single pool of resources.
You can create one or more clusters based on the purpose each cluster must fulfill, for example:

- Management
- Production
- Compute

A cluster can contain up to 64 ESXi hosts.

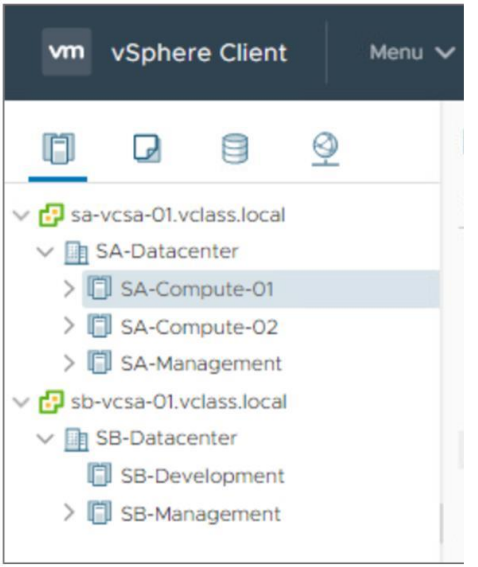

https://t.me/learningnets

9-8 Creating a vSphere Cluster and Enabling Cluster Features

When you create a cluster, you can enable one or more cluster features:

- vSphere DRS
- vSphere HA
- vSAN

You can also manage image setup and updates on all hosts collectively.

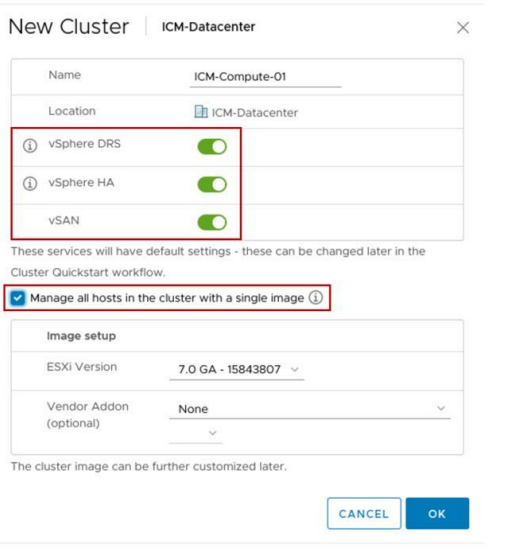

You can enable the following services in a vSphere cluster:

- vSphere HA, for high availability
- vSphere DRS, for VM placement and load balancing
- vSAN, for shared storage

You can also manage host updates using images. With vSphere Lifecycle Manager, you can update all hosts in the cluster collectively, using a specified ESXi image.

https://t.me/learningnets

9-9 Configuring the Cluster Using Quickstart

After you create a cluster, you can use the Cluster Quickstart workflow to configure the cluster. With Cluster Quickstart, you follow a step-by-step configuration wizard that makes it easy to expand the cluster as needed.

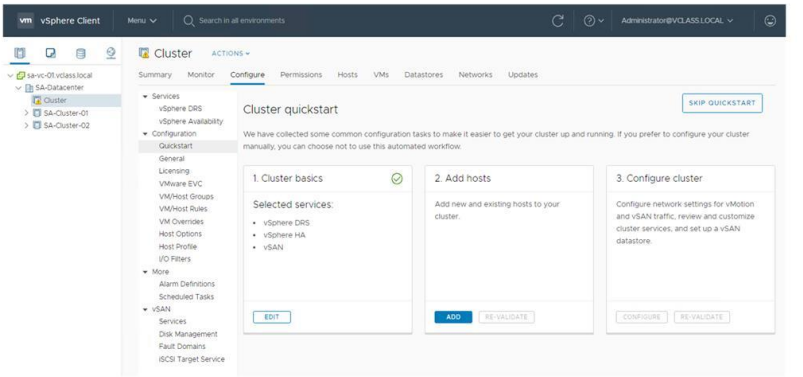

The Cluster Quickstart workflow guides you through the deployment process for clusters. It covers every aspect of the initial configuration, such as host, network, and vSphere settings. With Cluster Quickstart, you can also add additional hosts to a cluster as part of the ongoing expansion of clusters.

Cluster Quickstart reduces the time it takes to configure a cluster.

The workflow includes the following tasks:

- Setting up services such as vSphere HA and vSAN
- Verifying hardware and software compatibility
- Deploying Virtual Distributed Switches
- Configuring network settings for vSphere vMotion and vSAN

https://t.me/learningnets

- Creating a vSAN stretched cluster or vSAN fault domains
- Ensuring consistent NTP configuration across the cluster

The Cluster quickstart page provides workflow cards for configuring your new cluster:

- Cluster basics: Lists the services that you have already enabled and provides an option for editing the cluster's name.
- Add hosts: Adds ESXi hosts to the cluster. These hosts must already be present in the inventory. After hosts are added, the workflow shows the total number of hosts that are present in the cluster and provides health check validation for those hosts. At the start, this workflow is empty.
- Configure cluster: Informs you about what can be automatically configured, provides details on configuration mismatch, and reports cluster health results through the vSAN health service even after the cluster is configured.

For more information about creating clusters, see *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vcenterhost.doc/GUID-3B5AF2B1-C534-4426-B97A-D14019A8010F.html.

https://t.me/learningnets

9-10 Configuring the Cluster Manually

Alternatively, you can use the **Configure** tab to manually configure a cluster's settings.

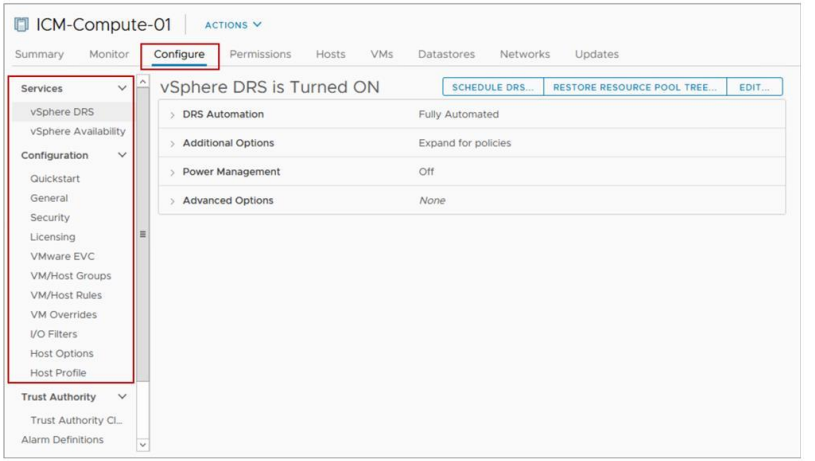

https://t.me/learningnets

9-11 Adding a Host to a Cluster

To add a host to a cluster, drag the host onto the cluster object in the inventory.

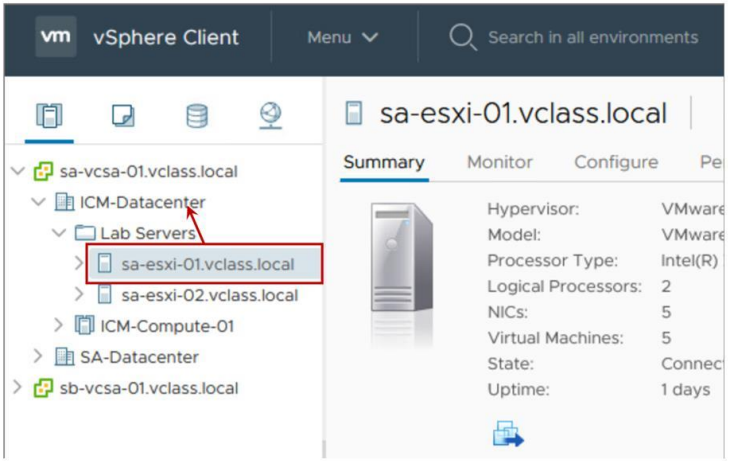

https://t.me/learningnets

9-12 Viewing Cluster Summary Information

For a quick view of your cluster configuration, the **Summary** tab provides general information about a cluster's resources and its consumers.

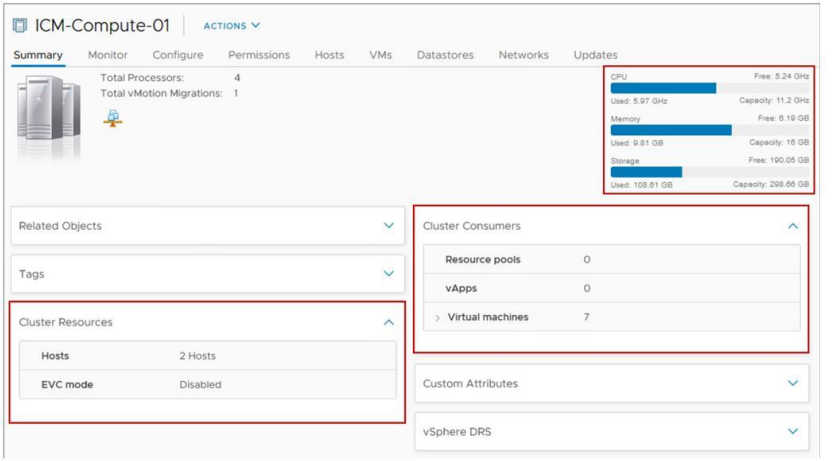

https://t.me/learningnets

9-13 Monitoring Cluster Resources

You can view a report of total cluster CPU, memory, memory overhead, storage capacity, the capacity reserved by VMs, and how much capacity remains available.

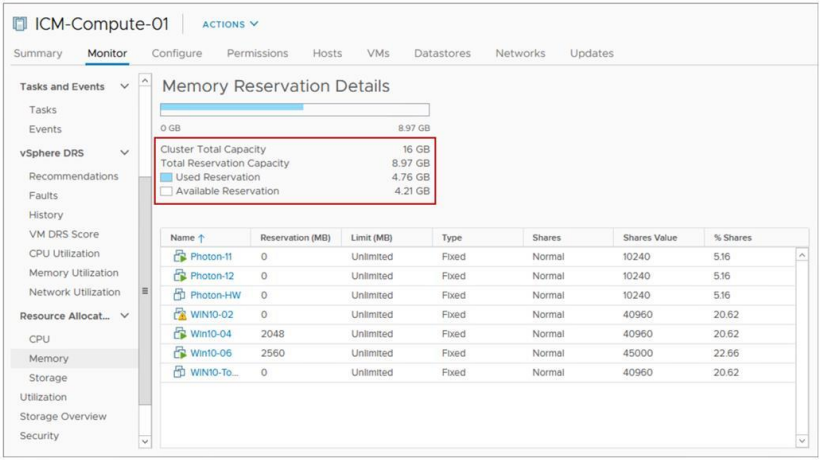

vCenter Server uses vSphere HA admission control to ensure that sufficient resources are available in a cluster to provide failover protection and to ensure that VM resource reservations are respected.

https://t.me/learningnets

9-14 Review of Learner Objectives

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

- Describe the benefits of vSphere clusters
- Create a vSphere cluster
- View information about a vSphere cluster

https://t.me/learningnets

https://t.me/learningnets

9-16 Learner Objectives

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

- Describe the functions of a vSphere DRS cluster
- Explain how vSphere DRS determines VM placement on hosts in the cluster
- Enable vSphere DRS in a cluster
- Monitor a vSphere DRS cluster

https://t.me/learningnets

9-17 About vSphere DRS

vSphere DRS is a cluster feature that helps improve resource allocation across all hosts in a cluster. It aggregates computing capacity across a collection of servers into logical resource pools. vSphere DRS is used in the following situations:

- Initial placement of a VM when it is powered on
- Load balancing
- Migrating VMs when an ESXi host is placed in maintenance mode

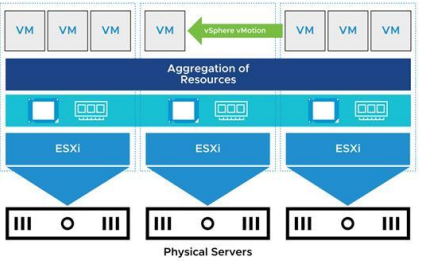

When you power on a VM in the cluster for the first time, vSphere DRS either places the VM on a particular host or makes a recommendation.

DRS attempts to improve resource use across the cluster by performing automatic migrations of VMs (vSphere vMotion) or by providing a recommendation for VM migrations.

Before an ESXi host enters maintenance mode, VMs running on the host must be migrated to another host (either manually or automatically by DRS) or shut down.

https://t.me/learningnets

9-18 vSphere DRS: VM Focused

vSphere DRS is VM focused:

- While the VM is powered on, vSphere DRS operates on an individual VM basis by ensuring that each VM's resource requirements are met.
- vSphere DRS calculates a score for each VM and gives recommendations (or migrates VMs) for meeting VM's resource requirements.

The DRS algorithm recommends where individual VMs should be moved for maximum efficiency. If the cluster is in fully automated mode, DRS executes the recommendations and migrates VMs to their optimal host based on the underlying calculations performed every minute.

https://t.me/learningnets

9-19 About the VM DRS Score

The VM DRS score is a metric that tracks a VM's execution efficiency on a given host. Execution efficiency is the frequency that the VM is reported as having its resources requirements met:

- Values closer to 0% indicate severe resource contention.
- Values closer to 100% indicate mild to no resource contention.

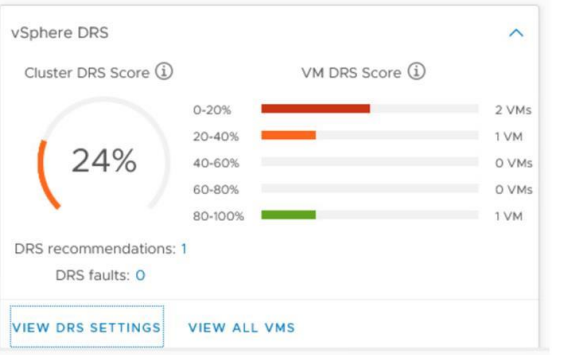

A VM DRS score is computed from an individual VM's CPU, memory, and network metrics. DRS uses these metrics to gauge the goodness or wellness of the VM.

In vSphere 7, the DRS algorithm runs every minute. The Cluster DRS Score is the last result of DRS running and is filed into one of five buckets. These buckets are simply 20 percent ranges: 0-20, 20-40, 40-60, 60-80 and 80-100 percent over the sample period.

https://t.me/learningnets

9-20 VM DRS Score List

The cluster's **Monitor** tab lists the VM DRS Score and more detailed metrics for all the VMs in the cluster.

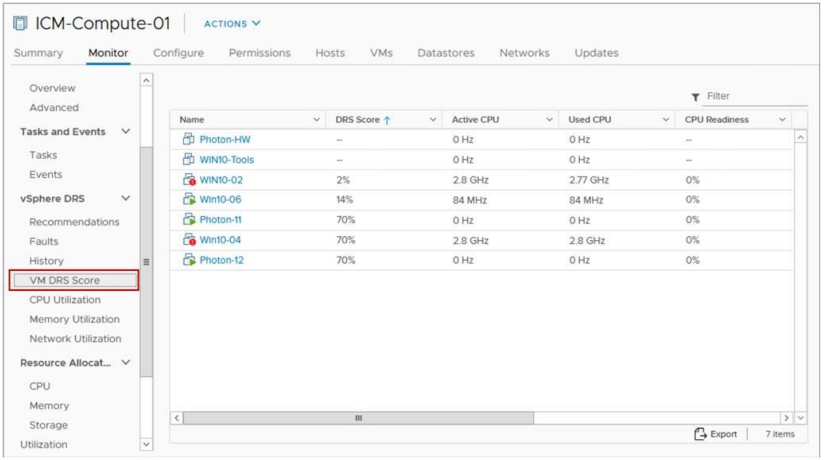

The VM DRS Score page shows the following values for VMs that are powered on:

- DRS Score
- Active CPU
- Used CPU
- CPU Readiness
- Granted Memory
- Swapped Memory
- Ballooned Memory

https://t.me/learningnets

9-21 Viewing VM DRS Scores Using Performance Charts (1)

The advanced performance chart for a cluster object provides the DRS Score counter.

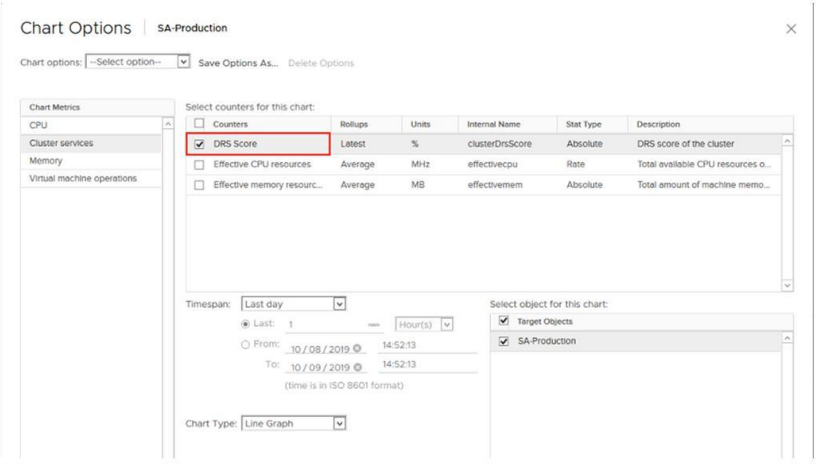

https://t.me/learningnets

9-22 Viewing VM DRS Scores Using Performance Charts (2)

The DRS Score counter displays the DRS scores for VMs in the cluster over the selected time period.

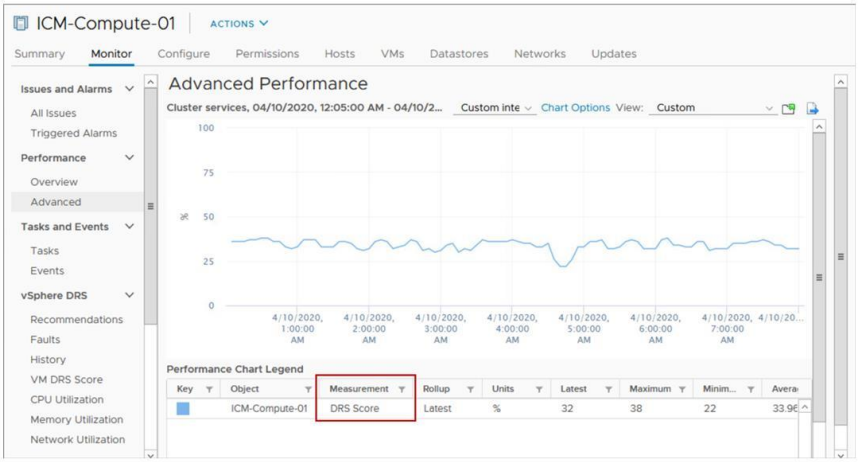

https://t.me/learningnets

9-23 Viewing vSphere DRS Settings

When you click **VIEW DRS SETTINGS**, the main vSphere DRS parameters and their current values are shown. vSphere DRS settings include:

- Automation level
- Migration threshold

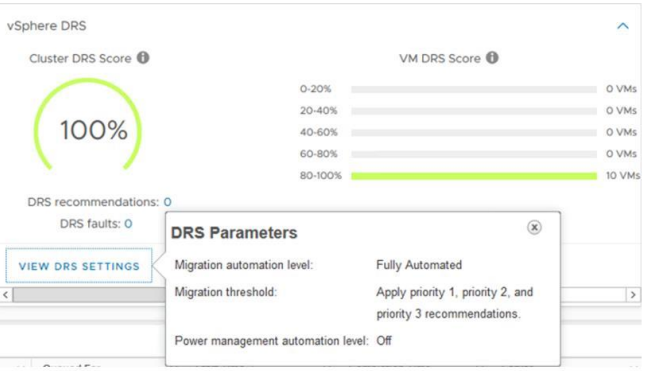

To view the vSphere DRS pane, go to the cluster's **Summary** tab.

https://t.me/learningnets

9-24 vSphere DRS Settings: Automation Level

You can configure the automation level for the initial placement of VMs and for dynamic balancing while VMs are running.

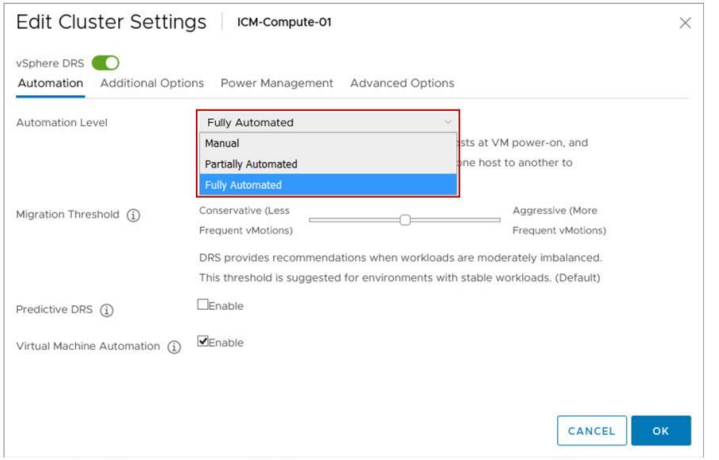

The automation level determines whether vSphere DRS makes migration recommendations or automatically places VMs on hosts. vSphere DRS makes placement decisions when a VM powers on and when VMs must be rebalanced across hosts in the cluster.

The following automation levels are available:

- **Manual**: When you power on a VM, vSphere DRS displays a list of recommended hosts on which to place the VM. When the cluster becomes imbalanced, vSphere DRS displays recommendations for VM migration.
- **Partially automated**: When you power on a VM, vSphere DRS places it on the best-suited host. When the cluster becomes imbalanced, vSphere DRS displays recommendations for manual VM migration.
- **Fully automated**: When you power on a VM, vSphere DRS places it on the best-suited host. When the cluster becomes imbalanced, vSphere DRS migrates VMs from overused hosts to underused hosts to ensure balanced use of cluster resources.

https://t.me/learningnets

9-25 vSphere DRS Settings: Migration Threshold

The migration threshold determines how aggressively vSphere DRS selects to migrate VMs.

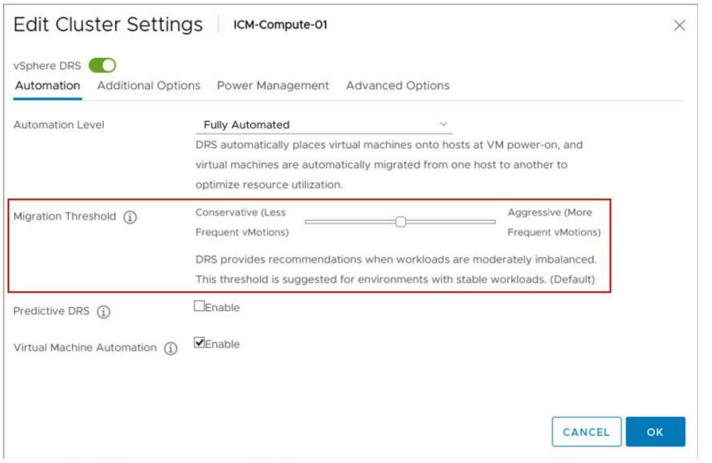

The following migration threshold settings are available:

- Level 1 (Conservative): Applies only priority 1 recommendations. vCenter Server applies only recommendations that must be taken to satisfy cluster constraints, such as affinity rules and host maintenance.
- Level 2: Apply priority 1 and priority 2 recommendations. vCenter Server applies recommendations that promise a significant improvement to the cluster's load balance.
- Level 3 (default): Apply priority 1, priority 2, and priority 3 recommendations. vCenter Server applies recommendations that promise at least good improvement to the cluster's load balance.

https://t.me/learningnets

- Level 4: Apply priority 1, priority 2, priority 3, and priority 4 recommendations. vCenter Server applies recommendations that promise even a moderate improvement to the cluster's load balance.
- Level 5 (Aggressive): Apply all recommendations. vCenter Server applies recommendations that promise even a slight improvement to the cluster's load balance.

https://t.me/learningnets

9-26 vSphere DRS Settings: Predictive DRS

vSphere DRS and vRealize Operations Manager combine data to predict future demand and determine when and where high resource utilization occurs.

To make predictive decisions, the vSphere DRS data collector retrieves the following data:

- Resource usage statistics from ESXi hosts
- Predicted usage statistics from the vRealize Operations Manager server

Predicted usage statistics always take precedence over current usage statistics.

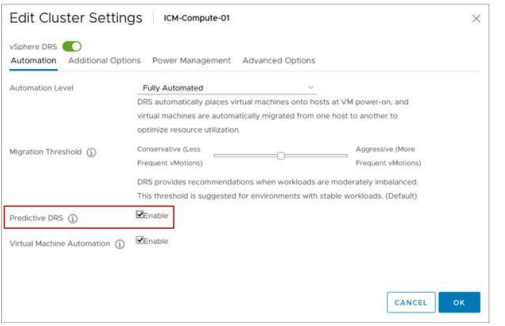

https://t.me/learningnets

9-27 vSphere DRS Settings: VM Swap File Location

By default, swap files for a VM are on a datastore in the folder containing the other VM files.
For all VMs in the cluster, you can place VM swap files on an alternative datastore.
If vSphere DRS is enabled, you should place the VM swap file in the VM's directory.

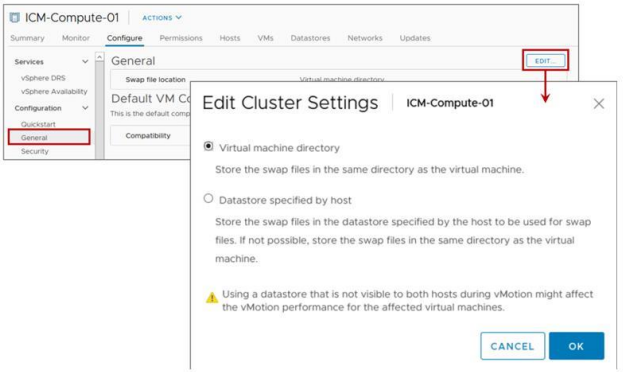

A VM's files can be on a VMFS datastore, an NFS datastore, a vSAN datastore, or a vSphere Virtual Volumes datastore. On a vSAN datastore or a vSphere Virtual Volumes datastore, the swap file is created as a separate vSAN or vSphere Virtual Volumes object.

A swap file is created by the ESXi host when a VM is powered on. If this file cannot be created, the VM cannot power on. Instead of accepting the default, you can also use the following options:

- Use per-VM configuration options to change the datastore to another shared storage location.
- Use host-local swap, which allows you to specify a datastore stored locally on the host. You can swap at a per-host level. However, it can lead to a slight degradation in performance for vSphere vMotion because pages swapped to a local swap file on the source host must be transferred across the network to the destination host. Currently, vSAN and vSphere Virtual Volumes datastores cannot be specified for host-local swap.

https://t.me/learningnets

9-28 vSphere DRS Settings: VM Affinity

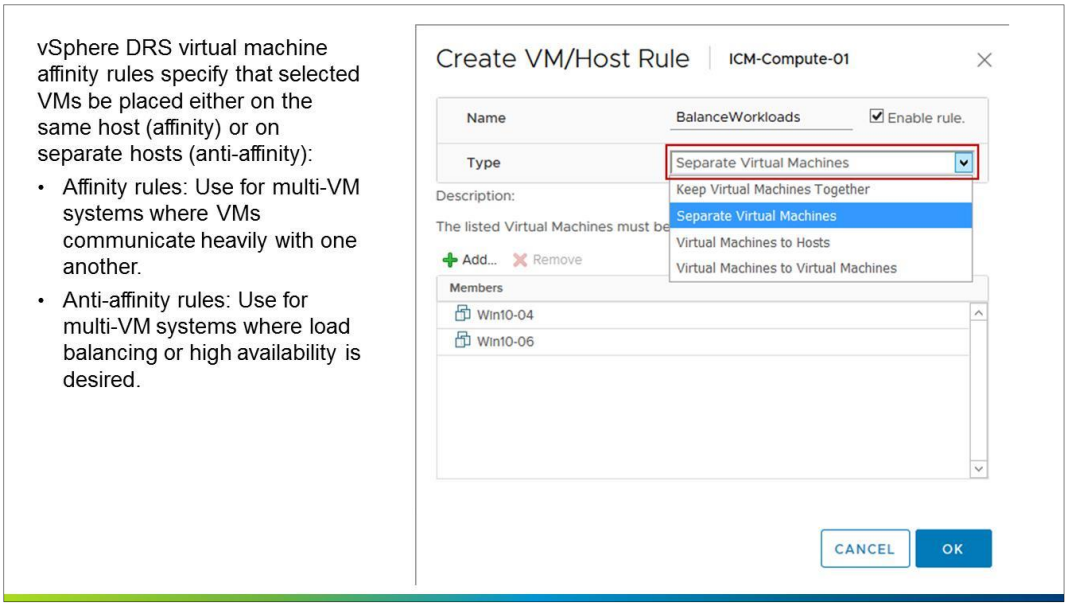

After a vSphere DRS cluster is created, you can edit its properties to create rules that specify affinity. The following types of rules can be created:

- Affinity rules: vSphere DRS keeps certain VMs together on the same host (for example, for performance reasons).
- Anti-affinity rules: vSphere DRS ensures that certain VMs are not together (for example, for availability reasons).

If two rules conflict, you are prevented from enabling both.

When you add or edit a rule, and the cluster is immediately in violation of that rule, the cluster continues to operate and tries to correct the violation.

For vSphere DRS clusters that have a default automation level of manual or partially automated, migration recommendations are based on both rule fulfillment and load balancing.

https://t.me/learningnets

9-29 vSphere DRS Settings: DRS Groups

VM groups and host groups are used in defining VM-Host affinity rules.

The VM-Host affinity rule specifies whether VMs can or cannot be run on a host.
Types of groups:

- VM group: One or more VMs
- Host group: One or more ESXi hosts

A VM can belong to multiple VM groups.

A host can belong to multiple host groups.

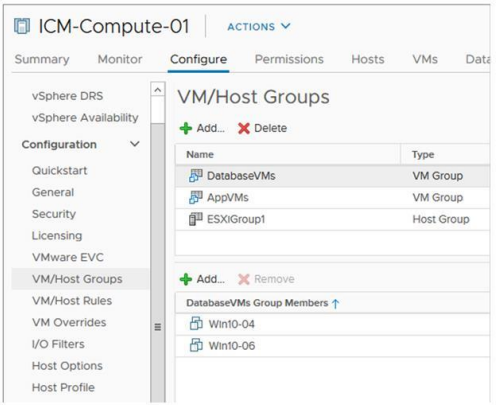

For ease of administration, virtual machines can be placed in VM or host groups. You can create one or more VM groups in a vSphere DRS cluster, each consisting of one or more VMs. A host group consists of one or more ESXi hosts.

The main use of VM groups and host groups is to help in defining the VM-Host affinity rules.

https://t.me/learningnets

9-30 vSphere DRS Settings: VM-Host Affinity Rules

A VM-Host affinity rule:

- Defines an affinity (or anti-affinity) relationship between a VM group and a host group
- Is either a required rule or a preferential rule

Rule options:

- Must run on hosts in group
- Should run on hosts in group
- Must not run on hosts in group
- Should not run on hosts in group

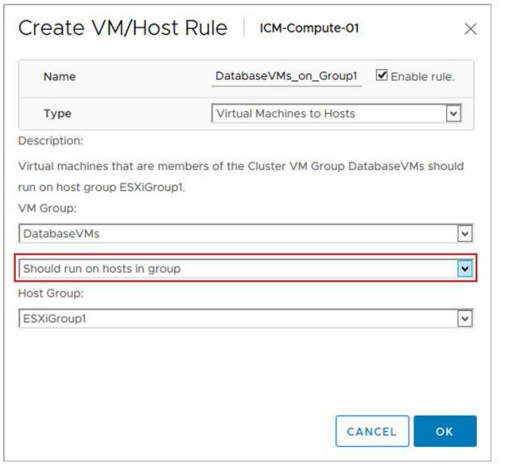

A VM-Host affinity or anti-affinity rule specifies whether the members of a selected VM group can run on the members of a specific host group.

Unlike an affinity rule for VMs, which specifies affinity (or anti-affinity) between individual VMs, a VM-Host affinity rule specifies an affinity relationship between a group of VMs and a group of hosts.

Because VM-Host affinity rules are cluster-based, the VMs and hosts that are included in a rule must all reside in the same cluster. If a VM is removed from the cluster, the VM loses its membership from all VM groups, even if it is later returned to the cluster.

https://t.me/learningnets

9-31 VM-Host Affinity Preferential Rules

A preferential rule is softly enforced and can be violated if necessary.
Example: Separate VMs on different blade systems for improved performance.

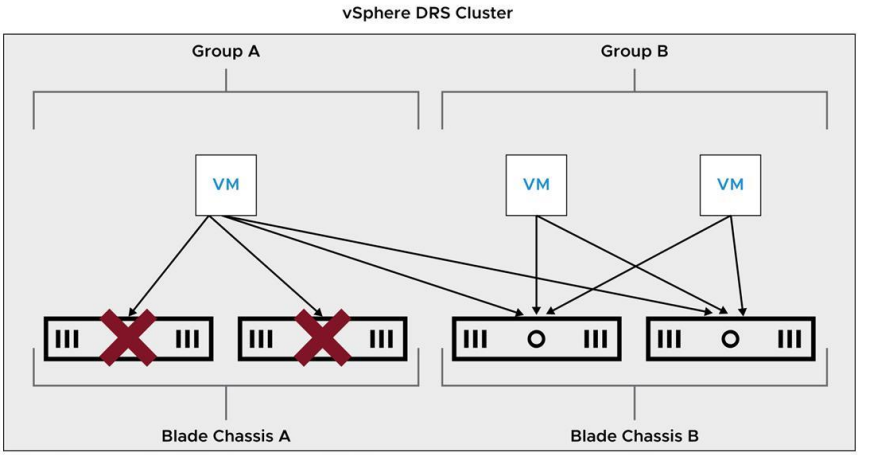

Preferential rules can be violated to allow the proper functioning of vSphere DRS, vSphere HA, and VMware vSphere DPM.

On the slide, Group A and Group B are VM groups. Blade Chassis A and Blade Chassis B are host groups. The goal is to force the VMs in Group A to run on the hosts in Blade Chassis A and to force the VMs in Group B to run on the hosts in Blade Chassis B. If the hosts fail, vSphere HA restarts the VMs on the other hosts in the cluster. If the hosts are put into maintenance mode or become overused, vSphere DRS moves the VMs to the other hosts in the cluster.

https://t.me/learningnets

9-32 VM-Host Affinity Required Rules

A required rule is strictly enforced and can never be violated.
Example: Enforce host-based ISV licensing.

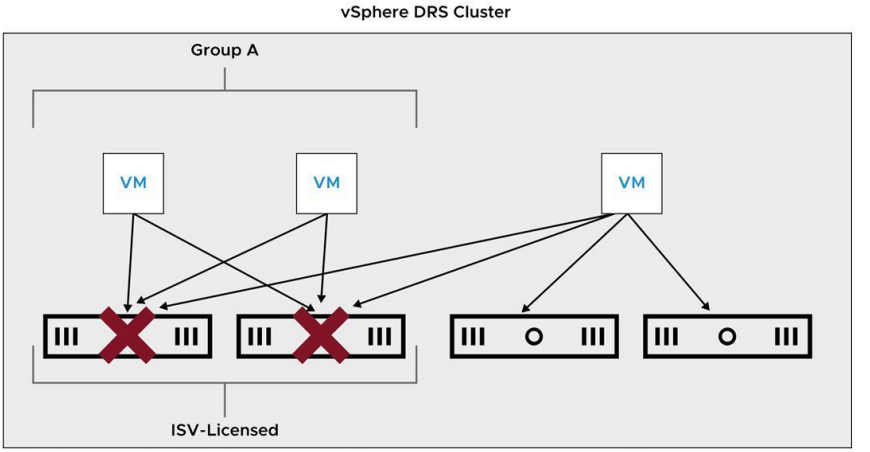

A VM-Host affinity rule that is required, instead of preferential, can be used when the software running in your VMs has licensing restrictions. You can enforce this rule when the software running in your VMs has licensing restrictions. You can place such VMs in a VM group. Then you can create a rule that requires the VMs to run on a host group, which contains hosts with the required licenses.

When you create a VM-Host affinity rule that is based on the licensing or hardware requirements of the software running in your VMs, you are responsible for ensuring that the groups are properly set up. The rule does not monitor the software running in the VMs. Nor does it know which third-party licenses are in place on which ESXi hosts.

On the slide, Group A is a VM group. You can force Group A to run on hosts in the ISV-Licensed group to ensure that the VMs in Group A run on hosts that have the required licenses. But if the hosts in the ISV-Licensed group fail, vSphere HA cannot restart the VMs in Group A on hosts that are not in the group. If the hosts in the ISV-Licensed group are put into maintenance mode or become overused, vSphere DRS cannot move the VMs in Group A to hosts that are not in the group.

https://t.me/learningnets

9-33 vSphere DRS Settings: VM-Level Automation

You can customize the automation level for individual VMs in a cluster to override the automation level set on the entire cluster.

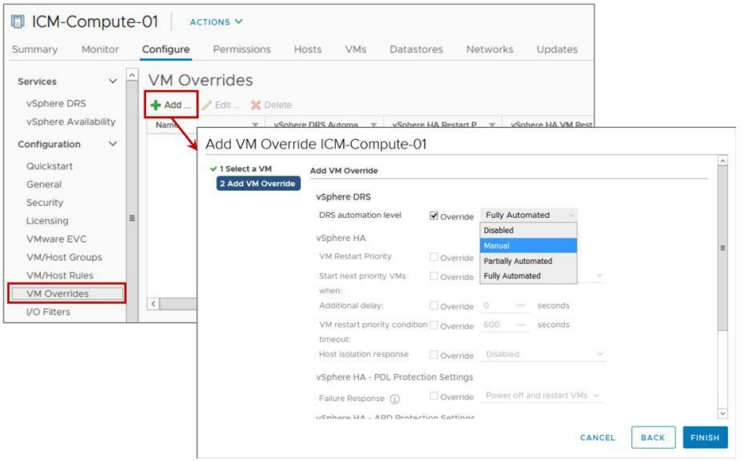

By setting the automation level for individual VMs, you can fine-tune automation to suit your needs. For example, you might have a VM that is especially critical to your business. You want more control over its placement so you set its automation level to **Manual**.

If a VM's automation level is set to disabled, vCenter Server does not migrate that VM or provide migration recommendations for it.

As a best practice, enable automation. Select the automation level based on your environment and level of comfort.

For example, if you are new to vSphere DRS clusters, you might select **Partially Automated** because you want control over the movement of VMs.

When you are comfortable with what vSphere DRS does and how it works, you might set the automation level to **Fully Automated**.

You can set the automation level to **Manual** on VMs over which you want more control, such as your business-critical VMs.

https://t.me/learningnets

9-34 vSphere DRS Cluster Requirements

ESXi hosts that are added to a vSphere DRS cluster must meet certain requirements to use cluster features successfully:

- To use vSphere DRS for load balancing, the hosts in your cluster must be part of a vSphere vMotion network:
  - If the hosts are not part of a vSphere vMotion network, vSphere DRS can still make initial placement recommendations.
  - vSphere DRS works best if the VMs meet vSphere vMotion requirements.
- Configure all managed hosts to use shared storage.

You can create vSphere DRS clusters, or you can enable vSphere DRS for existing vSphere HA or vSAN clusters.

https://t.me/learningnets

9-35 Viewing vSphere DRS Cluster Resource Utilization

From the cluster's **Monitor** tab, you can view CPU, memory, and network utilization per host.

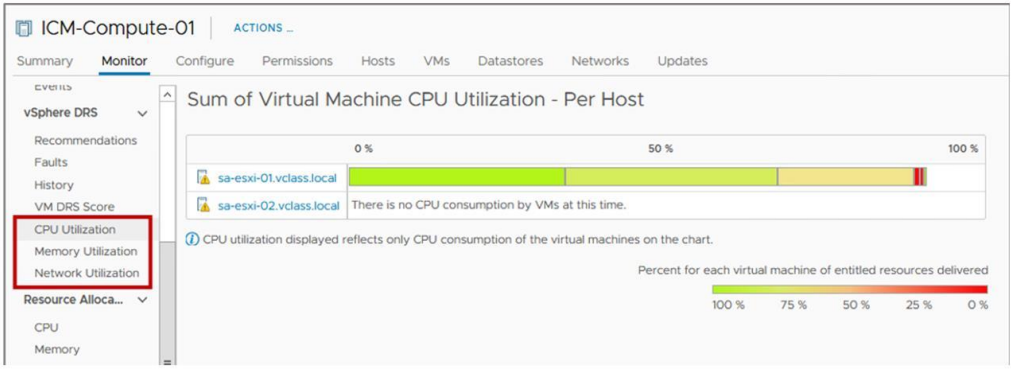

The CPU Utilization and Memory Utilization charts show all the hosts in the cluster and how their CPU and memory resources are allocated to each VM.

- For CPU usage, the VM information is represented by a colored box. If you point to the colored box, the VM's CPU usage information appears. If the VM is receiving the resources that it is entitled to, the box is green. Green means that 100 percent of the VM's entitled resources are delivered. If the box is not green (for example, entitled resources are 80 percent or less) for an extended time, you might want to investigate what is causing this shortfall (for example, unapplied recommendations).
- For memory usage, the VM boxes are not color-coded because the relationship between consumed memory and entitlement is often not easily categorized.

In the Network Utilization chart, the displayed network data reflects all traffic across physical network interfaces on the host.

https://t.me/learningnets

9-36 Viewing vSphere DRS Recommendations

The DRS Recommendations pane displays information about the vSphere DRS recommendations made for the cluster.

You can also view the faults that occurred when the recommendations were applied and the history of vSphere DRS actions.

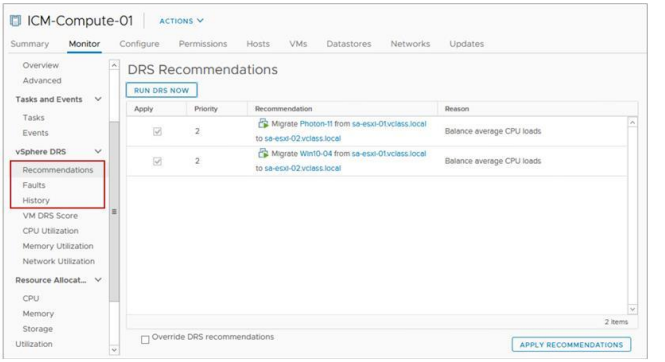

In the DRS Recommendations pane, you can see the current set of recommendations that are generated for optimizing resource use in the cluster through either migrations or power management. Only manual recommendations awaiting user confirmation appear in the list.

To refresh the recommendations, click **RUN DRS NOW**.

To apply all recommendations, click **APPLY RECOMMENDATIONS**.

To apply a subset of the recommendations, select the **Override DRS recommendations** check box. Select the check box next to each desired recommendation and click **APPLY RECOMMENDATIONS**.

https://t.me/learningnets

9-37 Maintenance Mode and Standby Mode

Maintenance mode:

- Removes a host's resources from a cluster, making those resources unavailable for use
- Is often used to service a host in a cluster

To place a host in maintenance mode:

- All running VMs on the host must be migrated to another host, shut down or suspended.
- When DRS is in fully automated mode, powered-on VMs are automatically migrated from a host that is placed in maintenance mode.

Standby mode:

- Is used by vSphere DPM to optimize power usage. When a host is placed in standby mode, it is powered off.

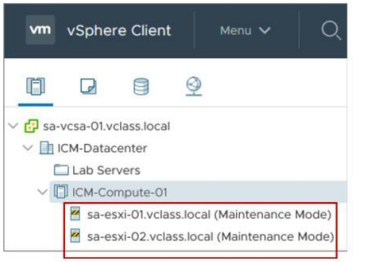

A host enters or leaves maintenance mode as the result of a user request. While in maintenance mode, the host does not allow you to deploy or power on a VM.

VMs that are running on a host entering maintenance mode must be shut down or migrated to another host, either manually (by a user) or automatically (by vSphere DRS). The host continues to run the Enter Maintenance Mode task until all VMs are powered down or moved away.

When no more running VMs are on the host, the host's icon indicates that it has entered maintenance mode. The host's **Summary** tab indicates the new state.

Place a host in maintenance mode before servicing the host, for example, when installing more memory or removing a host from a cluster.

You can place a host in standby mode manually. However, the next time that vSphere DRS runs, it might undo your change or recommend that you undo the changes. If you want a host to remain powered off, place it in maintenance mode and turn it off.

https://t.me/learningnets

9-38 Removing a Host from the vSphere DRS Cluster

To remove a host from a cluster:

1. Place the host in maintenance mode.
2. Drag the host to a different inventory location, for example, the data center or another cluster.

The resources available for the cluster decrease.

When a host is put into maintenance mode, all its running VMs must be shut down, suspended, or migrated to other hosts by using vSphere vMotion. VMs with disks on local storage must be powered off, suspended, or migrated to another host and datastore.

When you remove the host from the cluster, the VMs that are currently associated with the host are also removed from the cluster. If the cluster still has enough resources to satisfy the reservations of all VMs in the cluster, the cluster adjusts resource allocation to reflect the reduced amount of resources.

https://t.me/learningnets

9-39 vSphere DRS and Dynamic DirectPath I/O

Dynamic DirectPath I/O improves the vSphere DirectPath I/O functionality by adding a layer of abstraction between a VM and the physical PCI device:

A pool of PCI devices that are available in the cluster can be assigned to the VM.
vSphere DRS is aware of Dynamic DirectPath I/O devices:

- When the VM is powered on, vSphere DRS places the VM on any ESXi host that provides the assigned PCI device.
- vSphere DRS takes action only at VM power on and does not perform any load-balancing actions.

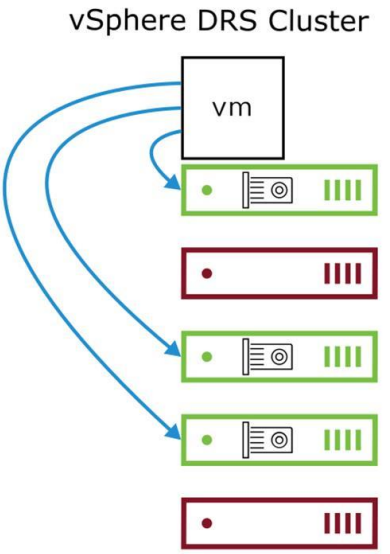

Dynamic DirectPath I/O is useful on hosts that have PCI passthrough devices and for virtualized devices that require a directly assigned hardware device to back it.

Dynamic DirectPath I/O is also called assignable hardware. The following devices can use assignable hardware:

- PCI passthrough devices
- Shared PCI passthrough devices, for example, NVIDIA GRID vGPUs

Full memory reservation is applied to the VMs with pass-through devices.

https://t.me/learningnets

9-40 Adding a Dynamic DirectPath I/O Device to a VM

You can add Dynamic DirectPath I/O devices to a VM by editing the VM's settings.

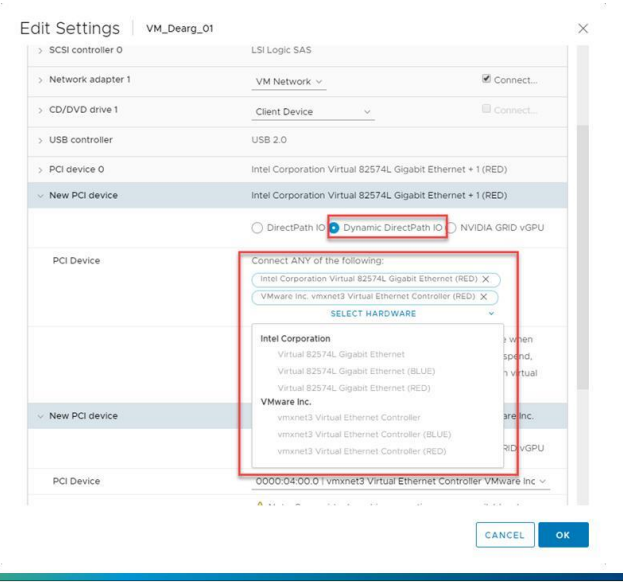

For New PCI device, click **Dynamic DirectPath IO**. Clicking **SELECT HARDWARE** displays a list of devices that can be attached to the VM. You can select one or more devices from the list. In the image, the VM can use either an Intel NIC with the RED hardware label or vmxnet3 NIC with the RED hardware label.

https://t.me/learningnets

9-41 Lab 25: Implementing vSphere DRS Clusters

Implement a vSphere DRS cluster and verify proper functionality:

1. Create a Cluster That Is Enabled for vSphere DRS
2. Modify vSphere DRS Settings
3. Add ESXi Hosts to the Cluster
4. Verify vSphere vMotion Configuration on the ESXi Hosts
5. Create a Load Imbalance
6. Verify Proper vSphere DRS Cluster Functionality

https://t.me/learningnets

Review of Learner Objectives

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

- Describe the functions of a vSphere DRS cluster
- Explain how vSphere DRS determines VM placement on hosts in the cluster
- Enable vSphere DRS in a cluster
- Monitor a vSphere DRS cluster

https://t.me/learningnets

https://t.me/learningnets

9-44 Learner Objectives

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

- Identify options for configuring a highly available vSphere environment
- Describe how vSphere HA responds when an ESXi host, a virtual machine, or an application fails

https://t.me/learningnets

9-45 Protection at Every Level

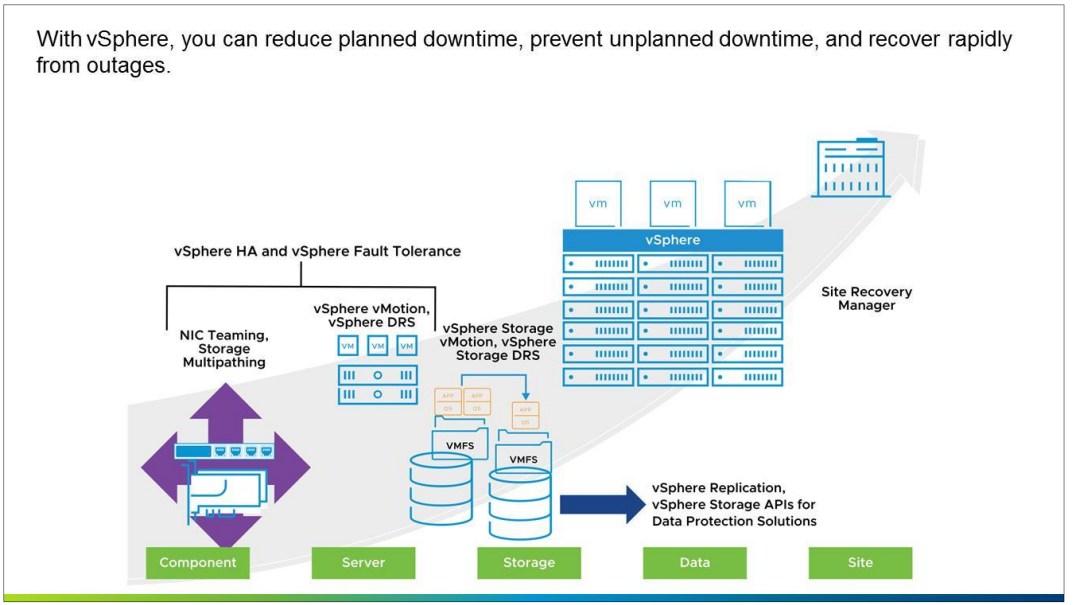

Whether planned or unplanned, downtime brings with it considerable costs. However, solutions to ensure higher levels of availability are traditionally costly, hard to implement, and difficult to manage.

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

- Provide higher availability independent of hardware, operating system, and applications.
- Reduce planned downtime for common maintenance operations.
- Provide automatic recovery in cases of failure.

vSphere HA provides a base level of protection for your VMs by restarting VMs if a host fails. vSphere Fault Tolerance provides a higher level of availability, allowing users to protect any VM

https://t.me/learningnets

from a host failure with no loss of data, transactions, or connections. vSphere Fault Tolerance provides continuous availability by ensuring that the states of the primary and secondary VMs are identical at any point in the instruction execution of the VM.

vSphere vMotion and vSphere Storage vMotion keep VMs available during a planned outage, for example, when hosts or storage must be taken offline for maintenance. System recovery from unexpected storage failures is simple, quick, and reliable with the encapsulation property of VMs. You can use vSphere Storage vMotion to support planned storage outages resulting from upgrades to storage arrays to newer firmware or technology and VMFS upgrades.

With vSphere Replication, a vSphere platform can protect VMs natively by copying their disk files to another location where they are ready to be recovered.

VM encapsulation is used by third-party backup applications that support file and image-level backups using vSphere Storage APIs - Data Protection. Backup solutions play prominent roles in recovering from deleted files or disks and corrupt or infected guest operating systems or file systems.

With Site Recovery Manager, you can quickly restore your organization's IT infrastructure, shortening the time that you experience a business outage. Site Recovery Manager automates setup, failover, and testing of disaster recovery plans. Site Recovery Manager requires that you install vCenter Server at the protected site and at the recovery site. Site Recovery Manager also requires either host-based replication through vSphere Replication or preconfigured array-based replication between the protected site and the recovery site.

https://t.me/learningnets

9-46 About vSphere HA

vSphere HA provides rapid recovery from outages and cost-effective high availability for applications running in VMs. vSphere HA protects application availability in several ways.

Protects Against	How Does vSphere HA Provide Protection?
ESXi host failure	By restarting the VMs on other hosts within the cluster
VM failure	By restarting the VM when a VMware Tools heartbeat is not received within a set time
Application failure	By restarting the VM when an application heartbeat is not received within a set time
Datastore accessibility failure	By restarting the affected VMs on other hosts that still can access the datastores.
Network isolation	By restarting VMs if their host becomes isolated on the management or vSAN network. This protection is provided even if the network becomes partitioned.

Unlike other clustering solutions, vSphere HA protects all workloads by using the infrastructure itself. After you configure vSphere HA, no actions are required to protect new VMs. All workloads are automatically protected by vSphere HA.

https://t.me/learningnets

9-47 vSphere HA Scenario: ESXi Host Failure

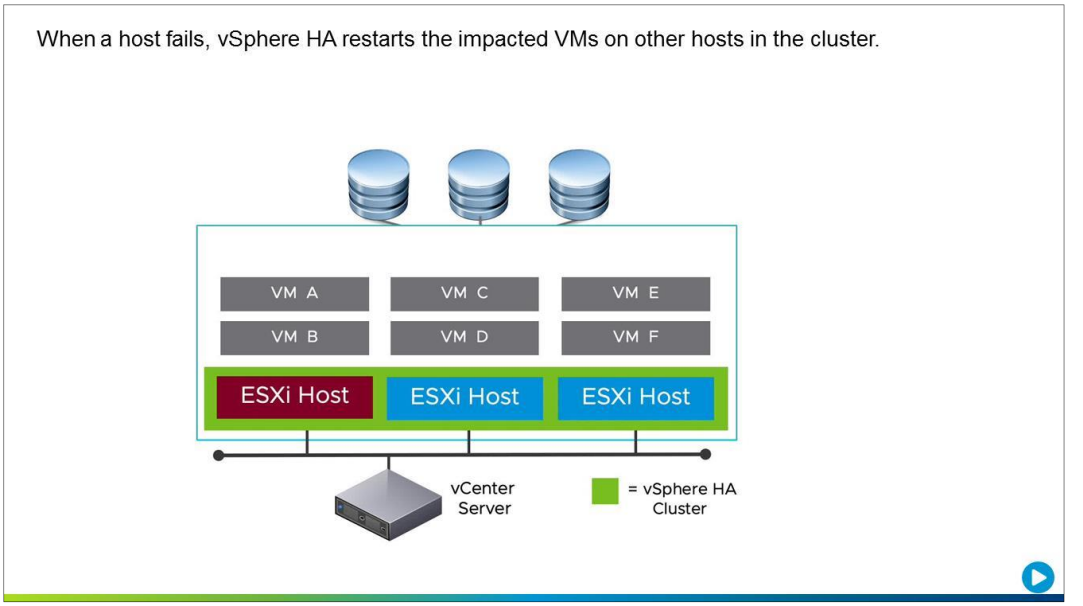

To play the animation, go to https://vmware.bravais.com/s/kvK76lswrsbmJq8kRueo.

vSphere HA can also determine whether a ESXi host is isolated or has failed. If an ESXi host fails, vSphere HA attempts to restart any VMs that were running on the failed host by using one of the remaining hosts in the cluster.

In every cluster, the time to recover depends on how long it takes your guest operating systems and applications to restart when the VM is failed over.

https://t.me/learningnets

9-48 vSphere HA Scenario: Guest Operating System Failure

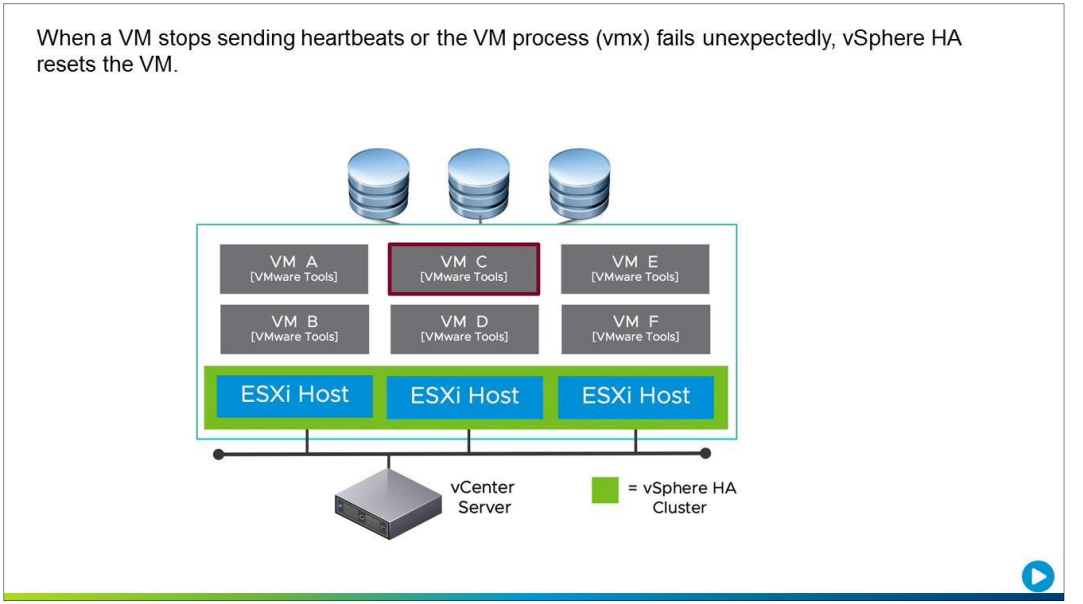

To play the animation, go to https://vmware.bravais.com/s/ikio4LtOkS6fPivIJpR6.

If VM monitoring is enabled, the vSphere HA agent on each individual host monitors VMware Tools in each VM running on the host. When a VM stops sending heartbeats, the guest operating system is reset. The VM stays on the same host.

https://t.me/learningnets

9-49 vSphere HA Scenario: Application Failure

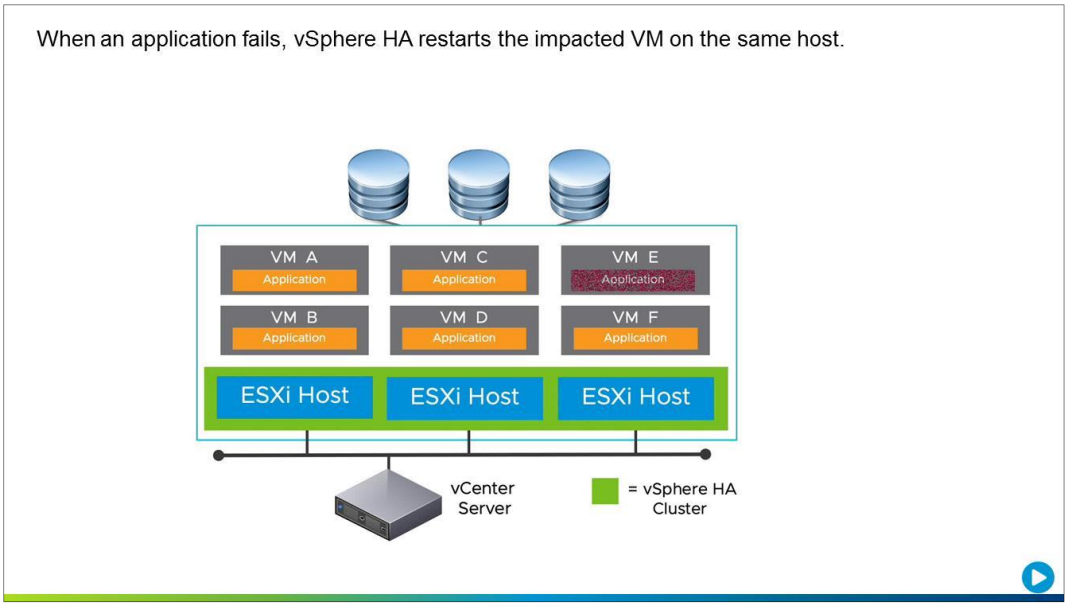

To play the animation, go to https://vmware.bravais.com/s/OgIz03mC2MiGVVPKCxdh.

The agent on each host can optionally monitor heartbeats of applications running in each VM. When an application fails, the VM on which the application was running is restarted on the same host. Application monitoring requires a third-party application monitoring agent designed to work with VM application monitoring.

https://t.me/learningnets

9-50 vSphere HA Scenario: Datastore Accessibility Failures

If VM Component Protection (VMCP) is enabled, vSphere HA can detect datastore accessibility failures and provide automated recovery for affected VMs.
You can determine the response that vSphere HA makes to such a failure, ranging from the creation of event alarms to VM restarts on other hosts:

- All paths down (APD):
  - Recoverable.
  - Represents a transient or unknown accessibility loss.
  - Response can be either **Issue events**, **Power off and restart VMs - Conservative restart policy**, or **Power off and restart VMs - Aggressive restart policy**.
- Permanent device loss (PDL):
  - Unrecoverable loss of accessibility.
  - Occurs when a storage device reports that the datastore is no longer accessible by the host.
  - Response can be either **Issue events** or **Power off and restart VMs**.

Power off and restart VMs - Conservative restart policy: vSphere HA does not attempt to restart the affected VMs unless vSphere HA determines that another host can restart the VMs. The host experiencing the all paths down (APD) communicates with the vSphere HA master host to determine whether the cluster has sufficient capacity to power on the affected VMs. If the master host determines that sufficient capacity is available, the host experiencing the APD stops the VMs so that the VMs can be restarted on a healthy host. If the host experiencing the APD cannot communicate with the master host, no action is taken.

Power off and restart VMs - Aggressive restart policy: vSphere HA stops the affected VMs even if it cannot determine that another host can restart the VMs. The host experiencing the APD attempts to communicate with the master host to determine whether the cluster has sufficient capacity to power on the affected VMs. If the master host is not reachable, sufficient capacity to restart the VMs is unknown. In this scenario, the host takes the risk and stops the VMs so they can be restarted on the remaining healthy hosts. However, if sufficient capacity is not available, vSphere HA might not be able to recover all the affected VMs. This result is common in a network

https://t.me/learningnets

partition scenario where a host cannot communicate with the master host to get a definitive response to the likelihood of a successful recovery.

https://t.me/learningnets

9-51 vSphere HA Scenario: Protecting VMs Against Network Isolation

vSphere HA restarts VMs if their host becomes isolated on the management or vSAN network.

Host network isolation occurs when a host is still running, but it can no longer observe traffic from vSphere HA agents on the management network:

- The host tries to ping the isolation addresses. An isolation address is an IP address or FQDN that can be manually specified (the default is the host's default gateway).
- If pinging fails, the host declares that it is isolated from the network.
- This protection is provided even if the network becomes partitioned.

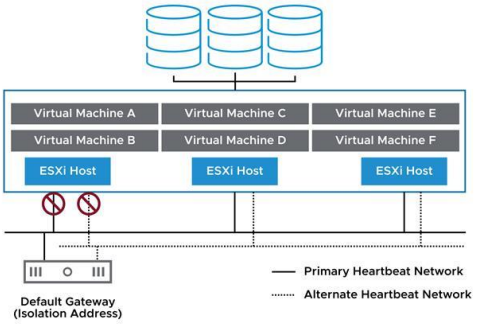

If you ensure that the network infrastructure is sufficiently redundant and that at least one network path is always available, host network isolation is less likely to occur.

https://t.me/learningnets

9-52 Importance of Redundant Heartbeat Networks

Redundant heartbeat networks ensure reliable failure detection and minimize the chance of host-isolation scenarios.
In a vSphere HA cluster, heartbeats have the following characteristics:

- They are sent between the master host and the subordinate hosts.
- They are used to determine whether a master host or a subordinate host has failed.
- They are sent over a heartbeat network.

Redundant heartbeat networking is the best approach for your vSphere HA cluster. When a master host's connection fails, a second connection is still available to send heartbeats to other hosts. If you do not provide redundancy, your failover setup has a single point of failure.

https://t.me/learningnets

9-53 Redundancy Using NIC Teaming

A heartbeat network is implemented in the following ways:

- By using a VMkernel port that is marked for management
- By using a VMkernel port that is marked for vSAN traffic when vSAN is in use

You can use NIC teaming to create a redundant heartbeat network on ESXi hosts.

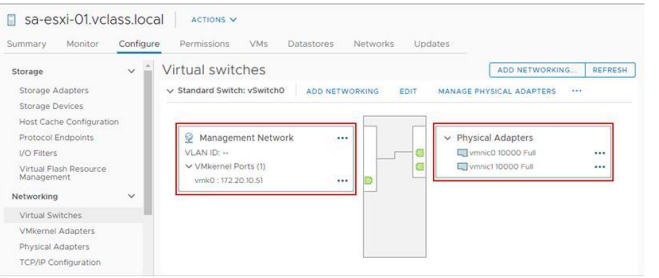

In this example, vmnic0 and vmnic1 form a NIC team in the Management network. The vmk0 VMkernel port is marked for management.

https://t.me/learningnets

9-54 Redundancy Using Additional Networks

You can create redundancy by configuring more heartbeat networks.

On each ESXi host, create a second VMkernel port on a separate virtual switch with its own physical adapter.

Redundant management networking supports the reliable detection of failures and prevents isolation or partition conditions from occurring, because heartbeats can be sent over multiple networks.

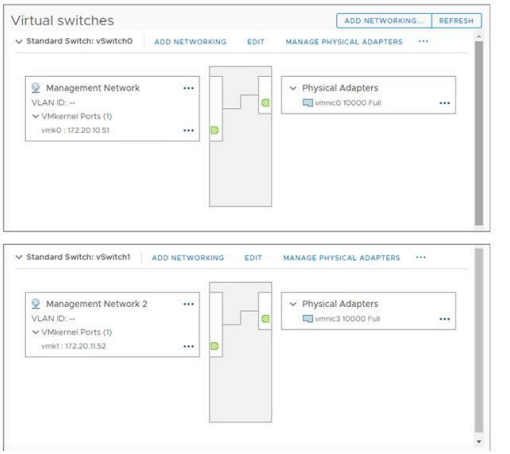

In most implementations, NIC teaming provides sufficient heartbeat redundancy, but as an alternative you can create a second management network connection attached to a separate virtual switch.

The original management network connection is used for network and management purposes. When the second management network connection is created, the master host sends heartbeats over both management network connections. If one path fails, the master host still sends and receives heartbeats over the other path.

https://t.me/learningnets

9-55 Review of Learner Objectives

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

- Identify options for configuring a highly available vSphere environment
- Describe how vSphere HA responds when an ESXi host, a virtual machine, or an application fails

https://t.me/learningnets

https://t.me/learningnets

9-57 Learner Objectives

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

- Identify the heartbeat mechanisms used by vSphere HA
- Describe failure scenarios
- Recognize vSphere HA design considerations

https://t.me/learningnets

9-58 vSphere HA Architecture: Agent Communication

When vSphere HA is enabled in a cluster, the Fault Domain Manager (FDM) service starts on the hosts in the cluster.

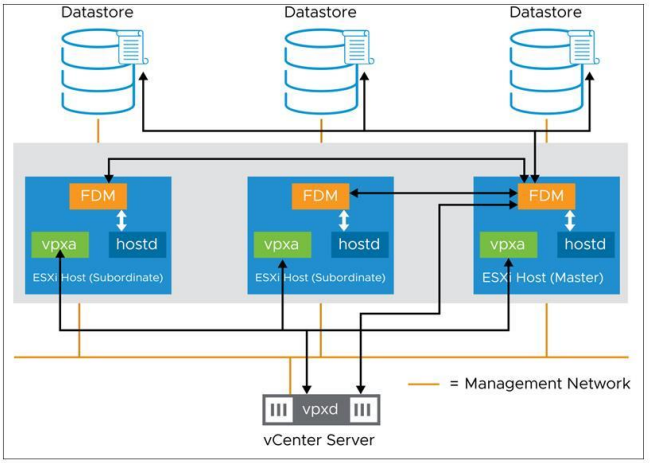

The vSphere HA cluster is managed by a master host. All other hosts are called subordinate hosts. Fault Domain Manager (FDM) services on subordinate hosts all communicate with FDM on the master host. Hosts cannot participate in a vSphere HA cluster if they are in maintenance mode, in standby mode, or disconnected from vCenter Server.

To determine which host is the master host, an election process takes place. The host that can access the greatest number of datastores is elected the master host. If more than one host sees the same number of datastores, the election process determines the master host by using the host Managed Object ID (MOID) assigned by vCenter Server.

https://t.me/learningnets

The election process for a new master host completes in approximately 15 seconds and occurs under these circumstances:

- vSphere HA is enabled.
- The master host encounters a system failure because of one of the following factors:
  - The master host is placed in maintenance mode.
  - The master host is placed in standby mode.
  - vSphere HA is reconfigured.
- The subordinate hosts cannot communicate with the master host because of a network problem.

During the election process, the candidate vSphere HA agents communicate with each other over the management network, or the vSAN network in a vSAN cluster, by using User Datagram Protocol (UDP). All network connections are point-to-point. After the master host is determined, the master host and subordinate hosts communicate using secure TCP. When vSphere HA is started, vCenter Server contacts the master host and sends a list of hosts with membership in the cluster with the cluster configuration. That information is saved to local storage on the master host and then pushed out to the subordinate hosts in the cluster. If additional hosts are added to the cluster during normal operation, the master host sends an update to all hosts in the cluster.

The master host provides an interface for vCenter Server to query the state of and report on the health of the fault domain and VM availability. vCenter Server tells the vSphere HA agent which VMs to protect with their VM-to-host compatibility list. The agent learns about state changes through hostd and vCenter Server learns through vpxa. The master host monitors the health of the subordinate hosts and takes responsibility for VMs that were running on a failed subordinate host.

A subordinate host monitors the health of VMs running locally and sends state changes to the master host. A subordinate host also monitors the health of the master host.

vSphere HA is configured, managed, and monitored through vCenter Server. The vpxd process, which runs on the vCenter Server system, maintains the cluster configuration data. The vpxd process reports cluster configuration changes to the master host. The master host advertises a new copy of the cluster configuration information and each subordinate host fetches an updated copy. Each subordinate host writes the updated configuration information to local storage. A list of protected VMs is stored on each datastore. The VM list is updated after each user-initiated power-on (protected) and power off (unprotected) operation. The VM list is updated after vCenter Server observes these operations.

https://t.me/learningnets

A VM becomes protected when an operation results in a power on. Reverting a VM to a snapshot with memory state causes the VM to power on and become protected. Similarly, a user action that causes the VM to power off, for example, reverting to a snapshot without memory state or a standby operation performed in the guest, causes the VM to become unprotected.

https://t.me/learningnets

9-59 vSphere HA Architecture: Network Heartbeats

The master host sends periodic heartbeats to the subordinate hosts.

In this way, the subordinate hosts know that the master host is alive and the master host knows that the subordinate hosts are alive.

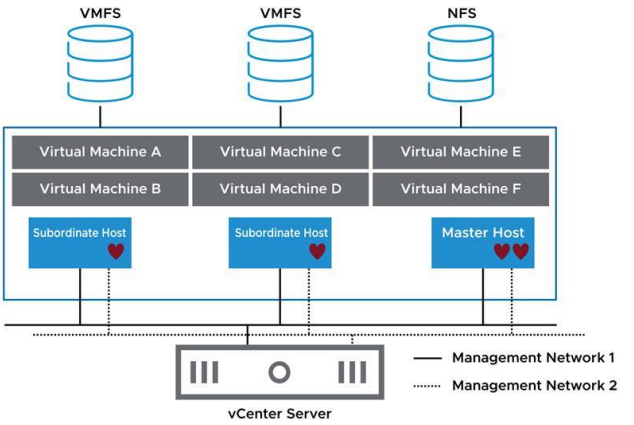

Heartbeats are sent to each subordinate host from the master host over all configured management networks. However, subordinate hosts use only one management network to communicate with the master host. If the management network used to communicate with the master host fails, the subordinate host switches to another management interface to communicate with the master host.

If the subordinate host does not respond within the predefined timeout period, the master host declares the subordinate host as agent unreachable. When a subordinate host is not responding, the master host attempts to determine the cause of the subordinate host's inability to respond. The master host must determine whether the subordinate host crashed, is not responding because of a network failure, or the vSphere HA agent is in an unreachable state.

https://t.me/learningnets

9-60 vSphere HA Architecture: Datastore Heartbeats

When the master host cannot communicate with a subordinate host over the management network, the master host uses datastore heartbeating to determine the cause:

- Subordinate host failure
- Network partition
- Network isolation

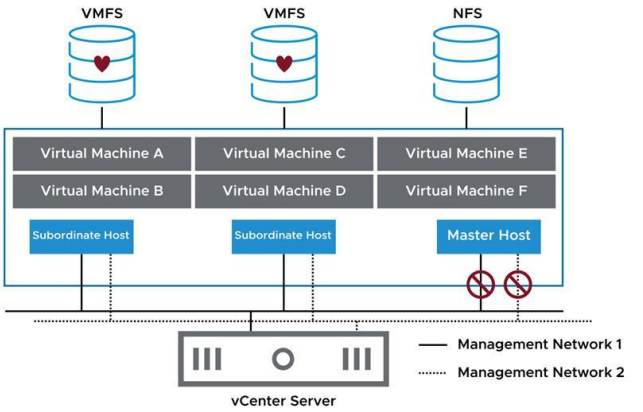

Using datastore heartbeating, the master host determines whether a host has failed or a network isolation has occurred. If datastore heartbeating from the host stops, the host is considered failed. In this case, the failed host's VMs are started on another host in the vSphere HA cluster.

https://t.me/learningnets

9-61 vSphere HA Failure Scenarios

vSphere HA can identify and respond to various types of failures:

- Subordinate host failure
- Master host failure
- Network failure (host isolation)

VMCP enables vSphere HA to detect and respond to datastore access failures:

- APD
- PDL

vSphere HA can also determine whether an ESXi host is isolated or has failed. Isolation refers to when an ESXi host cannot see traffic coming from the other hosts in the cluster and cannot ping its configured isolation address. If an ESXi host fails, vSphere HA attempts to restart the VMs that were running on the failed host on one of the remaining hosts in the cluster. If the ESXi host is isolated because it cannot ping its configured isolation address and sees no management network traffic, the host executes the Host Isolation Response.

https://t.me/learningnets

9-62 Failed Subordinate Hosts

When a subordinate host does not respond to the network heartbeat issued by the master host, the master host tries to identify the cause.

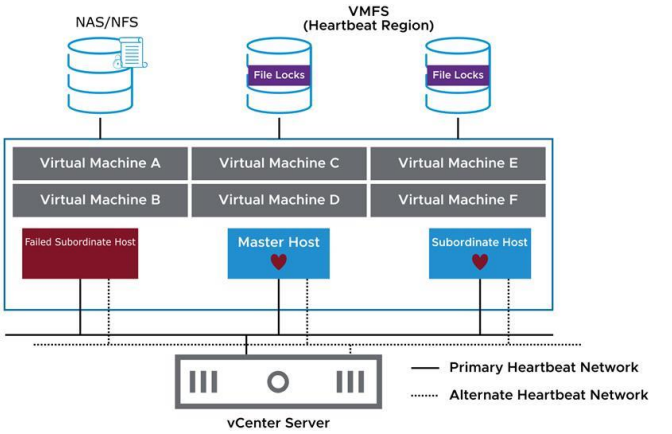

The master host must determine whether the subordinate host is isolated or has failed, for example, because of a misconfigured firewall rule or component failure. The type of failure dictates how vSphere HA responds.

When the master host cannot communicate with a subordinate host over the heartbeat network, the master host uses datastore heartbeating to determine whether the subordinate host failed, is in a network partition, or is network-isolated. If the subordinate host stops datastore heartbeating, the subordinate host is considered to have failed, and its virtual machines are restarted elsewhere.

For VMFS, a heartbeat region on the datastore is read to find out if the host is still heartbeating to it. For NFS datastores, vSphere HA reads the `host--hb` file, which is locked by the ESXi host accessing the datastore. The file guarantees that the VMkernel is heartbeating to the datastore and periodically updates the lock file.

The lock file time stamp is used by the master host to determine whether the subordinate host is isolated or has failed.

https://t.me/learningnets

In both storage examples, the vCenter Server instance selects a small subset of datastores for hosts to heartbeat to. The datastores that are accessed by the greatest number of hosts are selected as candidates. But two datastores are selected (by default) to keep the associated overhead and processing to a minimum.

https://t.me/learningnets

9-63 Failed Master Hosts

When the master host is placed in maintenance mode or fails, the subordinate hosts detect that the master host is no longer issuing heartbeats.

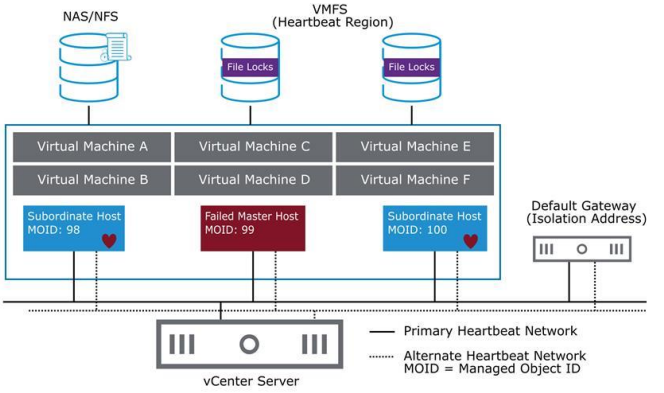

To determine which host is the master host, an election process takes place. The host that can access the greatest number of datastores is elected the master host. If more than one host sees the same number of datastores, the election process determines the master host by using the host Managed Object ID (MOID) assigned by vCenter Server. If the master host fails, is shut down, or is removed from the cluster a new election is held.

https://t.me/learningnets

9-64 Isolated Hosts

A host is declared isolated when the following conditions occur:

- The host is not receiving network heartbeats.
- The host cannot ping its isolation addresses.

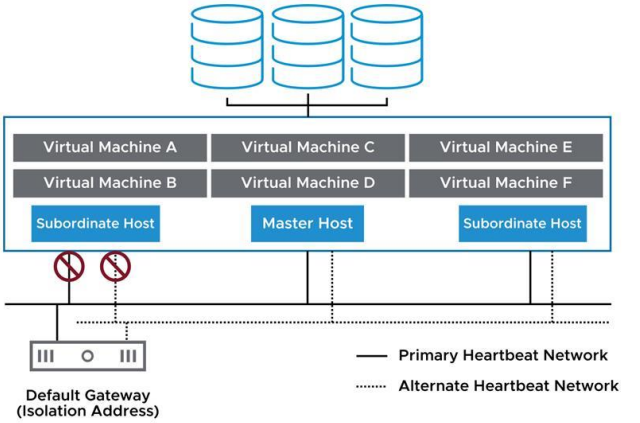

The slide illustrates one of several scenarios that might result in host isolation. If a host loses connectivity to both the primary heartbeat network and the alternate heartbeat network, the host no longer receives network heartbeats from the other hosts in the vSphere HA cluster. Furthermore, the slide depicts that this same host can no longer ping its isolation address.

If a host becomes isolated, the master host must determine if that host is still alive, and merely isolated, by checking for datastore heartbeats. Datastore heartbeats are used by vSphere HA only when a host becomes isolated or partitioned.

https://t.me/learningnets

9-65 VM Storage Failures

Storage connectivity problems might arise because of:

- Network or switch failure
- Array misconfiguration
- Power outage

Storage connectivity problems affect VM availability:

- VMs on affected hosts are difficult to manage.
- Applications with attached disks fail.

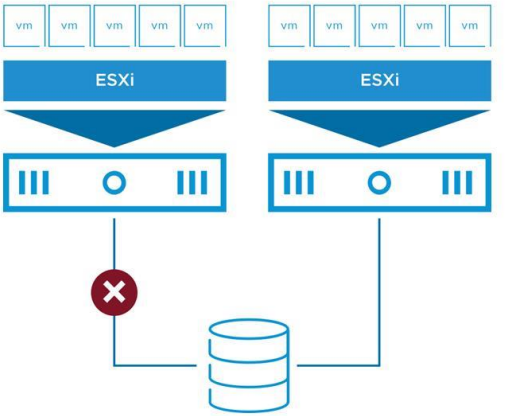

https://t.me/learningnets

Protecting Against Storage Failures with VMCP

VM Component Protection protects against storage failures on a VM.

- If VMCP is enabled, vSphere HA can detect datastore accessibility failures and provide automated recovery for affected VMs.
- VMCP is not supported with vSAN.

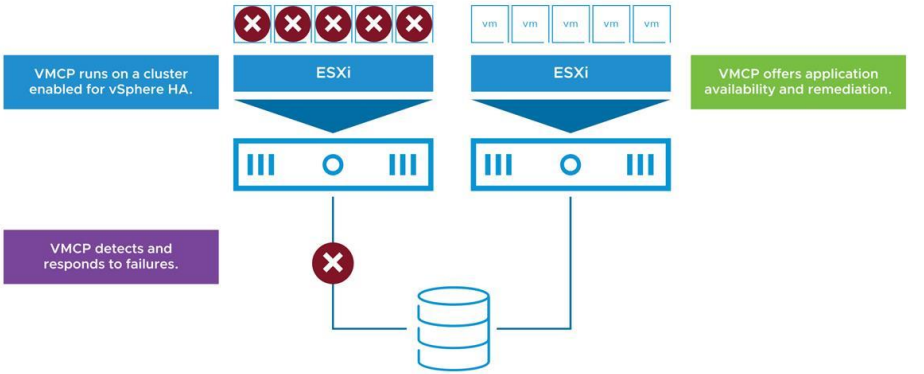

When a datastore accessibility failure occurs, the affected host can no longer access the storage path for a specific datastore. You can determine the response that vSphere HA gives to such a failure, ranging from the creation of event alarms to VM restarts on other hosts.

https://t.me/learningnets

9-67 vSphere HA Design Considerations

When designing your vSphere HA cluster, consider these guidelines:

- Implement redundant heartbeat networks and redundant isolation addresses:
  - Redundancy minimizes host isolation events.
- Physically separate VM networks from the heartbeat networks.
- Implement datastores so that they are separated from the management network by using one or both of the following approaches:
  - Use Fibre Channel over fiber optic for your datastores.
  - If you use IP storage, physically separate your IP storage network from the management network.

If a datastore is based on Fibre Channel, a network failure does not disrupt datastore access. When using datastores based on IP storage (for example, NFS, iSCSI, or Fibre Channel over Ethernet), you must physically separate the IP storage network and the management network (the heartbeat network). If physical separation is not possible, you can logically separate the networks.

https://t.me/learningnets

Review of Learner Objectives

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

- Identify the heartbeat mechanisms used by vSphere HA
- Describe failure scenarios
- Recognize vSphere HA design considerations

https://t.me/learningnets

https://t.me/learningnets

9-70 Learner Objectives

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

- Recognize the requirements for creating and using a vSphere HA cluster
- Configure a vSphere HA cluster

https://t.me/learningnets

9-71 vSphere HA Prerequisites

To create a vSphere HA cluster, you must meet several requirements:

- All hosts must be configured with static IP addresses. If you are using DHCP, you must ensure that the address for each host persists across reboots.
- All hosts must have at least one management network in common.
- For VM monitoring to work, VMware Tools must be installed in every VM.
- Only vSphere HA clusters that contain ESXi hosts 6.x and later can be used to enable VMCP.
- You must not exceed the maximum number of hosts that are allowed in a cluster.
  See VMware Configuration Maximums at https://configmax.vmware.com.

To determine the maximum number of hosts per cluster, see VMware Configuration Maximums at https://configmax.vmware.com.

https://t.me/learningnets

9-72 Configuring vSphere HA Settings

When you create or configure a vSphere HA cluster, you must configure settings that determine how the feature works.

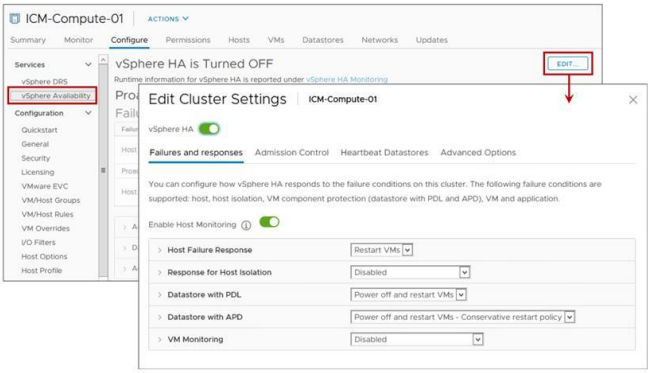

In the vSphere Client, you can configure the following vSphere HA settings:

- Availability failure conditions and responses: Provide settings for host failure responses, host isolation, VM monitoring, and VMCP.
- Admission control: Enable or disable admission control for the vSphere HA cluster and select a policy for how it is enforced.
- Heartbeat datastores: Specify preferences for the datastores that vSphere HA uses for datastore heart-eating.
- Advanced options: Customize vSphere HA behavior by setting advanced options.

https://t.me/learningnets

9-73 vSphere HA Settings: Failures and Responses

You use the Failures and responses pane to configure a cluster's response if a failure occurs.

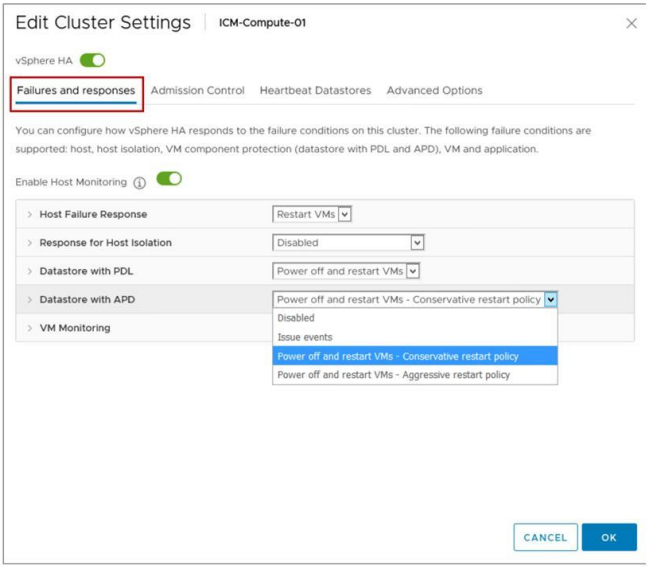

Using the Failures and Responses pane, you can configure how your cluster should function when problems are encountered. You can specify the vSphere HA cluster's response for host failures and isolation. You can also configure VMCP actions when permanent device loss and all paths down situations occur and enable VM monitoring.

If a datastore encounters an All Paths Down (APD) condition, the device state is unknown and might only be temporarily available. You can select the following options for a response to a datastore APD:

- **Disabled**: No action is taken for the affected VMs.
- **Issue events**: No action is taken against the affected VMs, however the administrator is notified when an APD event has occurred.
- **Power off and restart VMs - Conservative restart policy**: vSphere HA does not attempt to restart the affected VMs unless vSphere HA determines that another host can restart the VMs.
  .

https://t.me/learningnets

The host experiencing the APD communicates with the master host to determine whether sufficient capacity exists in the cluster to power on the affected VMs. If the master host determines sufficient capacity exists, the host experiencing the APD stops the VMs so that the VMs can be restarted on a healthy host. If the host experiencing the APD cannot communicate with the master host, no action is taken

- **Power off and restart VMs - Aggressive restart policy**: vSphere HA stops the affected VMs even if it cannot determine that another host can restart the VMs.

  The host experiencing the APD attempts to communicate with the master host to determine if sufficient capacity exists in the cluster to power on the affected VMs. If the master host is not reachable, sufficient capacity for restarting the VMs is unknown. In this scenario, the host takes the risk and stops the VMs so that they can be restarted on the remaining healthy hosts.

  However, if sufficient capacity is not available, vSphere HA might not be able to recover all the affected VMs. This result is common in a network partition scenario where a host cannot communicate with the master host to get a definitive response to the likelihood of a successful recovery.

For more information about VM Component Protection, see https://blogs.vmware.com/vsphere/.

https://t.me/learningnets

9-74 vSphere HA Settings: VM Monitoring

You use VM Monitoring settings to control the monitoring of VMs.

By default, VM Monitoring is set to Disabled.

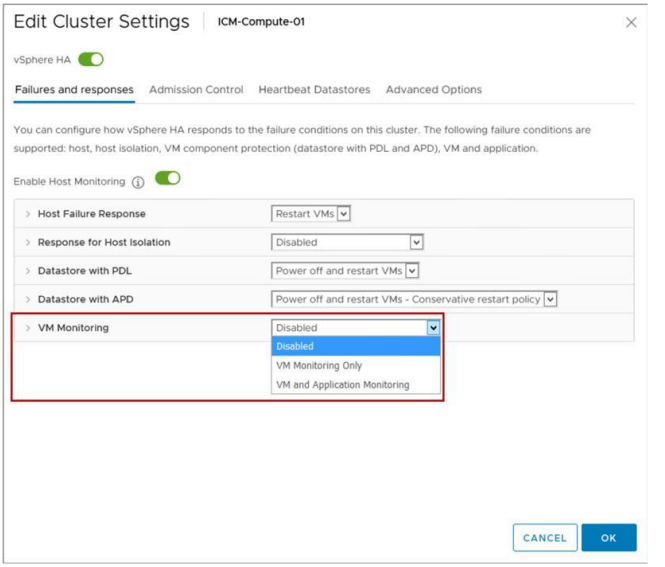

The VM monitoring service determines that the VM has failed if one of the following events occurs:

- VMware Tools heartbeats are not received.
- The guest operating system has not issued an I/O for the last 2 minutes (by default).

If the VM has failed, the VM monitoring service resets the VM to restore services.

You can configure the level of monitoring sensitivity. Highly sensitive monitoring results in a more rapid conclusion that a failure has occurred. Although unlikely, highly sensitive monitoring might lead to falsely identifying failures when the VM or application is still working but heartbeats have not been received because of factors like resource constraints. Low-sensitivity monitoring results in longer interruptions in service between actual failures and VMs being reset. Select an option that is an effective compromise for your needs.

You can select VM and Application Monitoring to enable application monitoring.

https://t.me/learningnets

9-75 vSphere HA Settings: Heartbeat Datastores

A heartbeat file is created on the selected datastores and is used if the management network fails.

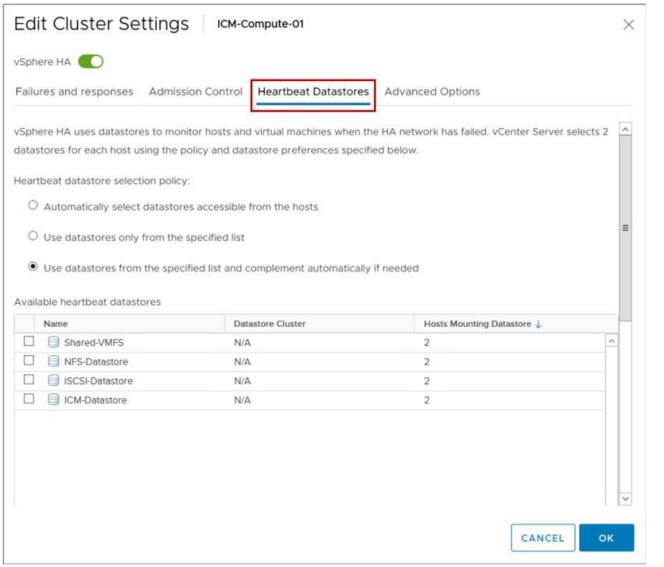

Datastore heartbeating takes checking the health of a host to another level by checking more than the management network to determine a host's health. You can configure a list of datastores to monitor for a particular host, or you can allow vSphere HA to decide. You can also combine both methods.

https://t.me/learningnets

9-76 vSphere HA Settings: Admission Control

vCenter Server uses admission control to ensure both that sufficient resources are available in a cluster to provide failover protection and that VM resource reservations are respected.

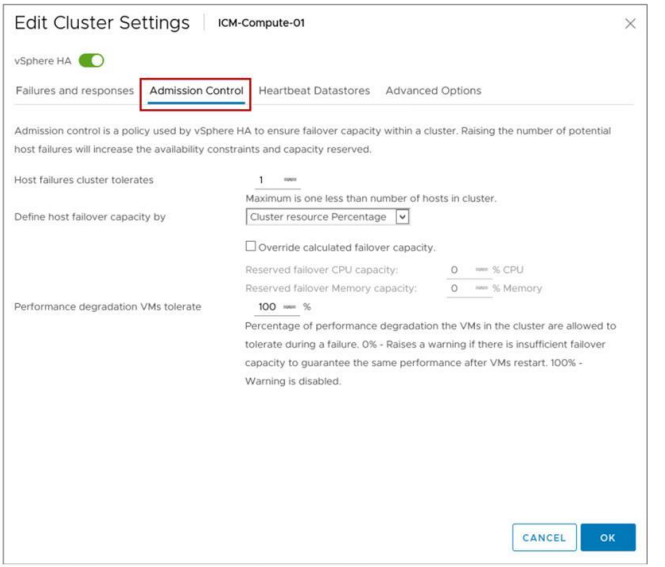

After you create a cluster, you can use admission control to specify whether VMs can be started if they violate availability constraints. The cluster reserves resources to allow failover for all running VMs for a specified number of host failures.

The admission control settings include:

- **Disabled**: (Not recommended) This option disables admission control, allowing the VMs violating availability constraints to power on.
- **Slot Policy**: A slot is a logical representation of memory and CPU resources. With the slot policy option, vSphere HA calculates the slot size, determines how many slots each host in the cluster can hold, and therefore determines the current failover capacity of the cluster.

https://t.me/learningnets

- **Cluster resource Percentage**: (Default) This value specifies a percentage of the cluster's CPU and Memory resources to be reserved as spare capacity to support failovers.
- **Dedicated failover hosts**: This option selects hosts to use for failover actions. If a default failover host does not have enough resources, failovers can still occur to other hosts in the cluster.

https://t.me/learningnets

9-77 Example: Admission Control Using Cluster Resources Percentage

Example of calculating total failover capacity using cluster resource percentages:

- Total cluster capacity:
  - CPU: 18 GHz
  - Memory: 24 GB
- Total VM reservations:
  - CPU: 7 GHz
  - Memory: 6 GB
- Current failover CPU capacity is 61%:
  ((18 GHz - 7 GHz)/18 GHz) = 61%
- Current failover memory capacity is 75%:
  ((24 GB - 6 GB)/24 GB) = 75%

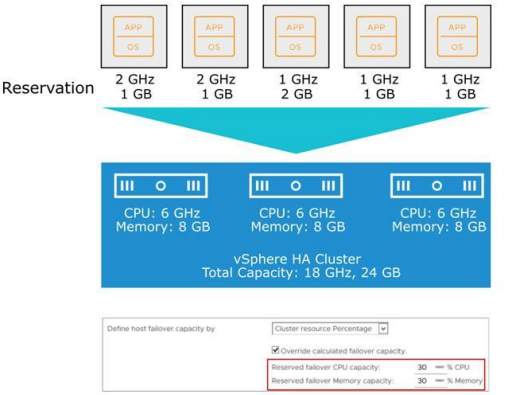

Cluster resource percentage is the default admission control policy. Recalculations occur automatically as the cluster's resources change, for example, when a host is added to or removed from the cluster.

https://t.me/learningnets

Example: Admission Control Using Slots (1)

A slot is calculated by combining the largest memory reservation and the largest CPU reservation of any running VM in the cluster. vSphere HA performs admission control by calculating the following values:

- Slot size:
  - In this example, the slot size is 2 GHz CPU and 2 GB memory.
- Number of slots each host in the cluster can hold:
  - Three
  - The cluster has a total of nine slots (3 + 3 + 3).

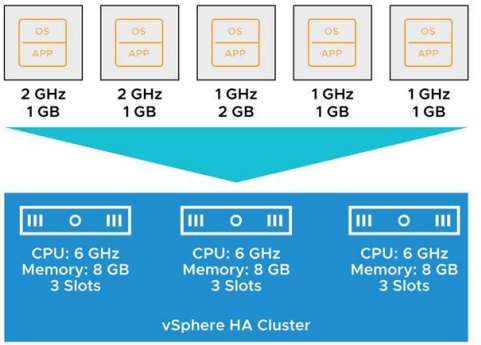

https://t.me/learningnets

Example: Admission Control Using Slots (2)

vSphere HA also calculates the current failover capacity. In this example, the failover capacity is one host:

- If the first host fails, six slots remain in the cluster, which is sufficient for all five of the powered-on VMs.
- If the first and second hosts fail, only three slots remain, which is insufficient for all five of the VMs.
- If the current failover capacity is less than the configured failover capacity, vSphere HA does not allow any more VMs to power on.

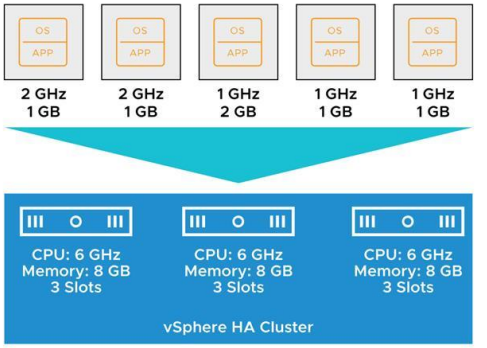

https://t.me/learningnets

9-80 vSphere HA Settings: Performance Degradation VMs Tolerate

The **Performance degradation VMs tolerate** threshold specifies the percentage of performance degradation that the VMs in the cluster are allowed to tolerate during a failure.

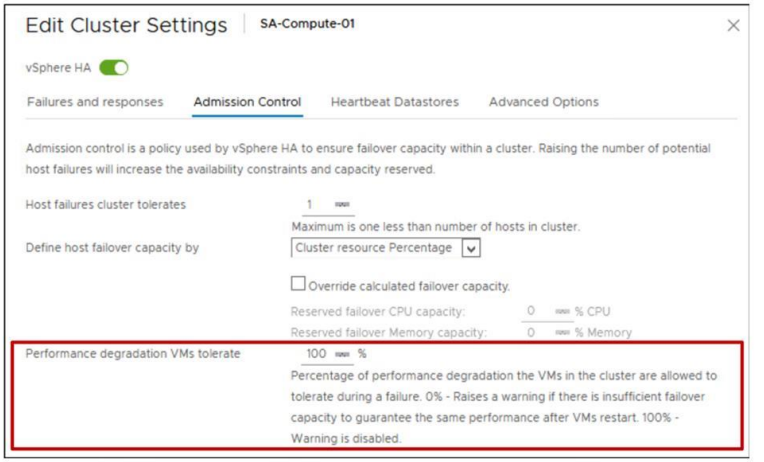

Admission control can also be configured to offer warnings when the actual use exceeds the failover capacity percentage. The resource reduction calculation takes into account a VM's reserved memory and memory overhead.

By setting the **Performance degradation VMs tolerate** threshold, you can specify when a configuration issue should generate a warning or notice. For example:

- The default value is 100 percent, which produces no warnings.
- If you reduce the threshold to 0 percent, a warning is generated when cluster use exceeds the available capacity.
- If you reduce the threshold to 20 percent, the performance reduction that can be tolerated is calculated as performance reduction = current use x 20 percent.

https://t.me/learningnets

When the current use minus the performance reduction exceeds the available capacity, a configuration notice is issued.

The **Performance degradation VMs tolerate** threshold is not available unless vSphere DRS is enabled.

https://t.me/learningnets

9-81 vSphere HA Setting: Default VM Restart Priority

The VM restart priority determines the order in which vSphere HA restarts VMs on a running host.

VMs are put in the Medium restart priority by default, unless the restart priority is explicitly set using VM overrides.

Exceptions:

- Agent VMs always start first, and the restart priority is nonconfigurable.
- vSphere Fault Tolerance secondary VMs fail over before regular VMs. Primary VMs follow the normal restart priority.

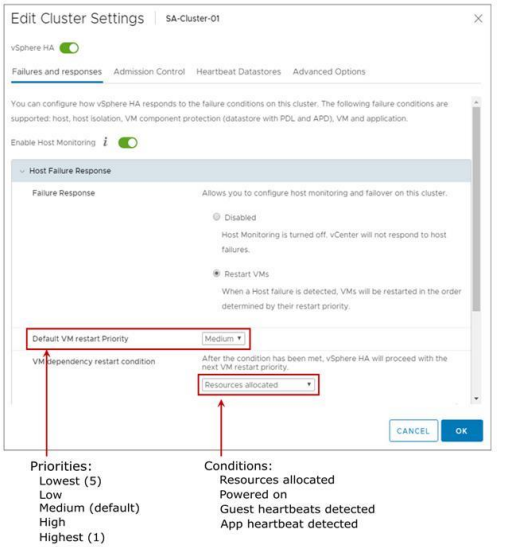

Optionally, you can configure a delay when a certain restart condition is met.

https://t.me/learningnets

9-82 vSphere HA Settings: Advanced Options

You can set advanced vSphere HA options to customize vSphere HA behavior.

Description	Option	Value
Force a cluster not to use the default isolation address (default gateway).	das.usedefaultisolationaddress	`false`
Force a cluster to ping alternate isolation addresses.	das.isolationaddressX	IP address or FQDN
Force a cluster to wait beyond the default 30-second isolation action window.	fdm.isolationpolicydelaysec	>=30 seconds
Force maximum bound on the memory slot size.	das.slotmeminmb	100
Force maximum bound on the CPU slot size.	das.slotcpuinmhz	32

You can set advanced options that affect the behavior of your vSphere HA cluster. For more details, see *vSphere Availability* at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-63F459B7-8884-4818-8872-C9753B2E0215.html.

https://t.me/learningnets

9-83 vSphere HA Settings: VM-Level Settings

You can customize the restart priority for individual VMs in a cluster to override the default level set for the entire cluster.

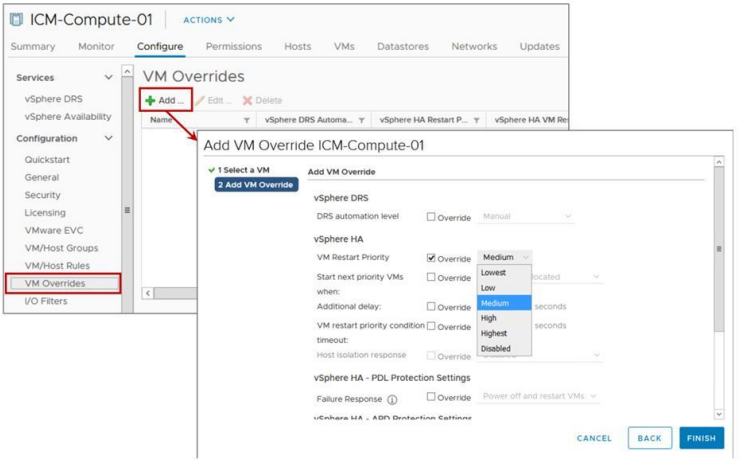

https://t.me/learningnets

9-84 About vSphere HA Orchestrated Restart

Orchestrated restart in vSphere HA enables five tiers for restarting VM and VM-VM dependencies. Choose from among the following conditions that must be met before a VM is considered ready:

- VM has resources secured.
- VM is powered on.
- VMware Tools heartbeat is detected.
- VMware Tools application heartbeat is detected.

VMs can be grouped into tiers indicating their startup priority:

- All VMs in the priority 1 tier receive their resources first and are powered on.
- After all the VMs in tier 1 have met their defined restart condition, vSphere HA continues to the VMs in the priority 2 tier, and so on.

After a host failure, VMs are assigned to other hosts with unreserved capacity, with the highest priority VMs placed first. The process continues to those VMs with lower priority until all have been placed or no more cluster capacity is available to meet the reservations or memory overhead of the VMs. A host then restarts the VMs assigned to it in priority order.

If insufficient resources exist, vSphere HA waits for more unreserved capacity to become available, for example, because of a host coming back online, and then retries the placement of these VMs. To reduce the chance of this situation occurring, configure vSphere HA admission control to reserve more resources for failures. With admission control, you can control the amount of cluster capacity that is reserved by VMs, which is unavailable to meet the reservations and memory overhead of other VMs if a failure occurs.

https://t.me/learningnets

9-85 VM Dependencies in Orchestrated Restart (1)

VMs can depend only on other VMs of the same or higher priority. Only direct dependencies are supported. VM-to-VM dependency is a hard rule. Creating cyclical dependencies causes VM restart to fail.

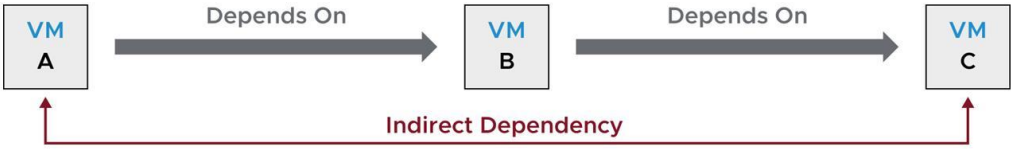

https://t.me/learningnets

9-86 VM Dependencies in Orchestrated Restart (2)

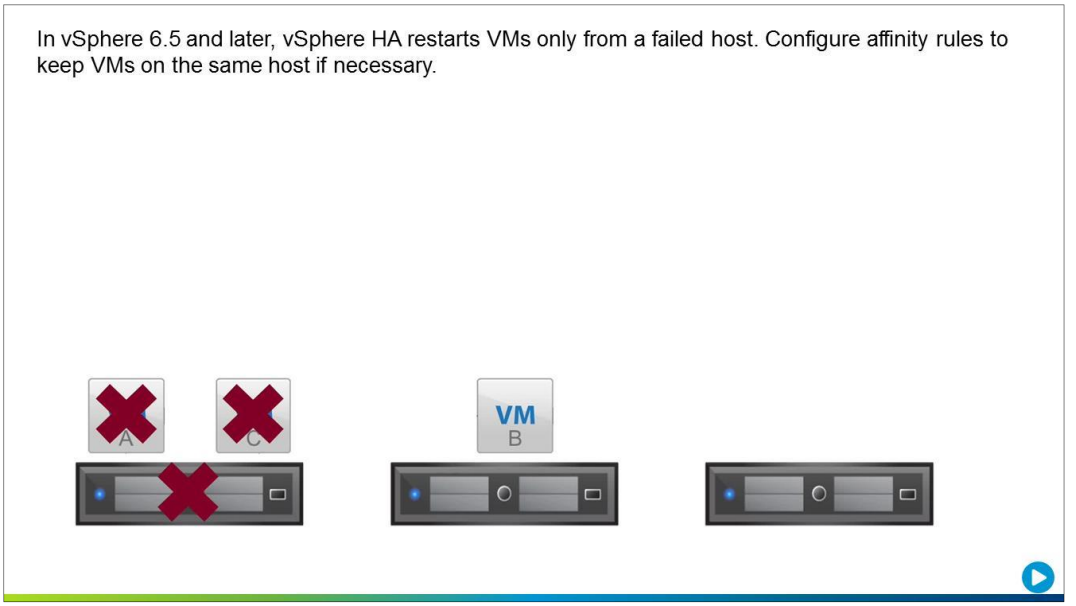

To play the animation, go to https://vmware.bravais.com/s/JDg7NJ3DjVli7r6FiMQ0.

https://t.me/learningnets

Disable host monitoring before modifying virtual networking components that involve the VMkernel ports configured for management or vSAN traffic.

This practice prevents unwanted attempts to fail over VMs.

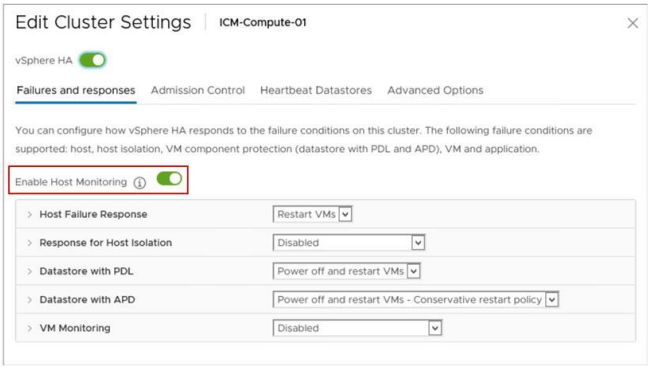

The following network maintenance suggestions can help you avoid the false detection of host failure and network isolation because of dropped vSphere HA heartbeats:

- Changing your network hardware or networking settings can interrupt the heartbeats used by vSphere HA to detect host failures, and might result in unwanted attempts to fail over VMs. When changing the management or vSAN networks of the hosts in the vSphere HA-enabled cluster, suspend host monitoring and place the host in maintenance mode.
- Disabling host monitoring is required only when modifying virtual networking components and properties that involve the VMkernel ports configured for the Management or vSAN traffic, which are used by the vSphere HA networking heartbeat service.
- After you change the networking configuration on ESXi hosts, for example, adding port groups, removing virtual switches, or suspending host monitoring, you must reconfigure vSphere HA on all hosts in the cluster. This reconfiguration causes the network information to be reinspected. Then, you must reenable host monitoring.

https://t.me/learningnets

9-88 Monitoring vSphere HA Cluster Status

You can monitor the status of a vSphere HA cluster on the Summary page of the **Monitor** tab.

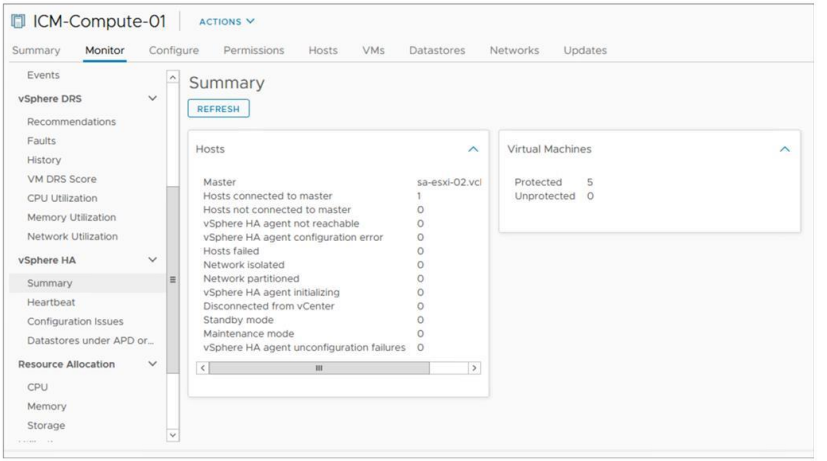

You cluster or its hosts can experience configuration issues and other errors that adversely affect proper vSphere HA operation. You can monitor these errors on the Configuration Issues page.

https://t.me/learningnets

9-89 Using vSphere HA with vSphere DRS

vSphere HA is closely integrated with vSphere DRS:

- When a failover occurs, vSphere HA checks whether resources are available on that host for the failover.
- If resources are not available, vSphere HA asks vSphere DRS to accommodate for the VMs where possible.

vSphere HA might not be able to fail over VMs for the following reasons:

- vSphere HA admission control is disabled, and resources are insufficient in the remaining hosts to power on all the failed VMs.
- Sufficient aggregated resources exist, but they are fragmented across hosts. In such cases, vSphere HA uses vSphere DRS to try to adjust the cluster by migrating VMs to defragment the resources.

When vSphere HA performs failover and restarts VMs on different hosts, its first priority is the immediate availability of all VMs. After the VMs are restarted, the hosts in which they were powered on are usually heavily loaded, and other hosts are comparatively lightly loaded. vSphere DRS helps to balance the load across hosts in the cluster.

https://t.me/learningnets

9-90 Lab 26: Using vSphere HA

Use vSphere HA functionality:

1. Enable vSphere HA in a Cluster
2. View Information About the vSphere HA Cluster
3. Configure Network Management Redundancy
4. Test the vSphere HA Functionality
5. View the vSphere HA Cluster Resource Usage
6. Configure the Percentage of Resource Degradation to Tolerate

https://t.me/learningnets

9-91 Review of Learner Objectives

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

- Recognize the requirements for creating and using a vSphere HA cluster
- Configure a vSphere HA cluster

https://t.me/learningnets

https://t.me/learningnets

9-93 Learner Objectives

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

- Describe the features and benefits of using vSphere Fault Tolerance
- Describe how vSphere Fault Tolerance works
- Describe how vSphere Fault Tolerance works with vSphere HA and vSphere DRS
- Enable vSphere Fault Tolerance using the vSphere Client

https://t.me/learningnets

9-94 About vSphere Fault Tolerance

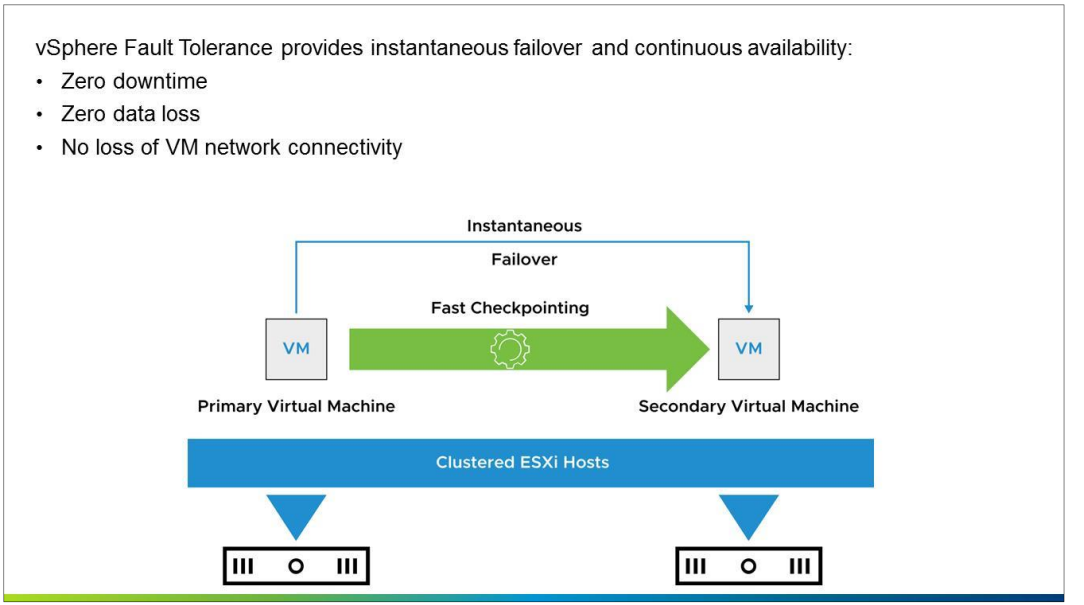

You can use vSphere Fault Tolerance for most mission-critical VMs. vSphere Fault Tolerance is built on the ESXi host platform.

The protected VM is called the primary VM. The duplicate VM is called the secondary VM. The secondary VM is created and runs on a different host to the primary VM. The secondary VM's execution is identical to that of the primary VM. The secondary VM can take over at any point without interruption and provide fault-tolerant protection.

The primary VM and the secondary VM continuously monitor the status of each other to ensure that fault tolerance is maintained. A transparent failover occurs if the host running the primary VM fails, in which case the secondary VM is immediately activated to replace the primary VM. A new secondary VM is created and started, and fault tolerance redundancy is reestablished automatically. If the host running the secondary VM fails, the secondary VM is also immediately replaced. In either case, users experience no interruption in service and no loss of data.

https://t.me/learningnets

9-95 vSphere Fault Tolerance Features

vSphere Fault Tolerance protects mission-critical, high-performance applications regardless of the operating system used.
vSphere Fault Tolerance:

- Supports VMs configured with up to 8 vCPUs and 128 GB memory
- Supports up to four fault-tolerant VMs per host with no more than eight vCPUs between them
- Supports vSphere vMotion migration for primary and secondary VMs
- Creates a secondary copy of all VM files and disks
- Provides fast checkpoint copying to keep primary and secondary VMs synchronized
- Supports multiple VM disk formats: thin provision, thick provision lazy-zeroed, and thick provision eager-zeroed
- Can be used with vSphere DRS only when Enhanced vMotion Compatibility is enabled
- Supports interoperability with vSAN

You can use vSphere Fault Tolerance with vSphere DRS only when the Enhanced vMotion Compatibility feature is enabled.

When you enable EVC mode on a cluster, vSphere DRS makes the initial placement recommendations for fault-tolerant VMs, and you can assign a vSphere DRS automation level to primary VMs. The secondary VM always assumes the same setting as its associated primary VM.

When vSphere Fault Tolerance is used for VMs in a cluster that has EVC mode disabled, the fault-tolerant VMs are given the disabled vSphere DRS automation level. In such a cluster, each primary VM is powered on only on its registered host, and its secondary VM is automatically placed.

https://t.me/learningnets

9-96 vSphere Fault Tolerance with vSphere HA and vSphere DRS

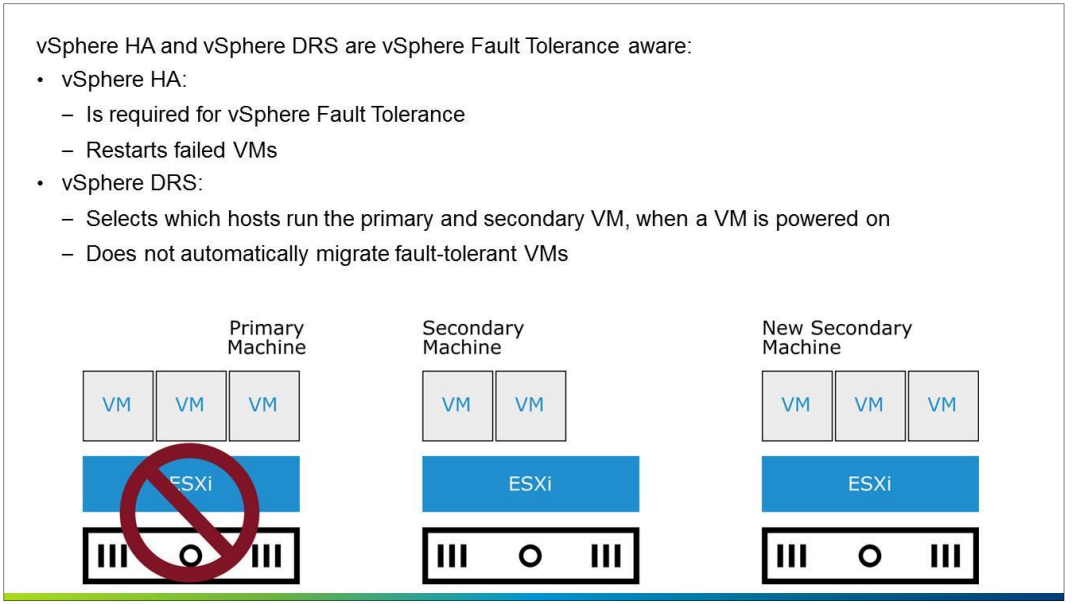

A fault-tolerant VM and its secondary copy are not allowed to run on the same host. This restriction ensures that a host failure cannot result in the loss of both VMs.

https://t.me/learningnets

9-97 Redundant VMDK Files

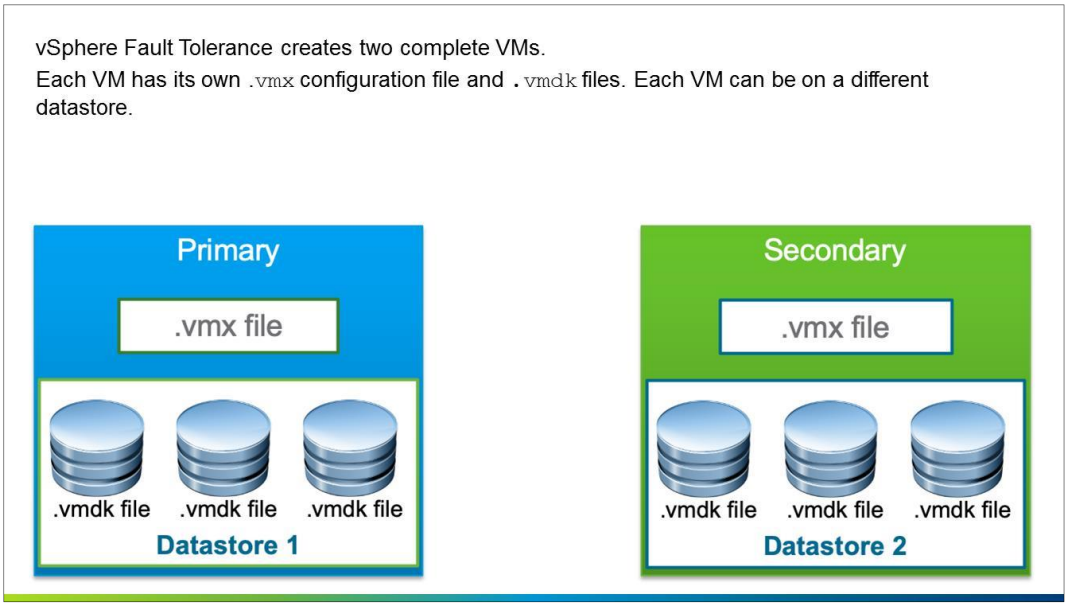

vSphere Fault Tolerance provides failover redundancy by creating two full VM copies. The VM files can be placed on the same datastore. However, VMware place these files on separate datastores to provide recovery from datastore failures.

https://t.me/learningnets

9-98 vSphere Fault Tolerance Checkpoint

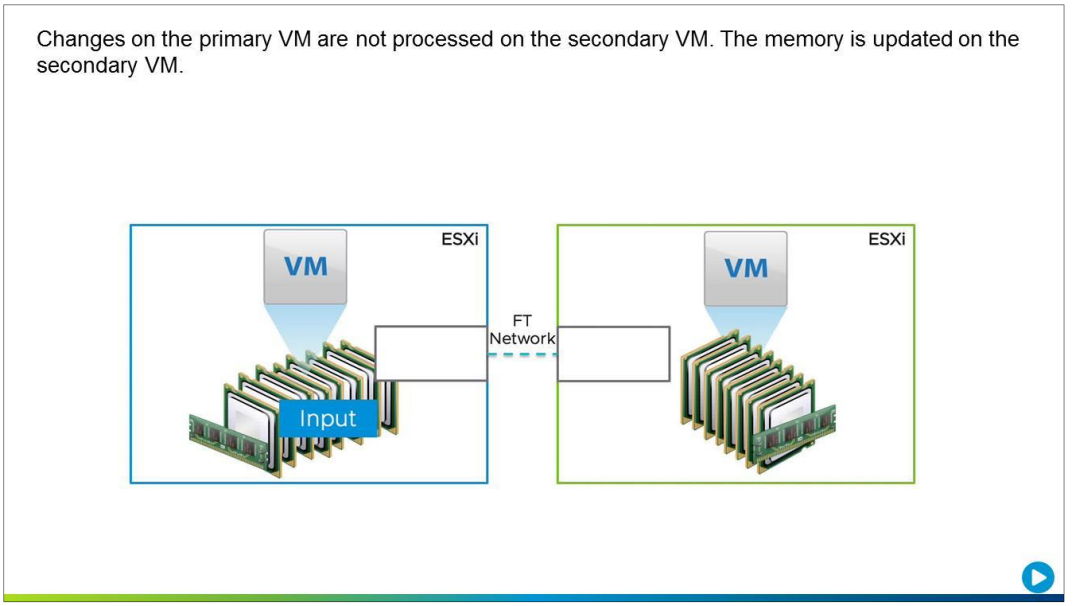

To play the animation, go to https://vmware.bravais.com/s/a8GAXMVDFHxWLstdhM1G.

https://t.me/learningnets

9-99 vSphere Fault Tolerance: Precopy

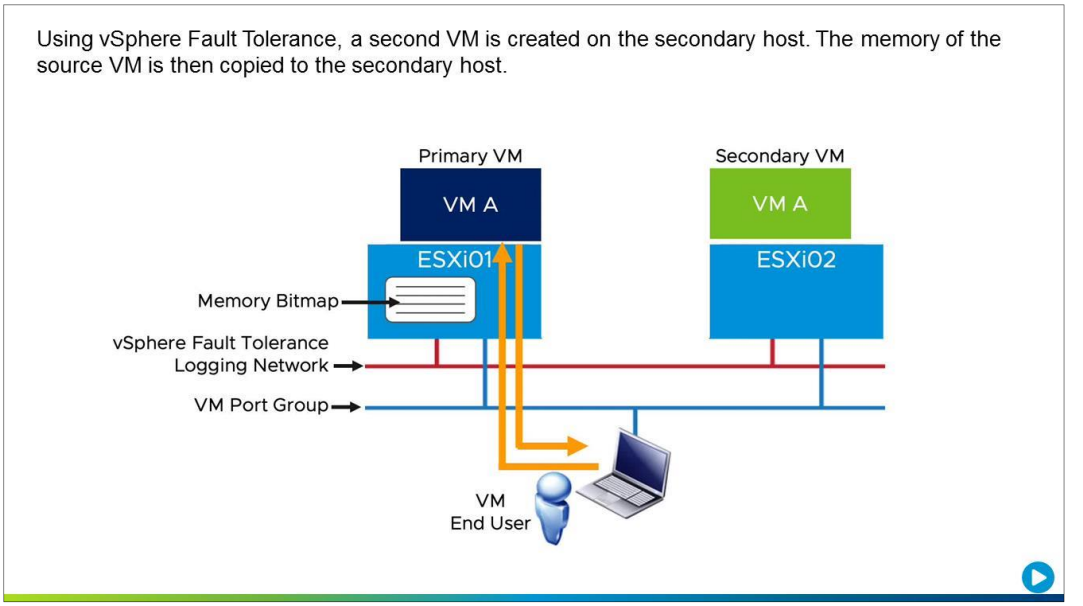

To play the animation, go to https://vmware.bravais.com/s/XM1mrfVGU5vPd6HVfBfv.

https://t.me/learningnets

9-100 vSphere Fault Tolerance Fast Checkpointing

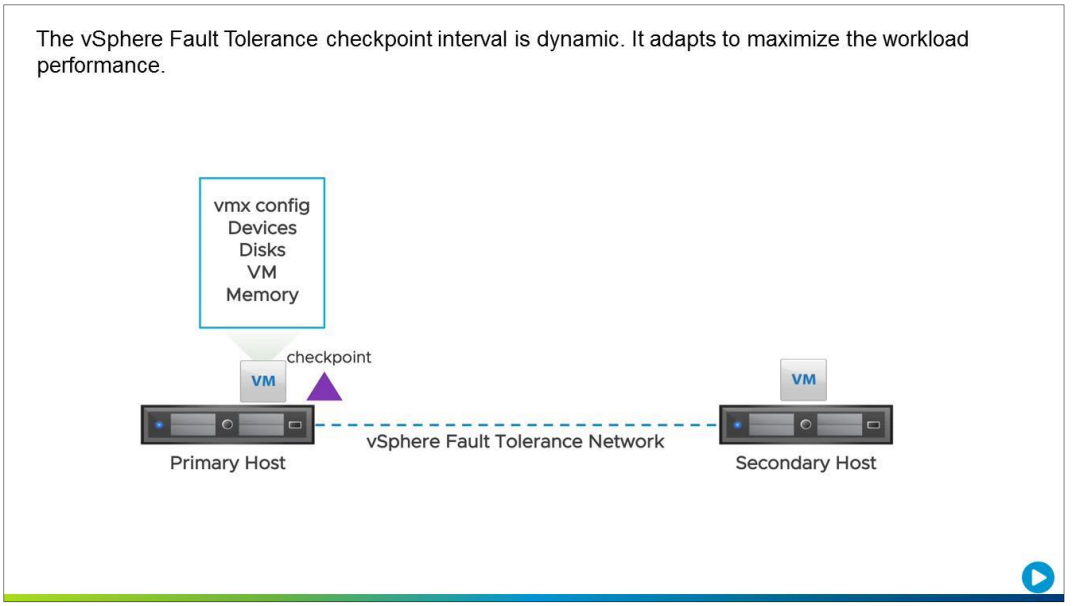

To play the animation, go to https://vmware.bravais.com/s/KafKVBJNsBpY7hn5bmGs.

vSphere Fault Tolerance uses an algorithm that provides fast, continuous copying (checkpointing) of the primary host VM. The primary VM is copied (checkpointed) periodically, and the copies are sent to a secondary host. If the primary host fails, the VM continues on the secondary host at the point of its last network send.

The goal is to take checkpoints of VMs at least every 10 milliseconds. The primary VM is continuously copied (checkpointed), and these copies (checkpoints) are sent to a secondary host.

The initial complete copy (checkpoint) is created using a modified form of vSphere vMotion migration to the secondary host. The primary VM holds each outgoing network packet until the following copy (checkpoint) has been sent to the secondary host.

In vSphere Fault Tolerance, checkpoint data makes up the last changed pages of memory. The source VM is paused to access this memory. This pause is typically under one second.

https://t.me/learningnets

9-101 vSphere Fault Tolerance Shared Files

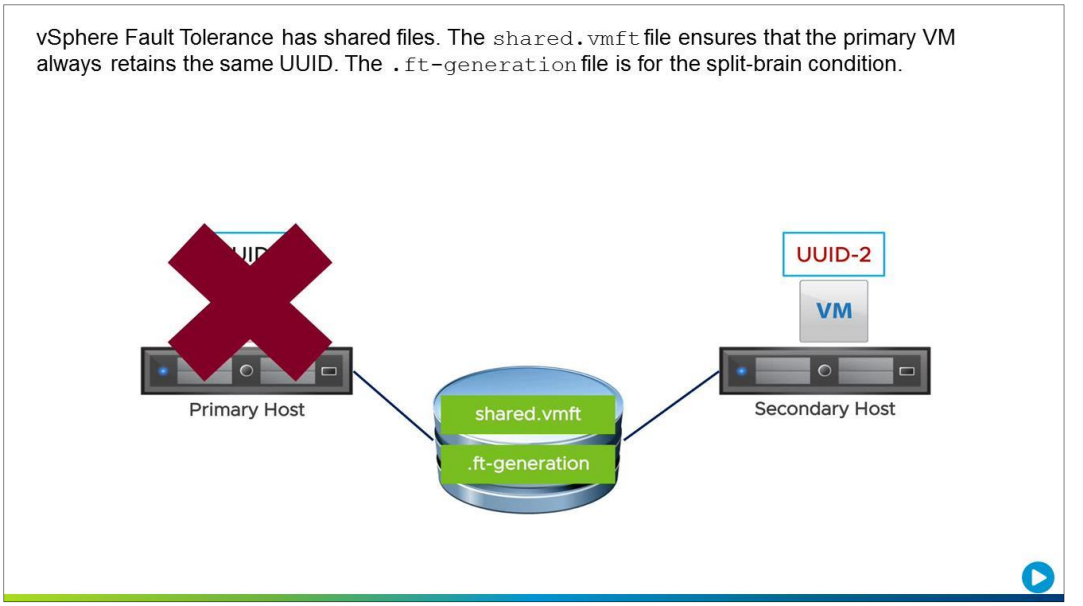

To play the animation, go to https://vmware.bravais.com/s/2c9Y6hQ4X4uFWWzSTeFk.

The `shared.vmft` file, which is found on a shared datastore, is the vSphere Fault Tolerance metadata file. This file contains the primary and secondary instance UUIDs and the primary and secondary vmx paths.

vSphere Fault Tolerance avoids split-brain situations, which can lead to two active copies of a virtual machine after recovery from a failure. The `.ftgeneration` file ensures that only one VM instance is designated as the primary VM.

https://t.me/learningnets

9-102 Enabling vSphere Fault Tolerance on a VM

You can turn on vSphere Fault Tolerance for a VM using the vSphere Client.

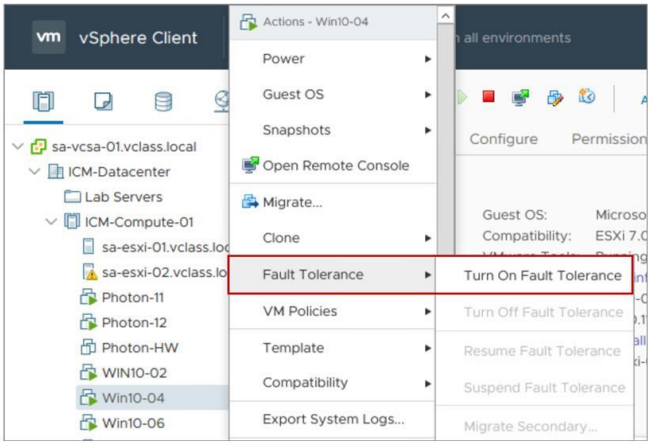

After you take all the required steps for enabling vSphere Fault Tolerance for your cluster, you can use the feature by turning it on for individual VMs.

Before vSphere Fault Tolerance can be turned on, validation checks are performed on a VM.

After these checks are passed, and you turn on vSphere Fault Tolerance for a VM, new options are added to the Fault Tolerance section of the VM's context menu. These options include turning off or disabling vSphere Fault Tolerance, migrating the secondary VM, testing failover, and testing restart of the secondary VM.

When vSphere Fault Tolerance is turned on, vCenter Server resets the VM's memory limit to the default (unlimited memory) and sets the memory reservation to the memory size of the VM. While vSphere Fault Tolerance is turned on, you cannot change the memory reservation, size, limit, number of virtual CPUs, or shares. You also cannot add or remove disks for the VM. When vSphere Fault Tolerance is turned off, any parameters that were changed are not reverted to their original values.

https://t.me/learningnets

9-103 Review of Learner Objectives

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

- Describe the features and benefits of using vSphere Fault Tolerance
- Describe how vSphere Fault Tolerance works
- Describe how vSphere Fault Tolerance works with vSphere HA and vSphere DRS
- Enable vSphere Fault Tolerance using the vSphere Client

https://t.me/learningnets

9-104 Activity: Virtual Beans Clusters (1)

As a Virtual Beans administrator, you want to place ESXi hosts in a vSphere cluster for a scalable and highly available infrastructure. Match the goal to the feature that helps achieve the goal.

Goal	vSphere Feature
○ Add ESXi hosts to the data center and let vSphere balance the load across the hosts.	○ vSphere HA
○ Make business-critical applications 99.99 percent available (downtime per year of 52.56 minutes).	○ VM scores
○ Identify VMs that are experiencing serious resource contention.	○ Cluster Quickstart
○ Improve the performance of certain VMs by ensuring that they always run together on the same host.	○ vSphere Fault Tolerance
	○ VM-Host affinity

https://t.me/learningnets

9-105 Activity: Virtual Beans Clusters (2)

As a Virtual Beans administrator, you want to place ESXi hosts in a vSphere cluster for a scalable and highly available infrastructure. Match the goal to the feature that helps achieve the goal.

Virtual Beans Goals	vSphere Feature
Add ESXi hosts to the data center and let vSphere balance the load across the hosts.	vSphere DRS
Make business-critical applications 99.99 percent available (downtime per year of 52.56 minutes).	vSphere HA or vSphere Fault Tolerance
Identify VMs that are experiencing serious resource contention.	VM scores
Improve the performance of certain VMs by ensuring that they always run together on the same host.	VM-Host affinity

https://t.me/learningnets

9-106 Key Points

- When you create a cluster, you can enable vSphere DRS, vSphere HA, vSAN, and the ability to manage image updates on all hosts collectively.
- vSphere DRS clusters provide automated resource management to ensure that a VM's resource requirements are satisfied.
- vSphere DRS works best when the VMs meet vSphere vMotion migration requirements.
- vSphere HA restarts VMs on the remaining hosts in the cluster.
- You implement redundant heartbeat networks either with NIC teaming or by creating additional heartbeat networks.
- vSphere Fault Tolerance provides zero downtime for applications that must always be available.

Questions?

https://t.me/learningnets

Module 10

vSphere Lifecycle Management

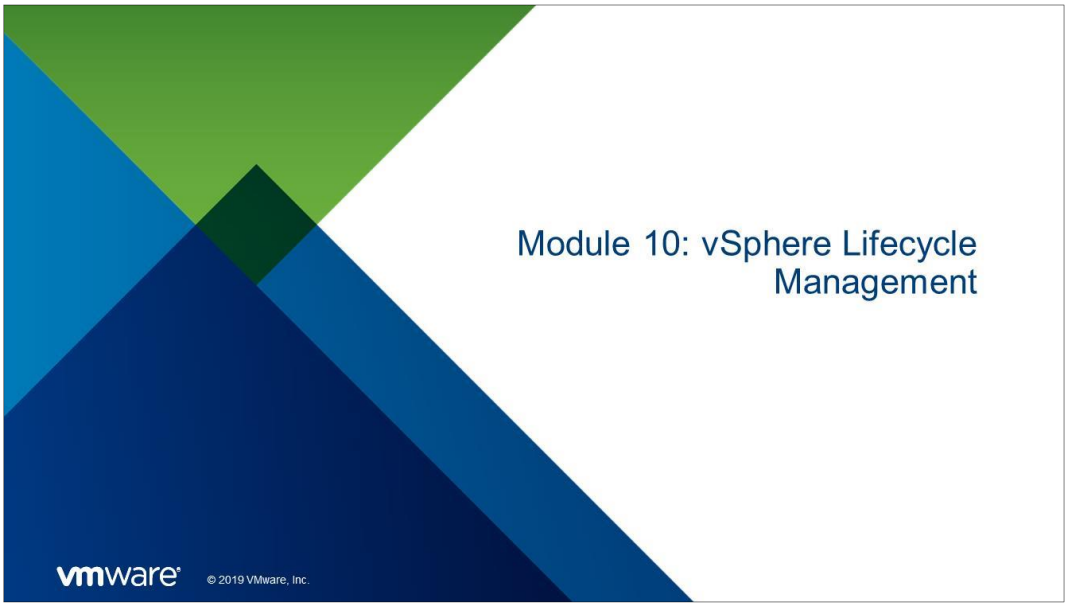

https://t.me/learningnets

10-2 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.

https://t.me/learningnets

10-3 Module Lessons

1. vCenter Server Update Planner
2. Overview of vSphere Lifecycle Manager
3. Working with Baselines
4. Working with Cluster Images
5. Managing the Life Cycle of VMware Tools and VM Hardware

https://t.me/learningnets

10-4 Virtual Beans: Lifecycle Management

Virtual Beans is struggling with its current lifecycle management process. The process is mostly manual and is error-prone and inefficient.
The company wants to use vSphere Lifecycle Manager. It hopes that this feature can provide a centralized, automated patch and version management system for keeping vSphere components up to date:

- vCenter Server
- ESXi hosts
- Virtual machines:
  - VM hardware
  - VMware Tools

As the vSphere administrator, you must implement vSphere Lifecycle Manager in the Virtual Beans data center.

https://t.me/learningnets

10-5 Lesson 1: vCenter Server Update Planner

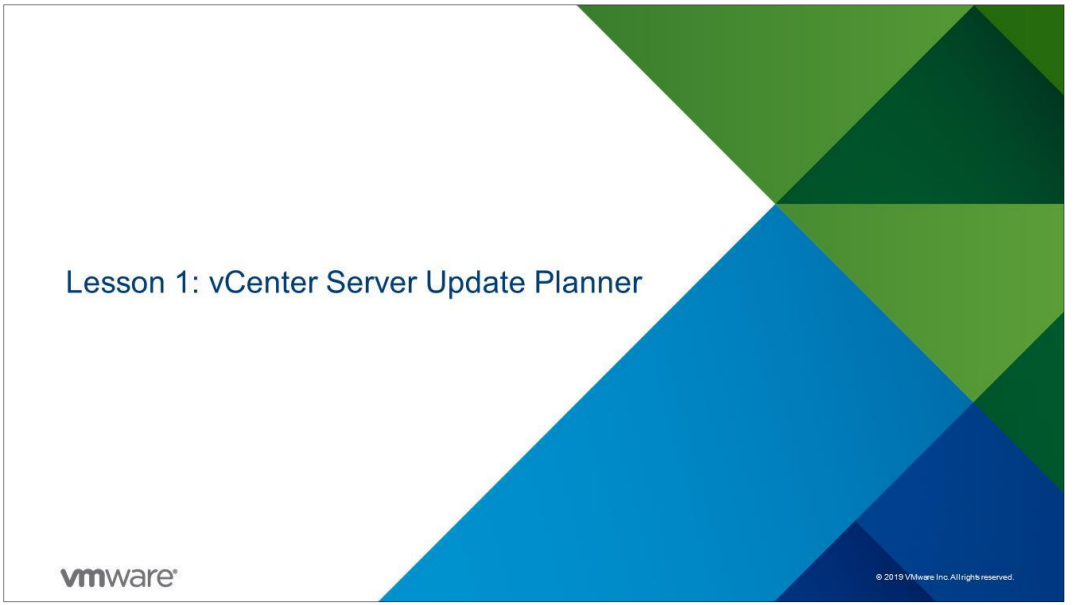

https://t.me/learningnets

10-6 Learner Objectives

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

- Describe features of the vCenter Server Update Planner
- Run vCenter Server upgrade prechecks and interoperability reports
- Export prechecks and interoperability report results

https://t.me/learningnets

10-7 Overview of vCenter Server Update Planner

In vSphere 7, you can use the Update Planner feature for planning updates to vCenter Server and other VMware products that are registered with it.
The Update Planner can perform the following tasks:

- Retrieve information about VMware products registered with vCenter Server.
- List available vCenter Server updates and upgrades.
- Create interoperability reports.
- Perform a precheck to verify that your system meets the minimum software and hardware requirements for a successful upgrade of vCenter Server.

https://t.me/learningnets

10-8 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.

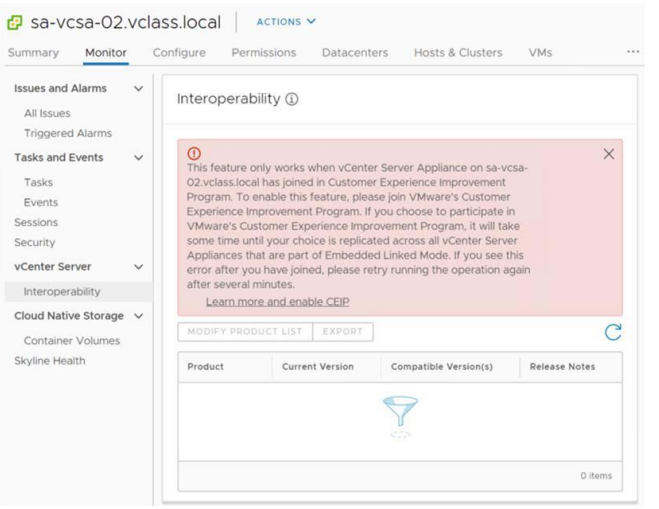

When generating reports, if the Customer Experience Improvement Program (CEIP) is not yet accepted, a prompt describing CEIP appears. Reports are not generated if you do not join CEIP.

https://t.me/learningnets

10-9 Update Planner View in the vSphere Client

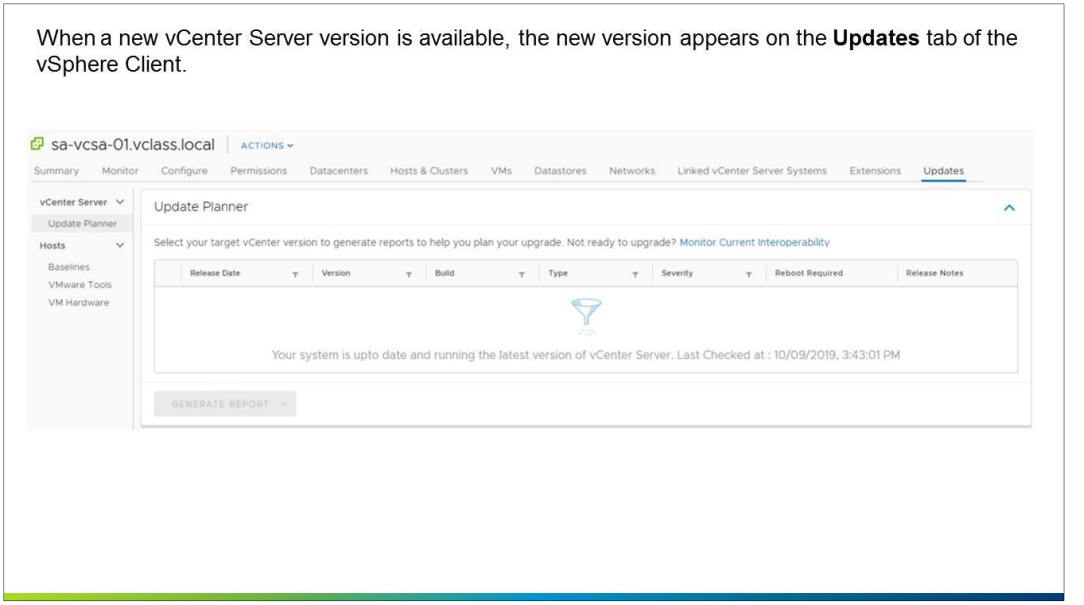

When new vCenter Server updates are released, the vSphere Client shows a notification in the **Summary** tab. Clicking the notification directs you to the **Updates** tab.

The **Updates** tab has an 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.

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.

https://t.me/learningnets

10-10 Interoperability View in vSphere Client

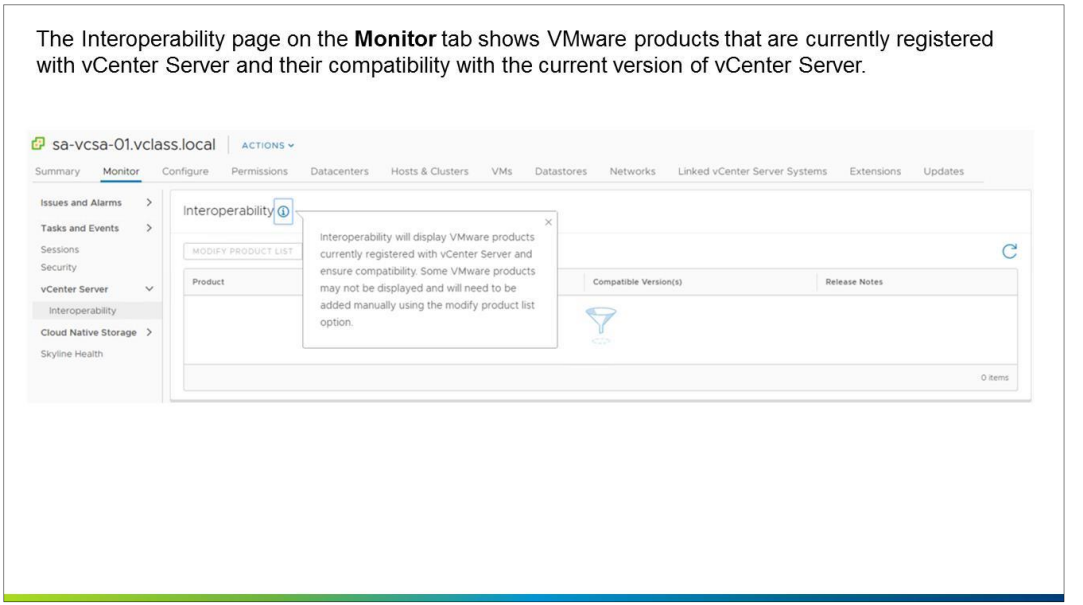

In the vSphere Client, the Interoperability page appears on the **Monitor** tab of vCenter Server. This page displays VMware products currently registered with vCenter Server.

Columns show the name, current version, compatible version, and release notes of each detected product.

If you do not see your registered VMware products, you can manually modify the list and add the appropriate names and versions.

https://t.me/learningnets

10-11 Exporting Report Results

You can export report results in CSV format and use the report as a guide to prepare for an update. Both product interoperability and precheck reports can be exported.

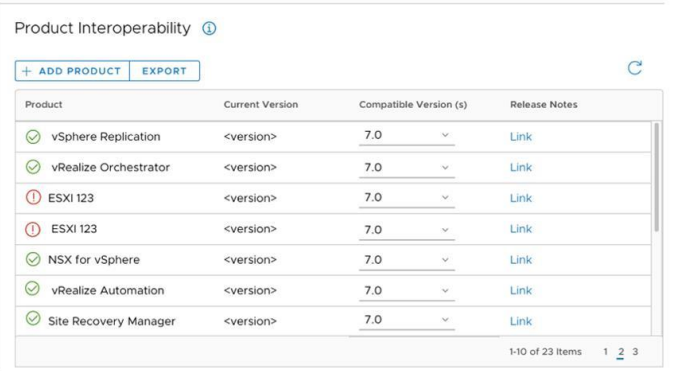

https://t.me/learningnets

10-12 Managing the vCenter Server Life Cycle

To manage the life cycle of vCenter Server, use the vCenter Server Management Interface (VAMI) to update and patch, and use the vCenter Server installer to upgrade.

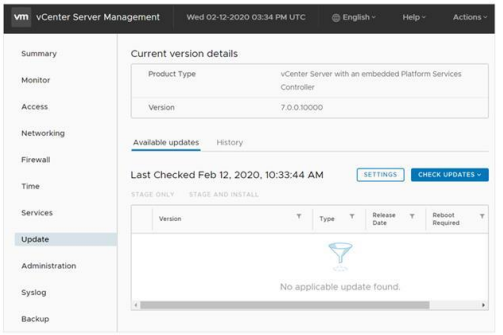

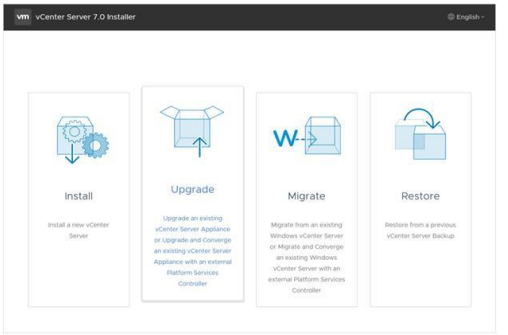

https://t.me/learningnets

10-13 Review of Learner Objectives

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

- Describe features of the vCenter Server Update Planner
- Run vCenter Server upgrade prechecks and interoperability reports
- Export prechecks and interoperability report results

https://t.me/learningnets

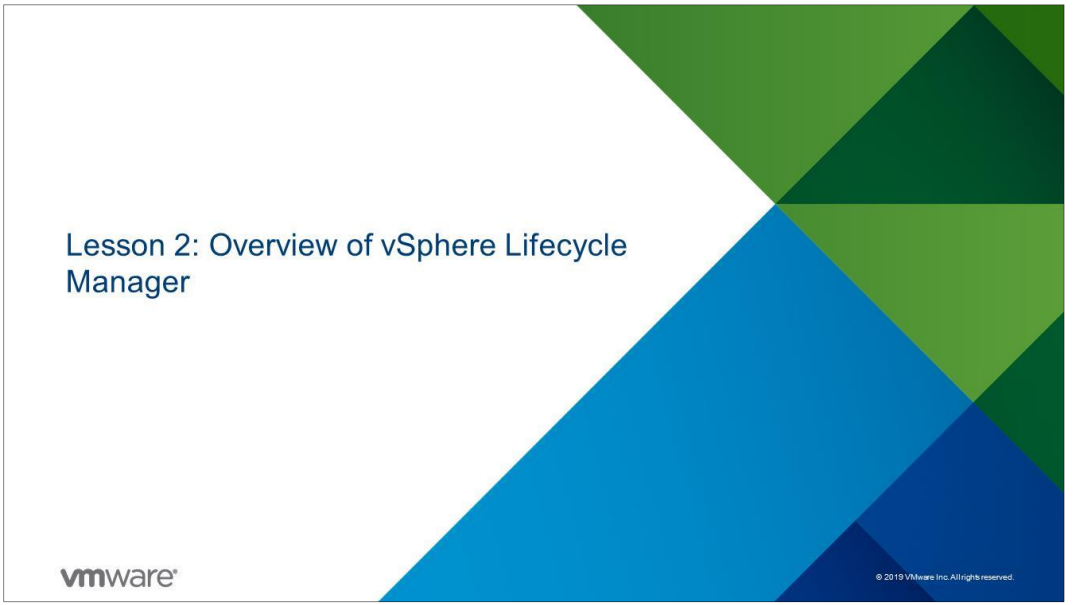

https://t.me/learningnets

10-15 Learner Objectives

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

- Recognize features of vSphere Lifecycle Manager
- Distinguish between managing hosts using baselines and managing hosts using images
- Change the patch download source

https://t.me/learningnets

10-16 Introduction to vSphere Lifecycle Manager

vSphere Lifecycle Manager centralizes automated patch and version management for clusters, ESXi, drivers and firmware, VM hardware, and VMware Tools.
vSphere Lifecycle Manager features include:

- Upgrading and patching ESXi hosts
- Installing and updating third-party software on ESXi hosts
- Standardizing images across hosts in a cluster
- Installing and updating ESXi drivers and firmware
- Managing VMware Tools and VM hardware upgrades

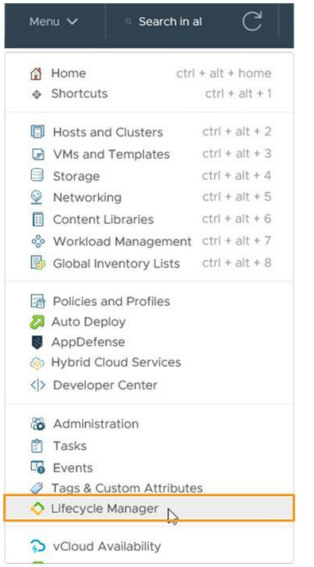

https://t.me/learningnets

10-17 Baselines and Images

vSphere Lifecycle Manager supports two methods for updating and upgrading ESXi hosts.
Only one method is supported at a time.
If you switch from managing using baselines to managing using images, you cannot switch back.

Managing Using Baselines	Managing Using Images
Compares ESXi hosts against an ESXi major version, group of patches, or set of extensions.	Compares ESXi hosts against a customized image that includes a base ESXi image, one or more add-on components, one or more vendor add-on components, firmware and drivers.
Supports all versions of ESXi from 6.5 and later.	Supports ESXi version 7.0 and later.
Baselines attach to individual ESXi hosts.	Hosts in a cluster are managed collectively, with one ESXi host image per cluster.
ESXi upgrades through ISO images	ESXi upgrades through image depots (ZIP files).
ESXi updates or patches are bundled into baselines.	ESXi updates or patches are bundled and distributed as new ESXi versions.

https://t.me/learningnets

10-18 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.

To access the vSphere Lifecycle Manager home view in the vSphere Client, select **Menu** > **Lifecycle Manager**.

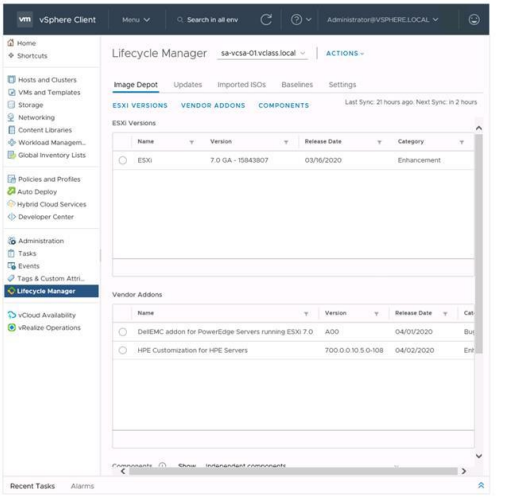

You do not require special privileges to access the vSphere Lifecycle Manager home view.

In the Lifecycle Manager pane, you can access the following tabs: **Image Depot**, **Updates**, **Imported ISOs**, **Baselines**, and **Settings**.

https://t.me/learningnets

10-19 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.

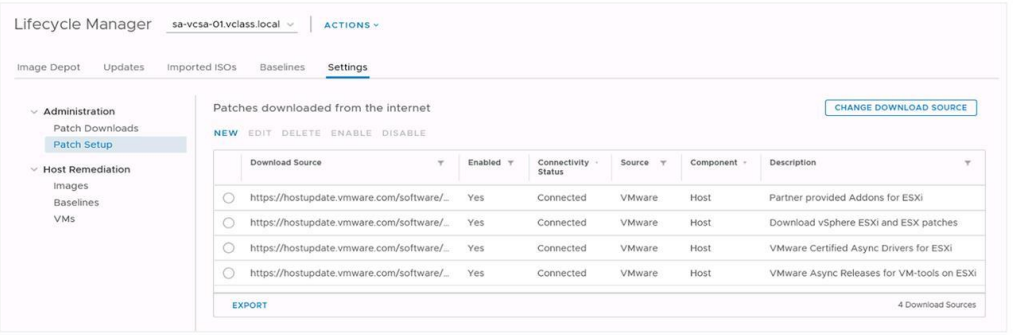

https://t.me/learningnets

10-20 vSphere Lifecycle Manager Integration with vSphere DRS

When performing remediation operations on a cluster that is enabled with vSphere DRS, vSphere Lifecycle Manager automatically integrates with vSphere DRS:

- When vSphere Lifecycle Manager places hosts into maintenance mode, vSphere DRS evacuates each host before the host is patched.
- When vSphere Lifecycle Manager attempts to place a host into maintenance mode, certain prechecks are performed to ensure that the ESXi host can enter maintenance mode.
- The vSphere Client reports any configuration issues that might prevent an ESXi host from entering maintenance mode.

https://t.me/learningnets

10-21 Review of Learner Objectives

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

- Recognize features of vSphere Lifecycle Manager
- Distinguish between managing hosts using baselines and managing hosts using images
- Change the patch download source

https://t.me/learningnets

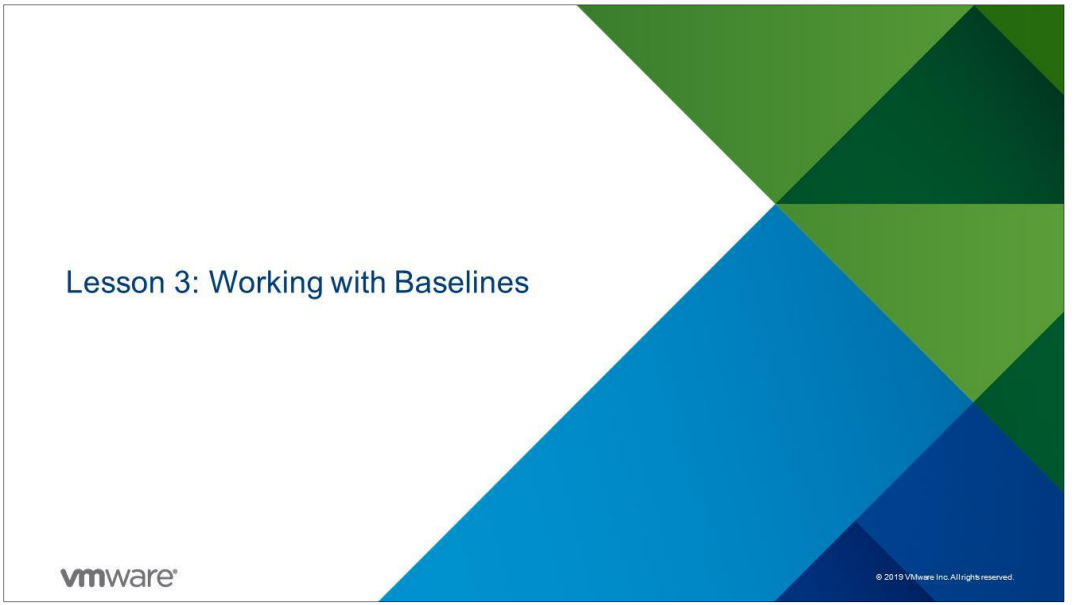

https://t.me/learningnets

10-23 Learner Objectives

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

- Identify types of baselines and baseline groups
- Recognize how to create baselines
- Describe how to update hosts using baselines

https://t.me/learningnets

10-24 Baselines and Baseline Groups

A baseline includes one or more patches, extensions, or upgrades.
vSphere Lifecycle Manager includes the following dynamic baselines by default:

- Critical Host Patches
- Non-Critical Host Patches
- Host Security Patches

A baseline group includes multiple baselines.

Baseline groups can contain one upgrade baseline and one or more patch and extension baselines.

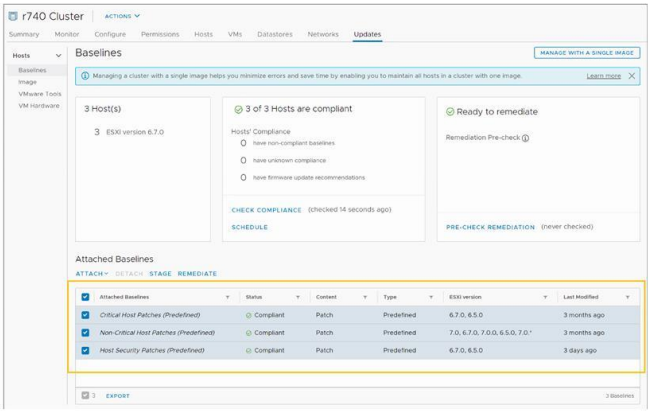

https://t.me/learningnets

10-25 Creating and Editing Patch or Extension Baselines

Using the New Baseline wizard, you can create baselines to meet the needs of your deployment:

- Fixed patch baseline: Set of patches that do not change as patch availability changes.
- Dynamic patch baseline: Set of patches that meet certain criteria.
- Host extension baseline: Contains additional software for ESXi hosts. This additional software might be VMware or third-party software.

When you create a patch or extension baseline, you can filter the patches and extensions available in the vSphere Lifecycle Manager repository to find specific patches and extensions to include in the baseline.

https://t.me/learningnets

10-26 Creating a Baseline

To create a baseline, select **Lifecycle Manager** from the **Menu** drop-down menu. Click **NEW** > **Baseline**.

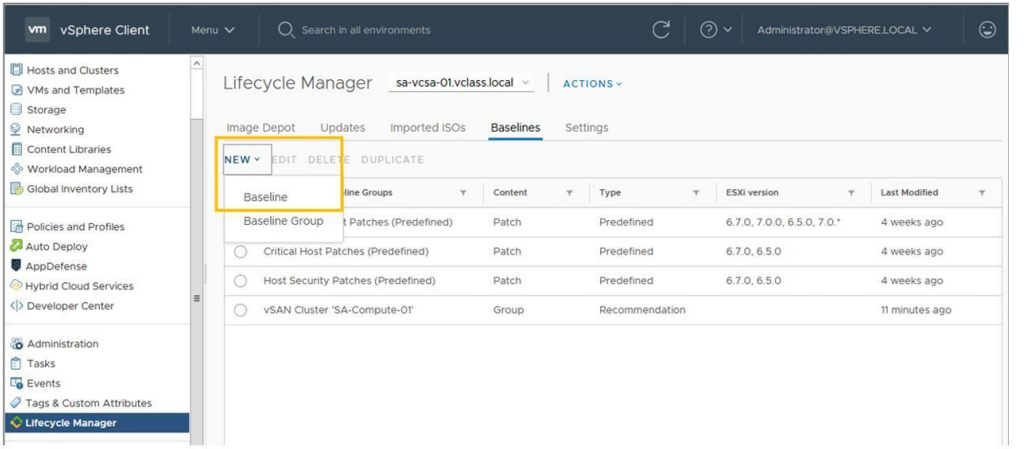

https://t.me/learningnets

10-27 Creating a Baseline: Name and Description

Provide the name, a description, the content of the baseline, and the ESXi version that this baseline applies to.

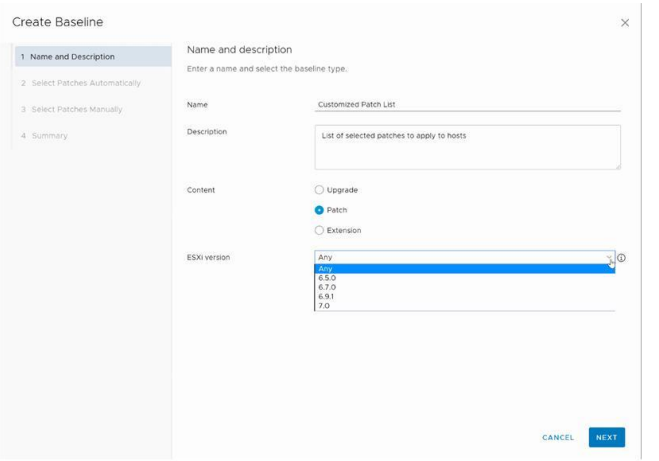

https://t.me/learningnets

10-28 Creating a Baseline: Select Patches Automatically

To create a dynamic baseline, set the criteria for adding patches to the baseline and select the check box for automatic updating of the baseline.

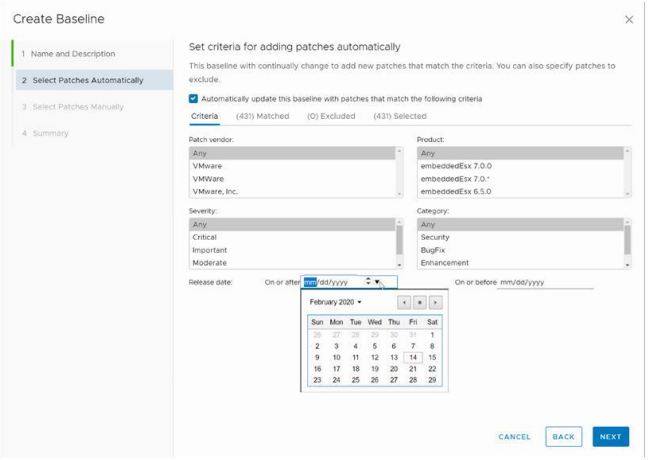

A dynamic baseline is a set of patches that meet certain criteria. The content of a dynamic baseline changes as the available patches change. You can manually exclude or add specific patches to the baseline.

https://t.me/learningnets

10-29 Creating a Baseline: Select Patches Manually

To create a fixed baseline, select the patches that you want to include in the baseline.

You must also disable the automatic updates by deselecting the check box on the Select Patches Automatically page.

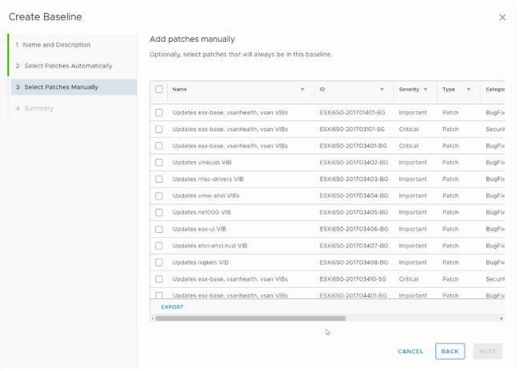

A fixed baseline is a set of patches that does not change as patch availability changes.

https://t.me/learningnets

10-30 Updating Your Host or Cluster with Baselines

Managing the life cycle of a standalone host or cluster of hosts is a five-step process:

1. Select your host or cluster and select the **Updates** tab.

   The Baselines window is the default view.
2. Attach one or more baselines.
3. Check compliance of your host or cluster with the attached baselines.
4. Perform a precheck before remediating.
5. Remediate the host or cluster.

   Optionally, stage your patches to copy them to hosts for remediation later.

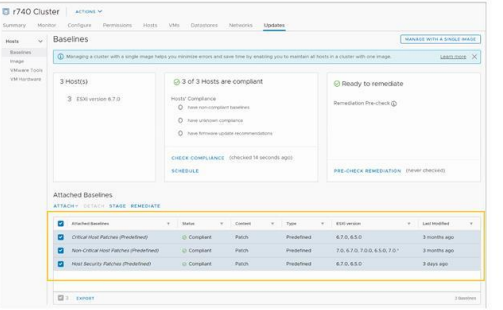

https://t.me/learningnets

10-31 Remediation Precheck

The Remediation Pre-check in vSphere Lifecycle Manager helps to verify that your remediation is successful.

vSphere Lifecycle Manager notifies you about any actions that it takes before the remediation and recommends actions for your attention.

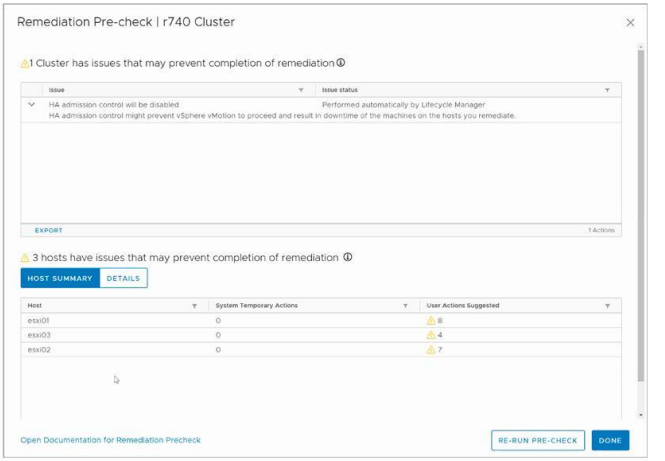

https://t.me/learningnets

10-32 Remediating Hosts

During the remediating process, the upgrades, updates, and patches from the compliance check are applied to your hosts:

- You can perform the remediation immediately or schedule it for a later date.
- Host remediation runs in different ways, depending on the types of baselines that you attach and whether the host is in a cluster.
- For ESXi hosts in a cluster, the process is sequential by default.
- The remediation of hosts in a cluster temporarily disables cluster features such as vSphere HA admission control.

https://t.me/learningnets

10-33 Review of Learner Objectives

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

- Identify types of baselines and baseline groups
- Recognize how to create baselines
- Describe how to update hosts using baselines

https://t.me/learningnets

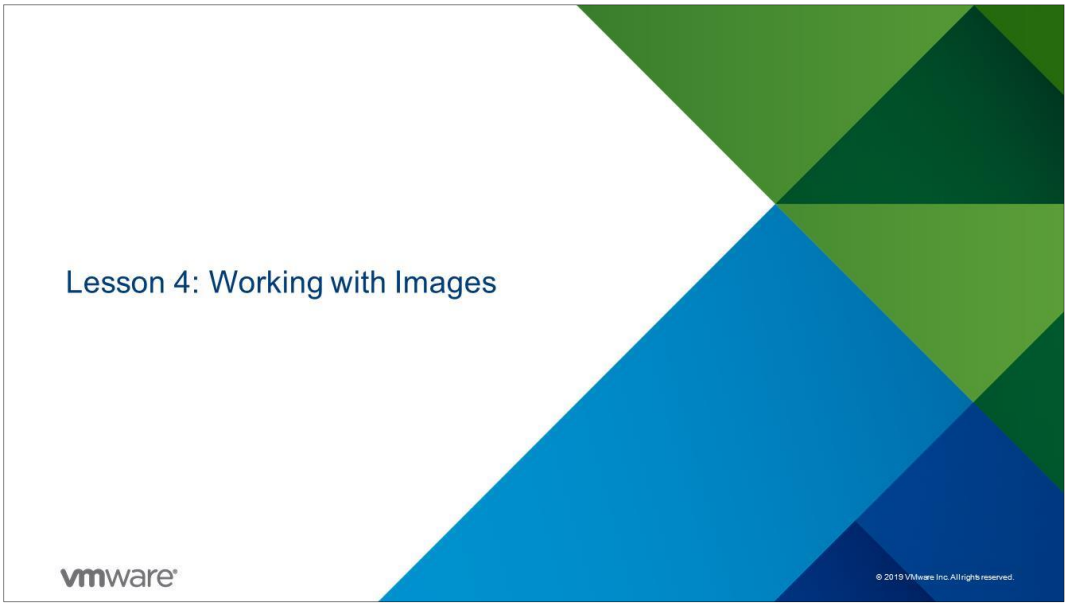

https://t.me/learningnets

10-35 Learner Objectives

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

- Describe ESXi images
- Import ESXi updates into the vSphere Client
- Enable vSphere Lifecycle Manager in a cluster
- Define a cluster image using vSphere Lifecycle Manager
- Validate ESXi host compliance against a cluster image
- Update ESXi hosts using vSphere Lifecycle Manager
- Apply a recommended image to the hosts in a cluster

https://t.me/learningnets

10-36 Elements of ESXi Images

Managing clusters with images helps to standardize the software running on your ESXi hosts.
An ESXi image consists of several elements:

- ESXi base image: An update that provides software fixes and enhancements
- Components: A logical grouping of one or more VIBs (vSphere Installation Bundles) that encapsulates a functionality in ESXi
- Vendor add-ons: Sets of components that OEMs bundle together with an ESXi base image
- Firmware and Drivers Add-On: Firmware and driver bundles that you can define for your cluster image

To maintain consistency, you apply a single ESXi image to all hosts in a cluster.

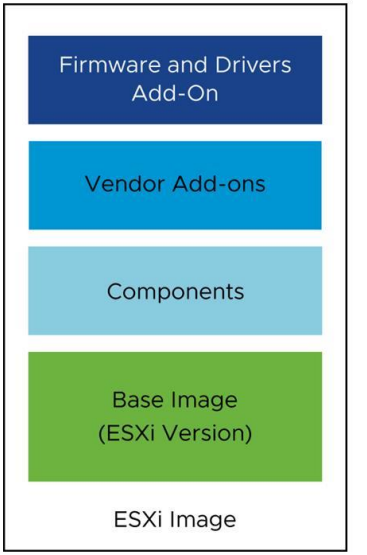

The ESXi base image is a complete ESXi installation package and is enough to start an ESXi host. Only VMware creates and releases ESXi base images.

The ESXi base image is a grouping of components. You must select at least the base image or vSphere version when creating a cluster image.

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. Components are the basic packaging for VIBs and metadata. The metadata provides the name and version of the component.

On installation, a component provides you with a visible feature. For example, vSphere HA is provided as a component. Components are optional elements to add to a cluster image.

Vendor add-ons are custom OEM images. Each add-on is a collection of components customized for a family of servers. OEMs can add, update, or remove components from a base image to create an add-on. Selecting an add-on is optional.

https://t.me/learningnets

The firmware and drivers add-on is a vendor-provided add-on. It contains the components that encapsulate firmware and driver update packages for a specific server type. 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.

https://t.me/learningnets

10-37 Image Depots

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

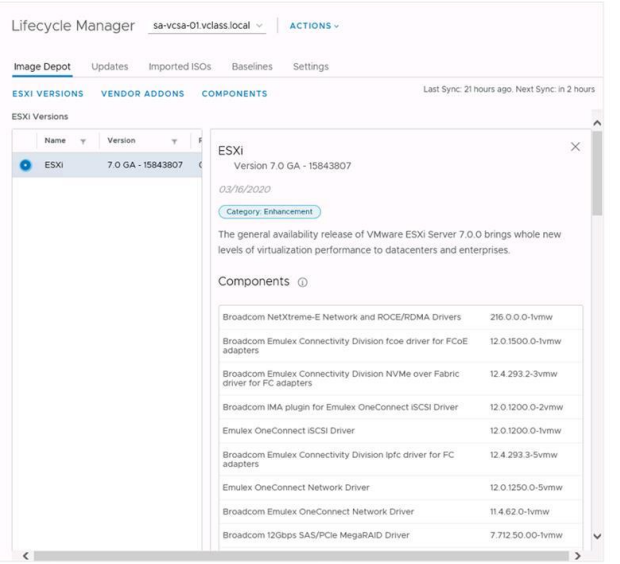

When you select a downloaded file, the details appear to the right:

- 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).

https://t.me/learningnets

10-38 Importing Updates

To use ESXi updates from a configured online depot, select **Sync Updates** from the **Actions** drop-down menu in the Lifecycle Manager pane.
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.

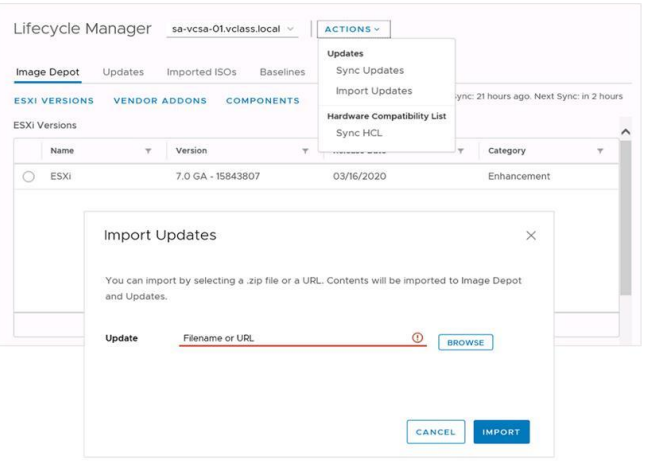

https://t.me/learningnets

10-39 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.

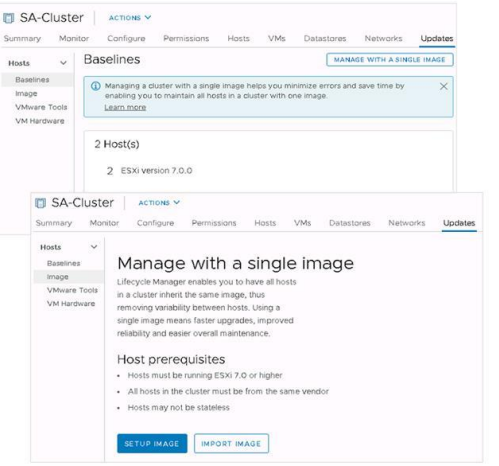

https://t.me/learningnets

10-40 Creating an ESXi Image for a New Cluster

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.

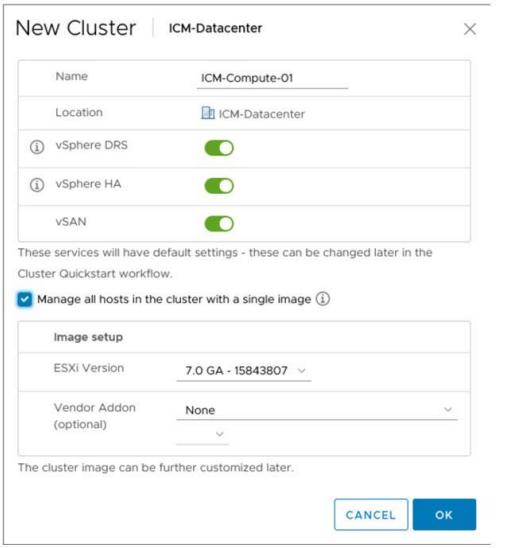

After your cluster is created, add ESXi hosts to it.

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.

https://t.me/learningnets

10-41 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.

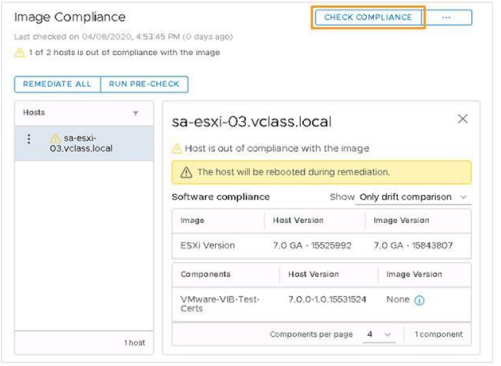

You can check the image compliance at the level of various vCenter Server objects:

- 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 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.

https://t.me/learningnets

10-42 Running a Remediation Precheck

To ensure that the cluster's health is good and that no problems occur during the remediation process of your ESXi hosts, you can perform a remediation precheck.

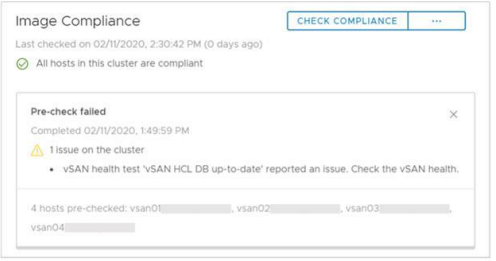

The procedure for a remediation precheck is as follows:

- In the vSphere Client, click **Hosts and Clusters** and select a cluster that is managed by an image.
- In the **Updates** tab, click **Image**.
- In the Image Compliance pane, run a remediation precheck.

https://t.me/learningnets

10-43 Hardware Compatibility

The hardware compatibility check verifies the underlying hardware of the ESXi host in the cluster against the vSAN Hardware Compatibility List (HCL).

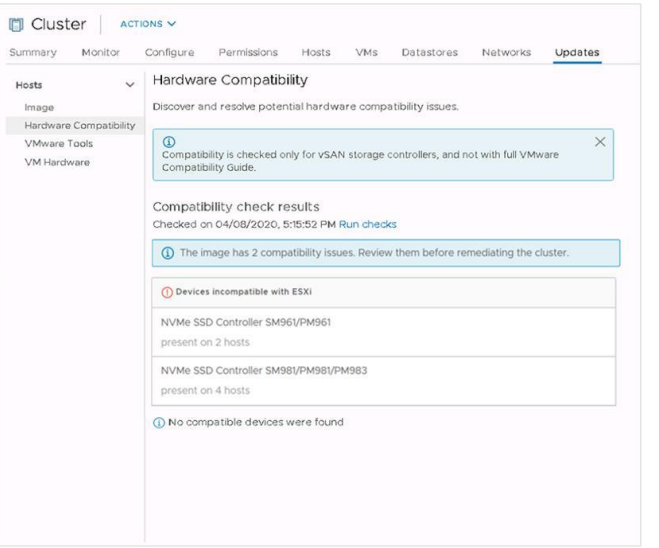

Hardware compatibility is checked only for vSAN storage controllers and not with the full VMware Compatibility Guide.

https://t.me/learningnets

10-44 Standalone VIBs

When you convert a cluster to use vSphere Lifecycle Manager, ESXi hosts are scanned.

During this scan, any VIB that is not part of an identified 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 let the update process remove the VIB from the host.

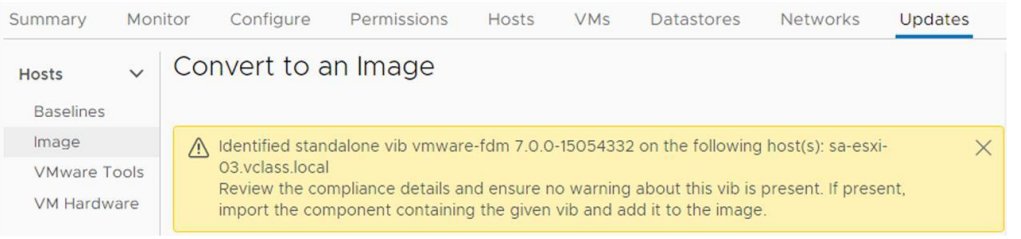

A warning about a standalone VIB does not block the process of converting the cluster to use vSphere Lifecycle Manager. If you continue to update ESXi, the VIB is uninstalled from the host as part of the process.

You cannot include standalone VIBs in a cluster image.

https://t.me/learningnets

10-45 Remediating a Cluster Against an Image

When you remediate a cluster that you manage with an image, vSphere Lifecycle Manager applies the following elements to the ESXi hosts:

- ESXi image version
- Optional: vendor addon
- Optional: firmware and driver addons
- Optional: user specified components

Remediation makes the selected hosts compliant with the desired image.

You can remediate a single ESXi host or an entire cluster, or simply pre-check hosts without updating them.

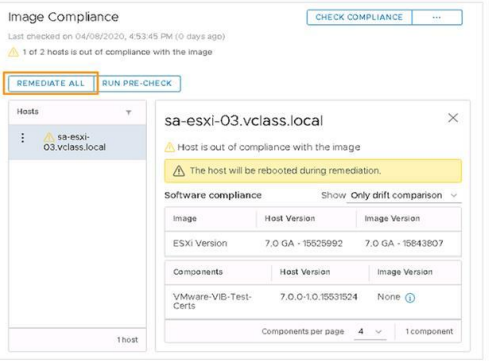

The Review Remediation Impact dialog box shows the impact summary, applicable remediation settings, End User License Agreement, and impact on specific hosts.

vSphere Lifecycle Manager performs a precheck on every remediation call. When the precheck is complete, vSphere Lifecycle Manager applies the latest saved cluster image to the hosts.

During each step of a remediation process, vSphere Lifecycle Manager 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.

https://t.me/learningnets

10-46 Reviewing Remediation Impact

The Review Remediation Impact dialog box includes the following information:

- Impact summary
- Applicable remediation settings
- End User License Agreement
- Impact to specific hosts

vSphere Lifecycle Manager performs a precheck before each remediation.

When the precheck is complete, vSphere Lifecycle Manager applies the latest saved cluster image to the hosts.

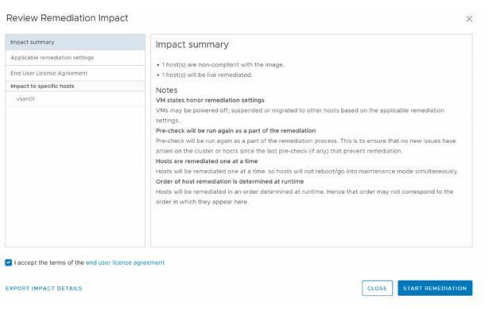

https://t.me/learningnets

10-47 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.

To check for recommended images:

1. In the **Updates** tab, select **Image** in the navigation pane.
2. Click the ellipses menu next to **EDIT** and select **Check for recommended images**.

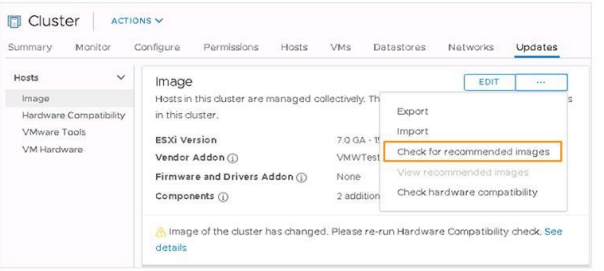

You check for image recommendations on demand and per cluster. You can check for recommendations for different clusters at the same time. When recommendation checks run concurrently with other checks, with compatibility scans, and with remediation operations, the checks are queued to run one at a time.

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 Checking for recommended images task is visible to all user sessions and cannot be canceled.

When the check completes, you can select **View recommended images**.

https://t.me/learningnets

10-48 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

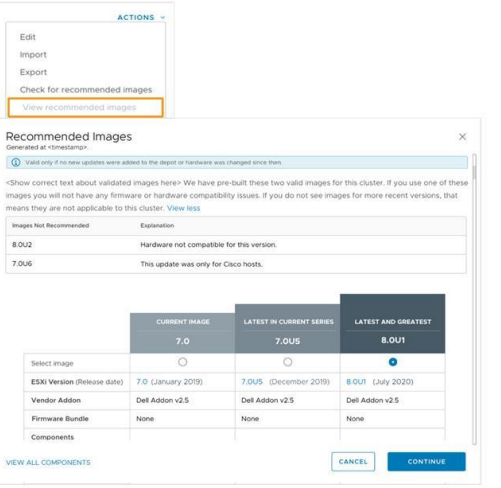

When you view recommended images, vSphere shows the following types of images:

- CURRENT IMAGE: The image specification that is being used to manage 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.

https://t.me/learningnets

vSphere might show one or more recommendations:

- If the latest release within the current series is the same as the latest major version released, only one recommendation appears.
- If the two releases are different, two recommendations appear.
- If the current image is the same as the latest release, no recommendations appear.

https://t.me/learningnets

10-49 Selecting a Recommended Image

You can select a recommended image and then validate and save it as the desired cluster image.

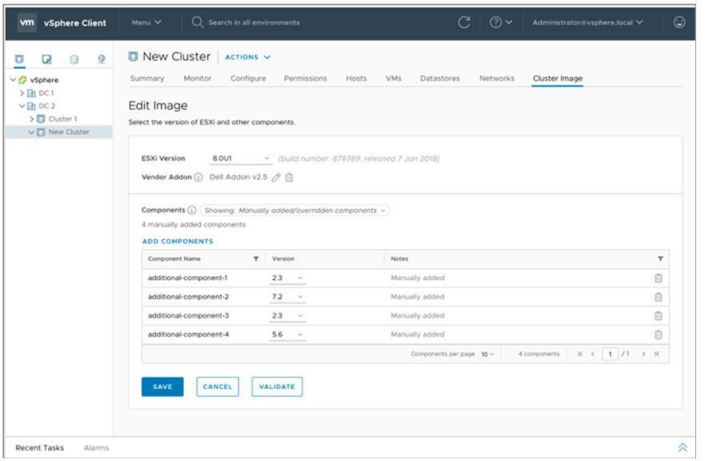

You can use a recommended image as a starting point to customize the cluster image. When you select a recommended image, the Edit Image workflow appears.

You can perform these actions:

- Add or remove image components.
- Validate and save the image.
- Scan the cluster for compatibility.
- Remediate the cluster.

https://t.me/learningnets

10-50 Customizing Cluster Images

After you start managing a cluster with an image, you can edit the image by changing, adding, or removing components, such as the ESXi image version, vendor add-ons, firmware and driver add-ons, and other components. 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

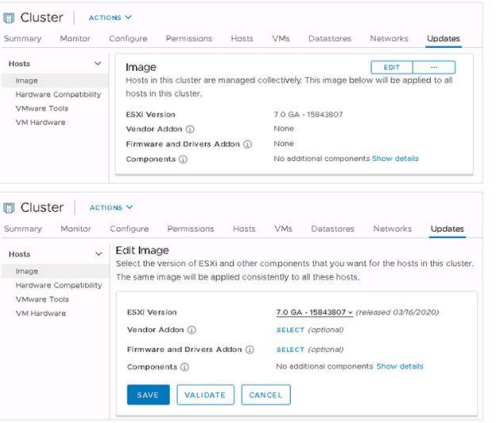

https://t.me/learningnets

10-51 Lab 27: Using vSphere Lifecycle Manager

Update ESXi hosts using vSphere Lifecycle Manager:

1. Import Update Files to the Image Depot
2. Create a Cluster with vSphere Lifecycle Manager Enabled
3. Add ESXi Hosts to the Cluster
4. Update ESXi Hosts Using the Cluster Image

https://t.me/learningnets

10-52 Review of Learner Objectives

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

- Describe ESXi images
- Import ESXi updates into the vSphere Client
- Enable vSphere Lifecycle Manager in a cluster
- Define a cluster image using vSphere Lifecycle Manager
- Validate ESXi host compliance against a cluster image
- Update ESXi hosts using vSphere Lifecycle Manager
- Apply a recommended image to the hosts in a cluster

https://t.me/learningnets

10-53 Lesson 5: Managing the Life Cycle of VMware Tools and VM Hardware

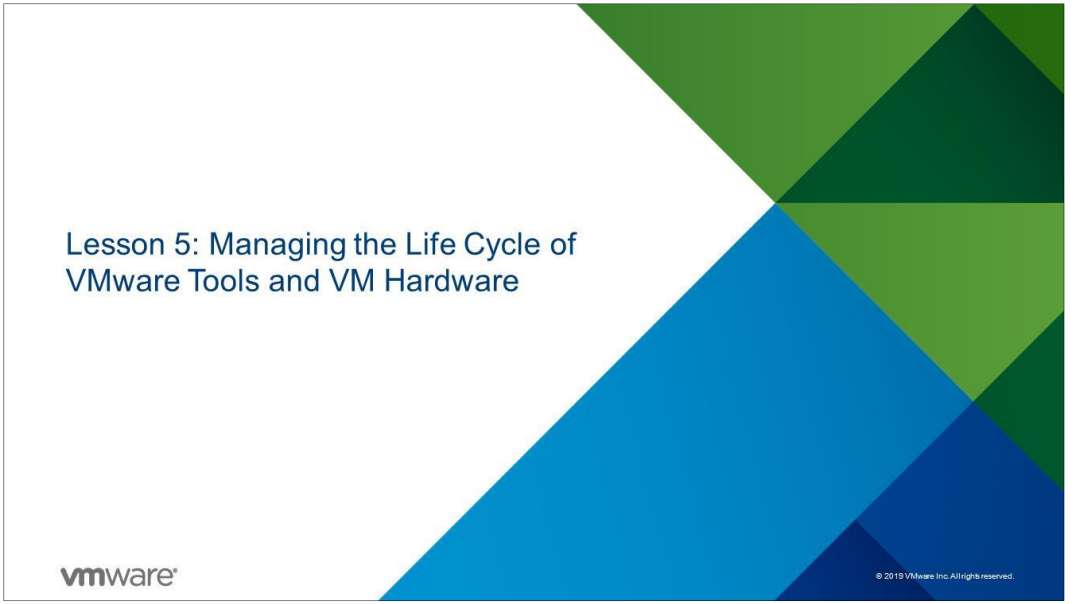

https://t.me/learningnets

10-54 Learner Objectives

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

- Use vSphere Lifecycle Manager to upgrade VMware Tools and VM hardware

https://t.me/learningnets

10-55 Keeping VMware Tools Up To Date

With each release of ESXi, VMware provides a new release of VMware Tools.
New releases include:

- Bug fixes
- Security patches
- New driver support for ESXi enhancements
- Performance enhancements for virtual devices

Keeping VMware Tools up to date is an important part of ongoing data center maintenance.

https://t.me/learningnets

10-56 Upgrading VMware Tools (1)

From a host or cluster's **Updates** tab, select VMware Tools to manage the life cycle of VMware Tools.
Step 1: Check the status of VMware Tools running in your VMs. A VM has one of the following status values:

- Upgrade Available
- Guest Managed
- Not Installed
- Unknown
- Up to Date

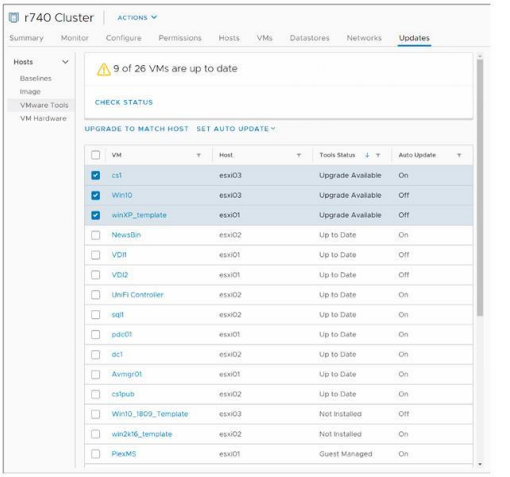

A VM has one of the following VMware Tools status values:

- Upgrade Available: You can upgrade VMware Tools to match the current version available for your ESXi hosts.
- Guest Managed: Your VM is running the Linux OpenVMTools package. Use native Linux package management tools to upgrade VMware Tools.
- Not Installed: Consider installing VMware Tools in this VM.
- Unknown: vSphere Lifecycle Manager has not yet checked the status of VMware Tools. Ensure that the VM is powered on before clicking the **CHECK STATUS** link.
- Up to Date: The version of VMware Tools running in the VM matches the latest available version for the ESXi host.

https://t.me/learningnets

10-57 Upgrading VMware Tools (2)

Select the VMs that use VMware Tools whose version you want to upgrade to a newer version.
Step 2: Click **UPGRADE TO MATCH HOST**.

1. Select the VMs to upgrade.
2. Schedule the upgrade.
   Plan the upgrade during your maintenance window.
3. Select rollback options.

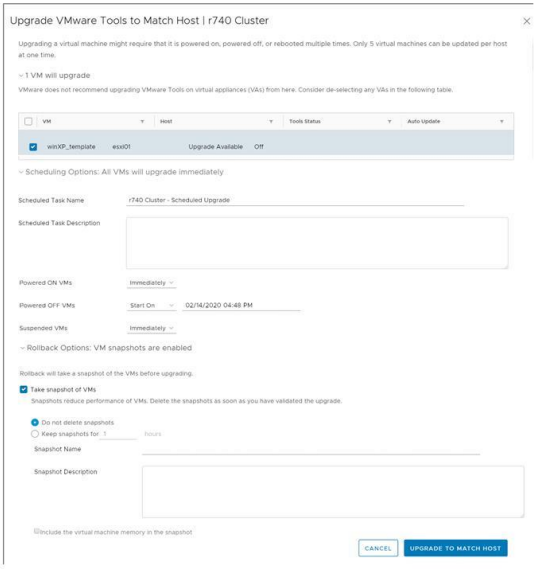

https://t.me/learningnets

10-58 Keeping VM Hardware Up To Date

With each subsequent release of ESXi, VMware provides a new release of VM hardware.

As ESXi improves its hardware support, VMware often carries that support into its VMs.
New releases include:

- Greater configuration maximums
- New types of hardware (for example, vGPU, vNVMe, vSGX, vTPM, and so on)

Consider upgrading VM hardware only when new features are required.

https://t.me/learningnets

10-59 Upgrading VM Hardware (1)

Select VM Hardware to upgrade your VMs' hardware.
Step 1: Check the status of the VM hardware running in your VMs. A VM has one of the following status values:

- Upgrade Available: You can choose to upgrade VM hardware to match the current version available for your ESXi hosts.
- Up to Date: The version of VM hardware running in the VM matches the latest available version for the ESXi host.

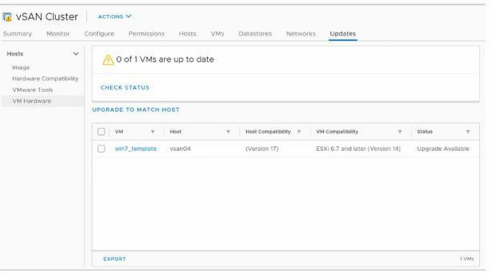

https://t.me/learningnets

10-60 Upgrading VM Hardware (2)

Select the VMs whose hardware version you want to upgrade to the latest version available on the ESXi host on which they run.
Step 2: Click **UPGRADE TO MATCH HOST**.

1. Select the VMs to upgrade.
2. Schedule the upgrade.

   Plan the upgrade during your maintenance window.
3. Select rollback options.

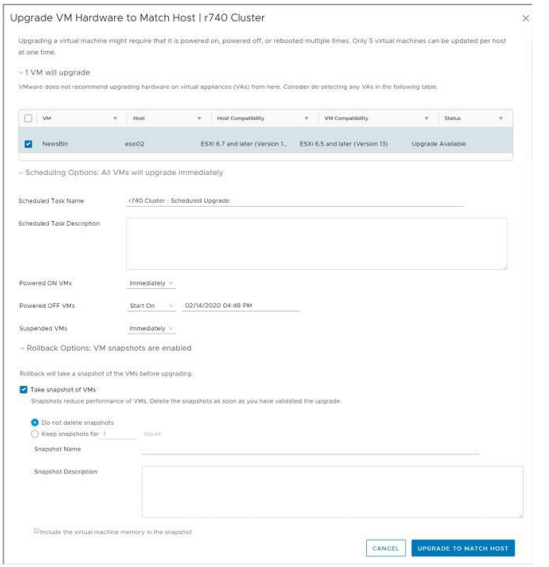

https://t.me/learningnets

10-61 Review of Learner Objectives

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

- Use vSphere Lifecycle Manager to upgrade VMware Tools and VM hardware

https://t.me/learningnets

10-62 Virtual Beans: Conclusion

By developing vSphere knowledge and skills and helping to create a modern data center at Virtual Beans, you help the company meet its expanding business demands.

Your manager recognizes your competence and assigns you as the lead vSphere administrator.

Thinking of the continuous company growth, your manager considers you for cross-training and additional responsibilities.

You think to yourself, "Now I must get that VMware certification!"

VMware Certified Professional - Data Center Virtualization:
https://www.vmware.com/education-services/certification/vcp-dcv.html

https://t.me/learningnets

10-63 Key Points

- With the Update Planner feature, you can perform prechecks to verify that your vCenter Server system meets the minimum requirements for a successful upgrade.
- vSphere Lifecycle Manager centralizes automated patch and version management for clusters, ESXi, drivers and firmware, VM hardware, and VMware Tools.
- In vSphere Lifecycle Manager, you can manage ESXi hosts by using baselines, or you can manage a cluster of ESXi hosts by using images.
- Keeping VMware Tools up to date is an important part of ongoing data center maintenance.
- Consider upgrading VM hardware only when new features are required.

Questions?

https://t.me/learningnets