VMware vSphere: Install, Configure, Manage

Lecture Manual
ESXi 6.7 and vCenter Server 6.7



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

https://t.me/learningnets

VMware vSphere:
Install, Configure, Manage
ESXi 6.7 and vCenter Server 6.7
Part Number EDU-EN-VSICM67-LECT (7/2018)
Lecture Manual

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

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

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

www.vmware.com/education

https://t.me/learningnets

Contents

Module 1 Course Introduction 1
1-2 Importance 2
1-3 Learner Objectives (1) 3
1-4 Learner Objectives (2) 4
1-5 You Are Here 5
1-6 Typographical Conventions 6
1-7 References 7
1-8 VMware Online Resources 8
1-9 VCP-Foundations Certification Alignment 9
1-10 VMware Education Overview 10
1-11 VMware Certification Overview 11

Module 2 Introduction to vSphere and the Software-Defined Data Center 13
2-2 You Are Here 14
2-3 Importance 15
2-4 Module Lessons 16
2-5 Lesson 1: Overview of vSphere and Virtual Machines 17
2-6 Learner Objectives 18
2-7 About Virtual Machines 19
2-8 Benefits of Using Virtual Machines 20
2-9 Types of Virtualization 22
2-10 About the Software-Defined Data Center 23
2-11 vSphere and Cloud Computing 24
2-12 Virtual Machine: Guest and Consumer of ESXi Host 26
2-13 Physical and Virtual Architecture 27
2-14 Physical Resource Sharing 28
2-15 CPU Virtualization 29
2-16 Physical and Virtualized Host Memory Usage 30
2-17 Physical and Virtual Networking 31
2-18 Comparison of Physical File Systems and Datastores 32
2-19 Virtual Machine Encapsulation 33
2-20 About vSphere Clients 34
2-21 About vSphere Web Client 35
2-22 About vSphere Client 36
2-23 About VMware Host Client 37
2-24 About Virtual Machine Files 38
2-25 About Virtual Machine Virtual Hardware 40
2-26 Virtual Hardware Versions 42
2-27 About CPU and Memory 43
2-28 About Virtual Storage 44
2-29 About Thick-Provisioned Virtual Disks 45
2-30 About Thin-Provisioned Virtual Disks 46

https://t.me/learningnets

2-31 About Virtual Networks . . . 48
2-32 About Virtual Network Adapters (1) . . . 49
2-33 About Virtual Network Adapters (2) . . . 51
2-34 About Miscellaneous Devices . . . 53
2-35 About the Virtual Machine Console . . . 54
2-36 Review of Learner Objectives . . . 55
2-37 Lesson 2: Shared Resources: Overview of ESXi . . . 56
2-38 Learner Objectives . . . 57
2-39 About ESXi . . . 58
2-40 Virtual Architecture . . . 60
2-41 Configuring an ESXi Host . . . 61
2-42 Configuring an ESXi Host: Root Access . . . 62
2-43 Configuring an ESXi Host: Management Network . . . 63
2-44 Configuring an ESXi Host: Other Settings . . . 64
2-45 Remote Access Settings: Security Profile . . . 65
2-46 Managing User Accounts: Best Practices . . . 66
2-47 ESXi Host as an NTP Client . . . 67
2-48 ESXi Quick Boot . . . 68
2-49 Demonstration: Installing and Configuring ESXi Hosts . . . 69
2-50 Review of Learner Objectives . . . 70
2-51 Key Points . . . 71

Module 3 Creating Virtual Machines . . . 73

3-2 You Are Here . . . 74
3-3 Importance . . . 75
3-4 Learner Objectives . . . 76
3-5 About Provisioning Virtual Machines . . . 77
3-6 Creating Virtual Machines with the New Virtual Machine Wizard (1) . . . 79
3-7 Creating Virtual Machines with the New Virtual Machine Wizard (2) . . . 80
3-8 New Virtual Machine Wizard Settings . . . 81
3-9 Installing the Guest Operating System . . . 82
3-10 Deploying OVF Templates . . . 83
3-11 About VMware Tools . . . 84
3-12 Managing VMware Tools . . . 86
3-13 VMware Tools: Supported ISO Images . . . 87
3-14 Troubleshooting an OS Installation in Virtual Machines . . . 88
3-15 Troubleshooting a VMware Tools Installation on a Guest Operating System . . . 89
3-16 Lab 1: Deploying and Configuring Virtual Machines . . . 90
3-17 Review of Learner Objectives . . . 91
3-18 Key Points . . . 92

https://t.me/learningnets

Module 4 vCenter Server 93
4-2 You Are Here 94
4-3 Importance 95
4-4 Module Lessons 96
4-5 Lesson 1: Centralized Management: Overview of vCenter Server 97
4-6 Learner Objectives 98
4-7 About the vCenter Server Management Platform 99
4-8 vCenter Server Architecture 100
4-9 vCenter Server Services 101
4-10 Monitoring the Health and Status of Services and Nodes Across vCenter Server Systems . 102
4-11 ESXi and vCenter Server Communication 103
4-12 Review of Learner Objectives 104
4-13 Lesson 2: vCenter Server Appliance 105
4-14 Learner Objectives 106
4-15 Overview of vCenter Server Appliance (1) 107
4-16 Overview of vCenter Server Appliance (2) 108
4-17 vCenter Server Appliance Scalability 109
4-18 About Platform Services Controller 110
4-19 About vCenter Single Sign-On 111
4-20 vCenter Single Sign-On Domains and Multiple Sites 112
4-21 About Enhanced Linked Mode 113
4-22 vCenter Server Deployment Options 114
4-23 Embedded Platform Services Controller 115
4-24 External vCenter Server System Linked to an Embedded vCenter Server System 116
4-25 Platform Services Controller with Load Balancer 117
4-26 vCenter Server APIs 118
4-27 High Availability for vCenter Server Appliance 119
4-28 Review of Learner Objectives 120
4-29 Lesson 3: Deploying vCenter Server Appliance 121
4-30 Learner Objectives 122
4-31 Preparing for vCenter Server Appliance Deployment 123
4-32 vCenter Server Appliance Native GUI Installer 124
4-33 vCenter Server Installation, Upgrade, Migration, and Restore 125
4-34 vCenter Server Appliance Deployment 126
4-35 Two-Stage Deployment of vCenter Server Appliance 127
4-36 vCenter Server Deployment Wizard 128
4-37 Configuring Custom Ports for a vCenter Server Appliance Deployment 129
4-38 Getting Started with vCenter Server 130
4-39 Adding License Keys to vCenter Server 131
4-40 Configuring vCenter Server Settings 132
4-41 Demonstration: Working with vCenter Server Appliance 133
4-42 Review of Learner Objectives 134
4-43 Lesson 4: Managing the vCenter Server Inventory 135

https://t.me/learningnets

4-44 Learner Objectives . . . 136
4-45 Accessing vSphere Clients . . . 137
4-46 vSphere Client Home Page . . . 138
4-47 Using the vSphere Client Navigation Pane . . . 139
4-48 vCenter Server Views: Hosts and Clusters, VMs and Templates . . . 140
4-49 vCenter Server Views: Storage and Networks . . . 141
4-50 Viewing Object Information . . . 142
4-51 Using Quick Filters in vSphere Client . . . 143
4-52 Dragging Objects in vSphere Client . . . 144
4-53 Using Pin and Unpin Functionality in vSphere Web Client . . . 145
4-54 About Data Center Objects . . . 146
4-55 Organizing Inventory Objects into Folders . . . 147
4-56 Adding a Data Center and Organizational Objects to vCenter Server . . . 149
4-57 Adding ESXi Hosts to vCenter Server . . . 150
4-58 Configuring Lockdown Mode . . . 151
4-59 Strict Lockdown Mode . . . 152
4-60 Creating Custom Tags for Inventory Objects . . . 153
4-61 Lab 2: Working with vCenter Server Appliance . . . 154
4-62 Review of Learner Objectives . . . 155
4-63 Lesson 5: vCenter Server Roles and Permissions . . . 156
4-64 Learner Objectives . . . 157
4-65 About vCenter Server Permissions . . . 158
4-66 Adding Permissions to the vCenter Server Inventory . . . 159
4-67 About Roles . . . 161
4-68 About Objects . . . 162
4-69 About Global Permissions . . . 163
4-70 Viewing Roles and Assignments . . . 164
4-71 Applying Permissions: Scenario 1 . . . 165
4-72 Applying Permissions: Scenario 2 . . . 166
4-73 Applying Permissions: Scenario 3 . . . 167
4-74 Applying Permissions: Scenario 4 . . . 168
4-75 Creating a Role . . . 169
4-76 Lab 3: Users, Groups, and Permissions . . . 170
4-77 Review of Learner Objectives . . . 171
4-78 Lesson 6: Backing Up and Restoring vCenter Server Appliance . . . 172
4-79 Learner Objectives . . . 173
4-80 Logging In to the vCenter Server Appliance Management Interface . . . 174
4-81 vCenter Server Appliance Management Interface Functions . . . 175
4-82 About Native vCenter Server Backup and Restore . . . 176
4-83 Methods for vCenter Server Appliance Backup and Restore . . . 177
4-84 File-Based Backup of vCenter Server Appliance . . . 178
4-85 File-Based Restore of vCenter Server Appliance . . . 179
4-86 Scheduling Backups . . . 180

https://t.me/learningnets

4-87 Creating a Backup Schedule . . . 181
4-88 Viewing the Backup Schedule . . . 182
4-89 Review of Learner Objectives . . . 183
4-90 Lesson 7: Monitoring vCenter Server Appliance . . . 184
4-91 Learner Objectives . . . 185
4-92 vCenter Server Events . . . 186
4-93 vCenter Server System Logs . . . 187
4-94 About Log Levels . . . 188
4-95 Setting Log Levels . . . 189
4-96 Forwarding vCenter Server Appliance Log Files to Another Host Machine . . . 190
4-97 vCenter Server Database Health . . . 191
4-98 Managing the vCenter Server Services . . . 192
4-99 Monitoring vCenter Server Appliance . . . 193
4-100 Monitoring vCenter Server Appliance Services . . . 194
4-101 vSphere Alarms for Detecting Resource Exhaustion and Service Failure . . . 195
4-102 Monthly Patch Updates for vCenter Server Appliance . . . 196
4-103 Review of Learner Objectives . . . 197
4-104 Key Points . . . 198

Module 5 Configuring and Managing Virtual Networks . . . 199

5-2 You Are Here . . . 200
5-3 Importance . . . 201
5-4 Module Lessons . . . 202
5-5 Lesson 1: Introduction to vSphere Standard Switches . . . 203
5-6 Learner Objectives . . . 204
5-7 Types of Virtual Switch Connections . . . 205
5-8 Adding ESXi Networking . . . 206
5-9 Virtual Switch Connection Examples . . . 207
5-10 About Standard Switches . . . 208
5-11 Viewing the Configuration of Standard Switches . . . 209
5-12 About VLANs . . . 210
5-13 Network Adapter Properties . . . 211
5-14 Types of Virtual Switches . . . 212
5-15 Distributed Switch Architecture . . . 213
5-16 Standard Switch and Distributed Switch Feature Comparison . . . 214
5-17 Review of Learner Objectives . . . 215
5-18 Lesson 2: Configuring Standard Switch Policies . . . 216
5-19 Learner Objectives . . . 217
5-20 Network Switch and Port Policies . . . 218
5-21 Configuring Security Policies . . . 219
5-22 Traffic-Shaping Policies . . . 220
5-23 Configuring Traffic Shaping . . . 221
5-24 NIC Teaming and Failover Policies . . . 222

https://t.me/learningnets

5-25 Load-Balancing Method: Originating Virtual Port ID . . . 224
5-26 Load-Balancing Method: Source MAC Hash . . . 225
5-27 Load-Balancing Method: Source and Destination IP Hash . . . 227
5-28 Detecting and Handling Network Failure . . . 229
5-29 Physical Network Considerations . . . 230
5-30 Lab 4: Using Standard Switches . . . 231
5-31 Review of Learner Objectives . . . 232
5-32 Key Points . . . 233

Module 6 Configuring and Managing Virtual Storage . . . 235

6-2 You Are Here . . . 236
6-3 Importance . . . 237
6-4 Module Lessons . . . 238
6-5 Lesson 1: Storage Concepts . . . 239
6-6 Learner Objectives . . . 240
6-7 About Datastores . . . 241
6-8 Basic Storage Overview . . . 242
6-9 Storage Protocol Overview . . . 244
6-10 About VMFS . . . 246
6-11 About NFS . . . 248
6-12 vSAN Overview . . . 249
6-13 About vSphere Virtual Volumes . . . 250
6-14 About Raw Device Mapping . . . 251
6-15 Physical Storage Considerations . . . 252
6-16 Review of Learner Objectives . . . 253
6-17 Lesson 2: Fibre Channel Storage . . . 254
6-18 Learner Objectives . . . 255
6-19 About Fibre Channel . . . 256
6-20 Fibre Channel SAN Components . . . 257
6-21 Fibre Channel Addressing and Access Control . . . 258
6-22 Multipathing with Fibre Channel . . . 259
6-23 FCoE Adapters . . . 260
6-24 Configuring Software FCoE: Creating VMkernel Ports . . . 261
6-25 Configuring Software FCoE: Activating Software FCoE Adapters . . . 262
6-26 Multipathing with Software FCoE . . . 263
6-27 Review of Learner Objectives . . . 264
6-28 Lesson 3: iSCSI Storage . . . 265
6-29 Learner Objectives . . . 266
6-30 iSCSI Components . . . 267
6-31 iSCSI Addressing . . . 268
6-32 Storage Device Naming Conventions . . . 269
6-33 iSCSI Adapters . . . 270
6-34 Setting Up iSCSI Adapters . . . 271

https://t.me/learningnets

6-35 ESXi Network Configuration for IP Storage . . . 272
6-36 Activating the Software iSCSI Adapter . . . 273
6-37 Discovering iSCSI Targets . . . 274
6-38 iSCSI Security: CHAP . . . 275
6-39 Multipathing with iSCSI Storage . . . 277
6-40 Lab 5: Accessing iSCSI Storage . . . 278
6-41 Review of Learner Objectives . . . 279
6-42 Lesson 4: VMFS Datastores . . . 280
6-43 Learner Objectives . . . 281
6-44 Using Datastores with ESXi Hosts . . . 282
6-45 Creating and Viewing VMFS Datastores . . . 283
6-46 Browsing Datastore Contents . . . 284
6-47 Managing Overcommitted Datastores . . . 285
6-48 Increasing the Size of VMFS Datastores . . . 286
6-49 Deleting or Unmounting VMFS Datastores . . . 287
6-50 Multipathing Algorithms . . . 288
6-51 Configuring Storage Load Balancing . . . 289
6-52 Lab 6: Managing VMFS Datastores . . . 290
6-53 Review of Learner Objectives . . . 291
6-54 Lesson 5: NFS Datastores . . . 292
6-55 Learner Objectives . . . 293
6-56 NFS Components . . . 294
6-57 Configuring NFS Datastores . . . 295
6-58 NFS v3 and NFS v4.1 . . . 296
6-59 NFS Version Compatibility with Other vSphere Technologies . . . 297
6-60 NFS Datastore Name and Configuration . . . 298
6-61 Configuring ESXi Host Authentication and NFS Kerberos Credentials . . . 299
6-62 Configuring Datastores to Use Kerberos . . . 300
6-63 Viewing IP Storage Information . . . 301
6-64 Unmounting NFS Datastores . . . 302
6-65 Multipathing and NFS Storage . . . 303
6-66 Enabling Session Trunking for NFS 4.1 . . . 305
6-67 Lab 7: Accessing NFS Storage . . . 306
6-68 Review of Learner Objectives . . . 307
6-69 Lesson 6: vSAN Datastores . . . 308
6-70 Learner Objectives . . . 309
6-71 About vSAN . . . 310
6-72 vSAN Requirements . . . 312
6-73 Configuring vSAN . . . 313
6-74 Disk Groups . . . 314
6-75 Viewing the vSAN Cluster Summary . . . 315
6-76 Using vSAN . . . 316
6-77 Objects in vSAN Datastores . . . 317

https://t.me/learningnets

6-78 Virtual Machine Storage Policies . . . 318
6-79 Configuring Virtual Machine Storage Policies . . . 319
6-80 Viewing Virtual Machine vSAN Datastores . . . 320
6-81 Review of Learner Objectives . . . 321
6-82 Key Points . . . 322

Module 7 Virtual Machine Management . . . 323
7-2 You Are Here . . . 324
7-3 Importance . . . 325
7-4 Module Lessons . . . 326
7-5 Lesson 1: Creating Templates and Clones . . . 327
7-6 Learner Objectives . . . 328
7-7 Using Templates . . . 329
7-8 Creating Templates . . . 330
7-9 Deploying Virtual Machines from a Template . . . 331
7-10 Updating Templates . . . 332
7-11 Cloning Virtual Machines . . . 333
7-12 Customizing the Guest Operating System . . . 334
7-13 About Instant Clones . . . 336
7-14 Removing a Virtual Machine . . . 337
7-15 Review of Learner Objectives . . . 338
7-16 Lesson 2: Working with Content Libraries . . . 339
7-17 Learner Objectives . . . 340
7-18 About Content Libraries . . . 341
7-19 Adding Templates to a Content Library . . . 342
7-20 Deploying VMs from Templates in a Content Library . . . 343
7-21 Benefits of Content Libraries . . . 344
7-22 Types of Content Libraries . . . 346
7-23 Labs . . . 347
7-24 Lab 8: Using Templates and Clones . . . 348
7-25 Lab 9: Using Content Libraries . . . 349
7-26 Review of Learner Objectives . . . 350
7-27 Lesson 3: Modifying Virtual Machines . . . 351
7-28 Learner Objectives . . . 352
7-29 Modifying Virtual Machine Settings . . . 353
7-30 Hot-Pluggable Devices . . . 354
7-31 Creating Raw Device Mapping Files . . . 356
7-32 Dynamically Increasing Virtual Disk Size . . . 357
7-33 Inflating Thin-Provisioned Disks . . . 358
7-34 Virtual Machine Options . . . 359
7-35 VMware Tools Options . . . 360
7-36 Boot Options . . . 361
7-37 Lab 10: Modifying Virtual Machines . . . 362

https://t.me/learningnets

7-38 Review of Learner Objectives . . . 363
7-39 Lesson 4: Migrating Virtual Machines . . . 364
7-40 Learner Objectives . . . 365
7-41 Migrating Virtual Machines . . . 366
7-42 Comparing Migration Types . . . 367
7-43 vSphere vMotion Migration . . . 368
7-44 vSphere vMotion Migration Workflow . . . 370
7-45 vSphere vMotion Migration Requirements . . . 371
7-46 Host Requirements for vSphere vMotion Migration . . . 373
7-47 CPU Constraints on vSphere vMotion Migration . . . 374
7-48 Other Cluster Settings: EVC for vSphere vMotion Migration . . . 375
7-49 CPU Baselines for an EVC Cluster . . . 376
7-50 EVC Cluster Requirements . . . 377
7-51 Virtual Machine EVC . . . 378
7-52 Change CPU Identification Mask Settings . . . 379
7-53 Checking vSphere vMotion Errors . . . 380
7-54 Migration with vSphere Storage vMotion . . . 381
7-55 vSphere Storage vMotion in Action . . . 382
7-56 vSphere Storage vMotion Guidelines and Limitations . . . 383
7-57 Shared-Nothing vSphere vMotion Migration . . . 384
7-58 Cross-vCenter migrations . . . 386
7-59 VMkernel Networking Layer and TCP/IP Stacks . . . 387
7-60 vSphere vMotion TCP/IP Stacks . . . 389
7-61 Long-Distance vSphere vMotion Migration . . . 390
7-62 Networking Prerequisite for Long-Distance vSphere vMotion Migration . . . 391
7-63 Network Checks for Migrations Between vCenter Server Instances . . . 392
7-64 Encrypted vSphere vMotion . . . 393
7-65 Lab 11: Migrating Virtual Machines . . . 394
7-66 Review of Learner Objectives . . . 395
7-67 Lesson 5: Creating Virtual Machine Snapshots . . . 396
7-68 Learner Objectives . . . 397
7-69 Virtual Machine Snapshots . . . 398
7-70 Taking Snapshots . . . 399
7-71 Types of Snapshots . . . 400
7-72 Virtual Machine Snapshot Files (1) . . . 401
7-73 Virtual Machine Snapshot Files (2) . . . 403
7-74 Managing Snapshots . . . 404
7-75 Deleting Virtual Machine Snapshots (1) . . . 405
7-76 Deleting Virtual Machine Snapshots (2) . . . 406
7-77 Deleting Virtual Machine Snapshots (3) . . . 407
7-78 Deleting All Virtual Machine Snapshots . . . 408
7-79 About Snapshot Consolidation . . . 409
7-80 Discovering When to Consolidate Snapshots . . . 410

https://t.me/learningnets

7-81 Performing Snapshot Consolidation . . . 411
7-82 Setting the Maximum Number of Snapshots Allowed for a VM . . . 412
7-83 Lab 12: Managing Virtual Machines . . . 413
7-84 Review of Learner Objectives . . . 414
7-85 Key Points . . . 415

Module 8 Resource Management and Monitoring . . . 417
8-2 You Are Here . . . 418
8-3 Importance . . . 419
8-4 Module Lessons . . . 420
8-5 Lesson 1: Virtual CPU and Memory Concepts . . . 421
8-6 Learner Objectives . . . 422
8-7 Memory Virtualization Basics . . . 423
8-8 Virtual Machine Memory Overcommitment . . . 424
8-9 Memory Reclamation Techniques . . . 425
8-10 Configuring Multicore Virtual CPUs . . . 427
8-11 About Hyperthreading . . . 428
8-12 CPU Load Balancing . . . 429
8-13 Review of Learner Objectives . . . 430
8-14 Lesson 2: Resource Controls and Resource Pools . . . 431
8-15 Learner Objectives . . . 432
8-16 Shares, Limits, and Reservations . . . 433
8-17 Defining Resource Settings for Individual VMs . . . 434
8-18 Resource Sharing by VMs . . . 435
8-19 About Resource Pools . . . 436
8-20 Resource Pool Attributes . . . 437
8-21 Reasons to Use Resource Pools . . . 438
8-22 Resource Pool Example . . . 439
8-23 Resource Pool Example: CPU Shares . . . 440
8-24 Resource Pool Example: CPU Contention . . . 441
8-25 About Expandable Reservations . . . 442
8-26 Example of Expandable Reservation (1) . . . 443
8-27 Example of Expandable Reservation (2) . . . 444
8-28 Admission Control for CPU and Memory Reservations . . . 445
8-29 Resource Pool Summary Tab . . . 446
8-30 Resource Reservation Tab . . . 447
8-31 Scheduling Changes to Resource Settings . . . 448
8-32 Managing Virtual Machines with vApps . . . 449
8-33 vApp Characteristics . . . 450
8-34 Exporting and Deploying vApps . . . 451
8-35 Lab 13: Managing Resource Pools . . . 452
8-36 Review of Learner Objectives . . . 453
8-37 Lesson 3: Monitoring Resource Use . . . 454

https://t.me/learningnets

8-38 Learner Objectives . . . 455
8-39 Performance-Tuning Methodology . . . 456
8-40 Resource-Monitoring Tools . . . 457
8-41 Guest Operating System Monitoring Tools . . . 458
8-42 Using Perfmon to Monitor VM Resources . . . 459
8-43 Using esxtop to Monitor VM Resources . . . 460
8-44 Monitoring Inventory Objects with Performance Charts . . . 461
8-45 Working with Overview Performance Charts . . . 462
8-46 Working with Advanced Performance Charts . . . 463
8-47 Chart Options: Real-Time and Historical . . . 464
8-48 About Chart Types . . . 465
8-49 Saving Charts . . . 466
8-50 About Objects and Counters . . . 467
8-51 About Statistics Types . . . 468
8-52 About Rollup . . . 469
8-53 Interpreting Data from Tools . . . 471
8-54 CPU-Constrained VMs . . . 472
8-55 Memory-Constrained VMs . . . 473
8-56 Memory-Constrained Hosts . . . 474
8-57 Monitoring VM Active Memory . . . 475
8-58 Disk-Constrained VMs . . . 476
8-59 Monitoring Disk Latency . . . 477
8-60 Network-Constrained VMs . . . 478
8-61 Lab 14: Monitoring Virtual Machine Performance . . . 479
8-62 Review of Learner Objectives . . . 480
8-63 Lesson 4: Using Alarms . . . 481
8-64 Learner Objectives . . . 482
8-65 About Alarms . . . 483
8-66 Alarm Settings . . . 484
8-67 Alarm Triggers . . . 485
8-68 Configuring Condition or State Triggers . . . 486
8-69 Configuring Event Triggers . . . 487
8-70 Configuring Actions . . . 488
8-71 Configuring vCenter Server Notifications . . . 490
8-72 Viewing and Acknowledging Triggered Alarms . . . 491
8-73 Lab 15: Using Alarms . . . 492
8-74 Review of Learner Objectives . . . 493
8-75 Key Points . . . 494

Module 9 vSphere HA, vSphere Fault Tolerance, and Protecting Data . 495

9-2 You Are Here . . . 496
9-3 Importance . . . 497
9-4 Module Lessons . . . 498

https://t.me/learningnets

9-5 Lesson 1: Introduction to vSphere HA . . . 499
9-6 Learner Objectives . . . 500
9-7 Protection at Every Level . . . 501
9-8 vCenter Server Availability: Recommendations . . . 503
9-9 About vSphere HA . . . 504
9-10 vSphere HA Scenario: ESXi Host Failure . . . 505
9-11 vSphere HA Scenario: Guest Operating System Failure . . . 506
9-12 vSphere HA Scenario: Application Failure . . . 507
9-13 vSphere HA Scenario: Datastore Accessibility Failures . . . 508
9-14 vSphere HA Scenario: Protecting VMs Against Network Isolation . . . 509
9-15 Importance of Redundant Heartbeat Networks . . . 510
9-16 Redundancy Using NIC Teaming . . . 511
9-17 Redundancy Using Additional Networks . . . 512
9-18 Review of Learner Objectives . . . 513
9-19 Lesson 2: vSphere HA Architecture . . . 514
9-20 Learner Objectives . . . 515
9-21 vSphere HA Architecture: Agent Communication . . . 516
9-22 vSphere HA Architecture: Network Heartbeats . . . 518
9-23 vSphere HA Architecture: Datastore Heartbeats . . . 519
9-24 vSphere HA Failure Scenarios . . . 520
9-25 Failed Slave Hosts . . . 521
9-26 Failed Master Hosts . . . 522
9-27 Isolated Hosts . . . 523
9-28 Design Considerations . . . 524
9-29 Virtual Machine Storage Failures . . . 525
9-30 About VMCP . . . 526
9-31 Review of Learner Objectives . . . 527
9-32 Lesson 3: Configuring vSphere HA . . . 528
9-33 Learner Objectives . . . 529
9-34 About Clusters . . . 530
9-35 vSphere HA Prerequisites . . . 531
9-36 Configuring vSphere HA Settings . . . 532
9-37 vSphere HA Settings: Failure and Responses . . . 533
9-38 vSphere HA Settings: Virtual Machine Monitoring . . . 534
9-39 vSphere HA Settings: Heartbeat Datastores . . . 535
9-40 vSphere HA Settings: Admission Control . . . 536
9-41 Example: Admission Control Using Cluster Resources Percentage . . . 537
9-42 Example: Admission Control Using Slots . . . 538
9-43 vSphere HA Settings: Performance Degradation VMs Tolerate . . . 539
9-44 vSphere HA Settings: Advanced Options . . . 540
9-45 About vSphere HA Orchestrated Restart . . . 541
9-46 Configuring vSphere HA Orchestrated Restart . . . 542
9-47 VM Dependencies in Orchestrated Restart . . . 543

https://t.me/learningnets

9-48 Network Configuration and Maintenance . . . 544
9-49 Cluster Resource Reservation . . . 545
9-50 Monitoring Cluster Status . . . 546
9-51 Lab 16: Using vSphere HA . . . 547
9-52 Review of Learner Objectives . . . 548
9-53 Lesson 4: Introduction to vSphere Fault Tolerance . . . 549
9-54 Learner Objectives . . . 550
9-55 About vSphere Fault Tolerance . . . 551
9-56 vSphere Fault Tolerance Features . . . 552
9-57 vSphere Fault Tolerance with vSphere HA and vSphere DRS . . . 553
9-58 Redundant VMDK Files . . . 554
9-59 vSphere Fault Tolerance Checkpoint . . . 555
9-60 vSphere Fault Tolerance: Precopy . . . 556
9-61 vSphere Fault Tolerance Fast Checkpointing . . . 557
9-62 vSphere Fault Tolerance Shared Files . . . 558
9-63 About the shared.vmft File . . . 559
9-64 Enabling vSphere Fault Tolerance on a Virtual Machine . . . 560
9-65 Review of Learner Objectives . . . 561
9-66 Lesson 5: vSphere Replication . . . 562
9-67 Learner Objectives . . . 563
9-68 About vSphere Replication . . . 564
9-69 About the vSphere Replication Appliance . . . 565
9-70 Replication Functions . . . 566
9-71 Deploying the vSphere Replication Virtual Appliance . . . 567
9-72 Configuring vSphere Replication for a Single VM . . . 568
9-73 Configuring Multiple Points in Time and the Recovery Point Objective . . . 569
9-74 Recovering Virtual Machines . . . 571
9-75 vSphere Storage APIs - Data Protection . . . 572
9-76 Offloaded Backup Processing . . . 573
9-77 About Changed-Block Tracking . . . 574
9-78 Review of Learner Objectives . . . 575
9-79 Key Points . . . 576

Module 10 vSphere DRS . . . 577

10-2 You Are Here . . . 578
10-3 Importance . . . 579
10-4 Learner Objectives . . . 580
10-5 About vSphere DRS . . . 581
10-6 vSphere DRS Cluster Prerequisites . . . 582
10-7 vSphere DRS Cluster Settings: Automation Level . . . 583
10-8 vSphere DRS Cluster Settings: Predictive DRS . . . 585
10-9 Other Cluster Settings: Swap File Location . . . 586
10-10 vSphere DRS Cluster Settings: Virtual Machine Affinity . . . 587

https://t.me/learningnets

10-11 vSphere DRS Cluster Settings: DRS Groups . . . 588
10-12 vSphere DRS Cluster Settings: VM-Host Affinity Rules . . . 589
10-13 VM-Host Affinity Rule: Preferential . . . 590
10-14 VM-Host Affinity Rule: Required . . . 591
10-15 vSphere DRS Cluster Settings: VM-Level Automation . . . 592
10-16 Adding a Host to a Cluster . . . 593
10-17 Viewing vSphere DRS Cluster Information . . . 594
10-18 Viewing vSphere DRS Recommendations . . . 595
10-19 Monitoring Cluster Status . . . 596
10-20 Maintenance Mode and Standby Mode . . . 597
10-21 Removing a Host from the vSphere DRS Cluster . . . 598
10-22 Disabling vSphere DRS and Restoring a Resource Pool Tree . . . 599
10-23 Improving Virtual Machine Performance Methods . . . 600
10-24 Using vSphere HA with vSphere DRS . . . 601
10-25 Lab 17: Implementing vSphere DRS Clusters . . . 602
10-26 Review of Learner Objectives . . . 603
10-27 Key Points . . . 604

Module 11 vSphere Update Manager . . . 605

11-2 You Are Here . . . 606
11-3 Importance . . . 607
11-4 Learner Objectives . . . 608
11-5 About vSphere Update Manager . . . 609
11-6 vSphere Update Manager Capabilities . . . 610
11-7 vSphere Update Manager Components . . . 611
11-8 Configuring vSphere Update Manager . . . 612
11-9 Baselines and Baseline Groups . . . 614
11-10 Creating and Editing Patch or Extension Baselines . . . 615
11-11 Attaching Baselines . . . 616
11-12 Scanning for Updates . . . 617
11-13 Reviewing vSphere Object Compliance . . . 618
11-14 Remediating Objects . . . 619
11-15 Patch Recall Notification . . . 621
11-16 EAM Integration with vSphere Update Manager . . . 622
11-17 vSAN Integration with vSphere Update Manager . . . 623
11-18 About System Baselines . . . 624
11-19 vSphere Update Manager Integration with vSphere DRS . . . 625
11-20 Lab 18: Using vSphere Update Manager . . . 626
11-21 Review of Learner Objectives . . . 627
11-22 Key Points . . . 628

https://t.me/learningnets

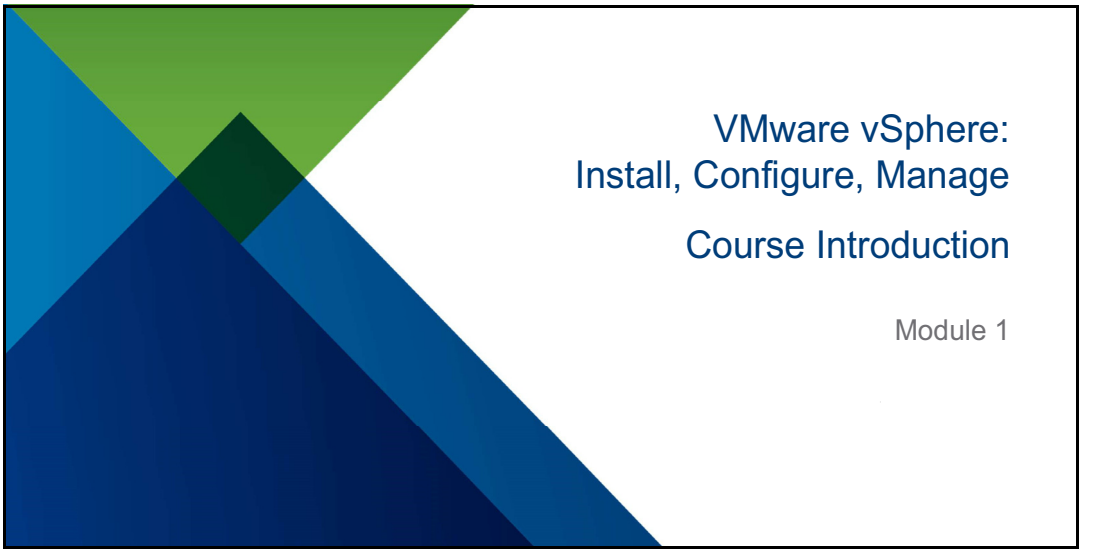
Module 12 vSphere Troubleshooting 629
12-2 You Are Here 630
12-3 Importance 631
12-4 Module Lessons 632
12-5 Lesson 1: Tools for Troubleshooting 633
12-6 Learner Objectives 634
12-7 About the VMware Knowledge Base 635
12-8 About the VMware Community 636
12-9 About VMware Host Client 637
12-10 Monitoring the Virtual Environment with Performance Charts 638
12-11 Viewing and Interpreting Diagnostic Information 639
12-12 Location of vCenter Server Logs 640
12-13 Common Logs 641
12-14 Management Node Logs 642
12-15 Platform Services Controller Logs 643
12-16 Important vCenter Server Log for Troubleshooting 644
12-17 Location of ESXi Host Logs 645
12-18 Useful ESXi Host Logs for Troubleshooting 646
12-19 Viewing vCenter Server Log Files in vSphere Web Client 647
12-20 Viewing Log Files in the DCUI 648
12-21 About vSphere Syslog Collector 649
12-22 Review of Learner Objectives 650
12-23 Lesson 2: Introduction to Troubleshooting 651
12-24 Learner Objectives 652
12-25 Troubleshooting Process 653
12-26 Definition of a System Problem 654
12-27 Collecting Symptoms of a Problem 655
12-28 Gathering Supplemental Information 656
12-29 Identifying Possible Causes and Taking the Appropriate Approach 657
12-30 Determining the Root Cause 658
12-31 Resolving the Problem 659
12-32 Example Scenario: Defining the Problem 660
12-33 Example: Gathering Information 661
12-34 Example: Identifying Possible Causes 662
12-35 Example: Resolving the Problem 663
12-36 VMware Severity Definitions and Response Targets 664
12-37 Review of Learner Objectives 665
12-38 Key Points 666

https://t.me/learningnets

https://t.me/learningnets

MODULE 1

Course Introduction



https://t.me/learningnets

1-2 Importance

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

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

https://t.me/learningnets

1-3 Learner Objectives (1)

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

- Install and configure ESXi hosts
- Deploy and configure VMware vCenter® Server Appliance™
- Create a backup schedule for the vCenter Server Appliance instance
- Use VMware 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 VMware vSphere® VMFS datastores
- Use vSphere Client to create virtual machines, templates, clones, and snapshots
- Create a content library for deploying virtual machines
- Manage virtual machine resource usage and manage resource pools

https://t.me/learningnets

1-4 Learner Objectives (2)

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

- Migrate virtual machines with VMware vSphere® vMotion® and VMware vSphere® Storage vMotion®
- Explain how to use a troubleshooting methodology to diagnose and solve problems in a vSphere environment
- Access the vCenter Server Appliance shell to diagnose problems and monitor the health of the vCenter Server database
- Use the command line and VMware vRealize® Log Insight™ to view log files
- Create and manage a vSphere cluster that is enabled with VMware vSphere® High Availability and VMware vSphere® Distributed Resource Scheduler™
- Describe the methods for protecting and recovering virtual machine data
- Use VMware vSphere® Update Manager™ to apply patches and perform upgrades to ESXi hosts and virtual machines
- Explain how to use a troubleshooting methodology to diagnose and solve problems in a vSphere

https://t.me/learningnets

1-5 You Are Here

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

https://t.me/learningnets

1-6 Typographical Conventions

The following typographical conventions are used in this course.

`Monospace`	Filenames, folder names, path names, and command names: Navigate to the `VMS` folder.
`Monospace bold`	What the user types: Enter **`ipconfig /release`**.
Boldface	User interface controls: Click the **Configuration** tab.
Italic	Book titles and placeholder variables: • *vSphere Virtual Machine Administration* • *ESXi_host_name*

https://t.me/learningnets

1-7 References

Title	Location
vSphere Installation and Setup	https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-vcenter-server-67-installation-guide.pdf
vCenter Server and Host Management	https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-67-host-management-guide.pdf
vSphere Virtual Machine Administration	https://docs.vmware.com/en/VMware-vSphere/6.0/vsphere-esxi-vcenter-server-601-virtual-machine-admin-guide.pdf
vSphere Networking	https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-67-networking-guide.pdf
vSphere Security	https://docs.vmware.com/en/VMware-vSphere/index.html
vSphere Resource Management	https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-67-resource-management-guide.pdf
vSphere Availability	https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-67-availability-guide.pdf
vSphere Monitoring and Performance	https://docs.vmware.com/en/VMware-vSphere/index.html

https://t.me/learningnets

1-8 VMware Online Resources

Documentation for vSphere: https://docs.vmware.com/en/VMware-vSphere/index.html

VMware Configuration Maximums: https://configmax.vmware.com/home

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

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

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

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

VMware Certification: https://www.vmware.com/education-services/certification.html

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

https://t.me/learningnets

1-9 VCP-Foundations Certification Alignment

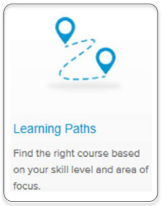
VMware vSphere: Install, Configure, Manage aligns with the VCP-Foundations certification:

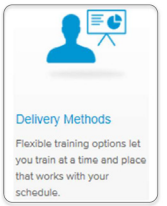
- The VCP-Foundations exam blueprint served as the basis for the design of this course.
- You should use the VCP-Foundations exam blueprint as a reference when preparing for the test.
- This course should not be used as the only resource for exam preparation.
- VMware certification details can be found at: https://mylearn.vmware.com/mgrReg/plan.cfm?plan=64178&ui=www_cert

https://t.me/learningnets

1-10 VMware Education Overview

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





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

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

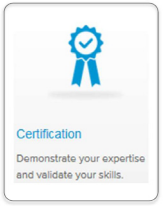
Examples of delivery methods:

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

https://t.me/learningnets

1-11 VMware Certification Overview

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





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



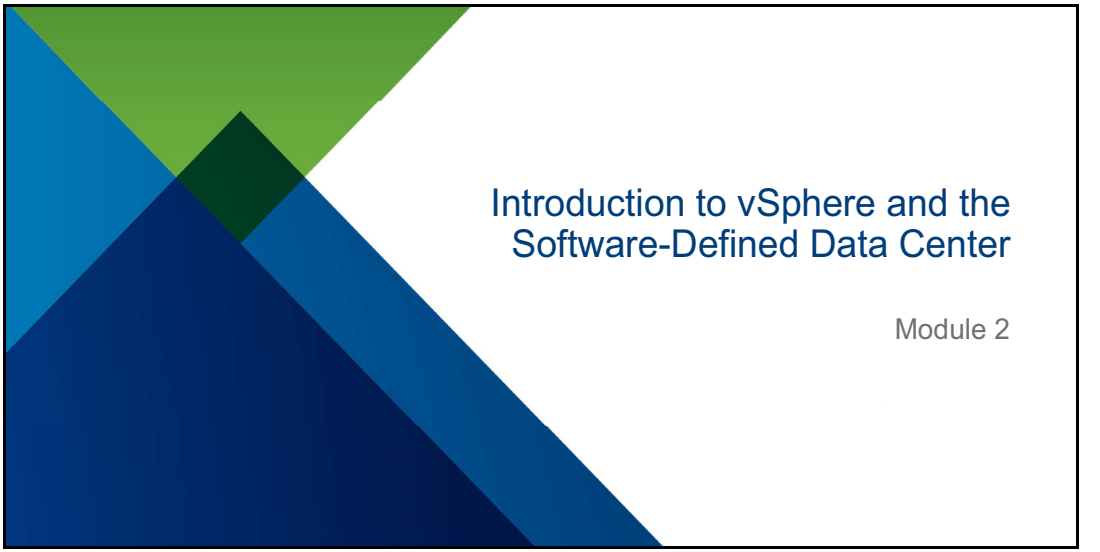
VMware Learning Zone provides a cloud-based video training library that can help you prepare efficiently and give you confidence to pass the certification exam.

https://t.me/learningnets

https://t.me/learningnets

MODULE 2

Introduction to vSphere and the Software-Defined Data Center



https://t.me/learningnets

2-2 You Are Here

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

https://t.me/learningnets

2-3 Importance

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

- 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

2-4 Module Lessons

Lesson 1:	Overview of vSphere and Virtual Machines
Lesson 2:	Shared Resources: Overview of ESXi

https://t.me/learningnets



https://t.me/learningnets

2-6 Learner Objectives

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

- Describe how vSphere fits into the software-defined data center and the cloud infrastructure
- Explain how vSphere interacts with CPUs, memory, networks, and storage
- Navigate vSphere Client and examine VM settings
- Identify the files and components that make up a virtual machine
- Identify the methods to access a virtual machine console
- Identify the virtual network adapters and the advantages of the VMXNET3 adapter
- Compare the types of virtual disk provisioning

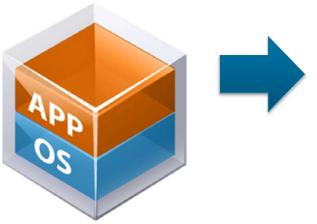
https://t.me/learningnets

2-7 About Virtual Machines

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

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

Virtual Machine



Virtual Machine Components

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

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

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

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

https://t.me/learningnets

2-8 Benefits of Using Virtual Machines

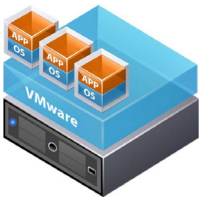
Physical Machines

- Difficult to move or copy
- Bound to a specific set of hardware components
- Often have a short lifecycle
- Require personal contact to upgrade hardware



Virtual Machines

- Easy to move or copy:
  - Encapsulated into files
  - Independent of physical hardware
- Easy to manage:
  - Isolated from other virtual machines running on same physical hardware
  - Insulated from physical hardware changes



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

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

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

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

https://t.me/learningnets

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

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

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

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

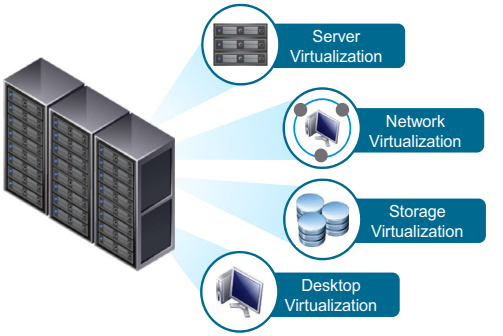
With virtual machines, you can support legacy applications and operating systems on newer hardware when maintenance contracts on the existing hardware expire.

https://t.me/learningnets

2-9 Types of Virtualization

Virtualization is the process of creating a software-based representation of something physical, such as a server, desktop, network, or storage device.

Virtualization is the single most effective way to reduce IT expenses while boosting efficiency and agility for all business sizes.



Server virtualization addresses these inefficiencies by allowing multiple operating systems to run on a single physical server as virtual machines, 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 the same as if on a physical network.

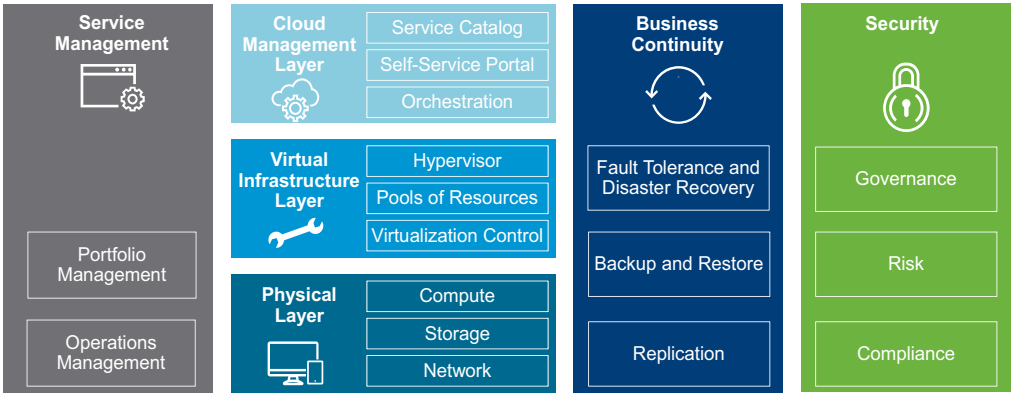
Storage virtualization is the process of creating a software-based representation of network storage devices into what appears to be a single unit.

Deploying desktops as a managed service enables you to respond more quickly to changing needs and opportunities.

https://t.me/learningnets

2-10 About the Software-Defined Data Center

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



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

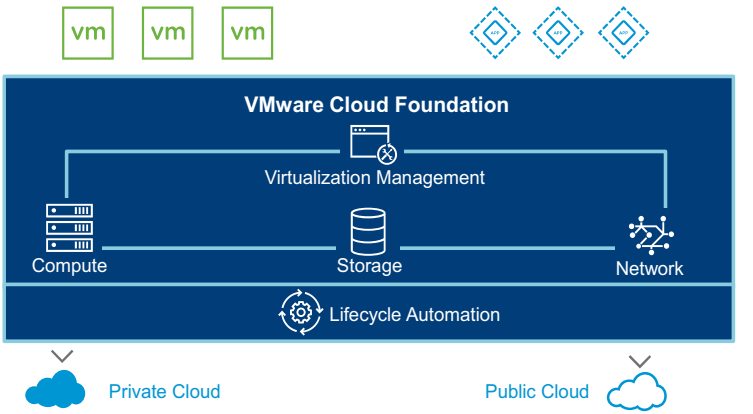
The software-defined data center can include the following components:

- Service Management: Use service management 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: The Cloud Management layer includes the service catalog, that houses the facilities to be deployed, orchestration, that provides the workflows to deploy catalog items, and the self-service portal that allows end users to use the SDDC.
- Virtual Infrastructure Layer: This layer establishes a robust virtualized environment into which all other solutions integrate. This layer includes the virtualization platform for the Hypervisor, Pools of Resources, and Virtualization Control. Additional processes and technologies build on the infrastructure to enable 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

2-11 vSphere and Cloud Computing

Cloud computing exploits the efficient pooling of an on-demand, self-managed, and virtual infrastructure.



As defined by the National Institute of Standards and Technology (NIST), cloud computing is a model for enabling ubiquitous, convenient, and on-demand network access to a shared pool of configurable computing resources. For example, networks, servers, storage, applications, and services can be rapidly provisioned and released with minimal management effort or little service provider interaction. vSphere is the foundation of the technology that enables shared and configurable resource pools. vSphere abstracts the physical resources of the data center to separate the workload from the physical hardware. A software user interface can provide the framework for managing and maintaining this abstraction and allocation.

VMware Cloud Foundation™ provides integrated cloud infrastructure (compute, storage, networking and security) and cloud management services to run enterprise applications in both private and public environments. Cloud Foundation extends the vSphere core hypervisor with integrated software-defined storage, networking, and security capabilities that can be consumed flexibly on premises or as a service in the public cloud (VMware Cloud™ on AWS or VMware cloud providers).

Use cases:

- Cloud infrastructure: Exploit the high performance, availability, and scalability of the VMware software-defined data center to run mission-critical applications such as databases, web apps, and virtual desktop infrastructure (VDI).
- IT automation: Automate infrastructure and application delivery with self-service capabilities.
- VDI: Provides a complete solution for VDI deployment at scale. It simplifies the planning and design with standardized and tested solutions fully optimized for VDI workloads.

https://t.me/learningnets

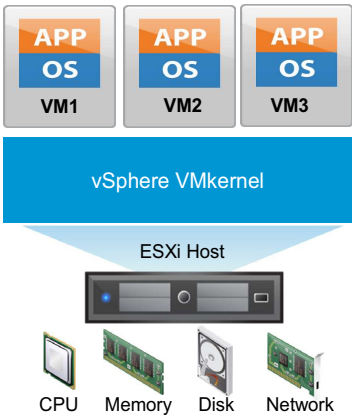
- Hybrid cloud: Build a true hybrid cloud with common infrastructure and consistent operational model, connecting your on-premise and off-premise data center that is compatible, stretched, and distributed.

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

https://t.me/learningnets

2-12 Virtual Machine: Guest and Consumer of ESXi Host

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



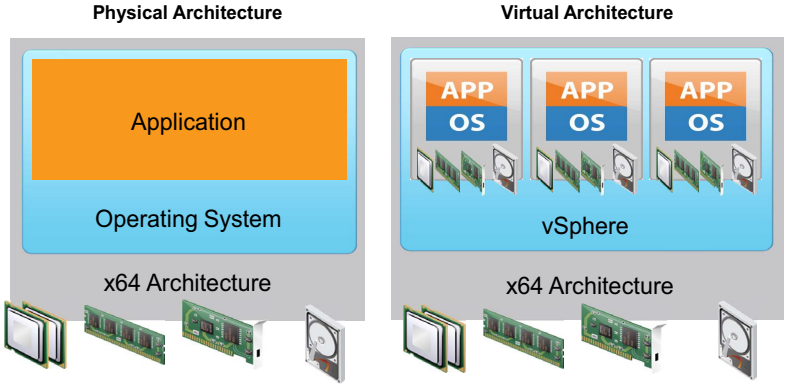
A virtual machine is an abstraction in software of a physical machine turning components into files that act like physical components.

For the list of all supported operating systems, see https://www.vmware.com/resources/compatibility/search.php.

https://t.me/learningnets

2-13 Physical and Virtual Architecture

Virtualization is a technology that abstracts physical components into software components and provides solutions for many problems that IT staff faces.



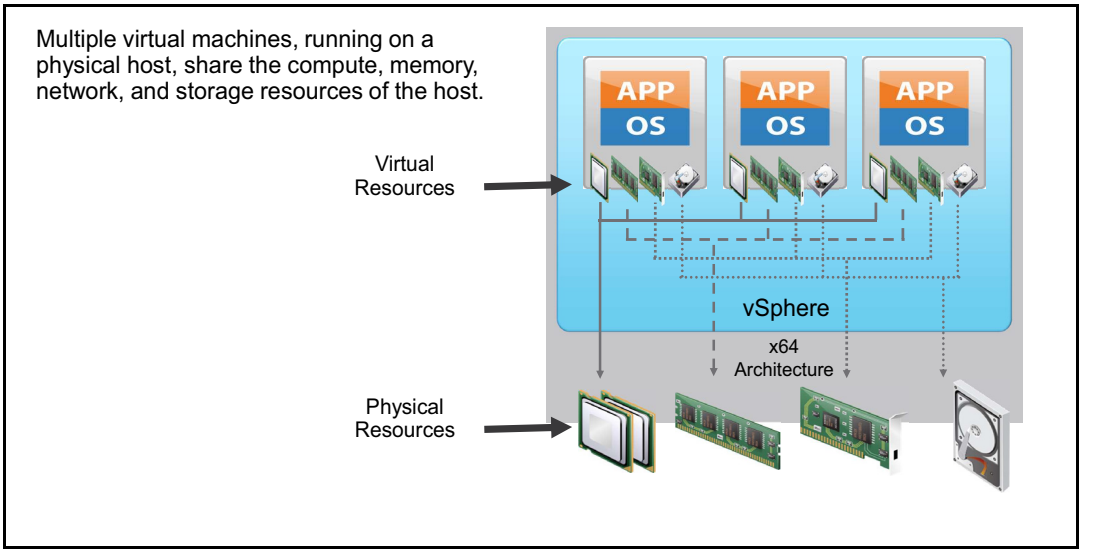
Virtualization enables you to consolidate and run multiple workloads as virtual machines on a single computer. A virtual machine is a computer that is created by software that, like a physical computer, runs an operating system and applications.

The slide 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 installed hardware through a thin layer of software called the virtualization layer or hypervisor.

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

https://t.me/learningnets

2-14 Physical Resource Sharing



With virtualization, you can run multiple virtual machines on a single physical host, with each virtual machine sharing the resources of one physical computer across multiple environments. Virtual machines share access to CPUs and are scheduled to run by the hypervisor. In addition, virtual machines are assigned their own region of memory to use and share access to the physical network cards and disk controllers. Different virtual machines can run different operating systems and applications on the same physical computer.

When multiple virtual machines run on an ESXi host, each virtual machine is allocated a portion of the physical resources. The hypervisor schedules virtual machines like a traditional operating system allocates memory and schedules applications. These virtual machines run on various CPUs.

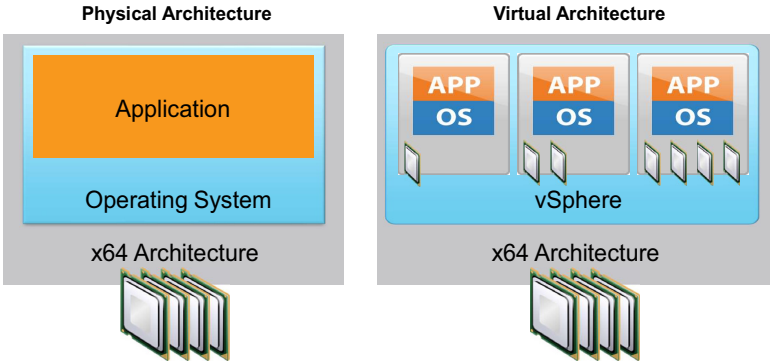
Virtual machines, like applications, use network and disk bandwidth. However, virtual machines are managed with elaborate control mechanisms to manage how much access is available for each virtual machine. With the default resource allocation settings, all virtual machines associated with the same ESXi host receive an equal share of available resources.

https://t.me/learningnets

2-15 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 virtual machines operate as if they were running directly on a physical machine. CPU virtualization is not emulation. A software emulator enables programs to run on a computer system other than the one for which they were originally written. Emulation provides portability, but might negatively affect performance. CPU virtualization is not emulation because the supported guest operating systems are those designed for x64 processors. The hypervisor allows the operating systems to run natively on the hosts' physical x64 processors.

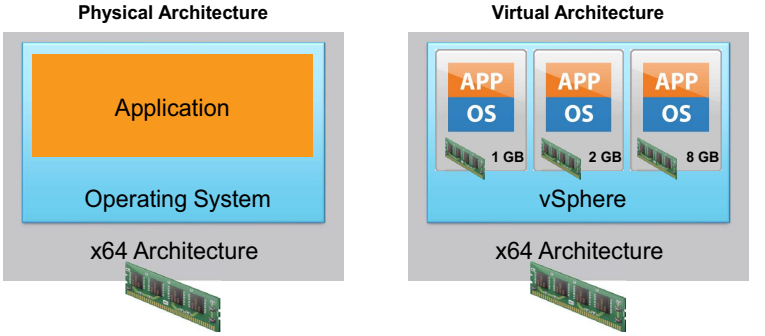
When many virtual machines are running on an ESXi host, those virtual machines might compete for CPU resources. When CPU contention occurs, the ESXi host time-slices the physical processors across all virtual machines so that each virtual machine runs as if it had a specified number of virtual processors.

https://t.me/learningnets

2-16 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.



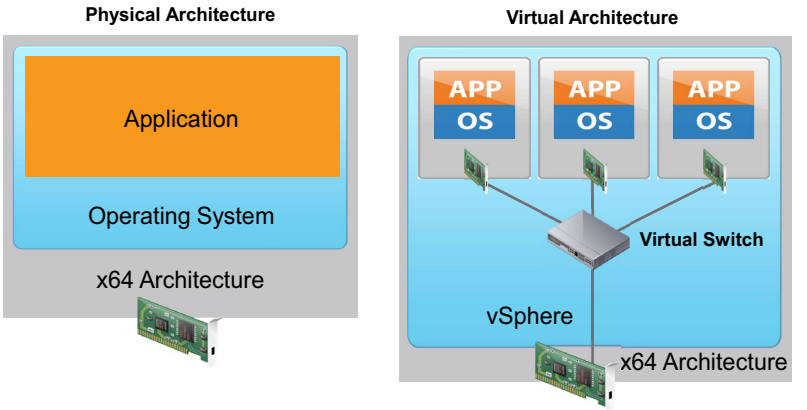
When an application starts, it uses the interfaces provided by the operating system to allocate or release virtual memory pages during the execution. Virtual memory is a decades-old technique used in most general-purpose operating systems. Virtual memory enables operating systems to present more memory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

Virtual memory creates a uniform virtual address space for applications and enables the operating system and hardware to handle the address translation between the virtual address space and the physical address space. This technique adapts the execution environment to support large address spaces, process protection, file mapping, and swapping in modern computer systems. In a virtualized environment, the VMware virtualization layer creates a contiguous addressable memory space for the virtual machine when it is started. The memory space allocated is configured when the virtual machine is created and has the same properties as the virtual address space. This configuration enables the hypervisor to run multiple virtual machines simultaneously while protecting the memory of each virtual machine from being accessed by others.

https://t.me/learningnets

2-17 Physical and Virtual Networking

Virtual Ethernet adapters and virtual switches are key virtual networking components.



A virtual machine can be configured with one or more virtual Ethernet adapters. Virtual switches enable virtual machines on the same ESXi host to communicate with one another by using the same protocols that are used over physical switches, without the need for additional hardware. Virtual switches also support VLANs that are compatible with standard VLAN implementations from other networking equipment vendors. With the VMware virtual networking, you can link local virtual machines together and link local virtual machines to the external network through a virtual switch.

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

Virtual switches are similar to modern physical Ethernet switches in many ways. Like a physical switch, each virtual switch is isolated and has its own forwarding table. So every destination that the switch looks up can match only ports on the same virtual switch where the frame originated. This feature improves security, making it difficult for hackers to break virtual switch isolation. 4

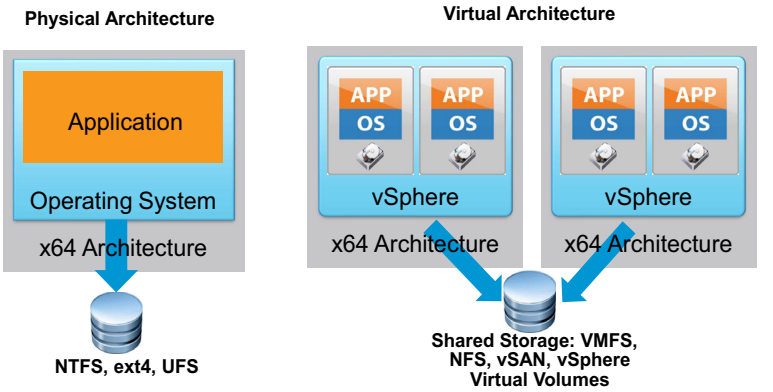
Virtual switches also support VLAN segmentation at the port level, so that each port can be configured as an access or trunk port, providing access to either single or multiple VLANs.

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

2-18 Comparison of Physical File Systems and Datastores

vSphere VMFS enables a distributed storage architecture, allowing multiple ESXi hosts to read or write to the shared storage concurrently.



To store virtual disks, ESXi uses datastores, which are logical containers that hide specifics of physical storage from virtual machines and provide a uniform model for storing virtual machine files. Datastores that you deploy on block storage devices use the VMware vSphere® 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. VMFS is a high-performance cluster file system designed for virtual machines. VMFS uses distributed journaling of its file system metadata changes to enable fast and resilient recovery if a hardware failure occurs. VMFS increases resource usage by providing multiple virtual machines with shared access to a consolidated pool of clustered storage. VMFS is also the foundation of distributed infrastructure services, such as live migration of virtual machines and virtual machine files, dynamically balanced workloads across available compute resources, automated restart of virtual machines, and fault tolerance.

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

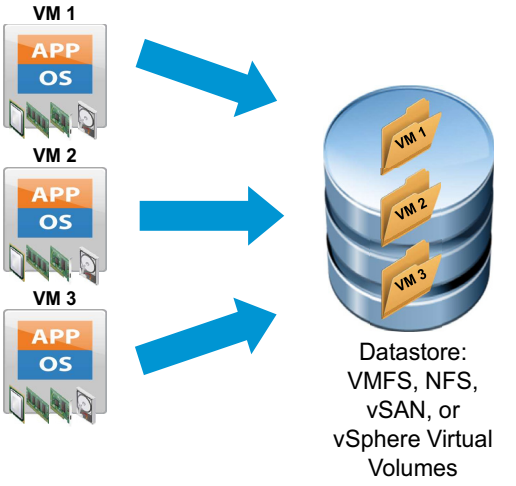
With the distributed locking methods, VMFS forges the link between the virtual machine and the underlying storage resources. The unique capabilities of VMFS enable virtual machines to join a cluster seamlessly, with no management overhead.

https://t.me/learningnets

2-19 Virtual Machine Encapsulation

vSphere encapsulates each virtual machine into a set of virtual machine files.

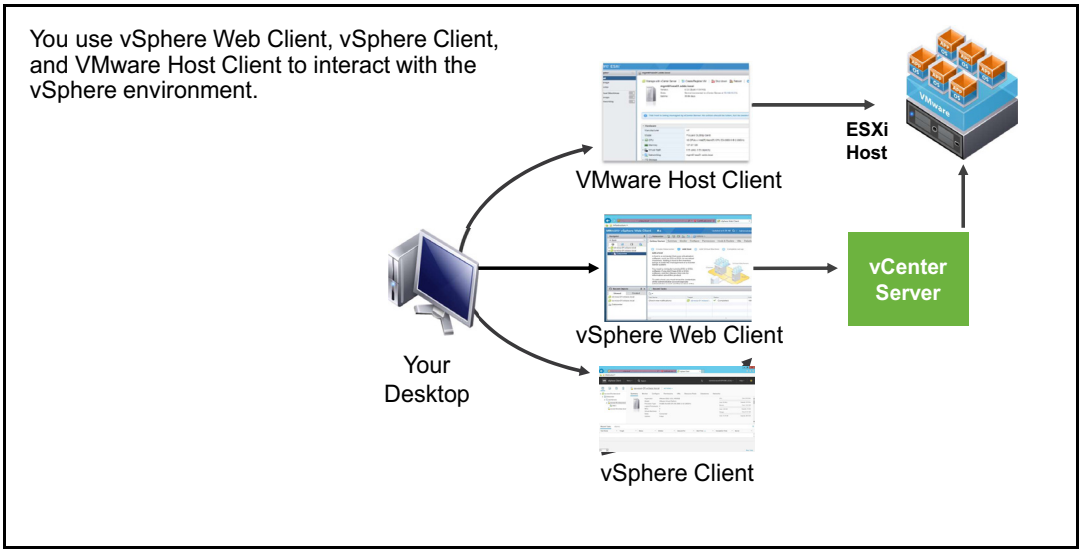
Virtual machine files are stored in directories on a VMFS, NFS, VMware vSAN™, or VMware vSphere® Virtual Volumes™ datastore.



vSphere encapsulates each virtual machine into a small number of files or objects, making virtual machines easier to manage and migrate. The files and objects for each virtual machine are stored in a separate folder on a datastore.

https://t.me/learningnets

2-20 About vSphere Clients



VMware Host Client™, introduced in vSphere 6.0 U2, provides direct management of individual ESXi hosts. In vSphere 6.5, it replaced Windows-based VMware vSphere® Client™. VMware Host Client is generally used only when management through VMware vCenter Server® is not possible.

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

The current version of vSphere Client is an HTML5-based client that can be used alongside vSphere Web Client.

VMware Host Client, vSphere Web Client, and 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

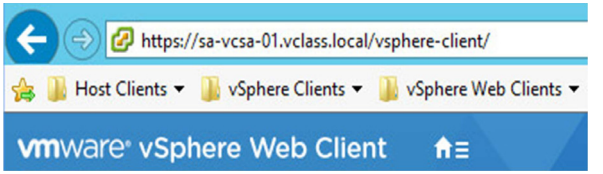
2-21 About vSphere Web Client

vSphere Web Client has the following components:

- Adobe Flex client application running in a browser
- Java server embedded in vCenter Server Appliance and vCenter Server
- A thin client, which requires only a browser with the Adobe Flash plug-in installed

You access vSphere Web Client at https*://your_vCenter_Server_Appliance/*vsphere-client.

An optional enhanced authentication plug-in provides Integrated Windows Authentication and Windows-based smart-card functionality.



The blue image on the slide is the unique label that you will see on the upper-left corner of the screen when using vSphere Web Client. This image helps you quickly differentiate vSphere Web Client from the other clients.

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

When you use https://*your_vCSA*/vsphere-client to access vSphere Web Client, the URL internally redirects to port 9443 on your vCenter Server system.

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

The tabbing structure of most vSphere objects is reorganized to be more familiar and easier to use.

The Client Integration Plug-In, though previously necessary for a certain set of functions in vSphere Web Client, is no longer required. vSphere Web Client no longer has any dependency on the following functions:

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

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

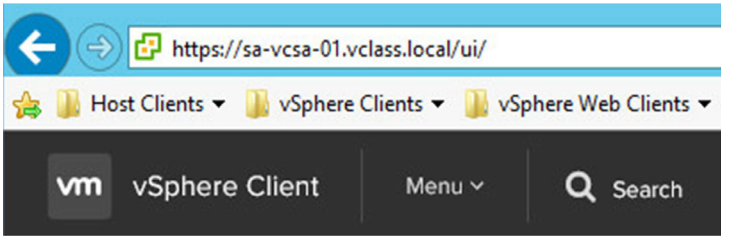
https://t.me/learningnets

2-22 About vSphere Client

vSphere Client is HTML5-based and has no dependence on Adobe Flex.

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

You access vSphere Client from a supported browser at https://*your_vCenter_Server_Appliance*/ui.



The dark image on the slide is the unique label that you will see on the upper-left corner of the screen when using vSphere Client. This image helps you quickly differentiate vSphere Client from the other clients.

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

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

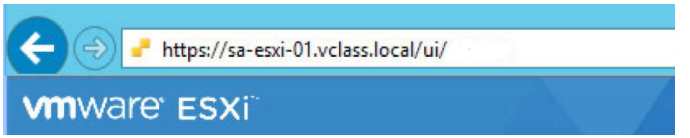
For the parity list of the features supported by vSphere Client compared to the features supported by vSphere Web Client, see https://docs.vmware.com/en/VMware-vSphere/6.5/rn/vsphere-client-65-html5-functionality-support.html.

https://t.me/learningnets

2-23 About VMware Host Client

VMware Host Client is an HTML5-based thin client that you can use to manage hosts directly when vCenter Server is unavailable.

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



VMware Host Client provides direct management access to individual ESXi hosts. It is used for scenarios when vCenter Server Appliance is unavailable. VMware Host Client is based on HTML5, and you access it through a browser.

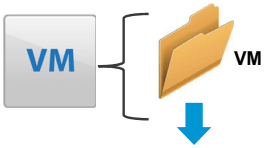
The blue image on the slide is the unique label that you will see on the upper-left corner of the screen when using VMware Host Client. This image helps you quickly differentiate VMware Host Client from the other clients.

VMware Host Client replaces the host client functionality of vSphere Client for Windows. The HTML5 UI layout is similar to that of vSphere Web Client to remain consistent in workflow navigation. The UI is simplified in areas where vCenter Server functionality is not required.

https://t.me/learningnets

2-24 About Virtual Machine Files

A virtual machine includes a set of related files.



Configuration file	*VM_name*.vmx
Swap files	*VM_name*.vswp
	vmx-*VM_name*.vswp
BIOS file	*VM_name*.nvram
Log files	vmware.log
Template configuration file	*VM_name*.vmtx
Disk descriptor file	*VM_name*.vmdk
Disk data file	*VM_name*-flat.vmdk
Raw device map file	*VM_name*-rdm(p).vmdk
Snapshot disk file	*VM_name*-######.vmdk
Snapshot data file	*VM_name*.vmsd
Snapshot state file	*VM_name-Snapshot#*.vmsn
Snapshot memory file	*VM_name-Snapshot#*.vmem
Suspend state file	*VM_name*-*.vmss
Suspended snapshot memory state	*VM_name*-*.vmem

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

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

NOTE

The list files shown on the slide is not comprehensive. For a complete list of all the types of Virtual machine files, see https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-CEFF6D89-8C19-4143-8C26-4B6D6734D2CB.html

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

https://t.me/learningnets

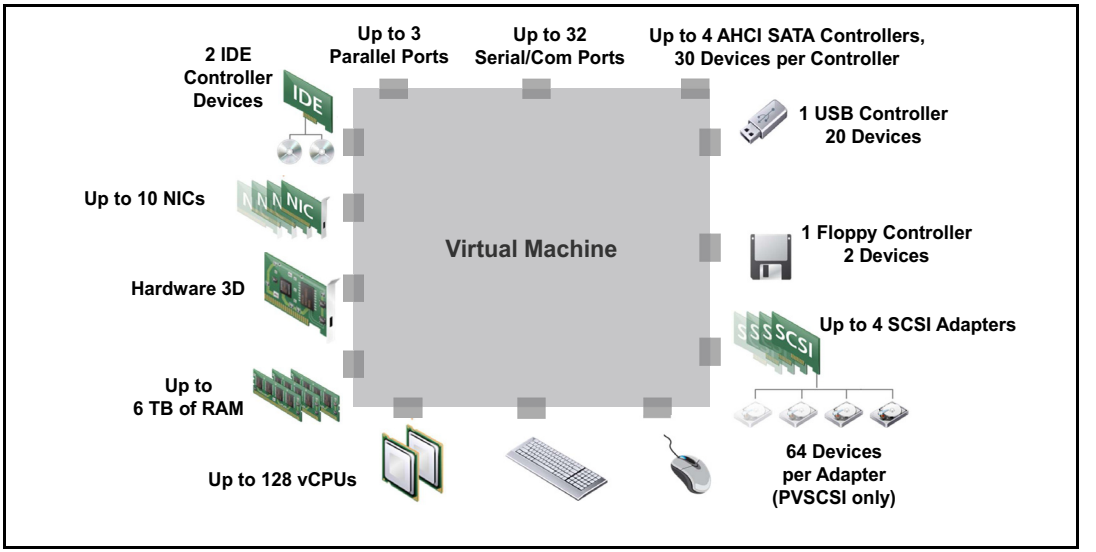
If the virtual machine has more than one disk file, the file pair for the second disk file and later is named *VM_name_#*.vmdk and V*M_name_#*-flat.vmdk, where # is the next number in the sequence, starting with 1. For example, if the virtual machine named Test01 has two virtual disks, this virtual machine has the Test01_0.vmdk, Test01_0-flat.vmdk, Test01_1.vmdk, and Test01_1-flat.vmdk files.

In addition to the current log file, vmware.log, up to six archive log files are maintained at one time. For example, -1.log to -6.log might exist at first. The next time an archive log file is created, for example, when the virtual machine is powered off and powered back on, the following actions occur: the -6.log is deleted, the -5.log is renamed to -6.log, and so on. Finally the previous vmware.log is renamed to the -1.log.

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

https://t.me/learningnets

2-25 About Virtual Machine Virtual Hardware



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

You can configure virtual machine memory and CPU settings. vSphere supports many of the latest CPU features, including virtual CPU performance counters. You can add virtual hard disks and NICs. You can also add and configure virtual hardware, such as CD/DVD drives, floppy drives, and SCSI devices. 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 virtual machine that resides on an ESXi host to which the devices are physically attached. When you attach a USB device to a physical host, the device is available only to virtual machines that reside on that host. Those virtual machines cannot connect to a device on another host in the data center. A USB device is available to only one virtual machine at a time. When you remove a device from a virtual machine, it becomes available to other virtual machines that reside on the host.

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

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

https://t.me/learningnets

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 VMCI is an infrastructure that provides fast and efficient communication between a virtual machine and the host operating system. The VMCI SDK facilitates the development of applications that use the VMCI infrastructure. Without VMCI, virtual machines communicate with the host using the network layer. Using the network layer adds overhead to the communication. With VMCI, communication overhead is minimal and tasks that require that communication can be optimized. VMCI can go up to nearly 10 Gbit/s with 128 K sized queue pairs.

The following types of communication exist:

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

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

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

https://t.me/learningnets

2-26 Virtual Hardware Versions

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

Compatibility	Hardware Version
ESXi 6.7	14
ESXi 6.5 and later	13
ESXi 6.0 and later	11
ESXi 5.5 and later	10

Each release of a VMware product has a corresponding virtual machine hardware version included. The table shows the highest hardware version level that each VMware ESX®/ESXi version supports. Each virtual machine compatibility level supports at least five major or minor vSphere releases. For example, a virtual machine with ESX/ESXi 4.0 and later compatibility can run on ESX/ESXi 4.0, ESX/ESXi 4.1, ESXi 5.0, ESXi 5.1, ESXi 5.5, ESXi 6.0, ESXi 6.5, and ESXi 6.7.

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

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

https://t.me/learningnets

2-27 About CPU and Memory

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

The maximum number of virtual CPUs (vCPUs) that you can assign to a virtual machine depends on the following factors:

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

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

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

The maximum memory size of a virtual machine
with ESXi 6.7 compatibility running on ESXi 6.7 is 6 TB.

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

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

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

You can set most of the memory parameters during virtual machine creation or after the guest operating system is installed. Some actions require that you power off the virtual machine before changing the settings. The memory resource settings for a virtual machine determine how much of the host's memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size. ESXi hosts limit the memory resource use to the maximum amount useful for the virtual machine so that you can accept the default of Unlimited memory resources. You can reconfigure the amount of memory allocated to a virtual machine to enhance performance. Maximum memory size for a virtual machine depends on the virtual machine's compatibility setting.

https://t.me/learningnets

2-28 About Virtual Storage

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

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

When you create a virtual machine, the wizard automatically selects a virtual SCSI adapter based on the choice of guest operating system. The virtual disk is stored in the same folder as the virtual machine configuration file. However, you can select to place a virtual disk in an alternate location, for example, when separating boot and data disks.You select a VMFS datastore to hold the new, blank virtual disk, and specify the disk's size. You can also site the disk at a specific virtual SCSI target ID and LUN.

https://t.me/learningnets

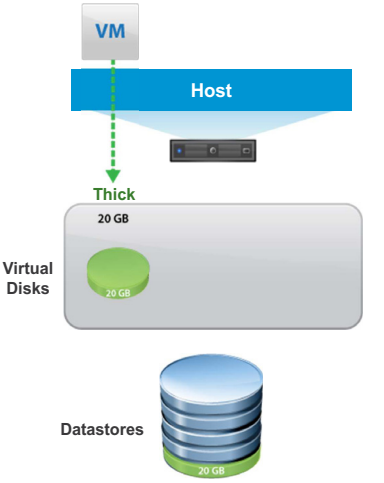
2-29 About Thick-Provisioned Virtual Disks

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

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

Eager-zeroed or lazy-zeroed:

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



When you create a virtual disk, the following virtual disk types are available:

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

https://t.me/learningnets

2-30 About Thin-Provisioned Virtual Disks

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

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

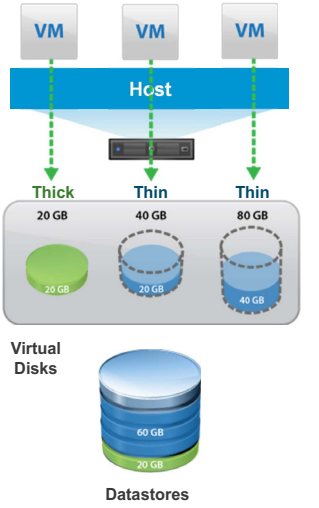
Run the unmap command to reclaim unused space from the array.

You can mix thick and thin formats.

Full reporting and alerts help manage allocations and capacity.

More efficient use of storage:

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



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

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

The following table identifies and compares the differences between the virtual disk options.

	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 upon first write to block.

https://t.me/learningnets

	Thick Provisioned Lazy-Zeroed	Thick Provisioned Eager-Zeroed	Thin Provisioned
Virtual disk layout	Higher chance of contiguous file blocks.	Highest chance of contiguous file blocks.	Layout varies according to dynamic state of the volume at time of block allocation.
Zeroing 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.

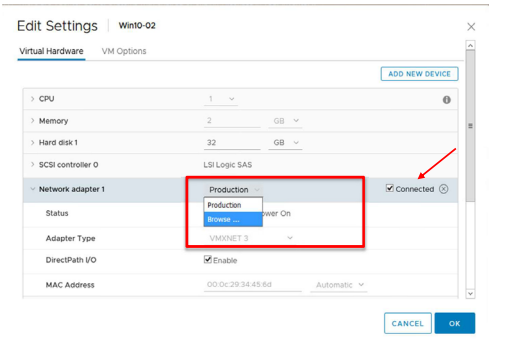
https://t.me/learningnets

2-31 About Virtual Networks

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

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

- The network adapter type
- The port group to connect to
- The network connection state
- Whether or not to connect to the network when the virtual machine powers on



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.
- AHCI 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. AHCI SATA is available only for virtual machines 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.

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

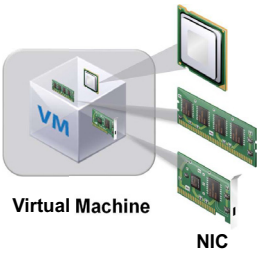
https://t.me/learningnets

2-32 About Virtual Network Adapters (1)

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

Supported network adapter types:

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



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

- The virtual machine compatibility level (or hardware version), which depends on the host that created or most recently updated it. For example, the VMXNET3 virtual NIC requires hardware version 7 (ESX/ESXi 4.0 or later).
- Whether or not the virtual machine compatibility is updated to the latest version for the current host.
- The guest operating system.

The following NIC types are supported:

- E1000E: Emulated version of the Intel 82574 Gigabit Ethernet NIC. E1000E is the default adapter for Windows 8 and Windows Server 2012.
- E1000: Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in most newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later.
- Flexible: Identifies itself as a Vlance adapter when a virtual machine boots, but initializes itself and functions as either a Vlance or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlance adapter to the higher performance VMXNET adapter.
- Vlance: Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in 32-bit legacy guest operating systems. A virtual machine configured with this network adapter can use its network immediately.

https://t.me/learningnets

- 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.

https://t.me/learningnets

2-33 About Virtual Network Adapters (2)

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

- SR-IOV pass-through: Representation of a virtual function on a physical NIC with SR-IOV support. This adapter type is suitable for virtual machines that require more CPU resources or where latency might cause failure. The virtual machine and the physical adapter exchange data without using the VMkernel as an intermediary. If certain virtual machines are sensitive to network delay, SR-IOV can provide direct access to the virtual functions of supported physical NICs, bypassing the virtual switches and hence reducing overhead.

  SR-IOV pass-through is available in ESXi 5.5 and later for Red Hat Enterprise Linux 6 and later, and Windows Server 2008 R2 with SP2. An operating system release might contain a default virtual function driver for certain NICs, while for others you must download and install it from a location provided by the NIC or host vendor.
- vSphere DirectPath I/O allows virtual machine access to physical PCI network functions on platforms with an I/O Memory Management Unit. Pass-through devices provide the means to more efficiently use resources and improve performance in your environment. You can configure a pass-through PCI device on a virtual machine in vSphere Web Client. The following features are unavailable for virtual machines configured with vSphere DirectPath I/O:
  - Hot adding and removing of virtual devices
  - Suspend and resume
  - Record and replay
  - Fault tolerance
  - High availability
  - VMware vSphere® Distributed Resource Scheduler™: Limited availability. The virtual machine can be part of a cluster, but cannot migrate across hosts.
  - Snapshots.

https://t.me/learningnets

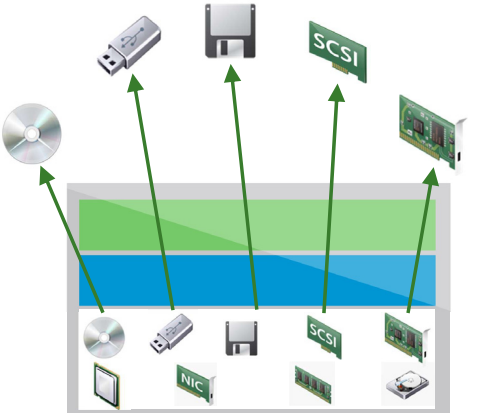
- vSphere 6.5 and later releases support remote direct memory access (RDMA) communication between virtual machines that have paravirtualized RDMA (PVRDMA) network adapters. PVRDMA provides an RDMA-like interface for vSphere guests. PVRDMA allows multiple guests to access the RDMA device by using Verbs API, an industry-standard interface. A set of these Verbs was implemented to expose an RDMA-capable guest device (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 usage and faster data access

https://t.me/learningnets

2-34 About Miscellaneous Devices

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

- CD/DVD drive: Connect to a CD, DVD, or ISO image.
- USB 3.0: Supported with host-connected and client-connected devices on Linux or Windows 8/Server 2012.
- Floppy drive: Connect a virtual machine to a floppy drive or a floppy image.
- Generic SCSI devices: A virtual machine can be connected to additional SCSI adapters.
- vGPUs: Enable a virtual machine to use GPUs on the physical host for high-computation activities.



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

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

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

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

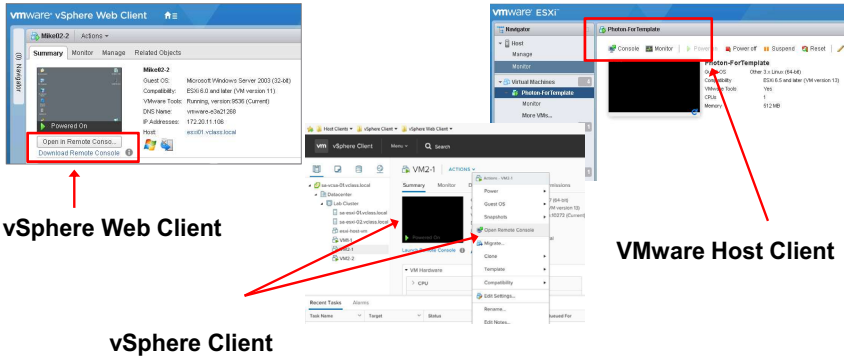
For information about creating a virtual machine, see *vSphere Virtual Machine Administration* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

2-35 About the Virtual Machine Console

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

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



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

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

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

The embedded VMware Host Client is an HTML5-based client that has a similar interface to the vSphere Web Client but is only used to manage single ESXi hosts. You use the VMware Host Client to conduct emergency management when vCenter Server is temporarily unavailable.

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

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

https://t.me/learningnets

2-36 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe how vSphere fits into the software-defined data center and the cloud infrastructure
- Explain how vSphere interacts with CPUs, memory, networks, and storage
- Navigate vSphere Client and examine VM settings
- Use various clients to access and manage your vCenter Server system and ESXi host
- Compare the uses of vSphere Client, VMware vSphere® Web Client, and VMware Host Client™
- Identify the files that make up a virtual machine
- Compare virtual machine hardware version 14 to other versions
- Describe the components of a virtual machine
- Identify the methods to access a virtual machine console
- Identify the virtual network adapters and the advantages of the VMXNET3 adapter
- Compare the types of virtual disk provisioning

https://t.me/learningnets



https://t.me/learningnets

2-38 Learner Objectives

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

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

https://t.me/learningnets

2-39 About ESXi

ESXi is a hypervisor with the following features:

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

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

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

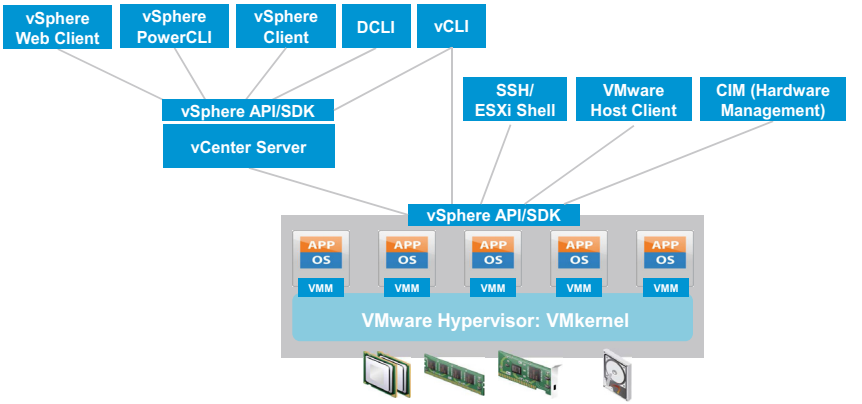
- Host-based firewall: To minimize the risk of an attack through the management interface, ESXi includes a firewall between the management interface and the network.
- Memory hardening: The ESXi kernel, user-mode applications, and executable components such as drivers and libraries are located at random, nonpredictable memory addresses. Combined with the non-executable memory protections made available by microprocessors, memory hardening provides protection that makes it difficult for malicious code to use memory exploits to take advantage of vulnerabilities.
- Kernel module integrity: Digital signing ensures the integrity and authenticity of modules, drivers, and applications as they are loaded by the VMkernel.
- Trusted Platform Module: A hardware element that creates a trusted platform. This element affirms that the boot process and all drivers loaded are genuine.
- UEFI secure boot: This feature is for systems that support UEFI secure boot firmware which contains a digital certificate that the VMware infrastructure bundles (VIBs) chain to. At boot time, a verifier is started before other processes that checks the VIB's chain to the certificate in the firmware.

https://t.me/learningnets

- Lockdown modes: Lockdown mode is a vSphere feature that disables login and API functions from being executed directly on an ESXi host. Before vSphere 6.0, only one lockdown mode existed. Feedback from customers indicated that this lockdown mode was inflexible in some use cases. With vSphere 6.0, the introduction of normal and strict lockdown modes aims to improve this feature.

https://t.me/learningnets

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



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

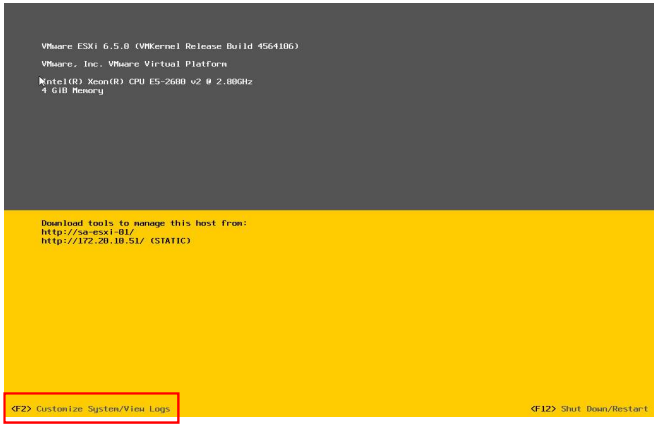
An ESXi host can be accessed through several interfaces:

- vSphere Web Client (connected to vCenter Server)
- vSphere Client (HTML5 UI for connecting to vCenter Server)
- VMware Host Client (HTML5 UI for connecting directly to a host)
- VMware vSphere® Command-Line Interface
- VMware vSphere® API and VMware vSphere® Management SDK
- Common Information Model

https://t.me/learningnets

2-41 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

2-42 Configuring an ESXi Host: Root Access

DCUI enables an administrator to configure root access settings:

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



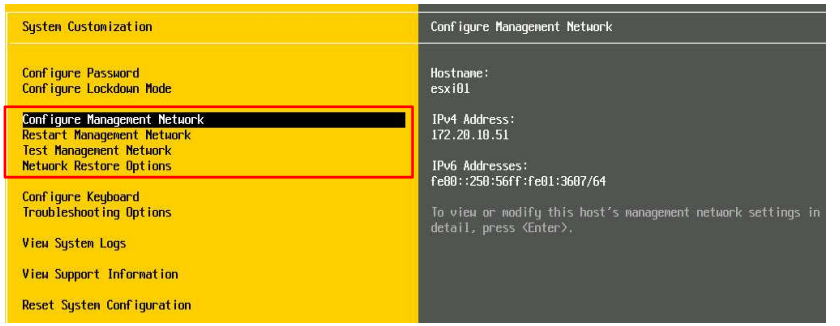
The administrative user name for the ESXi host is root. The root password must be configured during the ESXi installation process.

https://t.me/learningnets

2-43 Configuring an ESXi Host: Management Network

The DCUI enables you to modify network settings:

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



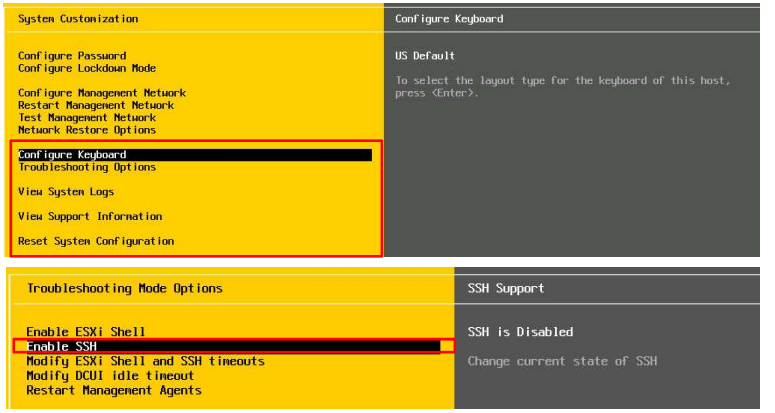
You must set up your IP address before your ESXi host is operational. By default, a DHCP-assigned address is configured for the ESXi host. To change or configure basic network settings, use the DCUI.

From the DCUI, you can change the host name, the IP settings (such as IP address, subnet mask, and default gateway), and the DNS servers. You can also modify the network adapter used for the management network, configure VLAN settings, configure IPv6 addressing, and set custom DNS suffixes. You can restart the management network (without rebooting the system), test the management network (using ping and DNS requests), and disable a management network.

https://t.me/learningnets

2-44 Configuring an ESXi Host: Other Settings

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



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

The troubleshooting options, which are disabled by default, allow you to enable or disable troubleshooting services:

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

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

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

https://t.me/learningnets

2-45 Remote Access Settings: Security Profile

The security profile controls remote access to an ESXi host:

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

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

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

https://t.me/learningnets

2-46 Managing User Accounts: Best Practices

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

- Strictly control root privileges to ESXi hosts.
- Create strong root account passwords that have at least eight characters. Use special characters, case changes, and numbers. Change passwords periodically.
- Manage ESXi hosts centrally through the vCenter Server system by using the appropriate vSphere client.
- 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.
  - If local users are created, manage them centrally using the `esxcli` command in vSphere CLI.

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

Whenever possible, use vSphere Web Client or 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 vSphere CLI or vSphere ESXi Shell, these access methods should be reserved for troubleshooting or configuration that cannot be accomplished by using VMware Host Client. If a host needs to be managed directly, avoid creating local users on the host. If possible, join the host to a Windows domain and log in with domain credentials instead.

https://t.me/learningnets

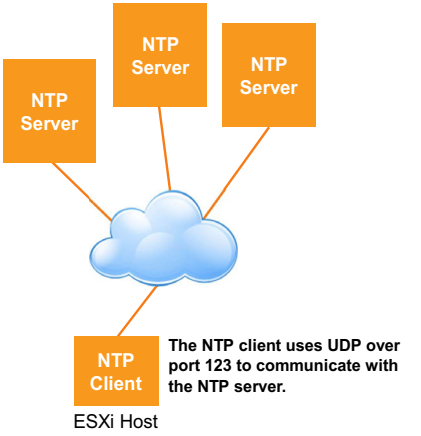
2-47 ESXi Host as an NTP Client

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

NTP is important:

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

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



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.
- Virtual machines can synchronize their time with the ESXi host. Time synchronization is beneficial to applications, such as database applications, running on the virtual machines.

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

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

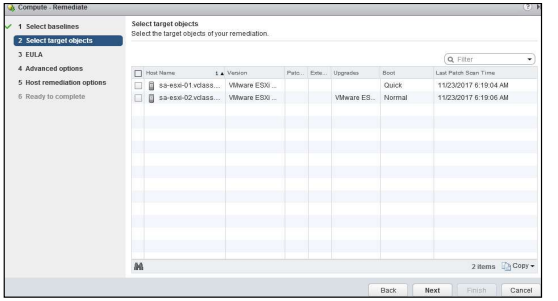
For more information about timekeeping, see VMware knowledge base articles 1318 at http://kb.vmware.com/kb/1318 and 1006427 at http://kb.vmware.com/kb/1006427.

https://t.me/learningnets

2-48 ESXi Quick Boot

In ESX 6.7, Quick Boot enables ESXi to reboot without reinitializing the physical server BIOS:

- Quick Boot reduces remediation time during host patch or host upgrade operations. Reboots can be performed in 1 to 2 minutes.
- Quick Boot is integrated into the vSphere Update Manager workflow during remediation operations.
- Quick Boot is enabled by default on supported hardware. It can be disabled if necessary.



The quick boot option is applicable only for the following hardware platforms:

- Dell PowerEdge R630 Rack Server
- Dell PowerEdge R640 Rack Server
- Dell PowerEdge R730 Rack Server
- Dell PowerEdge R740 Rack Server
- HPE ProLiant DL360 Gen9 Server
- HPE ProLiant DL360 Gen10 Server
- HPE ProLiant DL380 Gen9 Server
- HPE ProLiant DL380 Gen10 Server

https://t.me/learningnets

2-49 Demonstration: Installing and Configuring ESXi Hosts

Your instructor will run a demonstration.

https://t.me/learningnets

2-50 Review of Learner Objectives

You should be able to meet the following objectives:

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

https://t.me/learningnets

2-51 Key Points

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

Questions?

https://t.me/learningnets

https://t.me/learningnets

Creating Virtual Machines



https://t.me/learningnets

3-2 You Are Here

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

https://t.me/learningnets

3-3 Importance

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

https://t.me/learningnets

3-4 Learner Objectives

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

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

https://t.me/learningnets

3-5 About Provisioning Virtual Machines

You can create virtual machines in several ways:

- Use the New virtual machine wizard to create virtual machines.
- Deploy new virtual machines from existing templates/clones.
- Deploy virtual machines, virtual appliances, and vApps stored in OVF.

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

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

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

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

- The virtual machine name and folder. The resource on which the virtual machine will run: a host, a cluster, a vApp, or a resource pool. The virtual machine will have access to the resources of the selected object.
- The datastore on which to store the virtual machine's files. Each datastore might have a different size, speed, availability, and other properties. The available datastores are accessible from the destination resource that you selected.

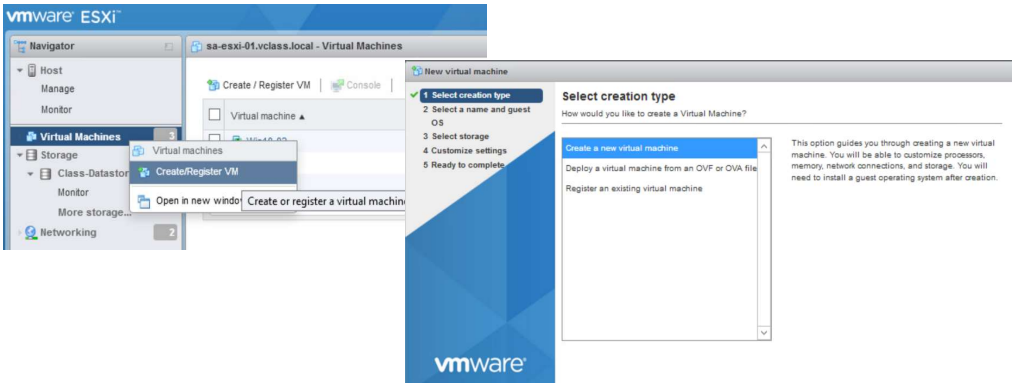
https://t.me/learningnets

- The guest operating system to be installed into the virtual machine.
- The number of NICs, the network to connect to, and the network adapter type.
- Virtual disk provisioning choice.

https://t.me/learningnets

3-6 Creating Virtual Machines with the New Virtual Machine Wizard (1)

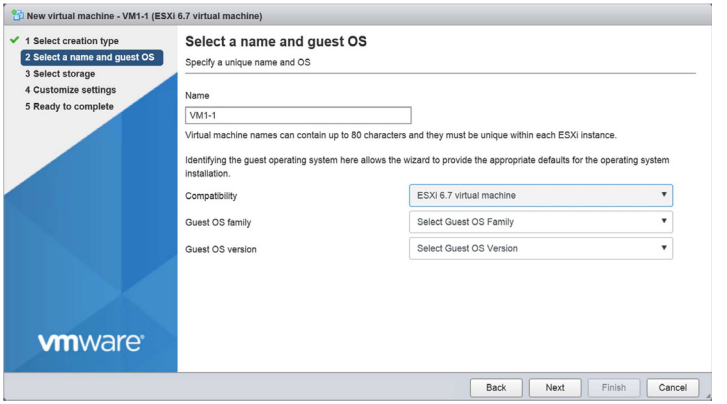
You can use the New virtual machine wizard in vSphere Web Client to create a virtual machine.



https://t.me/learningnets

3-7 Creating Virtual Machines with the New Virtual Machine Wizard (2)

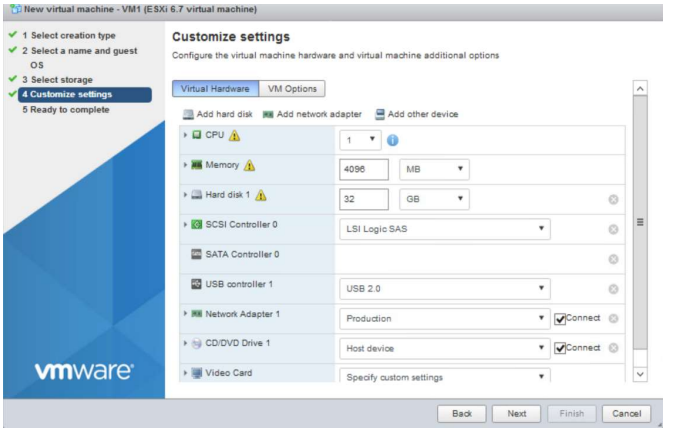
You can use the New virtual machine wizard in VMware Host Client to create a virtual machine.



https://t.me/learningnets

3-8 New Virtual Machine Wizard Settings

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



https://t.me/learningnets

3-9 Installing the Guest Operating System

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



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

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

For more information about installing guest operating systems, see *vSphere Virtual Machine Administration* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

For more about the supported guest operating systems, see *Hardware and Guest Operating System Compatibility Guides* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

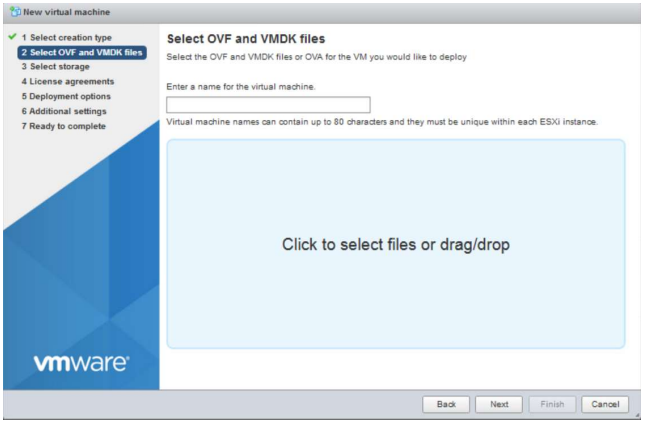
https://t.me/learningnets

3-10 Deploying OVF Templates

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

Virtual appliances are preconfigured virtual machines:

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



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

A virtual appliance can be added or imported to your vCenter Server system inventory or ESXi inventory. Virtual appliances can be imported from 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 virtual machines. OVF files are compressed, enabling faster downloads. vSphere Web Client validates an OVF file before importing it and ensures that it is compatible with the intended destination server. If the appliance is incompatible with the selected host, you cannot import it.

https://t.me/learningnets

3-11 About VMware Tools

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

VMware Tools benefits:

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

VMware Tools features:

- Guest OS heartbeat service
- Time synchronization
- Ability to shut down the virtual machine

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

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

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

- Significantly faster graphics performance and Windows Aero on operating systems that support Aero.
- Shared folders between host and guest file systems.
- Copying and pasting text, graphics, and files between the virtual machine and the host or client desktop.
- Improved mouse performance.
- Synchronization of the clock in the virtual machine with the clock on the host or client desktop.

https://t.me/learningnets

- Scripting that helps automate guest operating system operations.
- The ability to shut down the virtual machine.

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

https://t.me/learningnets

3-12 Managing VMware Tools

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

VMware Tools 10.1 provides the following features:

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

You can download ISO images for other operating systems from VMware.

vSphere 6.5 and later provides enhancements to VMware Tools by implementing security, reducing distribution media volume size that do not require as much space, offloading ISOs to an on-disk product locker, and providing access for other ISO images needed to be downloaded from the VMware website.

For more detailed information, see VMware knowledge base article 2129825 at https://kb.vmware.com/s/article/2129825.

https://t.me/learningnets

3-13 VMware Tools: Supported ISO Images

The following ISO files are included with vSphere 6.5 and later:

- `windows.iso`: For Vista and later guests
- `winPreVista.iso`: For Windows 2000, XP, and Server 2003 guests
- `linux.iso`: For Linux OS with glibc 2.5 or higher (for example, RHEL 5 or later, SLES 11 or later, Ubuntu 10.04 or later)

VMware Tools for other guest operating systems, such as Solaris, and macOS, can be downloaded from My VMware at https://download.vmware.com.

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

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

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

https://t.me/learningnets

3-14 Troubleshooting an OS Installation in Virtual Machines

Problem: The installation of a 64-bit operating system fails or hangs during installation.

Solution:

1. Verify that the guest operating system that you are trying to install is fully certified by VMware.
2. Verify that your VMware ESXi host meets the hardware and firmware requirements for running 64-bit virtual machines.

https://t.me/learningnets

3-15 Troubleshooting a VMware Tools Installation on a Guest Operating System

Problem: VMware Tools fails to install or hangs during installation.

Solution:

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

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

https://t.me/learningnets

3-16 Lab 1: Deploying and Configuring Virtual Machines

Create and prepare a virtual machine for use

1. Access Your Student Desktop
2. Create a Virtual Machine
3. Install VMware Tools
4. Copy Files to the Desktop

https://t.me/learningnets

3-17 Review of Learner Objectives

You should be able to meet the following objectives:

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

https://t.me/learningnets

3-18 Key Points

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

Questions?

https://t.me/learningnets

MODULE 4

vCenter Server



https://t.me/learningnets

4-2 You Are Here

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

https://t.me/learningnets

4-3 Importance

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

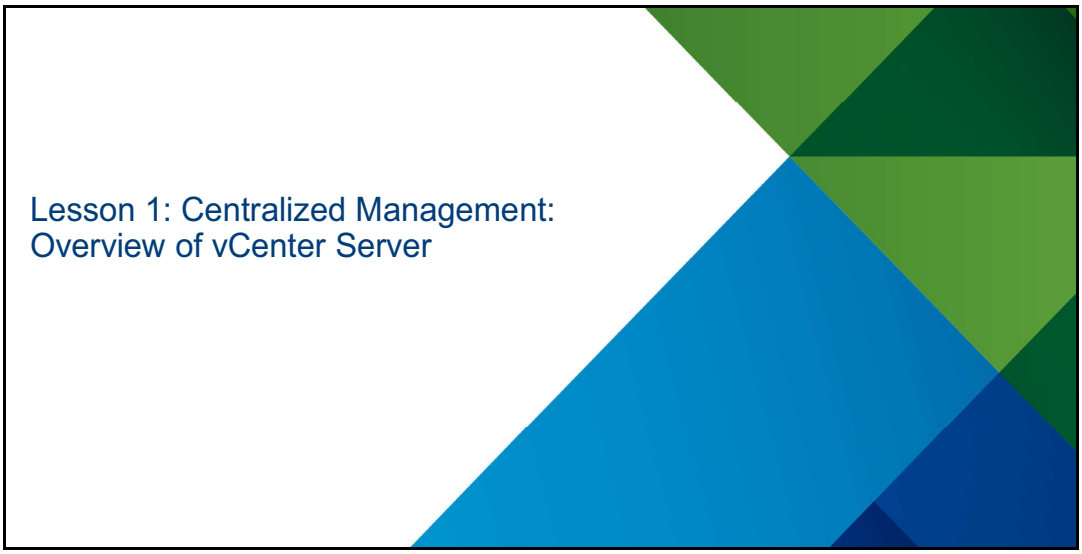
https://t.me/learningnets

4-4 Module Lessons

Lesson 1:	Centralized Management: Overview of vCenter Server
Lesson 2:	vCenter Server Appliance
Lesson 3:	Deploying vCenter Server Appliance
Lesson 4:	Managing the vCenter Server Inventory
Lesson 5:	vCenter Server Roles and Permissions
Lesson 6:	Backing Up and Restoring vCenter Server Appliance
Lesson 7:	Monitoring vCenter Server Appliance

https://t.me/learningnets

4-5 Lesson 1: Centralized Management: Overview of vCenter Server



https://t.me/learningnets

4-6 Learner Objectives

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

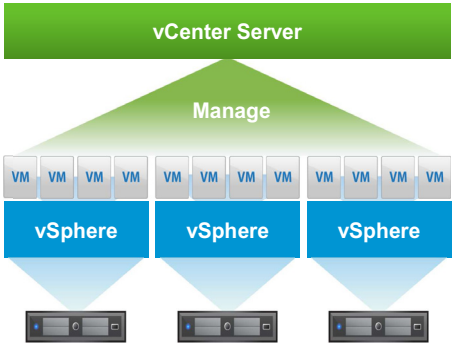
- Describe the vCenter Server architecture
- Discuss how ESXi hosts communicate with vCenter Server
- Identify the vCenter Server services, components, and modules
- Explain VMware Platform Services Controller™

https://t.me/learningnets

4-7 About the vCenter Server Management Platform

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

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



VMware has announced that 6.7 is the final release of vCenter Server on Windows.

vCenter Server is a service that acts as a central administrator for ESXi hosts connected in a network. vCenter Server enables you to pool and manage the resources of multiple hosts.

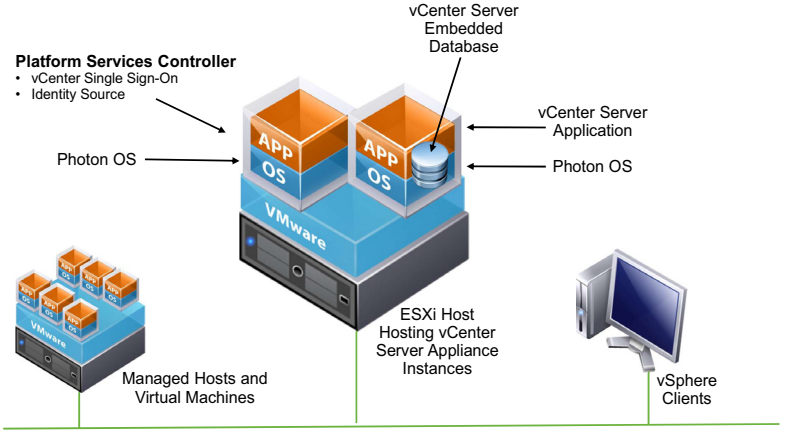
You can deploy vCenter Server Appliance on an ESXi host in your infrastructure. vCenter Server Appliance is a preconfigured Linux-based virtual machine that is optimized for running vCenter Server and the vCenter Server components.

vCenter Server for Windows and vCenter Server Appliance provide advanced features, such as VMware vSphere® Distributed Resource Scheduler™, VMware vSphere® High Availability, VMware vSphere® Fault Tolerance, VMware vSphere® vMotion®, and VMware vSphere® Storage vMotion®.

https://t.me/learningnets

4-8 vCenter Server Architecture

The diagram shows the supporting components for vCenter Server.



The vCenter Server architecture relies on the following components:

- vSphere Web Client and vSphere Client: These clients are used to connect to vCenter Server in order to centrally manage your ESXi hosts. When an ESXi host is managed by vCenter Server, administrators should always use vCenter Server and the vSphere clients to manage that host.
- vCenter Server database: The vCenter Server database is the most important component. The database stores inventory items, security roles, resource pools, performance data, and other critical information for vCenter Server.
- Managed hosts: vCenter Server enables you to manage ESXi hosts and the virtual machines that run on them. vCenter Server provides access to the ESXi host through a vCenter Server agent, which is started on the host when it is added to the vCenter Server inventory.
- VMware vCenter® Single Sign-On: vCenter Single Sign-On is one of the services running on the Platform Services Controller. It provides a security domain defined in your vSphere environment. Authentication is performed by the vCenter Single Sign-On service. The vCenter Single Sign-On server can be configured to authenticate against multiple user repositories, also called identity sources, such as an Active Directory domain.

Distributed services include vSphere vMotion, vSphere DRS, and vSphere HA, which are installed with vCenter Server.

Additional services are packaged as part of the base product, for example, VMware vSphere® Update Manager™ and VMware vRealize® Orchestrator™. No additional license is necessary.

https://t.me/learningnets

4-9 vCenter Server Services

The vCenter Server group of services includes:

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



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

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

- vSphere Web Client
- vCenter Server

vSphere Web Client is not a client application. It is a web server that is based on the Apache Tomcat web server.

In addition, the vCenter Server component also includes the following services:

- VMware vSphere® ESXi™ Dump Collector
- vSphere Update Manager
- vSphere Auto Deploy

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

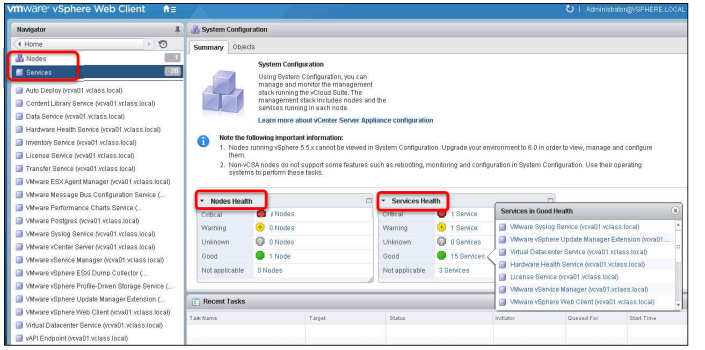
You can use vSphere Update Manager with either vCenter Server on Windows or with vCenter Server Appliance.

https://t.me/learningnets

4-10 Monitoring the Health and Status of Services and Nodes Across vCenter Server Systems

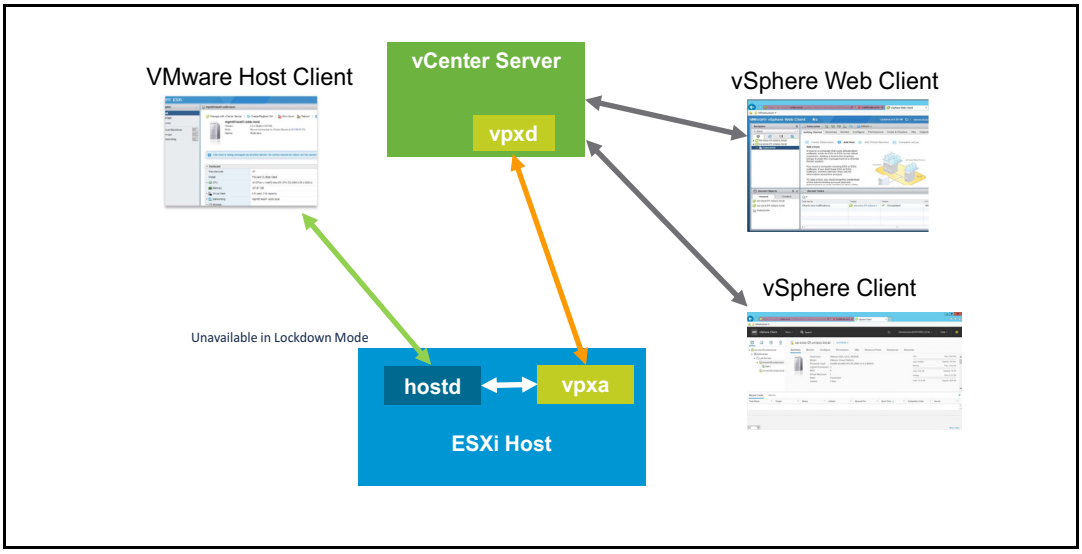
vSphere Web Client enables you to monitor the status of all manageable services and nodes across vCenter Server systems at **Administration > Deployment > System Configuration > Services**.

A list of default services is available in each vCenter Server instance.



https://t.me/learningnets

4-11 ESXi and vCenter Server Communication



vSphere Web Client and vSphere Client communicate directly with vCenter Server. To communicate directly with an ESXi host, you use VMware Host Client.

vCenter Server provides direct access to the ESXi host through a vCenter Server agent named virtual provisioning X agent (vpxa). The vpxa process is automatically installed onto 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).

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

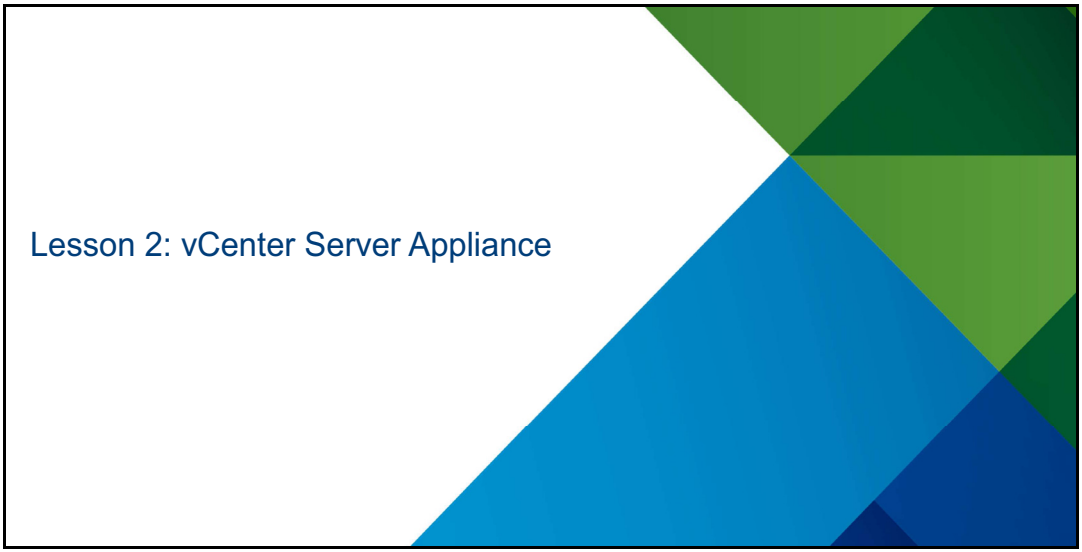
https://t.me/learningnets

4-12 Review of Learner Objectives

You should be able to meet the following objectives:

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

https://t.me/learningnets



https://t.me/learningnets

4-14 Learner Objectives

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

- Identify the purpose of vCenter Server Appliance
- Describe the vCenter Server architecture
- Discuss the REST-based API
- Describe VMware vCenter Server® High Availability

https://t.me/learningnets

4-15 Overview of vCenter Server Appliance (1)

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

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

The vCenter Server Appliance package contains the following software:

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

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

https://t.me/learningnets

4-16 Overview of vCenter Server Appliance (2)

vCenter Server Appliance uses the embedded PostgreSQL database, which can scale up to 2,000 hosts and 35,000 registered virtual machines or 25,000 powered-on virtual machines.

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

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

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

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

https://t.me/learningnets

4-17 vCenter Server Appliance Scalability

vCenter Server Appliance now scales to the same capacity as vCenter Server installed on a Windows server.

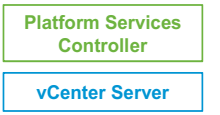
Metric	vCenter Server 6.7 on Windows	vCenter Server Appliance 6.7
Hosts per vCenter Server system	2,000	2,000
Powered-on virtual machines per vCenter Server System	25,000	25,000
Hosts per cluster	64	64
Virtual machines per cluster	8,000	8,000
Database	Must be Oracle or SQL for full scalability	PostgreSQL
Enhanced Linked Mode	Yes	Yes

https://t.me/learningnets

4-18 About Platform Services Controller

vCenter Server includes Platform Services Controller:

- Platform Services Controller includes a set of common infrastructure services:
  - vCenter Single Sign-On
  - VMware License Server
  - Lookup Service
  - VMware Certificate Authority
  - Certificate Store
  - VMware Directory Services



- Other features are installed under the vCenter Server component.
- You can install vCenter Server and Platform Services Controller on the same or different machines.
- When you install a Platform Services Controller instance, you are prompted to create a vCenter Single sign-on domain or join an existing domain.

vCenter Server has VMware Platform Services Controller. This system includes a set of common infrastructure services. Some of these, such as vCenter Single Sign-On, were present under vCenter Server in vSphere 5.x. Other functions, such as an internal certificate authority and certificate store, are present in vSphere 6.x and are subsystems designed to make vCenter Server more robust.

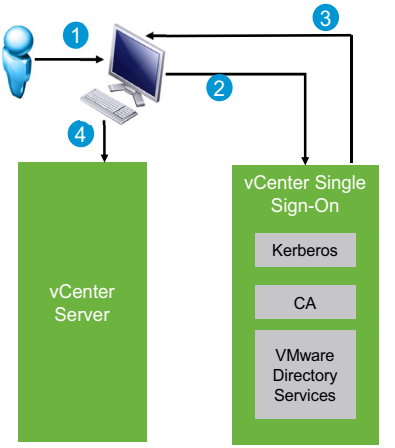
https://t.me/learningnets

4-19 About vCenter Single Sign-On

vCenter Single Sign-On enables vSphere components to communicate with each other through a secure token mechanism:

- Users authenticate to vCenter Single Sign-On by entering their credentials on the login page.
- vCenter Single Sign-On authenticates users' credentials with the configured identity source (for example, Active Directory).
- If successful, vCenter Single Sign-On returns a token that represents the user to vSphere Client.
- The vCenter Single Sign-On server returns the token to the vCenter Server system, using the vCenter Server Authorization Framework to allow user access.

After connecting to vCenter Server, authenticated users can view all vCenter Server instances or other vSphere objects for which their role gives them privileges. No further authentication is required.

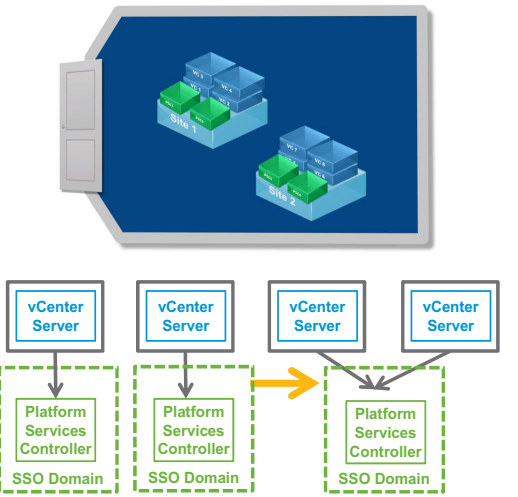


https://t.me/learningnets

4-20 vCenter Single Sign-On Domains and Multiple Sites

A vCenter Single Sign-On domain can be divided into multiple sites:

- For high availability or for multisite environments, deploy multiple Platform Services Controller instances, all of which synchronize their configurations and vCenter Single Sign-On domains.
- Each Platform Services Controller instance and vCenter Server instance can be associated with only one site.
- Cross vCenter Single Sign-On domain repointing allows a vCenter Server Appliance instance to move to a new vCenter Single Sign-On domain.
- vCenter Server instances that exist in separate vCenter Single Sign-On domains can be consolidated into one vCenter Single Sign-On domain.

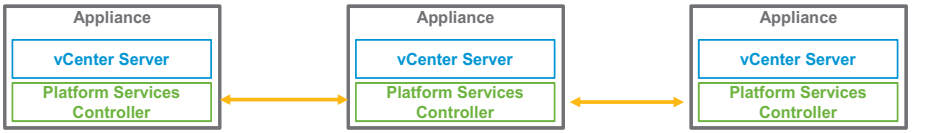


https://t.me/learningnets

4-21 About Enhanced Linked Mode

Enhanced Linked Mode allows you to log in to any single instance of vCenter Server Appliance or 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. Embedded Enhanced Linked Mode supports up to 10 vCenter Server instances.
- An Enhanced Linked Mode group can be created only during the deployment of vCenter Server Appliance or installation of vCenter Server.
- Enhanced Link Mode supports the following deployment configurations:
  - vCenter Server with embedded Platform Services Controller
  - vCenter Server with external Platform Services Controller



You can create an Enhanced Linked Mode group only during the deployment of vCenter Server Appliance or installation of vCenter Server. You cannot create an Enhanced Linked Mode group after you install vCenter Server or after you deploy vCenter Server Appliance.

Enhanced Linked Mode provides the following features for both vCenter Server on Windows and vCenter Server Appliance systems:

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

To join vCenter Server systems in Enhanced Linked Mode, connect them to the same Platform Services Controller instances, or to the Platform Services Controller instances that share the same vCenter Single Sign-On domain.

Enhanced Linked Mode requires the vCenter Server Standard licensing level, and it is not supported with VMware vCenter® Server Foundation™ or VMware vCenter® Server for Essentials™.

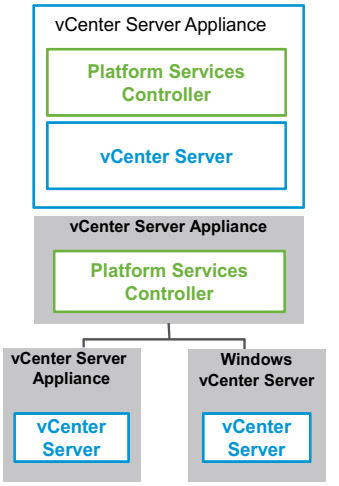
https://t.me/learningnets

4-22 vCenter Server Deployment Options

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

vCenter Server Appliance can be configured in the following ways:

- As an embedded system with an internal Platform Services Controller instance
- As a distributed system with an external Platform Services Controller instance



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

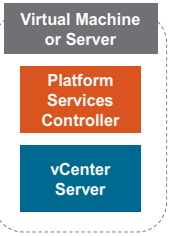
You can download the vCenter Server Appliance ISO file and run the installer. Native installers for Windows, Linux, and macOS are available.

https://t.me/learningnets

4-23 Embedded Platform Services Controller

Using an embedded Platform Services Controller instance results in a standalone deployment that has its own vCenter Single Sign-On domain with a single site (same virtual machine or physical server):

- The connection between vCenter Server and the embedded Platform Services Controller instance is not over the network, and vCenter Server is not prone to outages caused by connectivity and name-resolution issues between vCenter Server and Platform Services Controller.
- Embedded Platform Services Controller is suitable for small environments.
- Embedded Platform Services Controller cannot join other vCenter Server or Platform Services Controller instances.



vCenter Server with embedded Platform Services Controller:

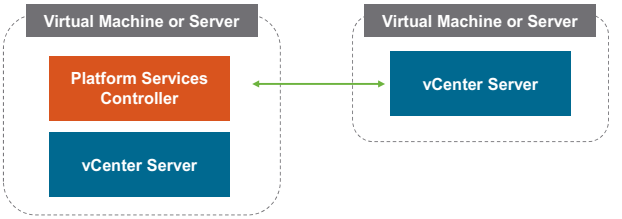
- Single virtual machine containing vCenter Server and Platform Services Controller services.
- Embedded Platform Services Controller supports VMware vCenter® Server High Availability.
- Each vCenter Server with embedded Platform Services Controller contains an isolated vCenter Single Sign-On domain.

https://t.me/learningnets

4-24 External vCenter Server System Linked to an Embedded vCenter Server System

All vCenter Server instances that are registered with one common or different joined Platform Services Controller instances are connected in Enhanced Linked Mode:

- vCenter Server instances can be made highly available by using vCenter Server High Availability (appliance only).
- Platform Services Controller is a single point of failure.



For an embedded Platform Services Controller and vCenter Server system linked with an external, standalone vCenter Server system, linking a second vCenter Server system to an existing embedded vCenter Server system and Platform Services Controller is possible.

When you install vCenter Server or deploy vCenter Server Appliance with an external Platform Services Controller instance, you must first install Platform Services Controller. During installation of Platform Services Controller, you can select whether to create a vCenter Single Sign-On domain or join an existing domain.When you join an existing vCenter Single Sign-On domain, the infrastructure data between the existing Platform Services Controller instance and the new Platform Services Controller instance is replicated.

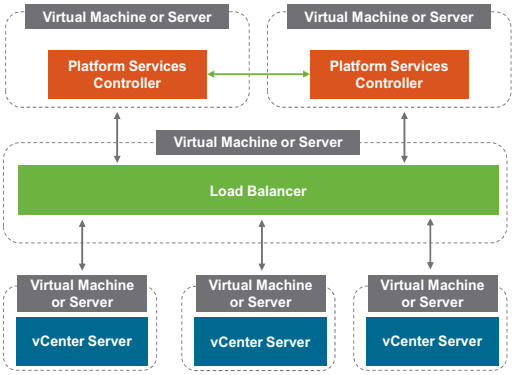
For more information about Platform Services Controller recommendations, see VMware knowledge base article 2108548 at http://kb.vmware.com/kb/2108548.

For more information about backing up and restoring vCenter Server 6.x external deployment models, see the VMware knowledge base article 2147672 at http://kb.vmware.com/kb/2147672.

https://t.me/learningnets

4-25 Platform Services Controller with Load Balancer

You can use a third-party load balancer per site to configure Platform Services Controller high availability with automatic failover for this site.

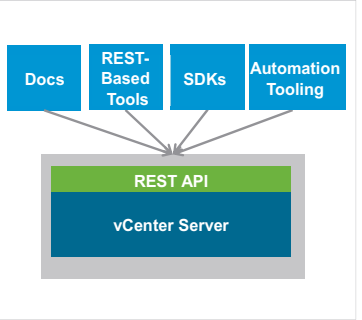


To provide higher levels of availability, two or more Platform Services Controller instances can be configured to use the same vCenter Single Sign-On domain and be placed behind a load balancer. Each of the vCenter Server instances is configured to communicate with the load balancer, which uses the Platform Services Controller instances in an active-passive configuration. If the active Platform Services Controller instance fails, the load balancer redirects to the second Platform Services Controller instance, and the vCenter Server instances continue to function.

https://t.me/learningnets

4-26 vCenter Server APIs

vSphere 6.x includes a developer-friendly, automation-friendly REST-based API and interfaces that simplify automation and development.



Overview

- REST API for VM management
- Simplified and modern API design
- Enables easy access to automation from standard tooling
- Designed with automation and DevOps in mind
- Full featured SDKs, CLI, and workflow access

Benefits

Key use case: VM automation, development, and vCenter Server Appliance services health check

Benefits:

- Reduces API development complexity and time
- Access through modern automation and development processes for consistency and repeatability
- Simplified API model
- Single point of access for all API samples

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

The following SDKs are included in this release:

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

Each of these SDKs contains:

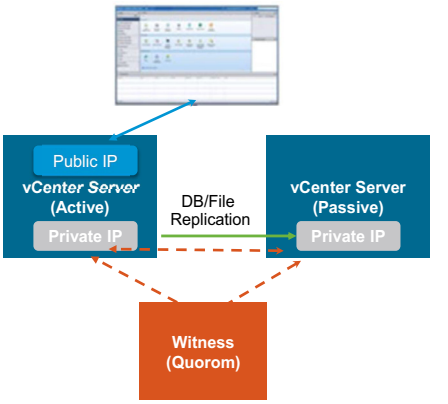
- Bindings to the vSphere Automation API
- vSphere Web Services API
- API reference
- Programming guide
- Samples

https://t.me/learningnets

4-27 High Availability for vCenter Server Appliance

vCenter Server High Availability protects against both hardware and software failures and ensures that your implementation can recover quickly:

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



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

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

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

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

Before you can configure vCenter Server High Availability, you must consider several factors. A vCenter Server Appliance deployment can use an internal or external Platform Services Controller instance. A brownfield deployment with components that use different versions of vSphere requires different considerations than a greenfield deployment that includes only vSphere 6.x components. Resource and software requirements and the networking setup must also be considered carefully.

https://t.me/learningnets

4-28 Review of Learner Objectives

You should be able to meet the following objectives:

- Identify the purpose of vCenter Server Appliance
- Describe the vCenter Server architecture
- Discuss the REST-based API
- Describe VMware vCenter Server® High Availability

https://t.me/learningnets



https://t.me/learningnets

4-30 Learner Objectives

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

- Deploy vCenter Server Appliance into an infrastructure
- Add license keys to vCenter Server
- Configure vCenter Server settings

https://t.me/learningnets

4-31 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.
- For the first installation of vCenter Server Appliance, Platform Services Controller must be deployed before vCenter Server:
  - If you deploy vCenter Server Appliance with an embedded Platform Services Controller instance, this operation occurs automatically.
  - If you install vCenter Server Appliance with an external Platform Services Controller instance, you must first install Platform Services Controller and then install vCenter Server.
  - You must provide the fully qualified domain name (FQDN) or the static IP of the host machine on which you are performing the install or upgrade. VMware recommends using the FQDN.
  - You must ensure that clocks on all machines on the vSphere network are synchronized.

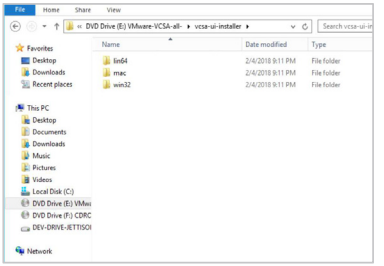
For more information, see *vSphere Installation and Setup* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

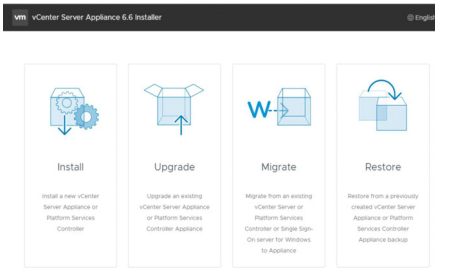
https://t.me/learningnets

4-32 vCenter Server Appliance Native GUI Installer

With vSphere 6.x, a native application facilitates the deployment of vCenter Server Appliance 6.x:

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



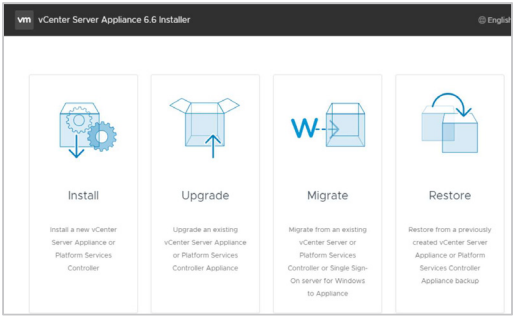


https://t.me/learningnets

4-33 vCenter Server Installation, Upgrade, Migration, and Restore

The GUI deployment tool has the following featured options:

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



https://t.me/learningnets

4-34 vCenter Server Appliance Deployment

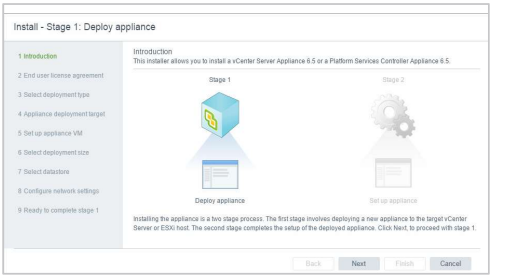
The installer is available for Windows, macOS, and Linux.

vSphere Update Manager is included.

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

- Stage 1: Deploy OVF
- Stage 2: Configuration

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



When you choose to deploy vCenter Server Appliance with an embedded Platform Services Controller instance, you deploy Platform Services Controller and vCenter Server as one appliance.

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

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

https://t.me/learningnets

4-35 Two-Stage Deployment of vCenter Server Appliance

Stage 1a, UI phase:

1. Accept the EULA.
2. Select the deployment type.
3. Connect to the target ESXi host or vCenter Server system to deploy vCenter Server Appliance.
4. Define the vCenter Server Appliance name and root password.
5. Select deployment size (memory, CPU), storage size, and data store location (thin disk).
6. Define networking settings.

Stage 1b, Deployment phase:

1. OVF is deployed to the ESXi host.
2. Disks and Networking are configured.

Stage 2, Configuration phase:

1. Create a vCenter Single Sign-On domain, or join an existing domain.
2. Join the Customer Experience Improvement Program (CEIP).

https://t.me/learningnets

4-36 vCenter Server Deployment Wizard

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

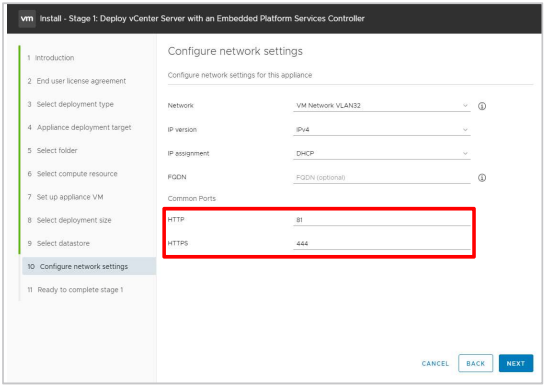
	Embedded	Platform Services Controller	vCenter Server
Stage 1			
Deployment target	✓	✓	✓
Deployment type	✓	✓	✓
Deployment size	✓	✗	✓
Define VM name	✓	✓	✓
Define root password	✓	✓	✓
Define datastore location	✓	✓	✓
Define networking	✓	✓	✓
Stage 2			
Create vCenter Single Sign-On domain	✓	✓	✗
Join vCenter Single Sign-On domain	✗	✓	✓
Configure CEIP	✓	✓	✓

https://t.me/learningnets

4-37 Configuring Custom Ports for a vCenter Server Appliance Deployment

You configure custom HTTP and HTTPS ports for vCenter Server Appliance during installation.

Custom ports might be a requirement for security reasons.



vCenter Server Appliance supports defining custom HTTP and HTTPS ports during installation. The default ports are 80 and 443. The ports cannot be changed after installation.

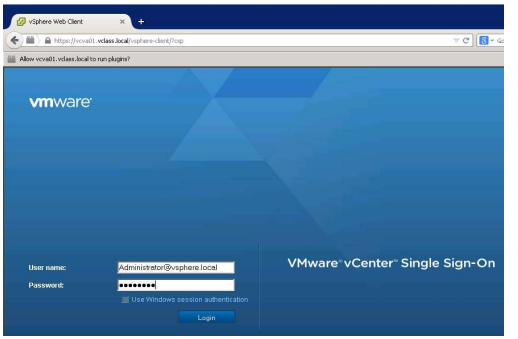
A customer might want to define custom HTTP/HTTPS ports if a corporate security policy disallows the use of default port values.

https://t.me/learningnets

4-38 Getting Started with vCenter Server

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

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

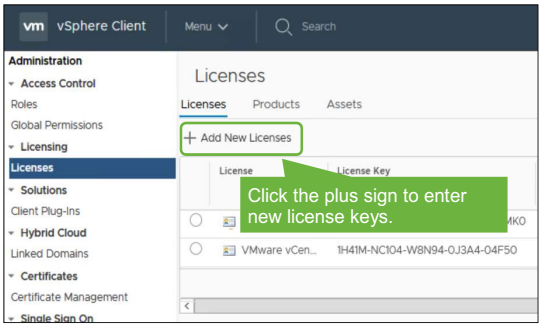


https://t.me/learningnets

4-39 Adding License Keys to vCenter Server

You must assign a license to vCenter Server before its 60-day evaluation period expires.

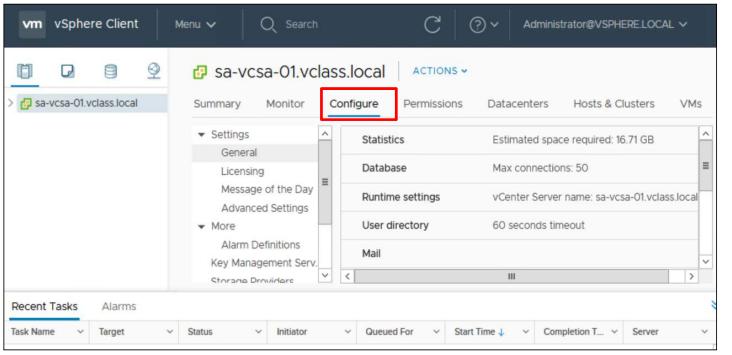
In vSphere Client, select **Home > Administration > Licenses**.



https://t.me/learningnets

4-40 Configuring vCenter Server Settings

You can configure your vCenter Server system from vSphere Client, including settings such as licensing, statistics collection, logging, and other settings. To access the vCenter Server system settings, navigate to the vCenter Server system in vSphere Client, click the **Configure** tab, and select **Settings**.



https://t.me/learningnets

4-41 Demonstration: Working with vCenter Server Appliance

Your instructor will run a demonstration.

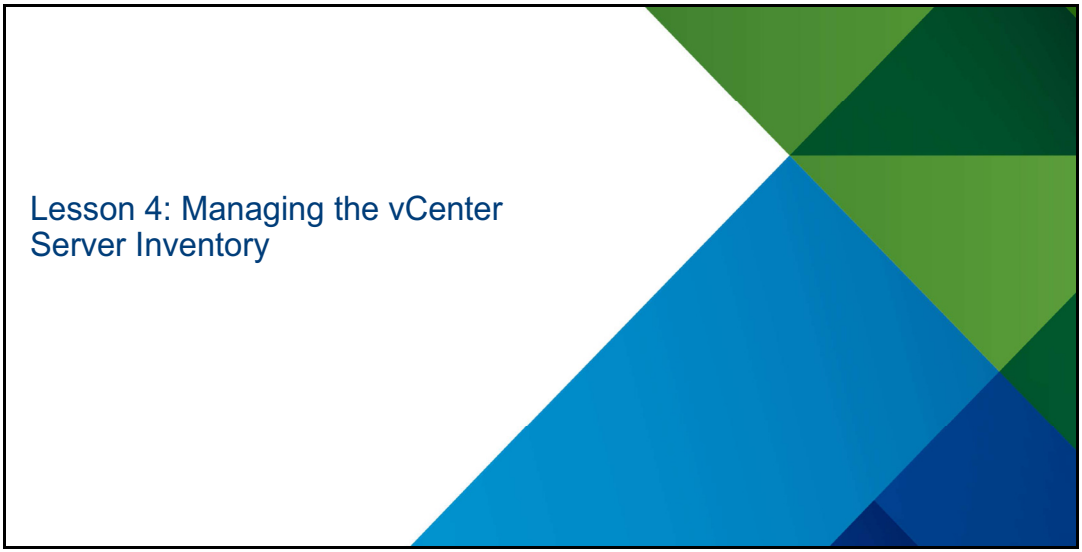
https://t.me/learningnets

4-42 Review of Learner Objectives

You should be able to meet the following objectives:

- Deploy vCenter Server Appliance into an infrastructure
- Add license keys to vCenter Server
- Configure vCenter Server settings

https://t.me/learningnets



https://t.me/learningnets

4-44 Learner Objectives

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

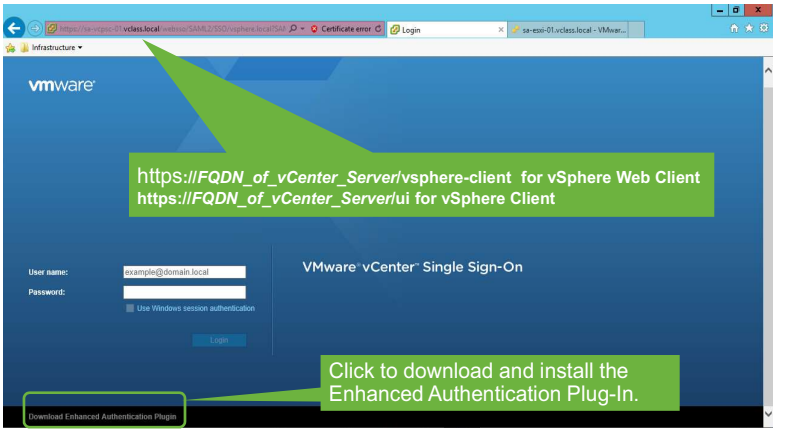
- Use vSphere Client to manage the vCenter Server inventory
- Create and organize vCenter Server inventory objects
- Add data center and organizational objects to vCenter Server
- Add hosts to vCenter Server
- Discuss how to create custom inventory tags for inventory objects

.

https://t.me/learningnets

4-45 Accessing vSphere Clients

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

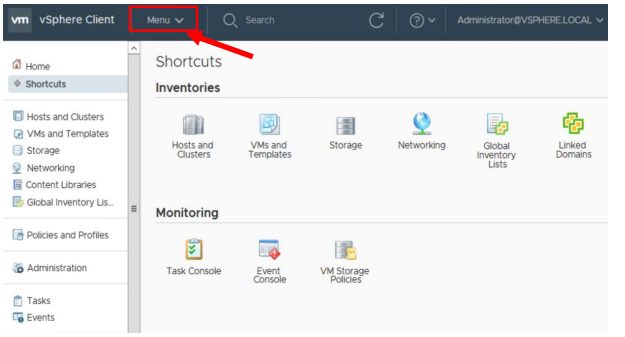


https://t.me/learningnets

4-46 vSphere Client Home Page

Select **Menu > Home** to go to the vSphere Client Home page.

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

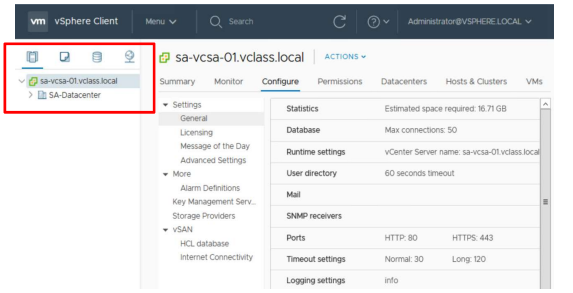


From the vSphere Client Home page, you can manage your vCenter Server system inventory, monitor your infrastructure environment, and complete system administration tasks.

https://t.me/learningnets

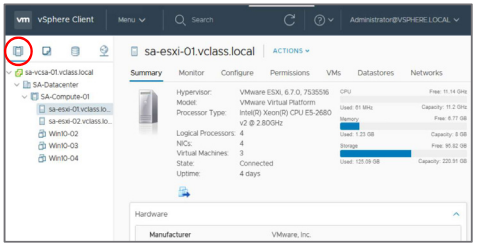
4-47 Using the vSphere Client Navigation Pane

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

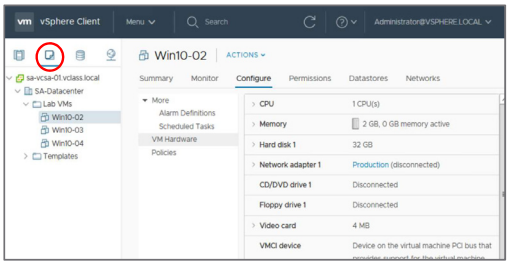


https://t.me/learningnets

4-48 vCenter Server Views: Hosts and Clusters, VMs and Templates



Hosts and Clusters Inventory View



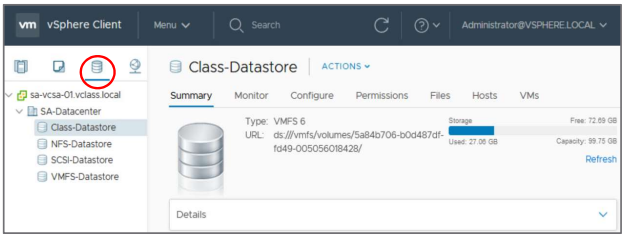
VMs and Templates Inventory View

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

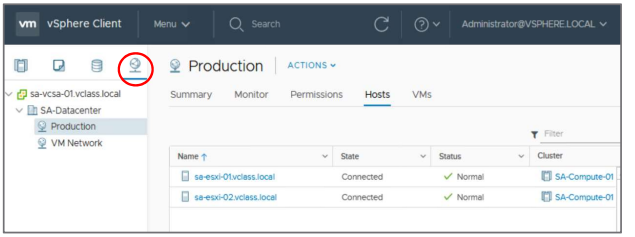
The VMs and Templates inventory view displays all virtual machine and template objects in a data center.

https://t.me/learningnets

4-49 vCenter Server Views: Storage and Networks



Storage Inventory View



Networks Inventory View

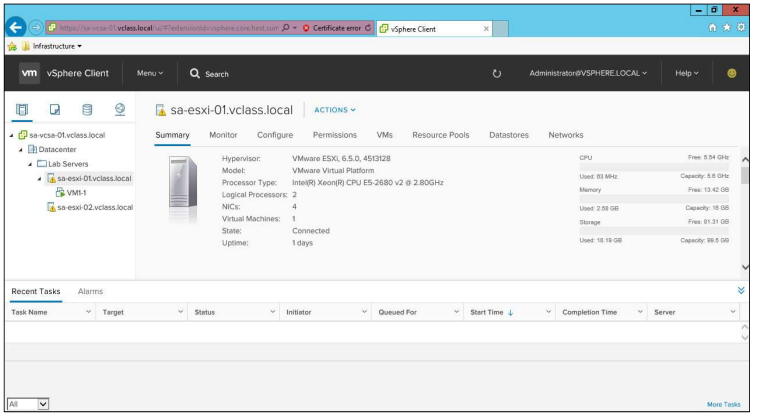
The Storage inventory view displays all the details for datastores in the data center. The Networking inventory view displays all virtual machine port groups and distributed switches.

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

https://t.me/learningnets

4-50 Viewing Object Information

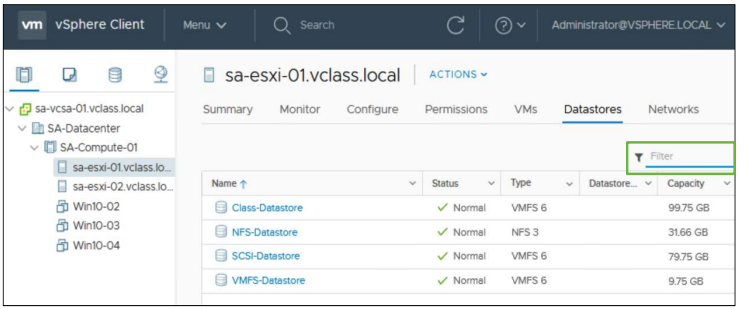
Because you can navigate to view object information and access related objects, monitoring and managing object properties is easy.



https://t.me/learningnets

4-51 Using Quick Filters in vSphere Client

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



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

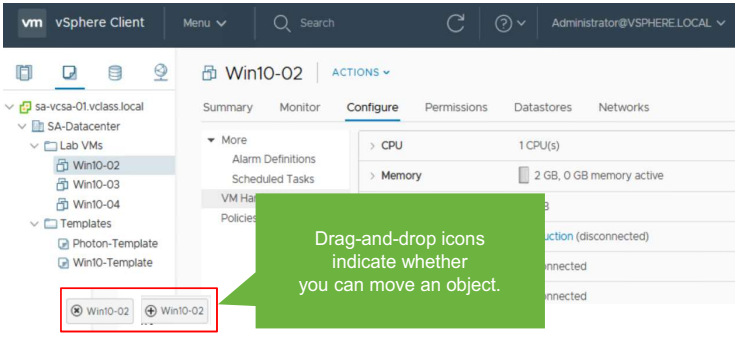
Quick filters are available in the list views, which appear on the **Objects** tab of an inventory list and in search results.

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

https://t.me/learningnets

4-52 Dragging Objects in vSphere Client

You can drag an inventory object to another location.

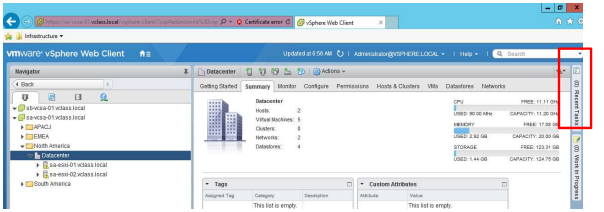


https://t.me/learningnets

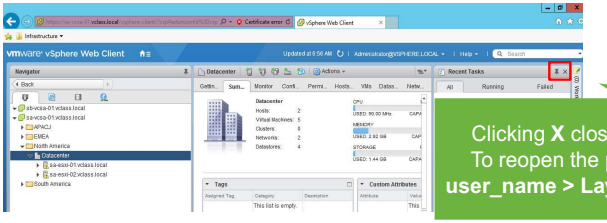
4-53 Using Pin and Unpin Functionality in vSphere Web Client

You can pin and unpin display panes in the UI.

Unpinned Recent Tasks



Pinned Recent Tasks

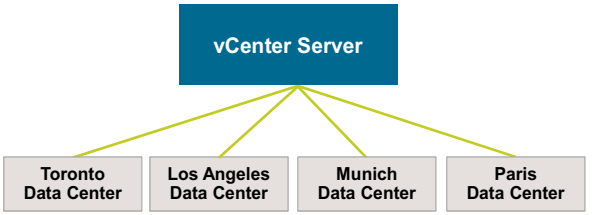


Clicking **X** closes the pane. To reopen the pane, select **user_name > Layout Settings.**

4-54 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 virtual machines:

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

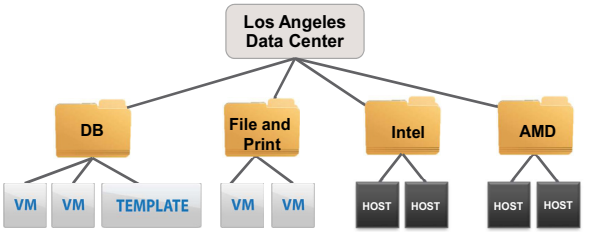


A virtual data center is a container for all the inventory objects required to complete a fully functional environment for operating virtual machines. You can create multiple data centers to organize sets of environments. For example, you might create a data center for each organizational unit in your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

https://t.me/learningnets

4-55 Organizing Inventory Objects into Folders

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



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

Populating and organizing your inventory involves the following tasks:

- Create data centers.
- Add hosts to the data centers.
- Organize inventory objects in folders.
- Set up networking by using vSphere standard switches or vSphere distributed switches. To use services such as vSphere vMotion, TCP/IP storage, VMware vSAN™, and vSphere Fault Tolerance, set up VMkernel networking for these services.

  For more information, see *vSphere Networking* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.
- Configure storage systems and create datastore inventory objects to provide logical containers for storage devices in your inventory.

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

https://t.me/learningnets

- Create clusters to consolidate the resources of multiple hosts and virtual machines. You can enable vSphere HA and vSphere DRS for increased availability and more flexible resource management.

  For information about configuring vSphere HA, see *vSphere Availability* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

  For information about configuring vSphere DRS, see *vSphere Resource Management* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

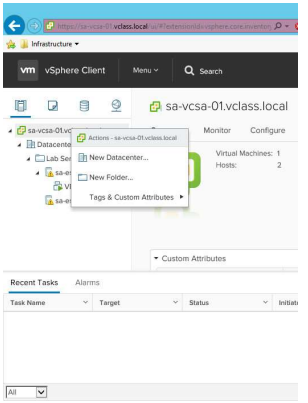
- Create resource pools to provide logical abstraction and flexible management of the resources in vSphere. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources.

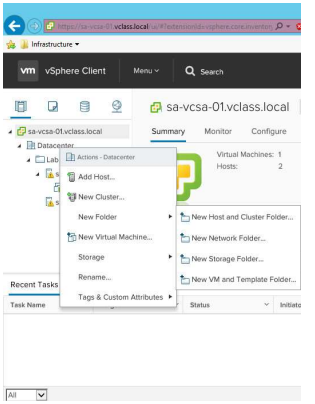
  For details, see *vSphere Resource Management* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

4-56 Adding a Data Center and Organizational Objects to vCenter Server

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

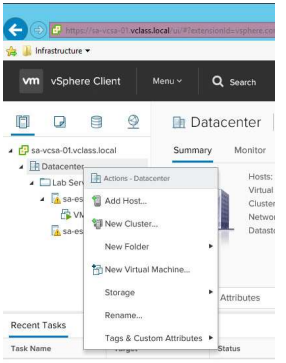


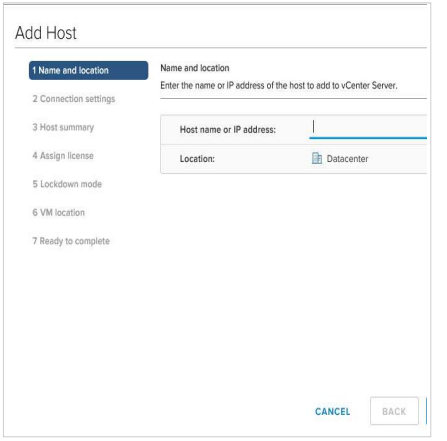


https://t.me/learningnets

4-57 Adding ESXi Hosts to vCenter Server

You can add an ESXi host to vCenter Server.





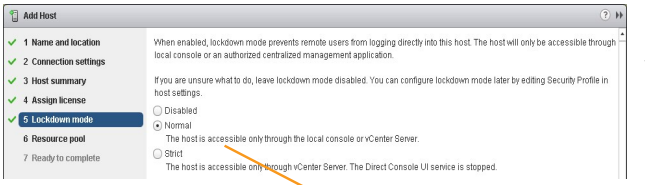
https://t.me/learningnets

4-58 Configuring Lockdown Mode

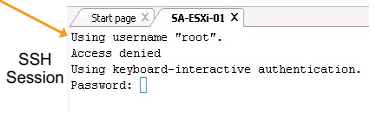
To increase the security of your ESXi hosts, you can put your hosts in lockdown mode. Two lockdown modes are available: normal and strict.

When you enable normal lockdown mode, only vpxuser has authentication permission. Also, users cannot perform operations against the host directly.

Only the root user can log in to the DCUI.



vSphere Web Client



SSH Session

Lockdown modes forces all host management to be performed through vCenter Server. When a host is in lockdown mode:

- You cannot manage it directly through VMware Host Client.
- You cannot run commands on it from vSphere CLI.
- External software or management tools that access the host directly cannot retrieve or modify information on the ESXi host, unless the credentials are added as exception users on the ESXi host.

If normal lockdown mode is enabled, only the root user is authorized to log in to the DCUI. If strict lockdown mode is enabled on a host, the DCUI service is automatically disabled. As result, even the root user cannot log in to the host.

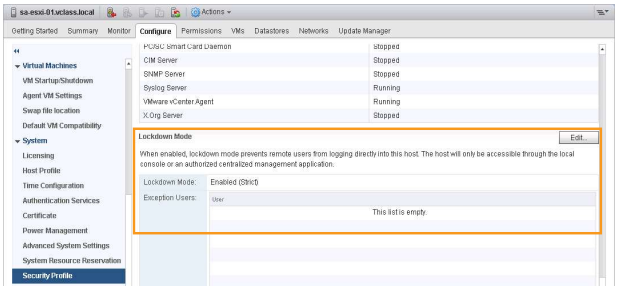
https://t.me/learningnets

4-59 Strict Lockdown Mode

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

If the connection to the vCenter Server system is lost and vSphere Client is no longer available, the ESXi host cannot be managed remotely.

In this situation, the host can be accessed only if VMware vSphere® ESXi™ Shell and SSH services are enabled and authorized service accounts are added to the exception users list.



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

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

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

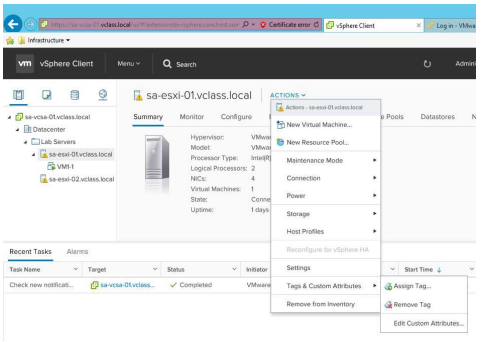
https://t.me/learningnets

4-60 Creating Custom Tags for Inventory Objects

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

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

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



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

https://t.me/learningnets

4-61 Lab 2: Working with vCenter Server Appliance

Configure and use vCenter Server Appliance

1. Access Your vCenter Server Appliance and Configure Licenses
2. Configure Single Sign-On and Create a Data Center Object
3. Add Your ESXi Hosts to the vCenter Server Inventory
4. Configure the ESXi Hosts as NTP Clients
5. Create a Host and Cluster Folder
6. Create Virtual Machine and Template Folders
7. Navigate vSphere Client

https://t.me/learningnets

4-62 Review of Learner Objectives

You should be able to meet the following objectives:

- Use 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
- Discuss how to create custom inventory tags for inventory objects

https://t.me/learningnets



https://t.me/learningnets

4-64 Learner Objectives

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

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

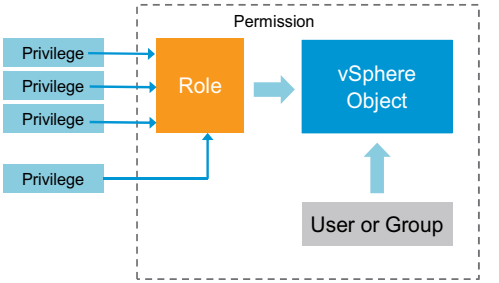
https://t.me/learningnets

4-65 About vCenter Server Permissions

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

The following concepts are important:

- Permissions: The permission model for vCenter Server systems relies on assigning permissions to objects in the vSphere object hierarchy. Each permission gives one user or group a set of privileges, that is, a role for the selected object.
- Privilege: An action that can be performed.
- Role: A set of privileges.
- Object: The target of the action.
- User or group: Indication of who can perform the action.



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

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

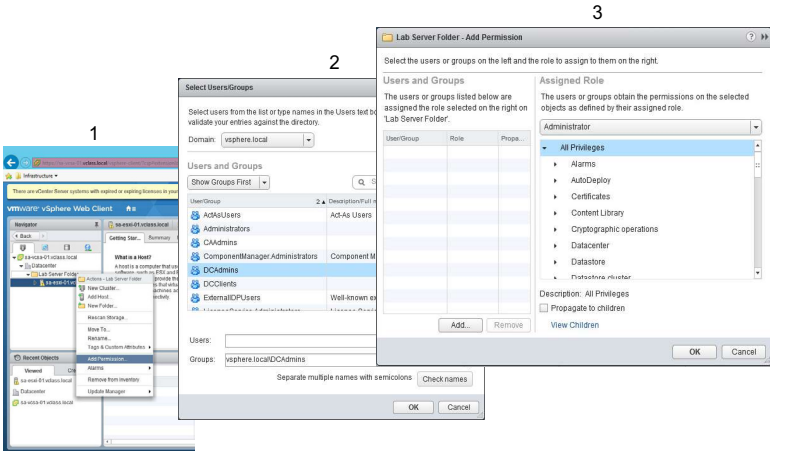
A permission is set on an object in the vCenter Server object hierarchy. Each permission associates the object with a group or user and the group or user access roles. For example, you can select a virtual machine object, add one permission that gives the ReadOnly role to group 1, and add a second permission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

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

https://t.me/learningnets

4-66 Adding Permissions to the vCenter Server Inventory

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



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

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

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

https://t.me/learningnets

You can also assign permissions to a global root object to apply the permissions to all objects in all solutions. For information about global permissions, see *vSphere Security* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

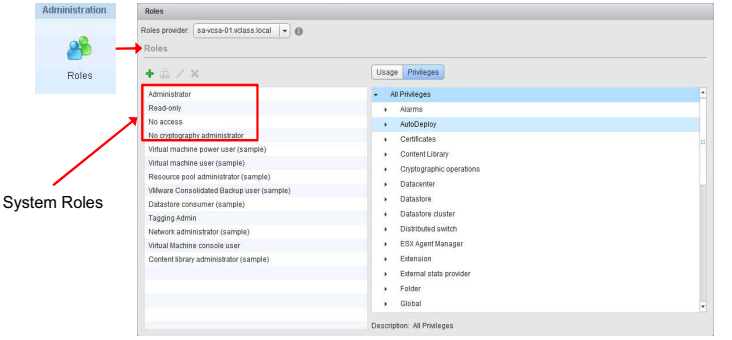
https://t.me/learningnets

4-67 About Roles

Privileges are grouped into roles:

- Roles allow users to perform tasks.
- Roles are grouped into categories to be simpler to configure.

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



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

You cannot change the privileges associated with the default roles:

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

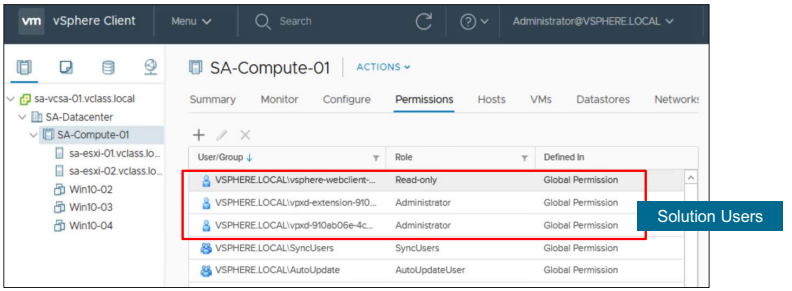
All roles are independent of each other. There is no hierarchy or inheritance between them.

https://t.me/learningnets

4-68 About Objects

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

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



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

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

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

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

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

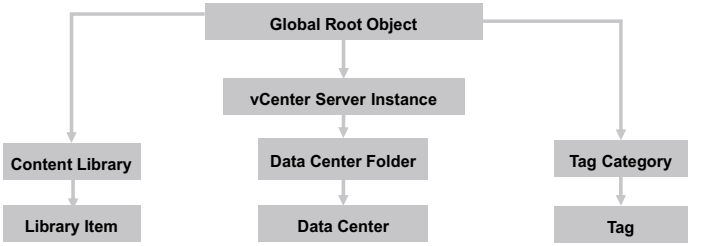
These solution users have the Administrator role and have global permissions.

https://t.me/learningnets

4-69 About Global Permissions

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

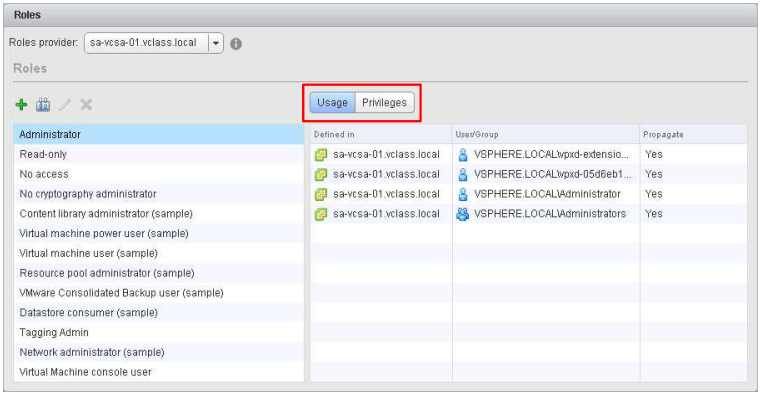
- Global permissions span solutions such as vCenter Server and VMware vRealize® Orchestrator™.
- Global permissions give a user or group privileges for all objects in all object hierarchies.



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

https://t.me/learningnets

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



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

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

https://t.me/learningnets

4-71 Applying Permissions: Scenario 1

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



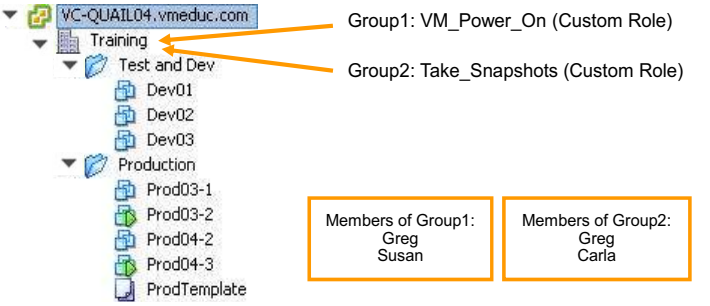
In addition to specifying whether or not permissions propagate downward, you can override permissions set at a higher level by explicitly setting different permissions for a lower-level object.

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

https://t.me/learningnets

4-72 Applying Permissions: Scenario 2

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



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

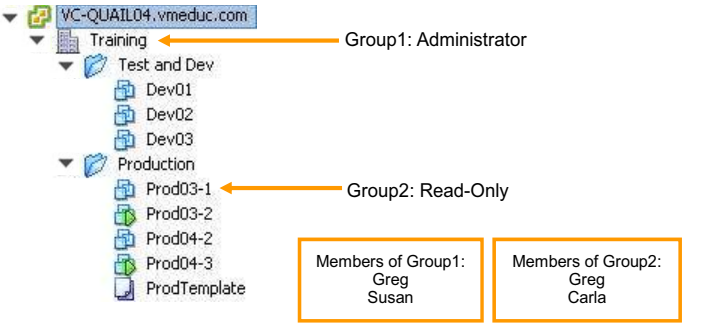
On the slide, Group1 is assigned the VM_Power_On role, a custom role that contains only one privilege: the ability to power on a virtual machine. Group2 is assigned the Take_Snapshots role, another custom role that contains the privileges to create and remove snapshots. Both roles propagate to the child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

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

4-73 Applying Permissions: Scenario 3

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



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

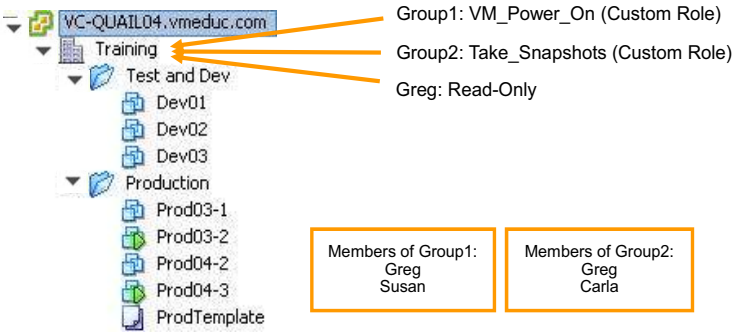
On the slide, Group1 is assigned the Administrator role at the Training data center and Group2 is assigned the Read-Only role on the virtual machine object, Prod03-1. The permission granted to Group1 is propagated to child objects.

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 virtual machine named Prod03-1 (the lower-level object) for which he gets read-only access.

https://t.me/learningnets

4-74 Applying Permissions: Scenario 4

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



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

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

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

https://t.me/learningnets

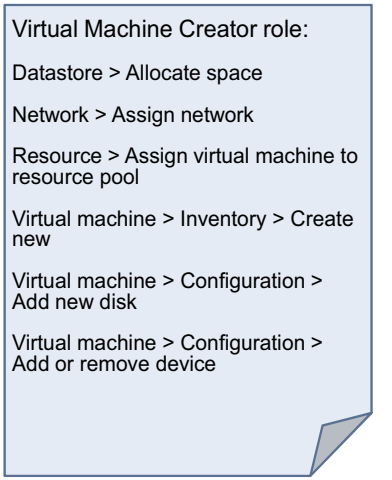
4-75 Creating a Role

Create roles that enable only the necessary tasks:

- For example, Virtual Machine Creator.

Use folders to contain the scope of permissions:

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



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

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

https://t.me/learningnets

4-76 Lab 3: Users, Groups, and Permissions

Assign roles and permissions to Active Directory users to perform functions in vCenter Server Appliance

1. Join vCenter Server Appliance to the vclass.local Domain
2. Add vclass.local as an Identity Source
3. View Active Directory Users
4. Assign Object Permissions to an Active Directory User
5. Assign Root-Level Global Permission
6. Log In with Windows Session Authentication
7. Use an Active Directory User to Manage a Virtual Machine

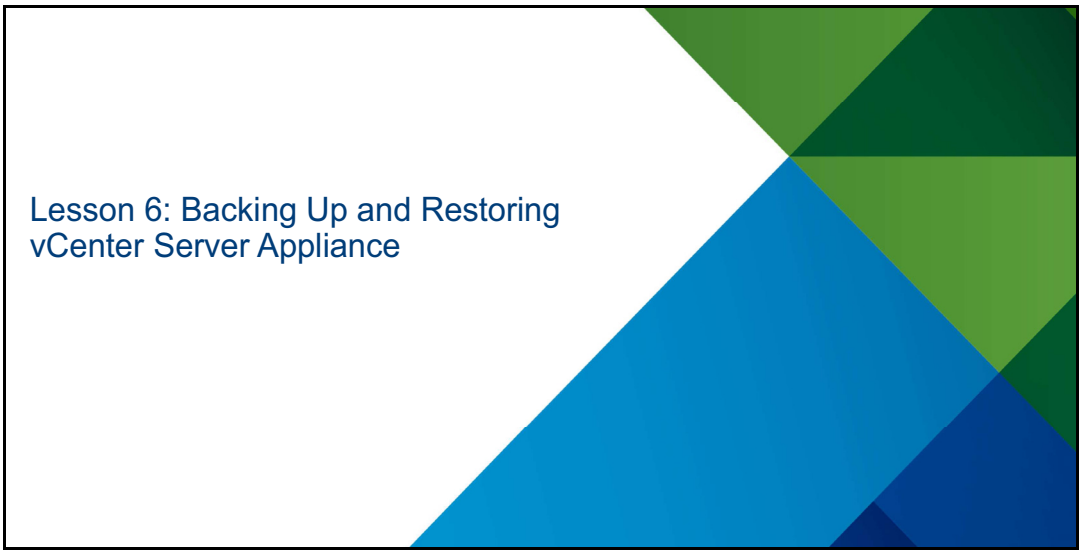
https://t.me/learningnets

4-77 Review of Learner Objectives

You should be able to meet the following objectives:

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

https://t.me/learningnets



https://t.me/learningnets

4-79 Learner Objectives

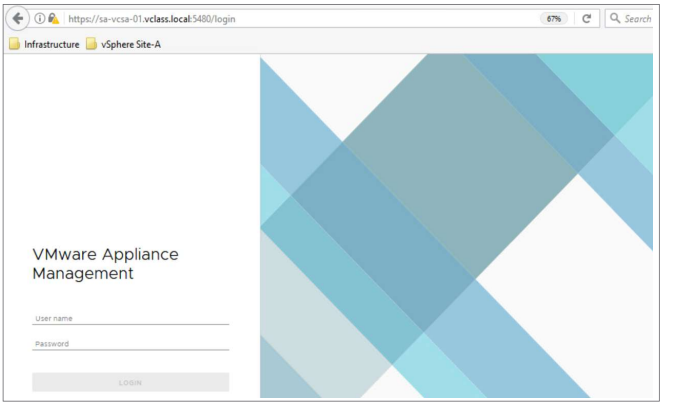
By the end of this lesson, you should be able to meet the following objectives:

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

https://t.me/learningnets

4-80 Logging In to the vCenter Server Appliance Management Interface

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



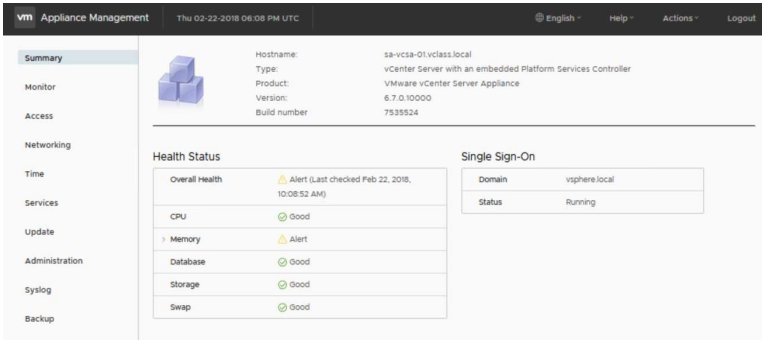
The vCenter Server Appliance Management Interface is an HTML client specifically designed to configure vCenter Server Appliance.

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

https://t.me/learningnets

4-81 vCenter Server Appliance Management Interface Functions

After logging in to the vCenter Server Appliance Management Interface, you see the functions that you can perform.

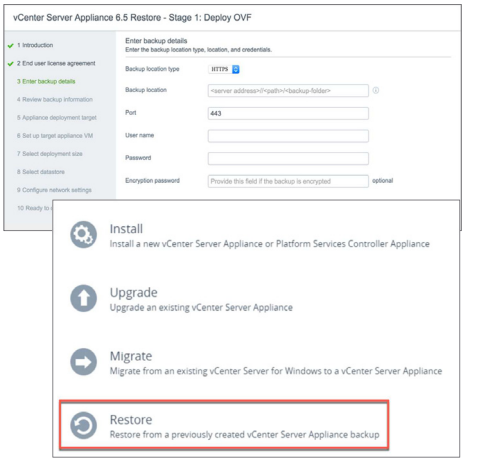


https://t.me/learningnets

4-82 About Native vCenter Server Backup and Restore

Native vCenter Server backup and restore offers several features:

- Removes dependency on third-party backup solutions
- Restores a vCenter Server instance to a brand-new appliance
- Supports backup or restore of vCenter Server Appliance and Platform Services Controller
- Supports protocols including HTTP/S, SCP, and FTP/S
- Includes option for encryption
- Restores directly from the vCenter Server Appliance ISO image



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 network load. To minimize storage impact, the transmission is streamed without caching in the appliance. To reduce the total time required to complete the backup operation, the backup process handles the different components in parallel.

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

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

https://t.me/learningnets

4-83 Methods for vCenter Server Appliance Backup and Restore

You can use the following 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 of the vCenter Server Appliance instance and the Platform Services Controller appliance.
  - After you create the backup, you can restore it through the GUI installer of the appliance.
  - File-based backup and restore can be scheduled.
- Image-based backup and restore:
  - Use a third-party product to back up and restore a virtual machine that contains vCenter Server, vCenter Server Appliance, or Platform Services Controller.
  - VMware vSphere® Storage APIs - Data Protection is a data protection framework that enables backup products to perform centralized, efficient, off-host, LAN-free backups of vSphere virtual machines.

https://t.me/learningnets

4-84 File-Based Backup of vCenter Server Appliance

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.

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 FTP, FTPS, HTTP, HTTPS, or SCP to a remote system. The backup is not stored on vCenter Server Appliance.

https://t.me/learningnets

4-85 File-Based Restore of vCenter Server Appliance

You can 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. The OVA file is deployed, which is included in the vCenter Server Appliance GUI installer.
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 have previously backed up by using the vCenter Server Appliance Management Interface. You can perform such 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

4-86 Scheduling Backups

vCenter Server Appliance provides the ability to schedule automatic file-based backups.

The backup scheduler supports:

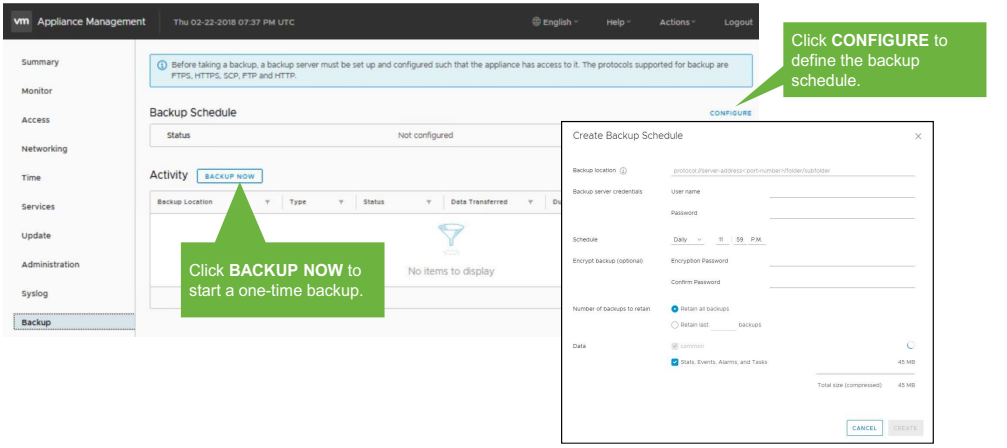
- FTPS, HTTPS, SCP, FTP and HTTP
- Encrypting the scheduled backups
- A retention policy to keep all backups or a defined number of backups
- Daily, weekly, or custom schedule

Failed backups trigger an alarm in vSphere Client.

https://t.me/learningnets

4-87 Creating a Backup Schedule

To create a backup schedule, go to the vCenter Server Appliance Management Interface and select **Backup**.



Creating a backup schedule requires the following information:

- Backup location: Backup server protocol, location, and directory
- Backup server credentials: User name and password of the backup server
- Schedule: Frequency of the backup schedule
- Encrypt backup: Encryption password (optional)
- Number of backups to retain: Retain all backups to define a set number
- Data: Choose to optionally backup SEAT data

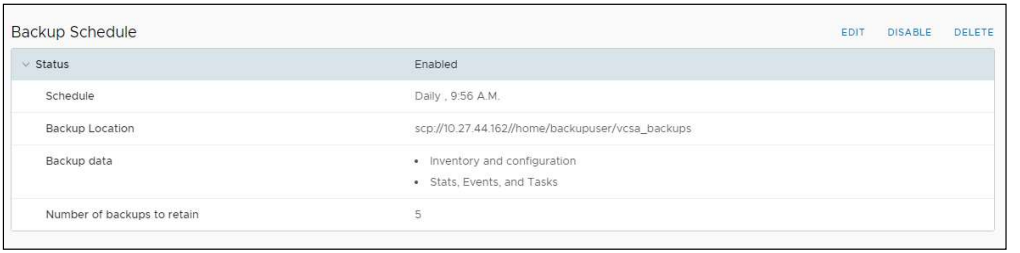
If a schedule is already defined, then clicking **BACKUP NOW** enables you to reuse the previously defined settings or overwrite the settings for the one-time backup.

https://t.me/learningnets

4-88 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.



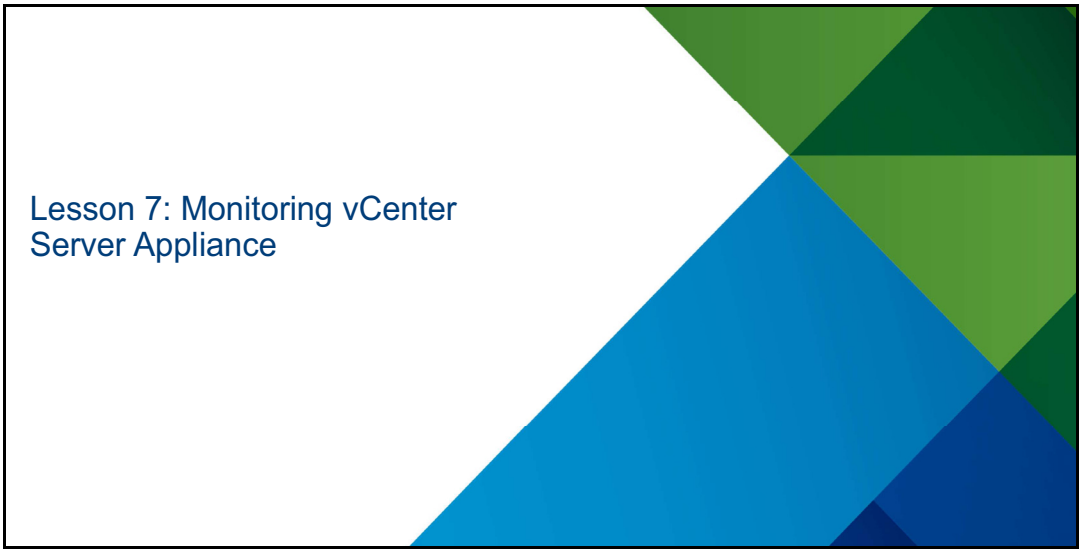
https://t.me/learningnets

4-89 Review of Learner Objectives

You should be able to meet the following objectives:

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

https://t.me/learningnets



https://t.me/learningnets

4-91 Learner Objectives

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

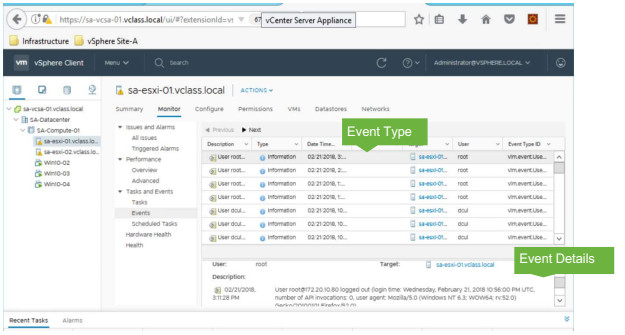
- Recognize how to view vCenter Server logs and events
- Manage the vCenter Server services
- Monitor vCenter Server Appliance
- Monitor vCenter Server Appliance for service and disk space usage
- Use vSphere alarms for resource exhaustion and service failures

https://t.me/learningnets

4-92 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.

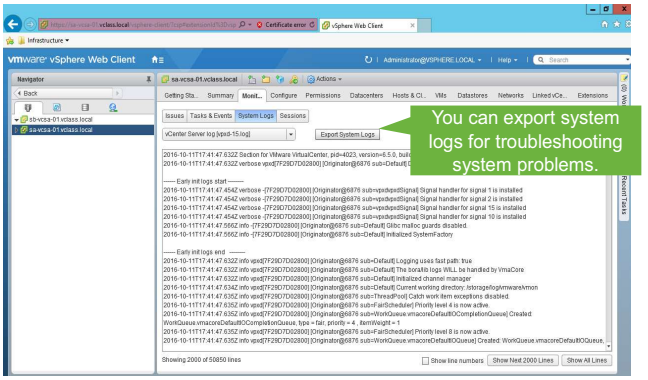


https://t.me/learningnets

4-93 vCenter Server System Logs

vSphere records events in the vCenter Server database.

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



https://t.me/learningnets

4-94 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.
- To control the amount of information being stored in the log files.

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

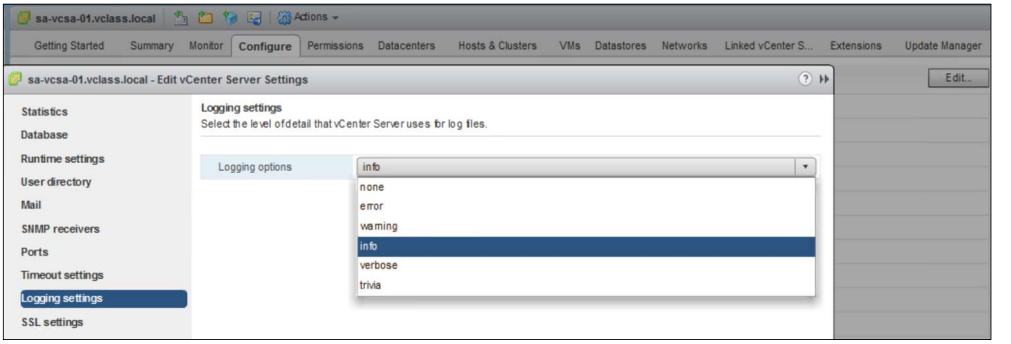
Changes to the logging settings take effect immediately. You do not have to restart the vCenter Server system.

https://t.me/learningnets

4-95 Setting Log Levels

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

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



To configure logging levels

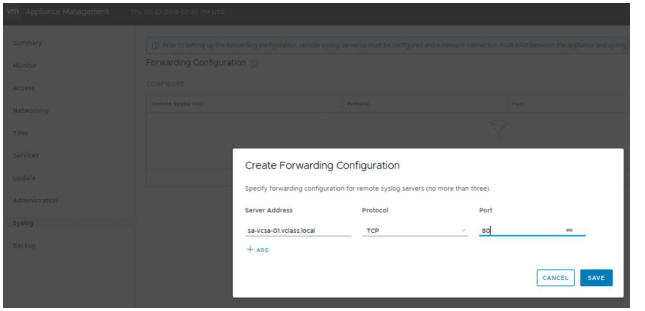
1. In vSphere Client, navigate to the vCenter Server instance.
2. Click the **Configure** tab.
3. Under Settings, select **General**.
4. Click **Edit**.
5. In the left pane, select **Logging settings**.
6. Select the logging options.

https://t.me/learningnets

4-96 Forwarding vCenter Server Appliance Log Files to Another Host Machine

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

- You can enable this feature in the vCenter Server Appliance Management Interface at https://*FQDN_of_vCenter_Server_Appliance*:5480.
- This feature can help prevent the vCenter Server Appliance file system from filling up with log files. It also enables further analysis of the vCenter Server Appliance log files with log analysis products, such as VMware vRealize® Log Insight™.



https://t.me/learningnets

4-97 vCenter Server Database Health

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

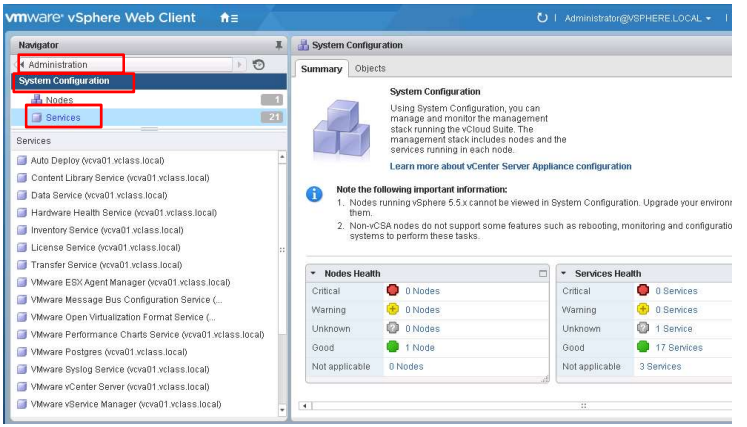
- By default, database health warnings trigger an alarm when the space used by the database reaches 80 percent.
- The alarm changes from warning to error when the space used by the database space reaches 95 percent. vCenter Server services shut down to enable the user to configure more disk space or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

These features are available for the PostgreSQL and Microsoft SQL Server databases. They are not available for Oracle databases.

https://t.me/learningnets

4-98 Managing the vCenter Server Services

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



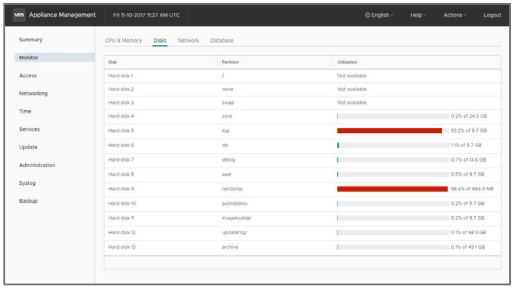
https://t.me/learningnets

4-99 Monitoring vCenter Server Appliance

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



CPU & Memory provides a historical view of CPU and memory use.

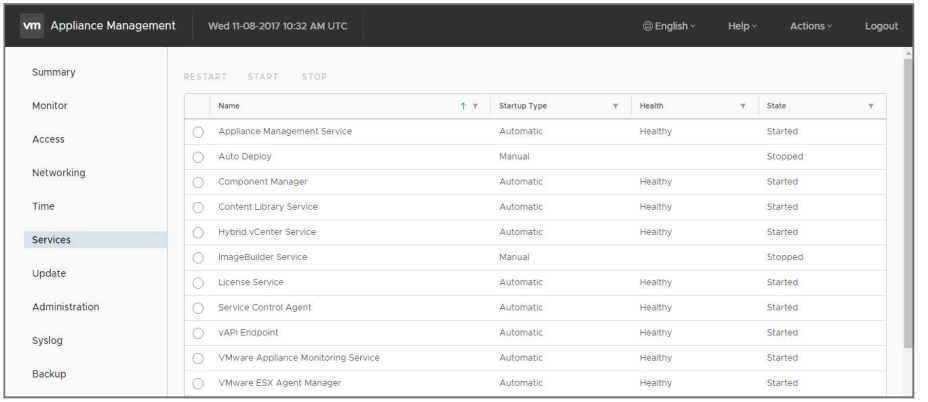


The vCenter Server Appliance disks can be monitored for available space.

https://t.me/learningnets

4-100 Monitoring vCenter Server Appliance Services

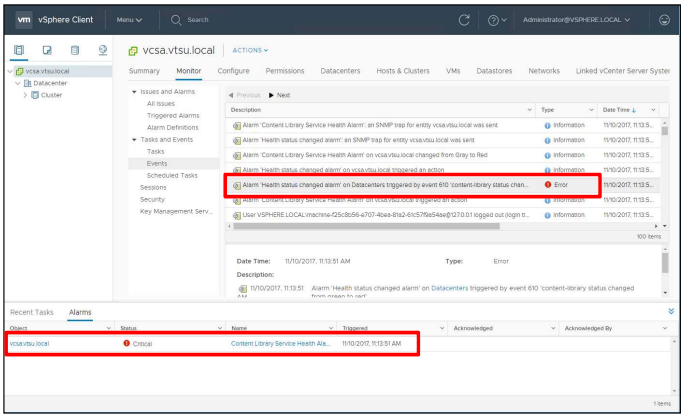
You can use the vCenter Server Appliance Management Interface to monitor the health and state of the vCenter Server Appliance services. You can restart, start, or stop services from the vCenter Server Appliance Management Interface.



https://t.me/learningnets

4-101 vSphere Alarms for Detecting Resource Exhaustion and Service Failure

vSphere Client displays events and alarms to alert the user to changes in the service health or when a service fails.



4-102 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 or low vulnerabilities are delivered with the next available vCenter Server patch or update.
- 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

4-103 Review of Learner Objectives

You should be able to meet the following objectives:

- Recognize how to view vCenter Server logs and events
- Manage the vCenter Server services
- Monitor vCenter Server Appliance
- Monitor vCenter Server Appliance for service and disk space usage
- Use vSphere alarms for resource exhaustion and service failures

https://t.me/learningnets

4-104 Key Points

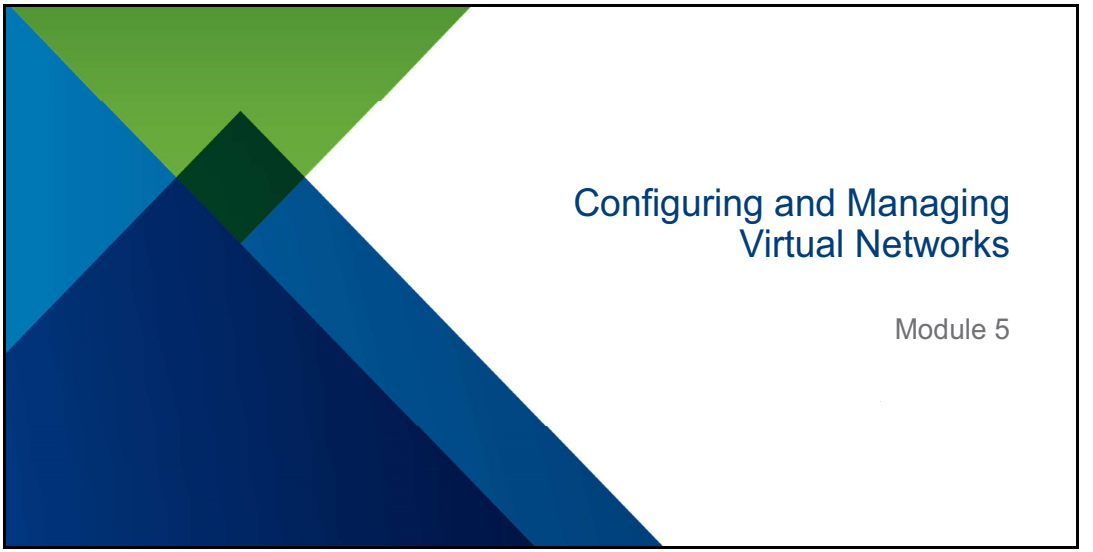
- The vCenter Server architecture includes the following components:
  - vCenter Server
  - vCenter Server database
  - Managed ESXi hosts
- vCenter Server has the following types of deployment models:
  - Embedded Platform Services Controller
  - External Platform Services Controller
- You use the vSphere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

Questions?

https://t.me/learningnets

MODULE 5

Configuring and Managing Virtual Networks



https://t.me/learningnets

5-2 You Are Here

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

https://t.me/learningnets

5-3 Importance

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

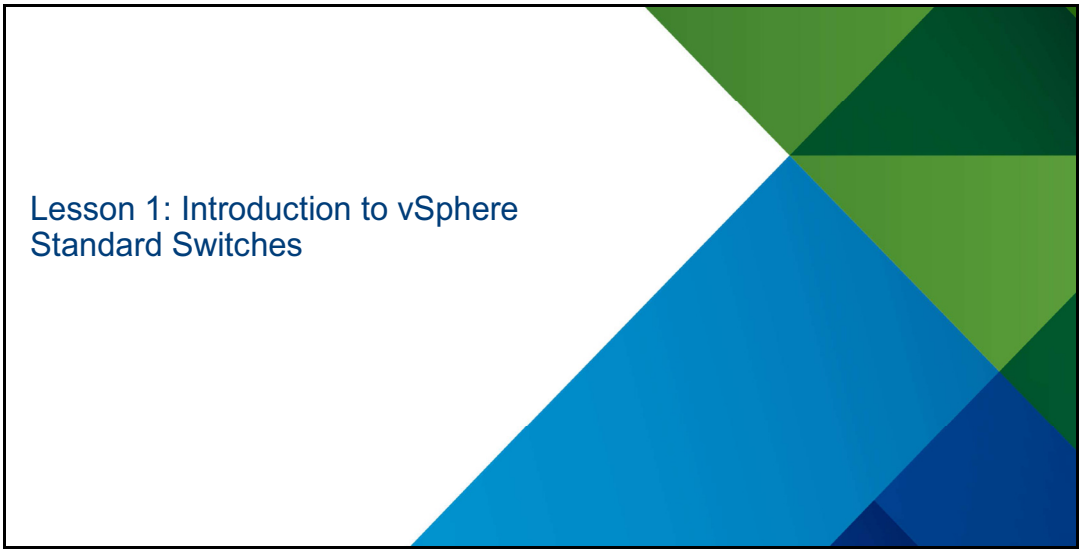
Failure to properly configure ESXi networking can negatively affect the operations of your virtual infrastructure.

https://t.me/learningnets

5-4 Module Lessons

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

https://t.me/learningnets



https://t.me/learningnets

5-6 Learner Objectives

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

- Describe the virtual switch connection types
- Configure and view standard switch configurations, such as virtual machine port groups, VMkernel ports, VLANs, and security features
- Compare the features of standard switches and distributed switches

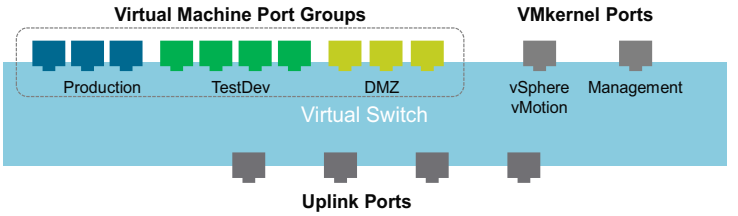
https://t.me/learningnets

5-7 Types of Virtual Switch Connections

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

A virtual switch has specific connection types:

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



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

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

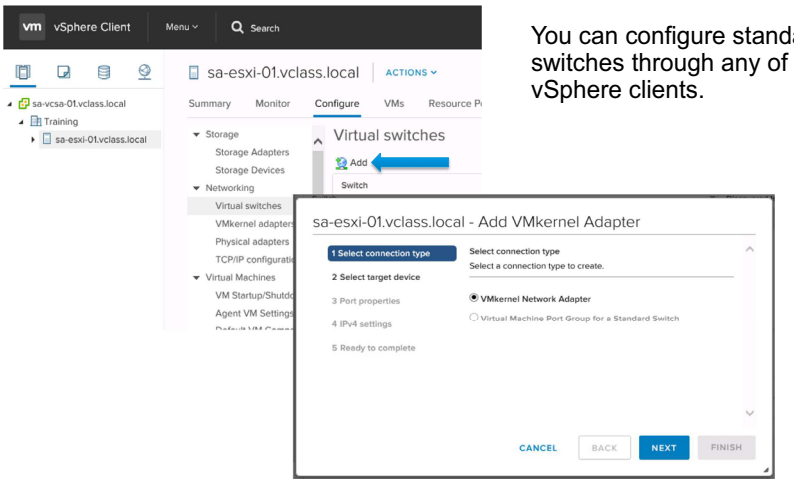
The ESXi management network port is used to connect to network or remote services, including vpxd on vCenter Server and VMware Host Client. Each ESXi management network port and each VMkernel port must be configured with its own IP address, netmask, and gateway.

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. Virtual machine port groups are used to connect virtual machines to one another with common networking properties.

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

https://t.me/learningnets

5-8 Adding ESXi Networking

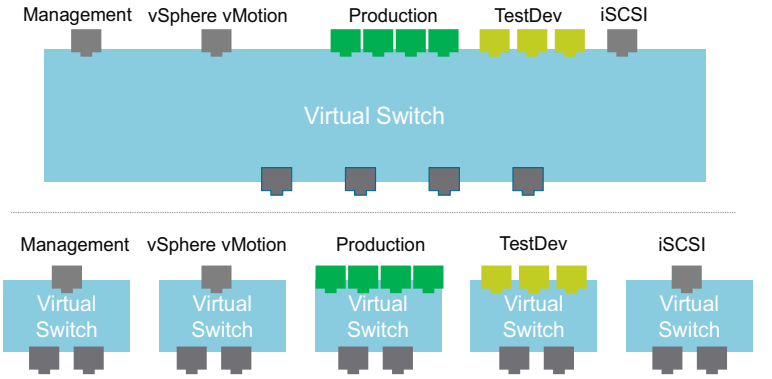


You can configure standard switches through any of the vSphere clients.

https://t.me/learningnets

5-9 Virtual Switch Connection Examples

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



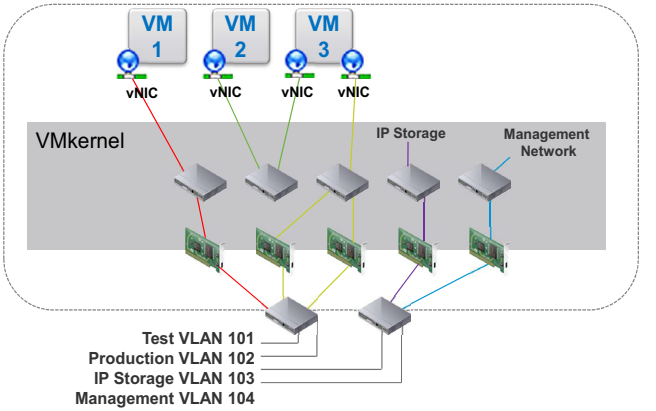
When you are designing your networking environment, vSphere enables you to place all your networks on a single virtual switch. Alternatively, you can opt for multiple virtual switches, each with a separate network. The decision partly depends on the layout of your physical networks. For example, you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

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

https://t.me/learningnets

5-10 About Standard Switches

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



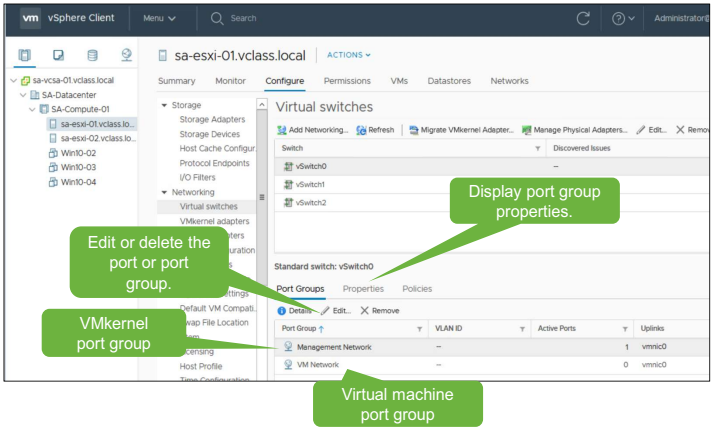
The slide shows five standard switches, each one devoted to a different purpose. From left to right, the switches are in numerical order:

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

https://t.me/learningnets

5-11 Viewing the Configuration of Standard Switches

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 an ESXi host. By default, the ESXi installation creates a virtual machine port group named VM Network and a VMkernel port named Management Network. It is good practice to remove the VM Network virtual machine port group and keep virtual machine networks and management networks separated, for performance and security reasons.

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

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

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

https://t.me/learningnets

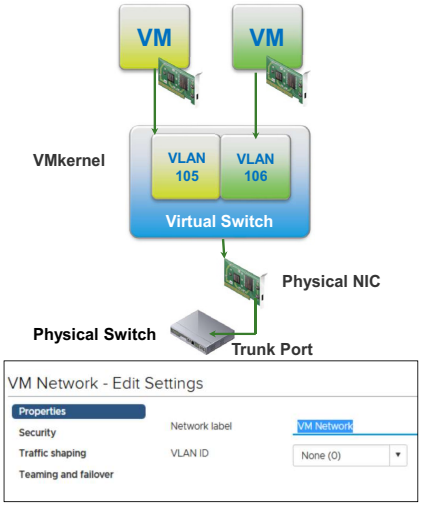
5-12 About VLANs

ESXi supports 802.1Q VLAN tagging.

Virtual switch tagging is one of the tagging policies supported:

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

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



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

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

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

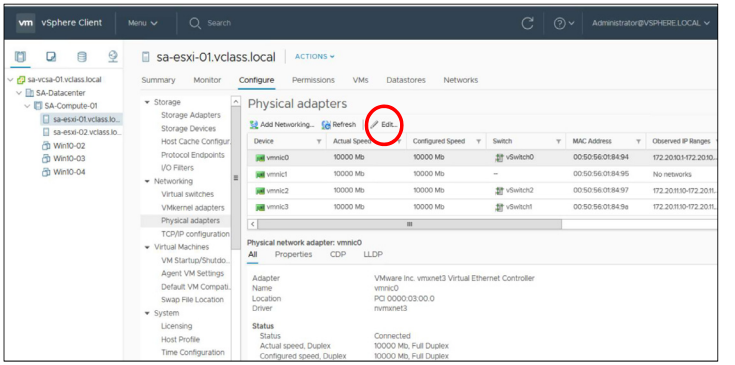
The port on a physical switch to which an ESXi host is connected must be defined as a static trunk port. A trunk port is a port on a physical Ethernet switch that is configured to send and receive packets tagged with a VLAN ID. No VLAN configuration is required in the virtual machine. In fact, the virtual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

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

https://t.me/learningnets

5-13 Network Adapter Properties

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



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

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

https://t.me/learningnets

5-14 Types of Virtual Switches

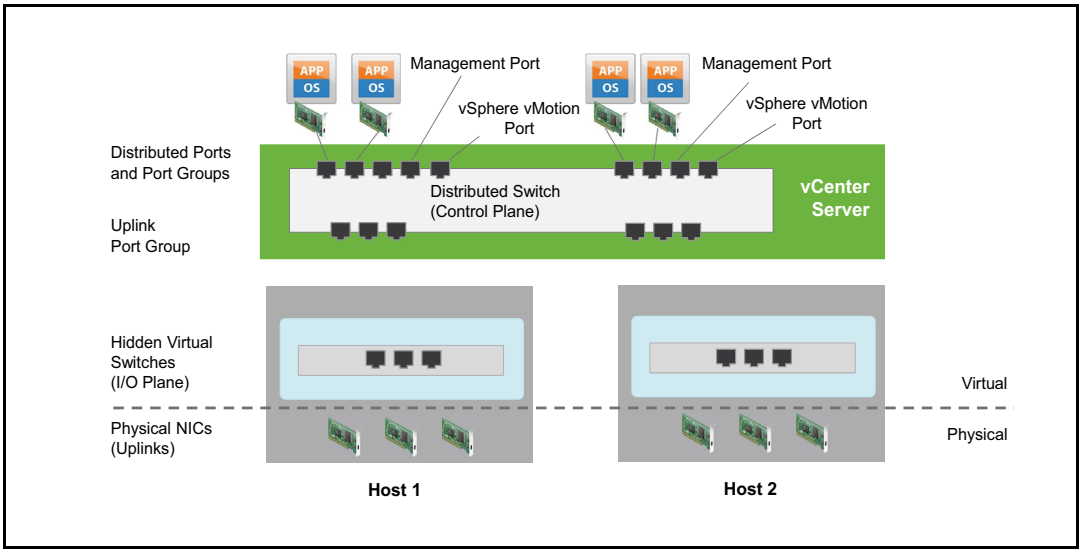
A virtual network supports the following types of virtual switches:

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

Both switch types are elastic: ports are created and removed automatically.

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 Switch and Distributed Switch Feature Comparison

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

https://t.me/learningnets

5-17 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe the virtual switch connection types
- Configure and view standard switch configurations, such as virtual machine port groups, VMkernel ports, VLANs, and security features
- Compare the features of standard switches and distributed switches

https://t.me/learningnets



https://t.me/learningnets

5-19 Learner Objectives

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

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

https://t.me/learningnets

5-20 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

Policies are defined at the following levels:

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

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

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

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

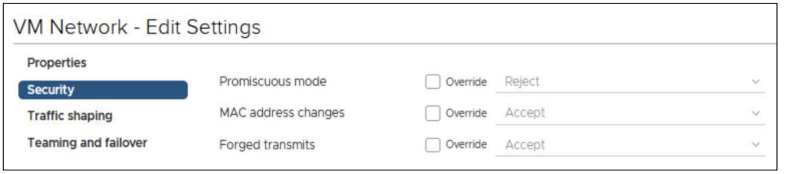
These policies are defined for the entire standard switch and can also be defined for a VMkernel port or a virtual machine port group. When a policy is defined for an individual port or port group, the policy at this level overrides the default policies defined for the standard switch.

https://t.me/learningnets

5-21 Configuring Security Policies

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

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



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

The network security policy contains the following exceptions:

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

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

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

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

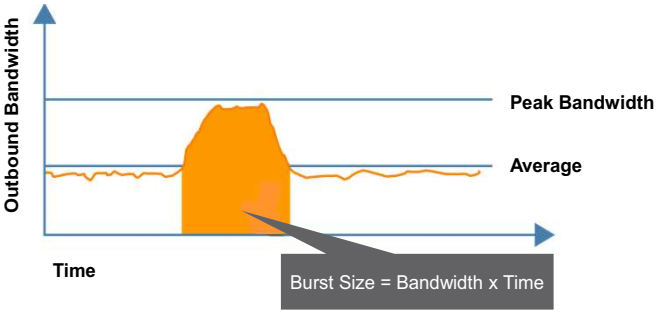
Leave **MAC Address Changes** and **Forged Transmits** at their default values (**Accept**) if your applications change the mapped MAC address, as do some guest operating system-based firewalls.

https://t.me/learningnets

5-22 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

5-23 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 travelling 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 **Average bandwidth**, the port might be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of kilobytes that have accumulated in the burst bonus and thus transfers at a higher speed.

Network traffic shaping is off by default.

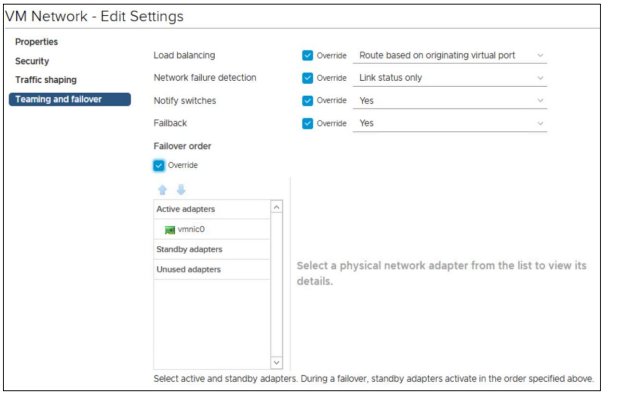
For information about configuring the traffic shaping policy in vSphere Web Client, see *VMware PowerCLI User's Guide* at https://www.vmware.com/support/developer/PowerCLI/.

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

https://t.me/learningnets

5-24 NIC Teaming and Failover Policies

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



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

- NIC teaming policy: You can use NIC teaming to connect a virtual switch to multiple physical NICs on a host to increase the network bandwidth of the switch and to provide redundancy. A NIC team can distribute the traffic between its members and provide passive failover in case of adapter failure or network outage. You set NIC teaming policies at the virtual switch or port group level for a vSphere standard switch and at the port group or port level for a vSphere distributed switch.
- Load balancing policy: The load balancing policy determines how network traffic is distributed between the network adapters in a NIC team. vSphere virtual switches load balance only the outgoing traffic. Incoming traffic is controlled by the load balancing policy on the physical switch.
- Failback policy: By default, a failback policy is enabled on a NIC team. If a failed physical NIC returns online, the virtual switch sets the NIC back to active by replacing the standby NIC that took over its slot.

  If the physical NIC that stands first in the failover order experiences intermittent failures, the failback policy might lead to frequent changes in the NIC that is used. The physical switch sees frequent changes in MAC addresses, and the physical switch port might not accept traffic

https://t.me/learningnets

immediately when an adapter becomes online. To minimize such delays, you might consider changing the following settings on the physical switch.

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

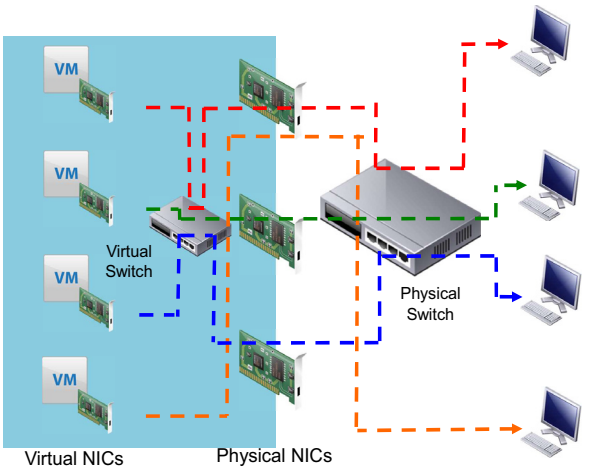
  NIC teaming and failover policies enable you to determine how network traffic is distributed between adapters and how to reroute traffic if an adapter fails. NIC teaming policies include load-balancing and failover settings. Default NIC teaming and failover policies are set for the entire standard switch.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

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

https://t.me/learningnets

5-25 Load-Balancing Method: Originating Virtual Port ID

Routing based on the originating port ID is called virtual port ID load balancing.



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

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

Advantages:

- An even distribution of traffic if the number virtual NICs is greater than the number of physical NICs in the team.
- Low resource consumption, because in most cases the virtual switch calculates uplinks for virtual machines only once.
- No changes on the physical switch are required.

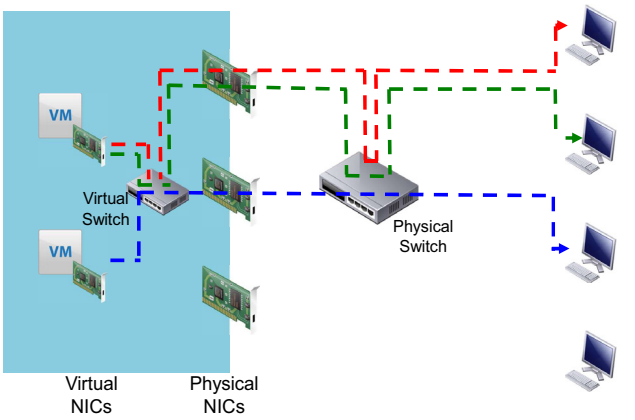
Disadvantages:

- 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 virtual machine is limited to the speed of the uplink that is associated with the relevant port ID, unless the virtual machine has more than one virtual NIC.

https://t.me/learningnets

5-26 Load-Balancing Method: Source MAC Hash

The diagram shows routing based on source MAC hash.



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

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

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

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

Advantages:

- A more even distribution of the traffic than Route Based on Originating Virtual Port, because the virtual switch calculates an uplink for every packet.
- Virtual machines use the same uplink because the MAC address is static. Powering a virtual machine on or off does not change the uplink that the virtual machine uses.
- No changes on the physical switch are required.

Disadvantages:

- The bandwidth that is available to a virtual machine is limited to the speed of the uplink that is associated with the relevant port ID, unless the virtual machine uses multiple source MAC addresses.

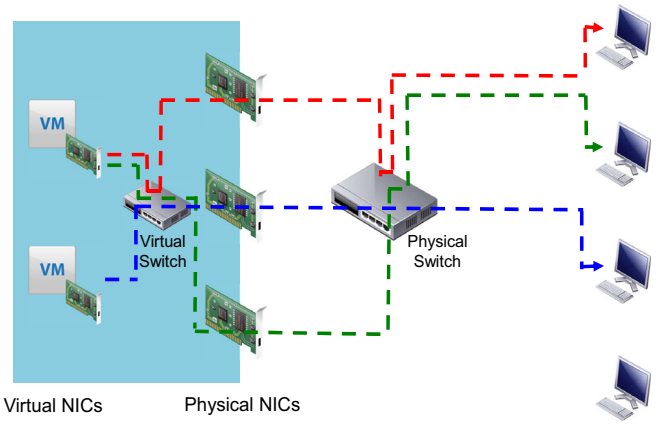
https://t.me/learningnets

- Higher resource consumption than Route Based on 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

5-27 Load-Balancing Method: Source and Destination IP Hash

The diagram shows routing based on IP hash.



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

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

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

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

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

Advantages:

- A more even distribution of the load compared to Route Based on Originating Virtual Port and Route Based on Source MAC Hash, as the virtual switch calculates the uplink for every packet.
- A potentially higher throughput for virtual machines that communicate with multiple IP addresses.

https://t.me/learningnets

Disadvantages:

- Highest resource consumption compared to the other load balancing algorithms.
- The virtual switch is not aware of the actual load of the uplinks.
- Requires changes on the physical network.
- Complex to troubleshoot.

https://t.me/learningnets

5-28 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 provided by the network adapter detects failures like cable pulls and physical switch power failures. This monitoring does not detect configuration errors, such as a physical switch port being blocked by the Spanning Tree Protocol or misconfigured VLAN membership. This method cannot detect upstream, nondirectly connected physical switch or cable failures.

Beaconing introduces a 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. See the VMware Knowledge Base article 1005577 at http://kb.vmware.com/kb/1005577.

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

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

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

https://t.me/learningnets

5-29 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

5-30 Lab 4: Using Standard Switches

Create a standard switch and a port group

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

https://t.me/learningnets

5-31 Review of Learner Objectives

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

5-32 Key Points

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

Questions?

https://t.me/learningnets

https://t.me/learningnets

MODULE 6

Configuring and Managing Virtual Storage



https://t.me/learningnets

6-2 You Are Here

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

https://t.me/learningnets

6-3 Importance

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

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

https://t.me/learningnets

6-4 Module Lessons

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

https://t.me/learningnets

6-5 Lesson 1: Storage Concepts



https://t.me/learningnets

6-6 Learner Objectives

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

- Describe vSphere storage technologies and datastores

https://t.me/learningnets

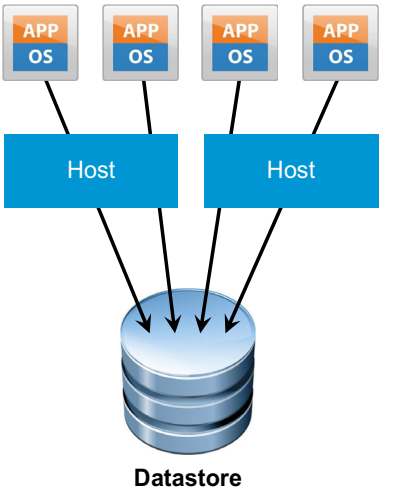
6-7 About Datastores

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 virtual machine files, templates, and ISO images.

vSphere supports the following types of datastores:

- VMFS
- NFS
- vSAN
- vSphere Virtual Volumes



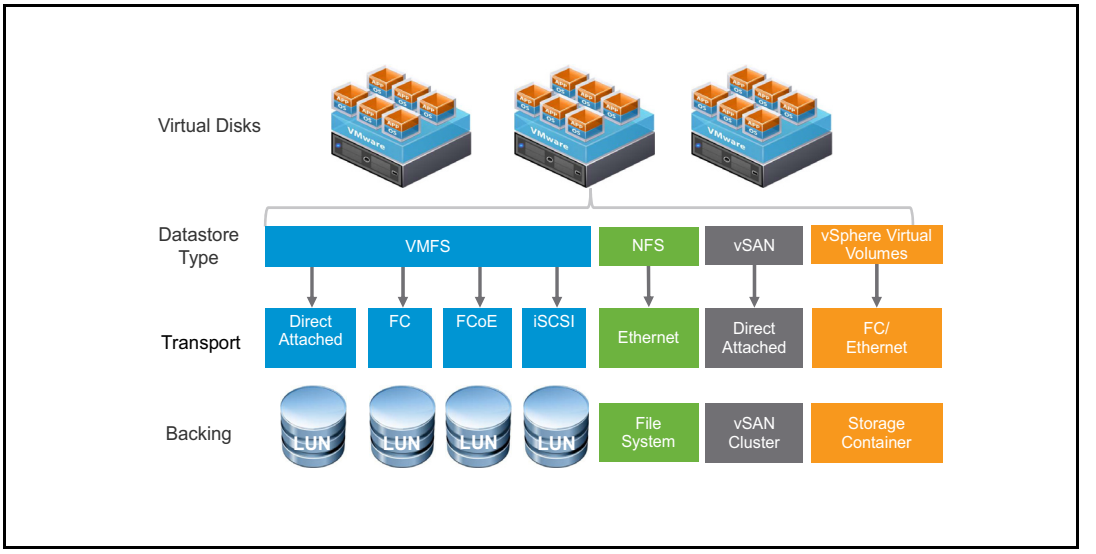
A datastore is a generic term for a container that holds files and objects. Datastores are logical containers, analogous to file systems, that hide the specifics of each storage device and provide a uniform model for storing virtual machine files.

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

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

https://t.me/learningnets

6-8 Basic Storage Overview



ESXi hosts should be configured so that they have shared access to datastores. Depending on the type of storage that you use, datastores can be formatted with VMFS or with a file system native to a storage device that is shared using the NFS protocol.

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

- Direct-attached storage: Internal or external storage disks or arrays attached to the host through a direct connection instead of a network connection.
- Fibre Channel: A high-speed transport protocol used for SANs. Fibre Channel encapsulates SCSI commands, which are transmitted between Fibre Channel nodes. In general, a Fibre Channel node is a server, a storage system, or a tape drive. A Fibre Channel switch interconnects multiple nodes, forming the fabric in a Fibre Channel network.
- FCoE: The Fibre Channel traffic is encapsulated into Fibre Channel over Ethernet (FCoE) frames. These FCoE frames are converged with other types of traffic on the Ethernet network. By carrying both Ethernet and Fibre Channel traffic on the same Ethernet link, usage of the physical infrastructure has greatly increased. FCoE also reduces the total number of network ports and cabling.
- iSCSI: A SCSI transport protocol, enabling access to storage devices and cabling over standard TCP/IP networks. iSCSI maps SCSI block-oriented storage over TCP/IP. Initiators, such as an iSCSI host bus adapter (HBA) in an ESXi host, send SCSI commands to targets, located in iSCSI storage systems.

https://t.me/learningnets

- NAS: Storage shared over standard TCP/IP networks at the file system level. NAS storage is used to hold NFS datastores. The NFS protocol does not support SCSI commands.

iSCSI, NAS, and FCoE can run over 1 Gigabit Ethernet or 10 Gigabit Ethernet. 10GigE provides increased storage performance levels and sufficient bandwidth that allows multiple types of high-bandwidth protocol traffic to coexist on the same network and run faster.

https://t.me/learningnets

6-9 Storage Protocol Overview

Datastore Type	Storage Protocol	Boot from SAN Support	vSphere vMotion Support	vSphere HA Support	vSphere DRS Support	Raw Device Mapping Support
VMFS	Fibre Channel	•	•	•	•	•
	FCoE	•	•	•	•	•
	iSCSI	•	•	•	•	•
	NFS		•	•	•	
Direct-Attached	DAS		•			•
vSphere Virtual Volumes	FC/ Ethernet		•	•	•	
	vSAN		•	•	•	

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

- CD/DVD ISO images
- Decommissioned virtual machines
- Virtual machine templates

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

- vSphere vMotion
- vSphere HA
- vSphere DRS

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

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

https://t.me/learningnets

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

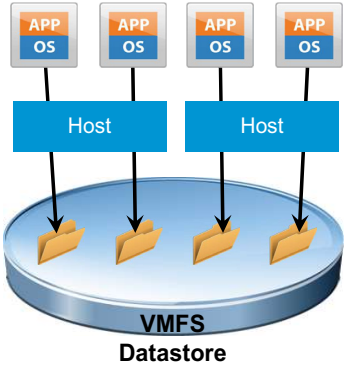
ESXi hosts allow booting from software iSCSI, a supported independent hardware SCSI adapter, and a supported dependent hardware iSCSI adapter. The network adapter must support only the iSCSI Boot Firmware Table (iBFT) format, which is a method of communicating parameters about the iSCSI boot device to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

https://t.me/learningnets

6-10 About VMFS

ESXi hosts version 6.5 and later support VMFS5 and VMFS6:

- Features supported by both VMFS5 and VMFS6:
  - They allow concurrent access to shared storage
  - They can be dynamically expanded
  - They use 1 MB and 512 MB block sizes that are good for storing large virtual disk files
  - The provide on-disk locking
- Features supported by VMFS6:
  - 4K native storage devices
  - Automatic space reclamation



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

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

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

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

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

On the slide, each ESXi host has two virtual machines running on it. The lines connecting the virtual machines to the virtual machine disks (VMDKs) are logical representations of the association and allocation of the larger VMFS datastore. The VMFS datastore includes one or more LUNs. The

https://t.me/learningnets

virtual machines see the assigned storage volume only as a SCSI target from within the guest operating system. The virtual machine contents are only files on the VMFS volume.

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

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

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

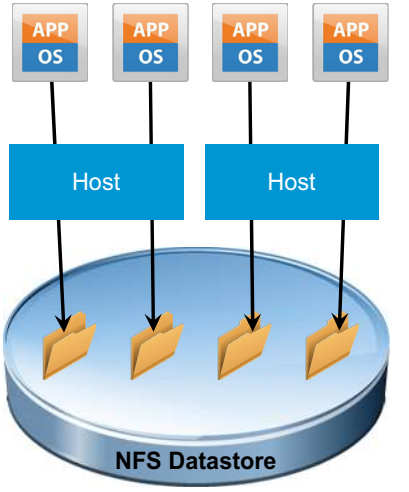
For the operating system in the virtual machine, VMFS preserves the internal file system semantics. As a result, the operating system running in the virtual machine sees a native file system, not VMFS. These semantics ensure correct behavior and data integrity for applications running on the virtual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

https://t.me/learningnets

6-11 About NFS

NFS is storage shared over the network at the file system level.

- NFS supports NFS version 3 and 4.1 over TCP/IP.



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

NFS datastores are treated like VMFS datastores because they can be used to hold virtual machine files, templates, and ISO images. In addition, like a VMFS datastore an NFS volume allows the vSphere vMotion migration of virtual machines 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/NFS servers.

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

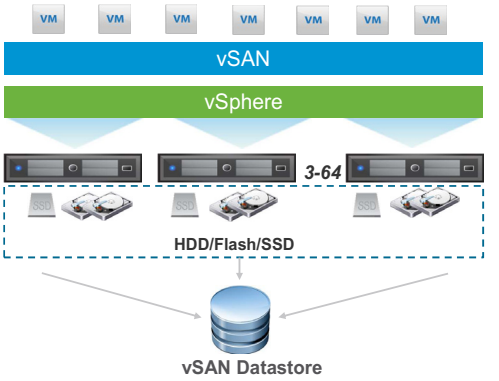
Because NFS 3 and NFS 4.1 clients do not use the same locking protocol, you cannot use different NFS versions to mount the same datastore on multiple hosts. Accessing the same virtual disks from two incompatible clients might result in incorrect behavior and cause data corruption.

https://t.me/learningnets

6-12 vSAN Overview

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

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



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

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

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

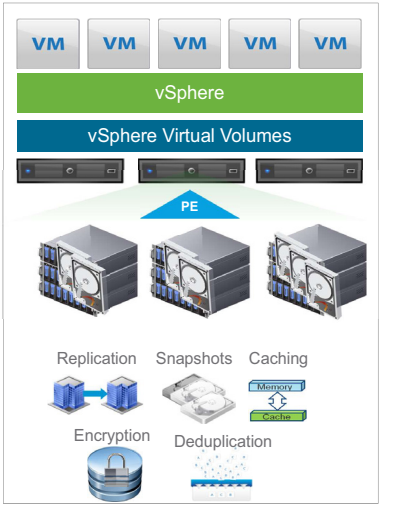
vSAN can be configured as hybrid or all-flash storage. In a hybrid storage architecture, vSAN pools server-attached HDDs and SSDs to create a distributed shared datastore that abstracts the storage hardware to provide a software-defined storage tier for virtual machines. Flash is used as a read cache/write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

https://t.me/learningnets

6-13 About vSphere Virtual Volumes

vSphere Virtual Volumes provides the following functionality:

- Native representation of VMDKs on SAN/NAS: No LUNs or volume management.
- Works with existing SAN/NAS systems.
- A new control path for data operations at the VM/VMDK level.
- Snapshots, replications, and other operations at the VM level on external storage.
- Automates control of per-VM service levels.
- VMware vSphere® API for Storage Awareness™ protocol endpoint provides standard protocol access to storage.
- Storage containers can span an entire array.



VMware vSphere® Virtual Volumes™ virtualizes SAN and NAS devices by abstracting physical hardware resources into logical pools of capacity. The vSphere Virtual Volumes storage paradigm is designed to address the requirements of next-generation storage in the software-defined data center.

vSphere Virtual Volumes provides the following benefits:

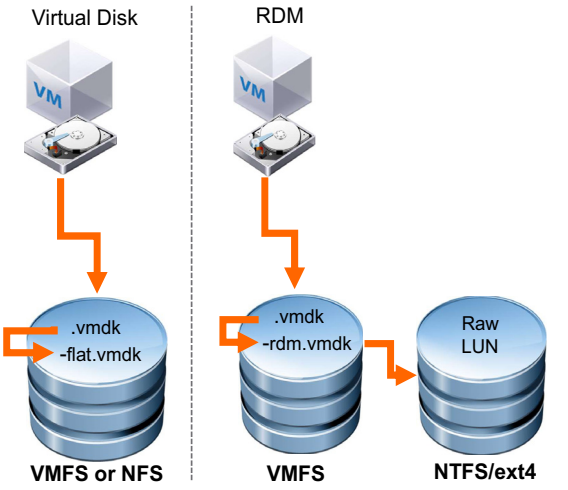
- Lower storage cost
- Reduced storage management overhead
- Greater scalability
- Better response to data access and analytical requirements

https://t.me/learningnets

6-14 About Raw Device Mapping

Though not a datastore, raw device mapping (RDM) allows a virtual machine direct access to a physical LUN.

The mapping file (`-rdm.vmdk`) that points a VM to a LUN must be stored on a VMFS datastore.



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

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

https://t.me/learningnets

6-15 Physical Storage Considerations

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

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

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

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

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

https://t.me/learningnets

6-16 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe vSphere storage technologies and datastores

https://t.me/learningnets



https://t.me/learningnets

6-18 Learner Objectives

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

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

https://t.me/learningnets

6-19 About Fibre Channel

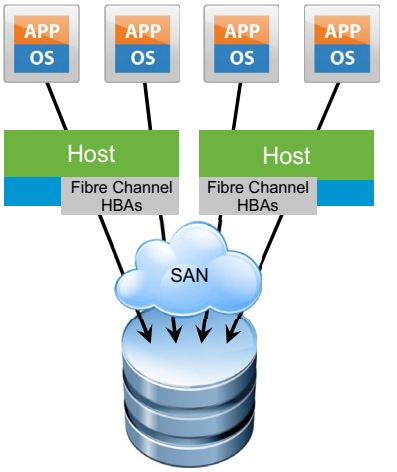
Fibre Channel stores virtual machine files remotely on a Fibre Channel SAN.

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

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

ESXi supports:

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



VMFS or vSphere Virtual Volumes Datastore on Fibre Array

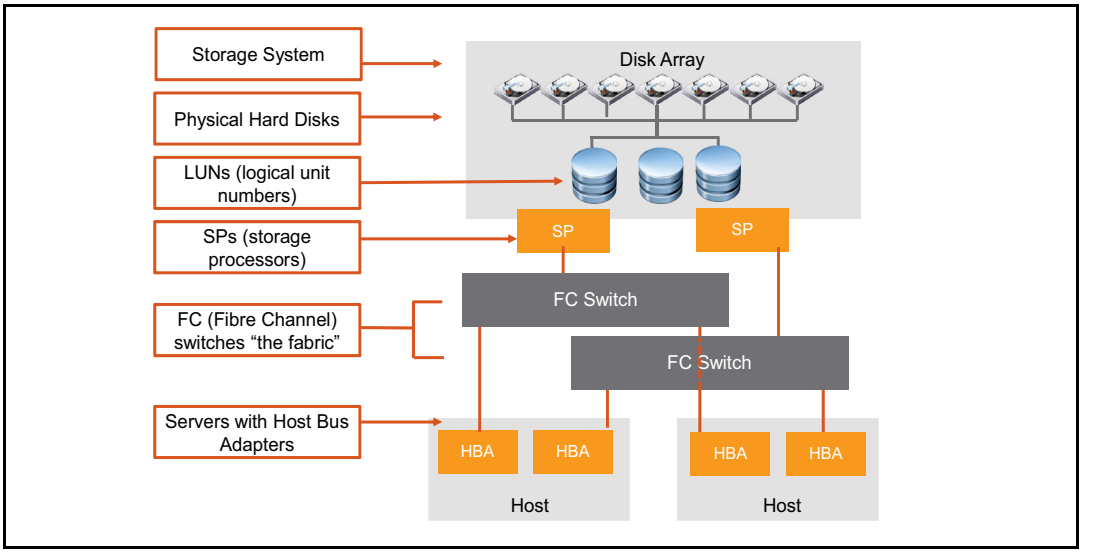
To connect to the Fibre Channel SAN, your host should be equipped with Fibre Channel host bus adapters (HBAs).

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

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



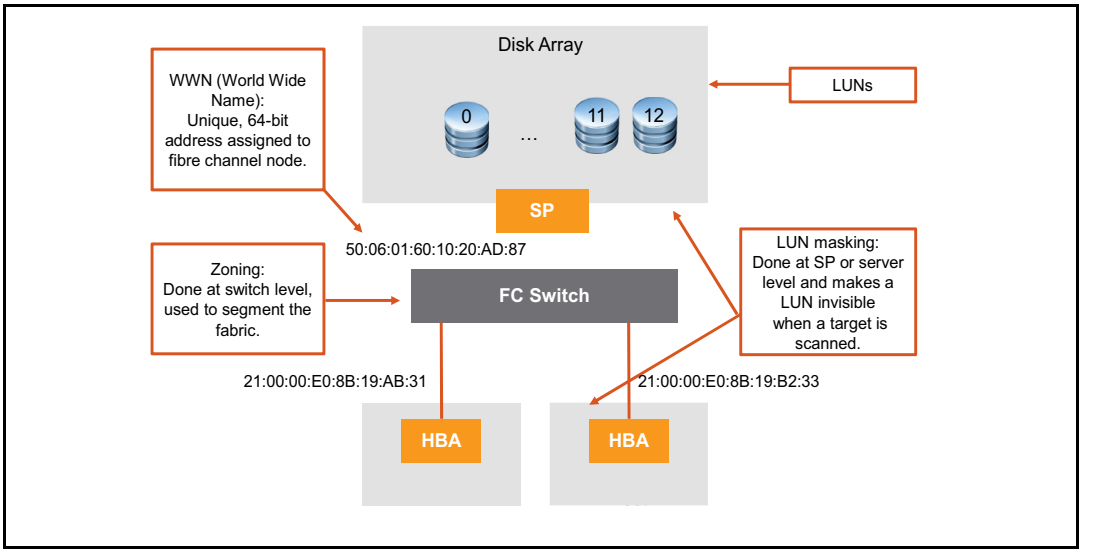
In its simplest form, a SAN consists of one or more servers attached to a storage array using one or more SAN switches. Each server might host numerous applications that require dedicated storage for applications processing.

The following components are involved:

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

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

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

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

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

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

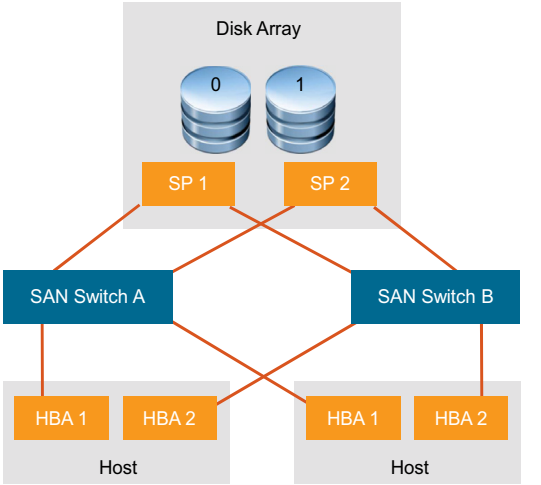
See your SAN vendor's documentation and *vSphere Storage* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

6-22 Multipathing with Fibre Channel

Multipathing enables continued access to SAN LUNs if hardware fails.

It also provides 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

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

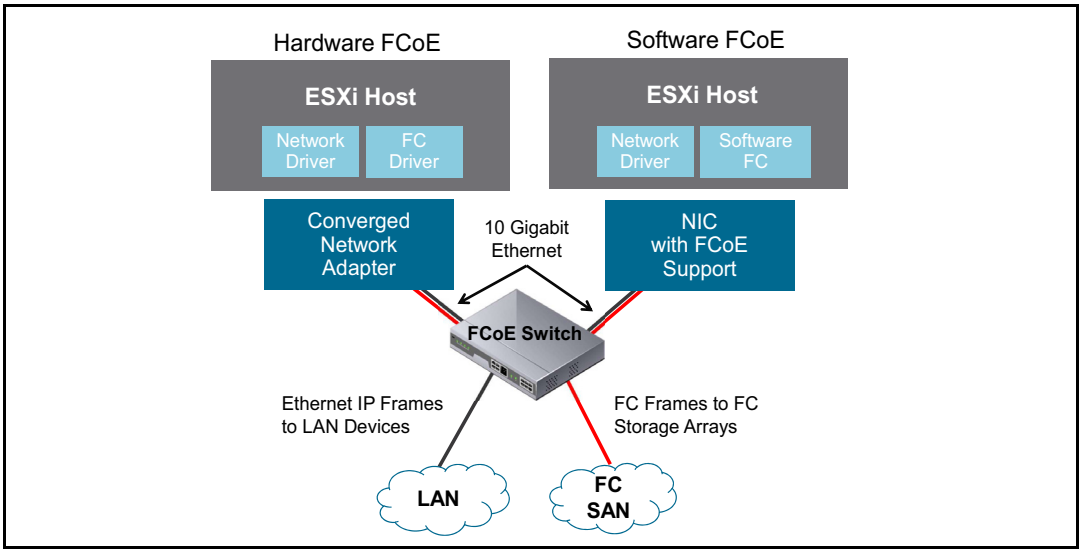
By default, ESXi hosts use only one path from a host to a given LUN at any time. If the path actively being used by the ESXi host fails, the server selects another of the available paths. The process of detecting a failed path and switching to another is called path failover. A path fails if any of the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

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

- An active-active disk array allows access to the LUNS simultaneously through all the storage processors that are available without significant performance degradation. All the paths are active at all times (unless a path fails).
- In an active-passive disk array, one storage processor is actively servicing a given LUN. The other storage processor acts as a backup for the LUN and might be actively servicing other LUN I/O. I/O can be sent only to an active processor. If the primary storage processor fails, one of the secondary storage processors becomes active, either automatically or through administrative intervention.

https://t.me/learningnets

6-23 FCoE Adapters



https://t.me/learningnets

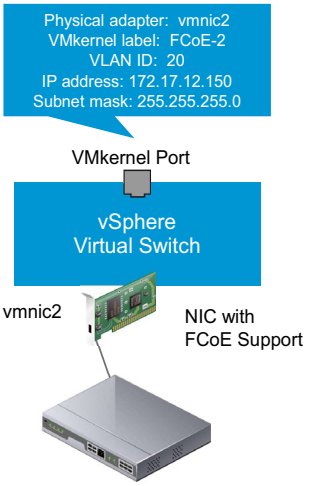
6-24 Configuring Software FCoE: Creating VMkernel Ports

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

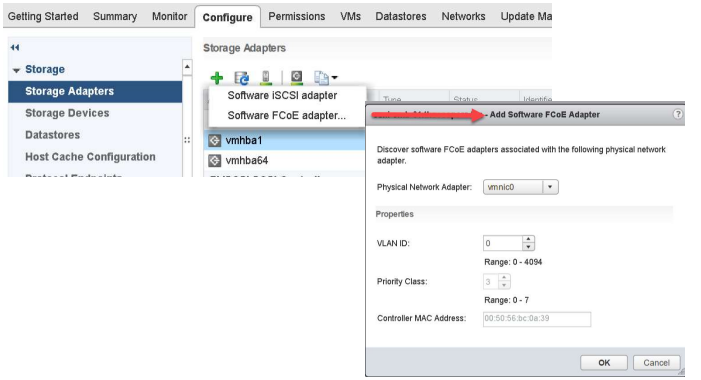


https://t.me/learningnets

6-25 Configuring Software FCoE: Activating Software FCoE Adapters

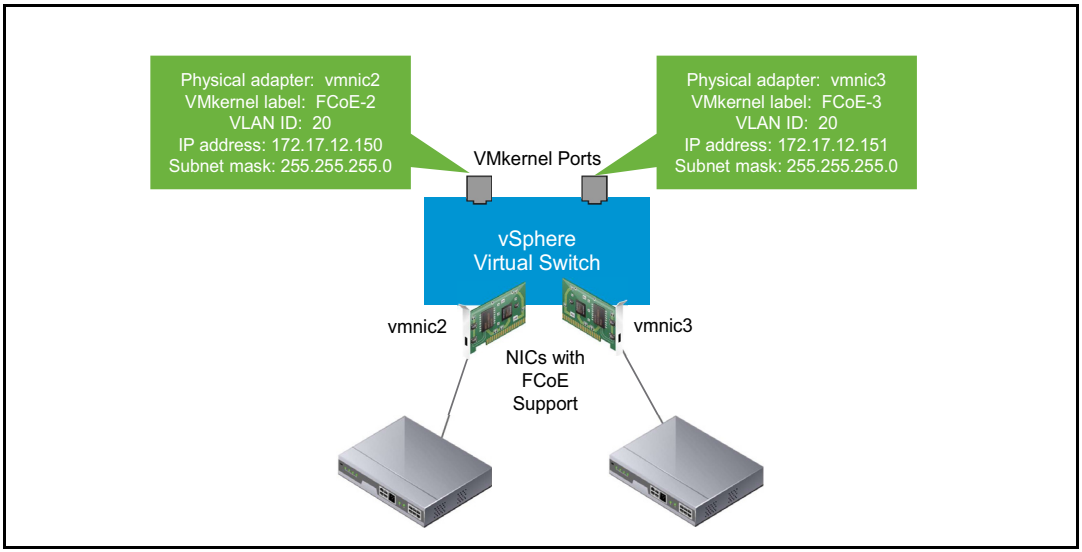
Step 2:

Add the software FCoE adapter. Select the host, click the **Configure** tab, select **Storage Adapters**, click the green plus sign, select **Software FCoE adapter**, and configure as needed.



https://t.me/learningnets

6-26 Multipathing with Software FCoE



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

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

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

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

https://t.me/learningnets

6-27 Review of Learner Objectives

You should be able to meet the following objectives:

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

https://t.me/learningnets



https://t.me/learningnets

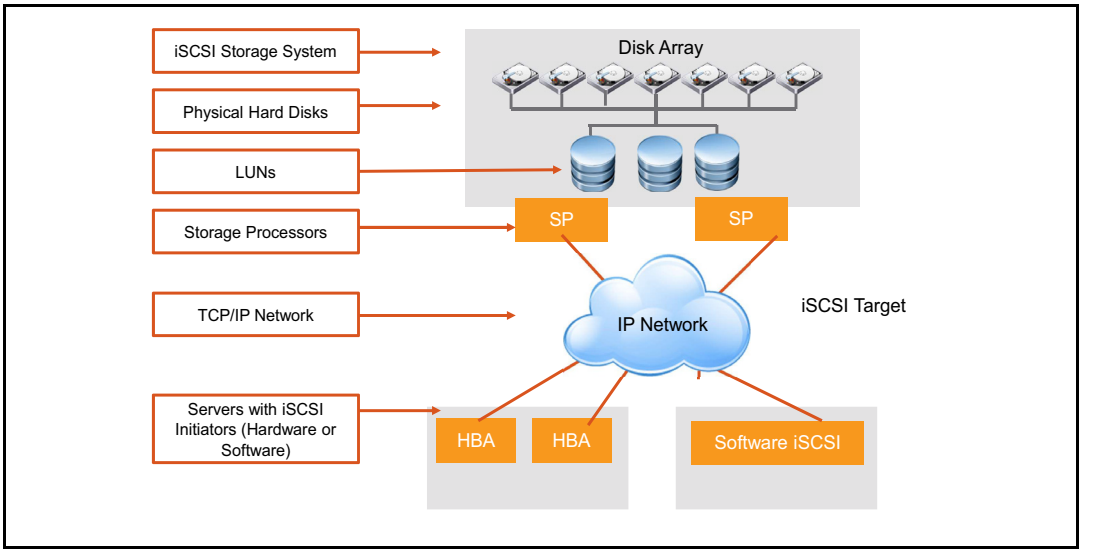
6-29 Learner Objectives

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

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

https://t.me/learningnets

6-30 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, in which case 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. You can have multiple initiators and targets in your iSCSI network. iSCSI is SAN-oriented because of the following reasons:

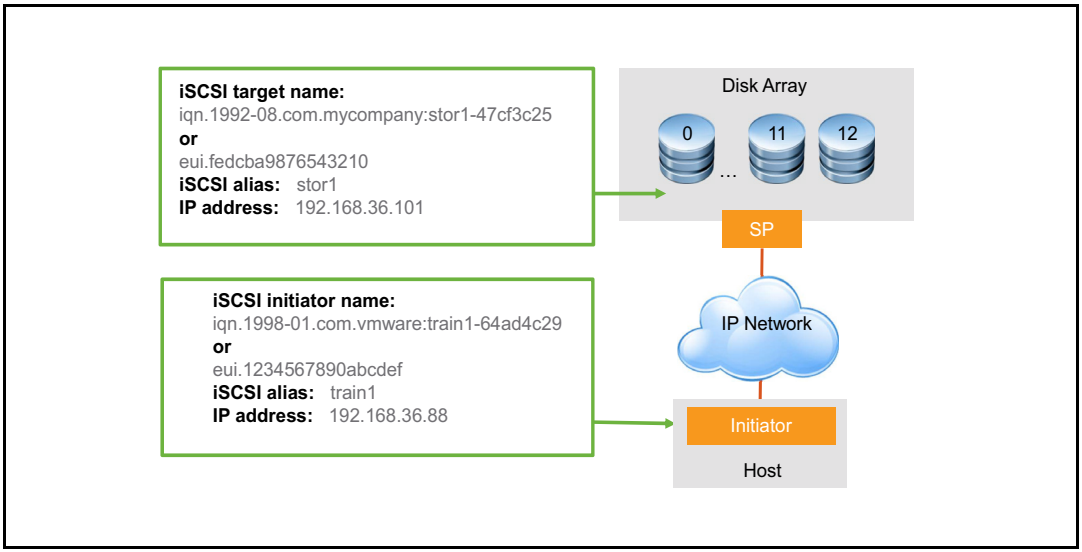
- The initiator finds one or more targets.
- A target presents LUNs to the initiator.
- The initiator sends SCSI commands to a target.

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

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

https://t.me/learningnets

6-31 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. The following naming convention is used:

- Prefix iqn
- Date code specifying the year and month in which the organization registered the domain or subdomain name 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

The following EUI naming convention is used:

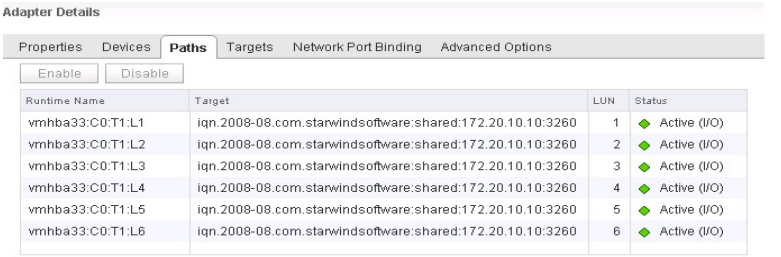
- Prefix eui, followed by a 16-character name. The name includes 24 bits for a company name that is assigned by the IEEE and 40 bits for a unique ID, such as a serial number.

https://t.me/learningnets

6-32 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 a unique identifier cannot be provided by the storage device, the VMkernel must generate a unique identifier to represent each LUN or disk.

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

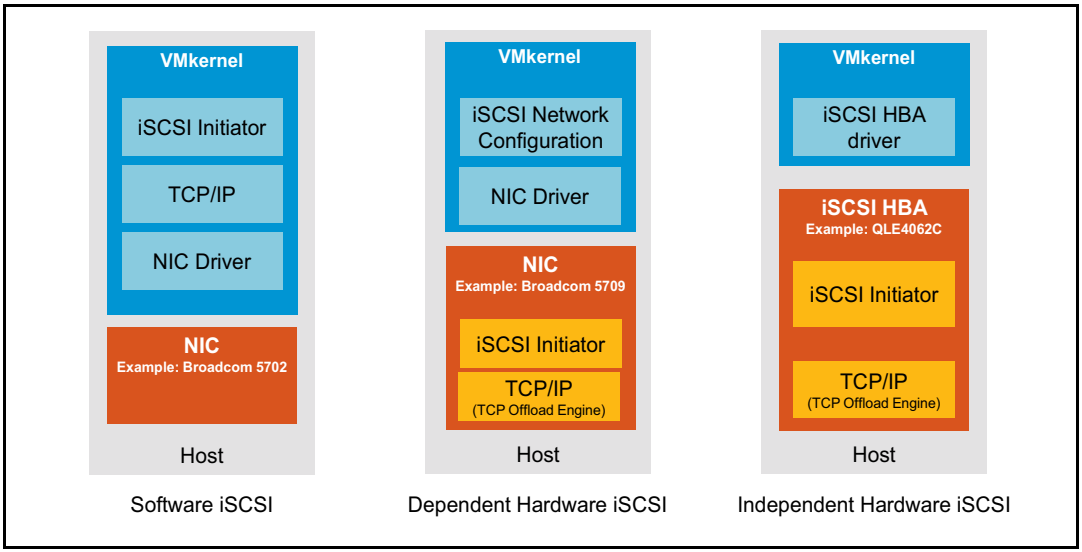
The following SCSI storage device identifiers are available:

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

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

https://t.me/learningnets

6-33 iSCSI Adapters



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

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

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

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

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

https://t.me/learningnets

6-34 Setting Up iSCSI Adapters

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

Supported iSCSI adapter types (vmhba):

- Software adapters
- Hardware adapters (independent or dependent)

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

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

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

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

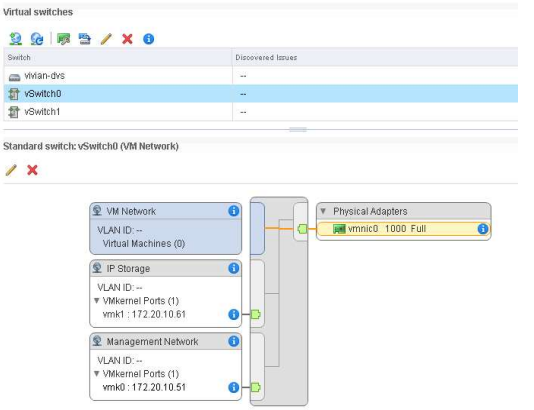
https://t.me/learningnets

6-35 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/NFS storage.

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

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



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

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

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

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

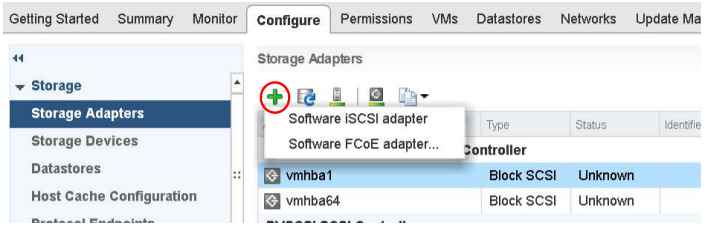
For the configuration steps for adding a VMkernel port to a virtual switch, see *vSphere Storage* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

6-36 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 the green plus sign.
3. Select **Software iSCSI 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

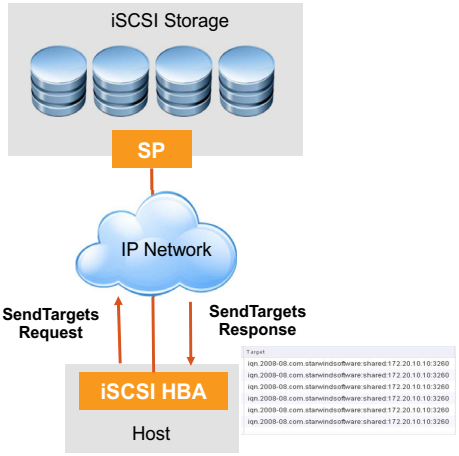
6-37 Discovering iSCSI Targets

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, called 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 all of the targets that it will contact in advance and uses their IP addresses and domain names to communicate with them.
- Dynamic discovery or SendTargets discovery: Each time the initiator contacts a specified iSCSI server, it sends the SendTargets request to the server. The server responds by supplying a list of available targets to the initiator. The names and IP addresses of these targets appear as static targets in vSphere Client. You can remove a static target that was added by dynamic discovery. If you remove the target, the target might be returned to the list during the next rescan operation. The target might also be returned to the list if the HBA is reset or the host is rebooted.

https://t.me/learningnets

6-38 iSCSI Security: CHAP

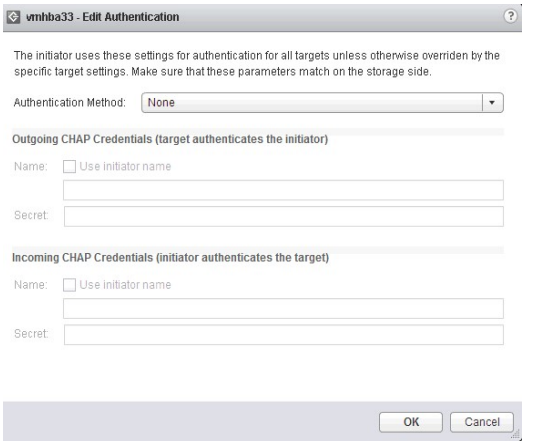
iSCSI initiators use CHAP for authentication purposes.

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: An additional level of security enables the initiator to authenticate the target. You must specify different target and initiator secrets.

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

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

Before configuring CHAP, check whether CHAP is enabled at the iSCSI storage system and check the CHAP authentication method that the system supports. If CHAP is enabled, you must enable it for your initiators, making sure that the CHAP authentication credentials match the credentials on the iSCSI storage. Although VMware recommends using CHAP in your iSCSI SAN implementation, consult with your storage vendor to ensure that best practices are followed.

https://t.me/learningnets

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

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

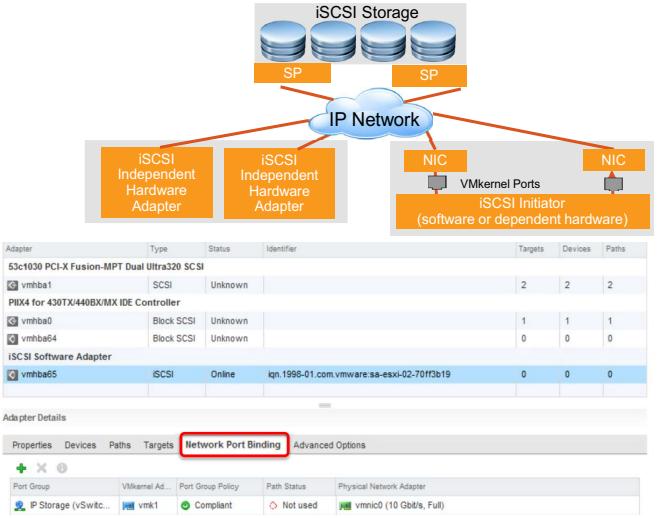
https://t.me/learningnets

6-39 Multipathing with iSCSI Storage

Software or dependent hardware iSCSI:

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

Independent hardware iSCSI uses two or more hardware iSCSI adapters



When setting up your ESXi host for multipathing and failover, you can use multiple hardware iSCSI adapters or multiple NICs. The choice depends on the type of iSCSI initiators on your host.

With software iSCSI and dependent hardware iSCSI, you can use multiple NICs that provide failover for iSCSI connections between your host and iSCSI storage systems. For this setup, because multipathing plug-ins do not have direct access to physical NICs on your host, you must first connect each physical NIC to a separate VMkernel port. Then you use a port-binding technique to associate all VMkernel ports with the iSCSI initiator. As a result, each VMkernel port connected to a separate NIC becomes a different path that the iSCSI storage stack and its multipathing plug-in can use. For dependent hardware iSCSI, ensure that the physical network card is correctly installed and appears on the host's **Configure** tab in the **Networking** > **Virtual Switches** view.

With independent hardware iSCSI, the host typically has two or more hardware iSCSI adapters available, from which the storage system can be reached using one or more switches. Alternatively, the setup might include one adapter and two storage processors so that the adapter can use a different path to reach the storage system.

After iSCSI multipathing is set up, each port on the ESXi system has its own IP address, but they all share the same iSCSI initiator IQN. When iSCSI multipathing is configured, the VMkernel routing table is not consulted when identifying the outbound NIC to use. Instead, iSCSI multipathing is managed using vSphere multipathing modules. Due to the latency that can be incurred, VMware does not recommend routing iSCSI traffic. For more about configuring iSCSI multipathing, see *vSphere Storage* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

6-40 Lab 5: Accessing iSCSI Storage

Configure access to an iSCSI datastore

1. Validate an Existing ESXi Host iSCSI Configuration
2. Add a VMkernel Port Group to a Standard Switch
3. Configure the iSCSI Software Adapter
4. Connect the iSCSI Software Adapters to Storage

https://t.me/learningnets

6-41 Review of Learner Objectives

You should be able to meet the following objectives:

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

https://t.me/learningnets



https://t.me/learningnets

6-43 Learner Objectives

By the end of 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

Use VMFS or vSphere Virtual Volumes datastores when accessing block-level storage:

- VMFS is optimized for storing and accessing large files.
- A VMFS datastore can have a maximum volume size of 64 TB.
- vSphere Virtual Volumes can use advanced storage services that include replication, encryption, deduplication, and compression on individual virtual disks.
- vSphere Virtual Volumes supports such vSphere features as VMware vSphere® vMotion™, vSphere Storage vMotion, snapshots, linked clones, Flash Read Cache, and vSphere DRS.

Use RDMs if any of the following conditions are true of your virtual machine:

- The VM is taking storage array-level snapshots.
- It is clustered to a physical machine.
- It has large amounts of data that you do not want to convert into a virtual disk.

VMFS datastores primarily serve as repositories for the files of virtual machines. A VMFS datastore is optimized for storing and accessing large files like virtual disks and memory images of suspended virtual machines. The maximum size of a VMFS datastore is 64 TB.

A Virtual Volumes datastore represents a storage container in vCenter Server and vSphere Web Client. From a vSphere administrator prospective, the Virtual Volumes datastore is similar to any other datastore and is used to hold virtual machines. The Virtual Volumes datastore capacity is configurable by the storage administrator outside of vSphere.

You can use an NFS datastore to store your virtual machines, templates, and ISO images, but not all functions are supported. For example, you cannot store an RDM on an NFS datastore. A VMFS datastore is required for an RDM to store the RDM mapping file (`*-rdm.vmdk`).

As for RDMs, choose RDMs over VMFS datastores if:

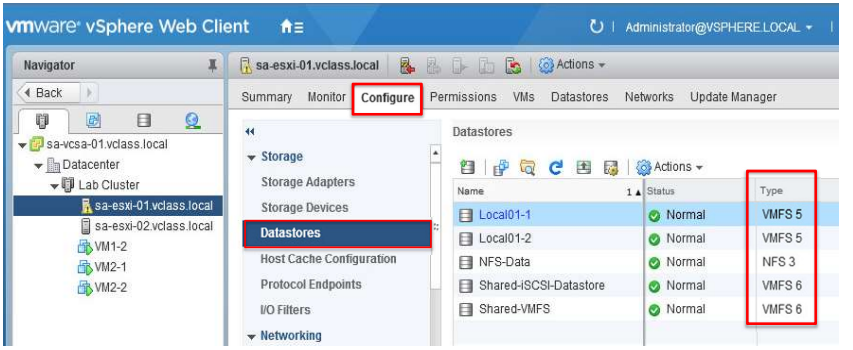
- A virtual machine is using storage array-level snapshot applications.
- A virtual machine is clustered with a physical machine.
- You want to keep the virtual machine's data on a raw disk instead of converting it to a virtual disk because, for example, the data disk is very large.

https://t.me/learningnets

6-45 Creating and Viewing VMFS Datastores

VMFS datastores serve as repositories for virtual machines.

Using the New Datastore wizard, you can create VMFS datastores on any SCSI-based storage devices that the host discovers, including Fibre Channel, iSCSI, and local storage devices.



By viewing datastores, you can determine the type of datastores in use. For example, you can determine whether your datastore is an NFS datastore, a VMFS datastore, or any other type of datastores.

ATS+SCSI mechanism: A VMFS datastore that supports the ATS+SCSI mechanism is configured to use ATS and attempts to use it when possible. If ATS fails, the VMFS datastore reverts to SCSI reservations. In contrast with the ATS locking, the SCSI reservations lock an entire storage device while an operation that requires metadata protection is performed. After the operation completes, VMFS releases the reservation and other operations can continue.

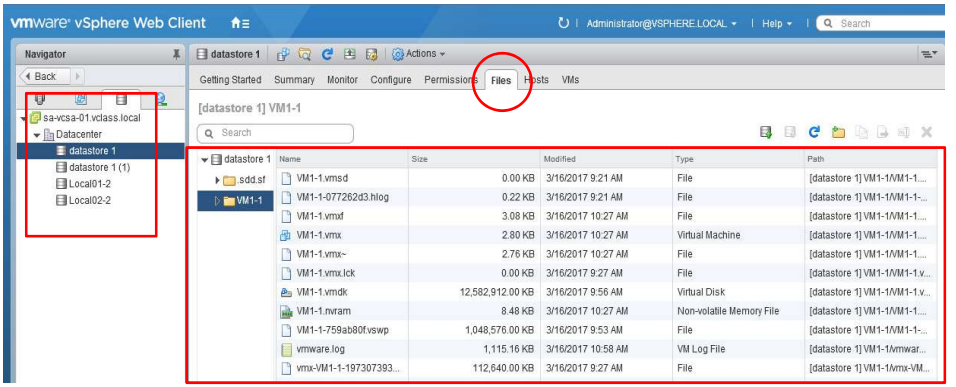
Datastores that use the ATS+SCSI mechanism include VMFS5 datastores that were upgraded from VMFS3. In addition, new VMFS5 or VMFS6 datastores on storage devices that do not support ATS use the ATS+SCSI mechanism.

If your VMFS datastore reverts to SCSI reservations, you might notice performance degradation caused by excessive SCSI reservations. For information about how to reduce SCSI reservations, see *vSphere Troubleshooting* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

6-46 Browsing Datastore Contents

You use the datastore file browser to manage the contents of your datastores.



The Datastores pane shows the datastores that are available in many views. The tab lists all datastores currently configured for the ESXi host.

The example shows the contents of the VMFS datastore named datastore1. The contents are the virtual machine folders. The files for each virtual machine are in the respective virtual machine folder.

https://t.me/learningnets

6-47 Managing Overcommitted Datastores

A datastore becomes overcommitted when the total provisioned space of thin-provisioned disks is greater than the size of the datastore.

Actively monitor your datastore capacity:

- Alarms assist through notifications:
  - Datastore disk overallocation
  - Virtual machine disk usage
- Use reporting to view space usage.

Actively manage your datastore capacity:

- Increase datastore capacity when necessary.
- Use vSphere Storage vMotion to mitigate space usage problems on a particular datastore.

Using thin-provisioned virtual disks for your virtual machines is a way to make the most of your datastore capacity. But if your datastore is not sized properly, it can become overcommitted. A datastore becomes overcommitted when the full capacity of its thin-provisioned virtual disks is greater than the datastore's capacity. If disk capacity is managed correctly, then you (the vSphere administrator) should not run out of space, despite external issues like budgeting.

When a datastore reaches capacity, vSphere Web Client prompts you to provide more space on the underlying VMFS datastore and freezes the virtual machine until you do so. Monitor your datastore capacity by using alarms to alert you of how much a datastore's disks are overallocated or how much disk space a virtual machine is using.

Manage your datastore capacity. Dynamically increase the size of your datastore when necessary. You can also use vSphere Storage vMotion to mitigate space usage issues. For example, with vSphere Storage vMotion, you can migrate a virtual machine off a datastore. The migration can be done with a change from virtual disks of thick format to thin format at the target datastore.

https://t.me/learningnets

6-48 Increasing the Size of VMFS Datastores

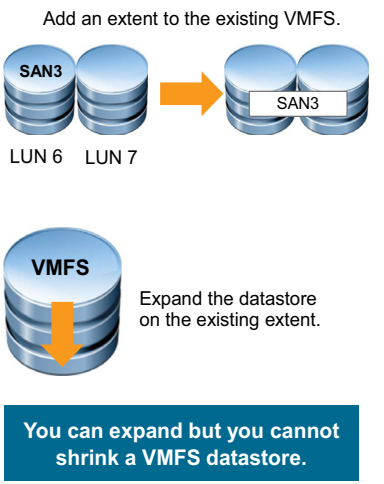
In general, before making any changes to your storage allocation:

- Perform a rescan to ensure that all hosts see the most current storage.
- Record the unique identifier.

Increase a VMFS datastore's size to give it more space or to possibly improve performance.

Ways to dynamically increase the size of a VMFS datastore:

- Add an extent (LUN).
- Expand the datastore within its extent.



In general, before you make storage allocation changes, a good practice is to rescan to ensure that all hosts see the current storage view. Also, record the unique identifier of the volume that you want to expand, for example, the NAA ID. You need this information to identify the VMFS datastore whose size you want to increase.

You can dynamically increase the capacity of a VMFS datastore if the datastore has insufficient disk space. Insufficient disk space is realized when you create a virtual machine or you try to add more disk space to a virtual machine.

Use one of the following methods:

- Add an extent to the VMFS datastore: An extent is a partition on a LUN. You can add an extent to any VMFS datastore. The datastore can stretch over multiple extents, up to 32.
- Expand the VMFS datastore: Increase the size of the VMFS datastore in its extent. Only extents with free space immediately after them are expandable. As a result, rather than adding the new extent, you can expand the existing extent so that it fills the available adjacent capacity.

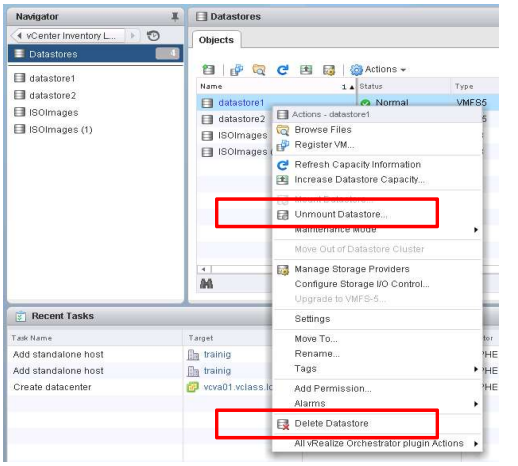
For more information about growing or expanding a VMFS volume or datastore, see VMware knowledge base article 1017662 at http://kb.vmware.com/kb/1017662.

https://t.me/learningnets

6-49 Deleting or Unmounting VMFS Datastores

An unmounted datastore remains intact but cannot be seen from the hosts that you specify. The datastore continues to appear on other hosts, where it remains mounted.

A deleted datastore is destroyed and disappears from all hosts that have access to it. All virtual machine files on the datastore are permanently removed.



When you unmount a datastore, it remains intact, but can no longer be seen from the hosts that you specify. Unmounting a VMFS datastore preserves the files on the datastore but makes the datastore inaccessible to the ESXi host. The datastore is still visible and continues to serve other hosts, where it remains mounted. Do not perform any configuration operations that might result in I/O to the datastore while the unmount is in progress.

You can delete any type of VMFS datastore, including copies that you have mounted without resignaturing. When you delete a datastore, it is permanently destroyed and disappears from all hosts that have access to the datastore. This operation permanently removes all files that are associated with virtual machines on the datastore. Although you can delete the datastore without unmounting, it is preferable that you unmount the datastore first. Deleting a VMFS datastore destroys the pointers to the files on the datastore, so the files disappear from all hosts that have access to the datastore.

Before you delete or unmount a VMFS datastore, power off all virtual machines whose disks reside on the datastore. If you do not power off the virtual machines and you try to continue, an error message tells you that the resource is busy. Before you unmount a VMFS datastore, use vSphere Web Client to verify the following:

- No virtual machines reside on the datastore.
- The datastore is not part of a datastore cluster.
- The datastore is not managed by VMware vSphere® Storage DRS™.
- VMware vSphere® Storage I/O Control is disabled.
- The datastore is not used for vSphere HA heartbeat.

To keep your data, back up the contents of your VMFS datastore before you delete the datastore.

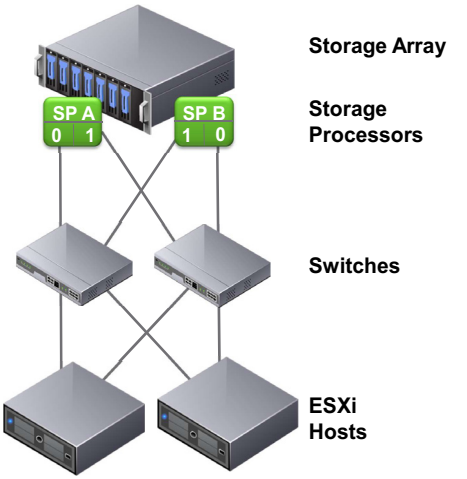
https://t.me/learningnets

6-50 Multipathing Algorithms

Arrays provide various features. Some offer active-active storage processors. Others offer active-passive storage processors.

vSphere offers native path selection, load-balancing, and failover mechanisms.

Third-party vendors can create their own software to be installed on ESXi hosts. This third-party software enables hosts to properly interact with the storage arrays.



The Pluggable Storage Architecture is a VMkernel layer responsible for managing multiple storage paths and providing load balancing. An ESXi host can be attached to storage arrays with the following storage processor configuration:

- Active-active
- Active-passive

VMware offers native load-balancing and failover mechanisms. Examples of VMware path selection policies include:

- Round Robin
- Most Recently Used (MRU)
- Fixed

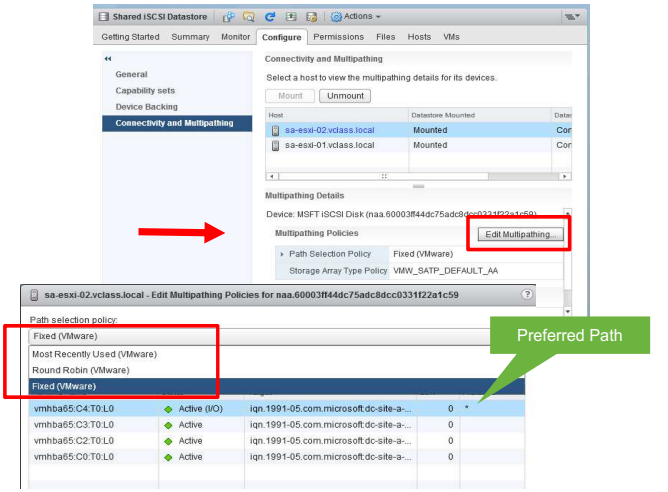
Third-party vendors are able to design their own load-balancing techniques and failover mechanisms for particular storage array types to add support for new arrays. Third-party vendors do not need to provide internal information or intellectual property about the array to VMware.

https://t.me/learningnets

6-51 Configuring Storage Load Balancing

Path selection policies exist for:

- Scalability:
  - Round Robin: A multipathing policy that performs load balancing across paths
- Availability:
  - MRU
  - Fixed



Multiple paths can exist to a datastore from an ESXi host. You can view these paths in the datastore properties by clicking the **Manage** tab and the **Settings** tab.

The following path selection policies are supported for multipathing with Fibre Channel or iSCSI:

- Fixed path selection policy: The host always uses the preferred path to the disk when that path is available. If the host cannot access the disk through the preferred path, it tries the alternative paths. This is the default policy for active-active storage devices.
- Most Recent Used (MRU) path selection policy: The host uses the most recent path to the disk until this path becomes unavailable. That is, the host does not revert back until this path becomes unavailable. A failover to a new path is performed. If the original path becomes available again, the host does not fail back to the original path. MRU is the default policy for active-passive storage devices and is required for those devices.
- Round Robin path selection policy: The host uses a path selection algorithm that rotates through all available paths. In addition to path failover, the Round Robin policy supports load balancing across the paths. Before using this policy, check with storage vendors to find out whether a Round Robin configuration is supported on their storage.

https://t.me/learningnets

6-52 Lab 6: 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 Shared VMFS Datastore Using iSCSI

https://t.me/learningnets

6-53 Review of Learner Objectives

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

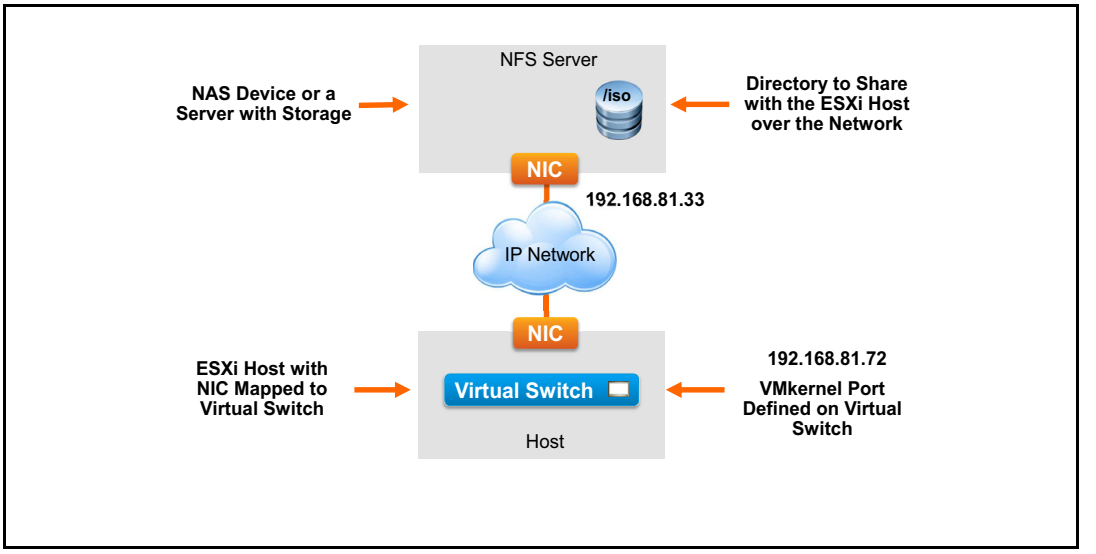
6-55 Learner Objectives

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

- Describe NFS components
- Describe the differences between NFS v3 and NFS v4.1
- Configure and manage NFS datastores

https://t.me/learningnets

6-56 NFS Components



An NFS file system is located on a NAS device that is called the NFS server. The NFS server contains one or more directories that are shared with the ESXi host over a TCP/IP network. An ESXi host accesses the NFS server through a VMkernel port that is defined on a virtual switch.

https://t.me/learningnets

6-57 Configuring NFS Datastores

Create a VMkernel port:

- For better performance and security, separate your NFS network from the iSCSI network.

Provide the following information:

- NFS version: v3 or v4.1
- Datastore name
- NFS server names or IP addresses
- Folder on the NFS server, for example, `/templates` and `/nfs_share`
- Select hosts that will mount the datastore
- Whether or not to mount the NFS file system read-only
- Authentication parameters

For each ESXi host to access an NFS datastore over the network, a VMkernel port must be configured on a virtual switch. The name of this port can be anything that you want. For performance and security reasons, a good practice is to isolate your NFS networks from the other networks, such as your iSCSI network and your virtual machine networks.

https://t.me/learningnets

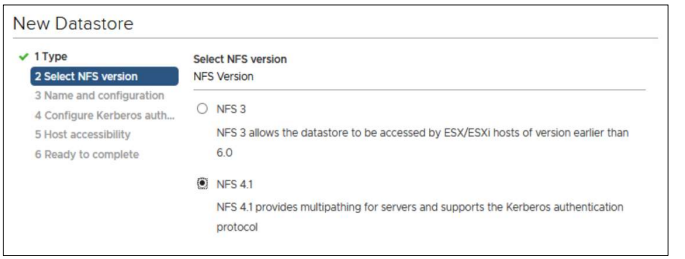
6-58 NFS v3 and NFS v4.1

NFS v3:

- ESXi managed multipathing
- AUTH_SYS (root) authentication
- VMware proprietary file locking
- Client-side error tracking

NFS v4.1:

- Native multipathing and session trunking
- Optional Kerberos authentication
- Built-in file locking
- Server-side error tracking



An NFS datastore can be created as either NFS v3 or NFS v4.1. Various compatibility issues between the two NFS versions preclude accessing datastores using both protocols at the same time from different hosts. If a datastore is configured as NFS v4.1, all hosts that access that datastore must mount the share as NFS 4.1. Data corruption can occur if hosts access a datastore with the wrong NFS version.

As of vSphere 6.5, NFS v4.1 is not compatible with vSphere Storage DRS, VMware vSphere® Storage I/O Control, and VMware Site Recovery Manager™, because of server protocol locking.

https://t.me/learningnets

6-59 NFS Version Compatibility with Other vSphere Technologies

	NFS v3	NFS v4.1
vSphere vMotion and vSphere Storage vMotion	Yes	Yes
vSphere HA	Yes	Yes
vSphere Fault Tolerance	Yes	Yes
vSphere DRS and vSphere DPM	Yes	Yes
Stateless ESXi and Host Profiles	Yes	Yes
vSphere Storage DRS and Storage I/O Control	Yes	No
Site Recovery Manager	Yes	No
vSphere Virtual Volumes	Yes	Yes

vSphere 6.x supports NFS v4.1 to overcome many limitations when using NFS v3. In vSphere 6.x, both NFS v3 and NFS v4.1 shares can be used, although some important constraints must be taken into consideration when designing a vSphere environment in which both versions will be used.

NFS v4.1 provides the following enhancements:

- Native multipathing and session trunking: NFS v4.1 provides multipathing for servers that support session trunking. When trunking is available, you can use multiple IP addresses to access a single NFS volume. Client ID trunking is not supported.
- Kerberos authentication: NFS v4.1 introduces Kerberos authentication in addition to the traditional AUTH_SYS method used by NFS v3.
- Improved built-in file locking.
- Enhanced error recovery using server-side tracking of open files and delegations.
- Many general efficiency improvements including session leases and less protocol overhead.

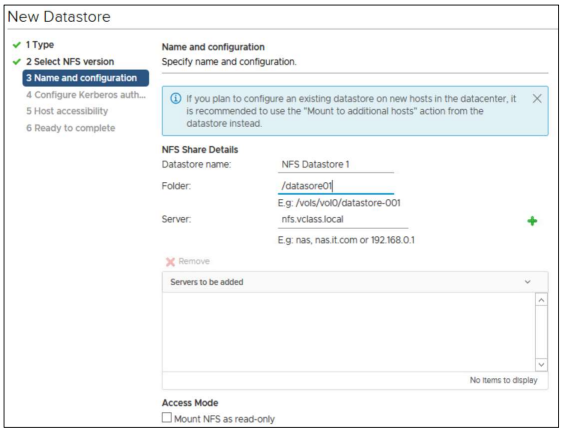
The NFS v4.1 client offers the following new features:

- Stateful locks with share reservation using a mandatory locking semantic
- Protocol integration, side-band (auxiliary) protocol no longer required to lock and mount
- Trunking (true NFS multipathing), where multiple paths (sessions) to the NAS array can be created and load distributed across those sessions
- Enhanced error recovery to mitigate server crash and loss of connectivity

https://t.me/learningnets

6-60 NFS Datastore Name and Configuration

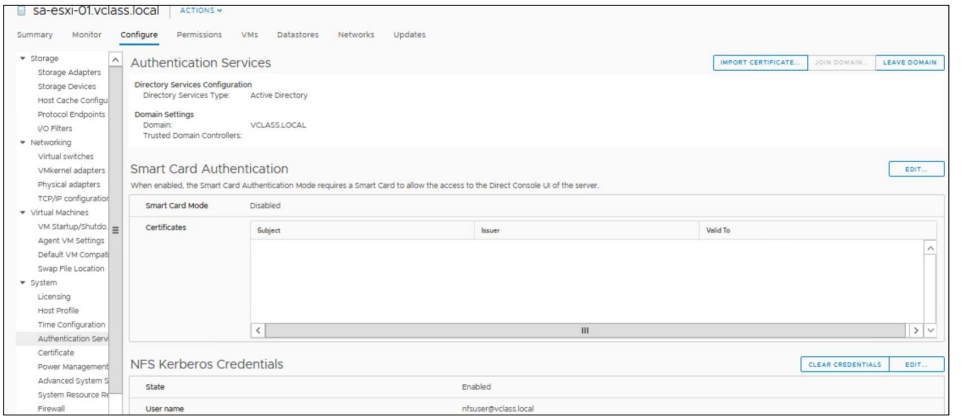
You assign the NFS datastore a name and edit datastore parameters.



https://t.me/learningnets

6-61 Configuring ESXi Host Authentication and NFS Kerberos Credentials

As a requirement of Kerberos authentication, you add each ESXi host to the AD domain. Then you configure NFS Kerberos credentials.



Each ESXi host must be added to the Active Directory domain.

A number of configuration steps must be taken to prepare each ESXi host to use Kerberos authentication. Kerberos authentication requires that all nodes involved (the Active Directory Server, the NFS servers, and the ESXi hosts) be synchronized so that little to no time drift exists. Kerberos authentication fails if any significant drift exists between the nodes. To prepare your ESXi host to use Kerberos authentication, configure the NTP client settings to reference a common NTP server (or the domain controller if applicable).

Administrators must consider the following when planning on using NFS Kerberos:

- NFS v3 and v4.1 use different authentication credentials, resulting in incompatible UID/GID on files.
- Using different Active Directory users on different hosts that access the same NFS share can cause the vSphere vMotion migration to fail.
- NFS Kerberos configuration can be automated using host profiles to reduce configuration conflicts.

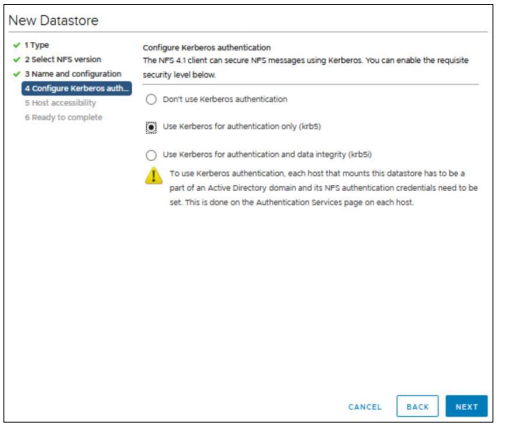
Time must be synchronized between all participating components.

https://t.me/learningnets

6-62 Configuring Datastores to Use Kerberos

Enable Kerberos authentication when creating each datastore and select the security mode:

- Kerberos5 authentication
- Kerberos5i authentication and data integrity



After the configuration steps are performed, you can configure the datastore to use Kerberos authentication.

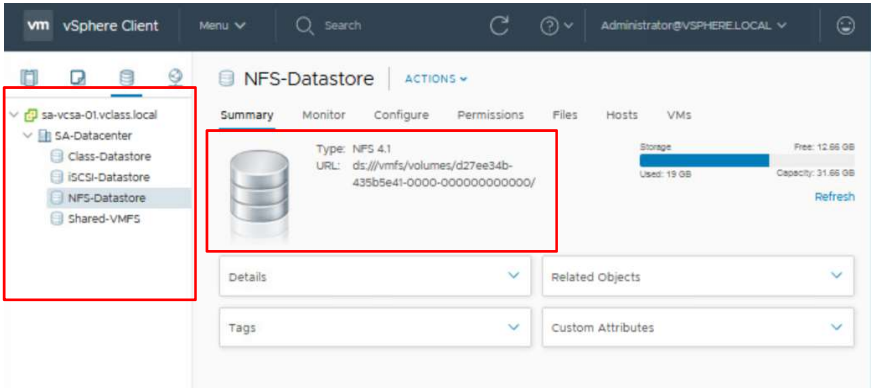
The screenshot shows a choice of Kerberos authentication only (krb5) or authentication with data integrity (krb5i). The difference is whether only the header or the header and the body of each NFS operation is signed using a secure checksum.

For more information about how to configure the ESXi hosts for Kerberos authentication, see https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.storage.doc/GUID-BDCB7500-72EC-4B6B-9574-CFAEAF95AE81.html.

https://t.me/learningnets

6-63 Viewing IP Storage Information

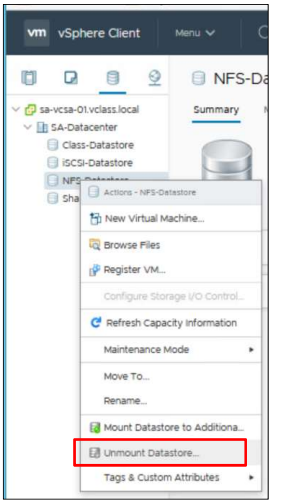
You can view the details of the VMFS or NFS datastores that you created.



https://t.me/learningnets

6-64 Unmounting NFS Datastores

Unmounting an NFS datastore causes the files on the datastore to become inaccessible to the ESXi host.



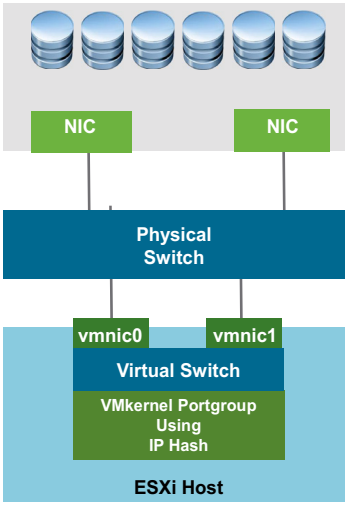
When an NFS datastore is unmounted, the files located on the NFS datastore become inaccessible to the ESXi host. Before unmounting an NFS datastore, you must stop all virtual machines whose disks reside on the datastore.

https://t.me/learningnets

6-65 Multipathing and NFS Storage

An example of a configuration for NFS multipathing:

- Configure one VMkernel port.
- Use adapters attached to the same physical switch to configure NIC teaming.
- Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.
- To better use multiple links, configure NIC teams with the IP hash load-balancing policy.



To create a highly available NAS architecture, you must avoid single points of failure. Examples of a single point of failure include the NIC card in an ESXi host, and the cable between the NIC card and the switch. One example is shown on the slide. To avoid single points of failure and to create a highly available NAS architecture, configure the ESXi host with redundant NIC cards and redundant physical switches.

The best approach is to install multiple NICs on an ESXi host and configure them in NIC teams. NIC teams should be configured on separate external switches, with each NIC pair configured as a team on the respective external switch. In addition, you may apply load-balancing algorithm, based on the link aggregation protocol type supported on the external switch, such as 802.3ad or EtherChannel.

As another example, an even higher level of performance and high availability can be achieved with cross-stack, EtherChannel-capable switches. With certain network switches, you can team ports across two or more separate physical switches that are managed as one logical switch. NIC teaming across virtual switches provides additional resilience as well as some performance optimization. Having more paths available to the ESXi host can improve performance by enabling distributed load sharing.

Only one active path is available for the connection between the ESXi host and a single storage target (LUN or mount point). Although alternative connections might be available for failover, the bandwidth for a single datastore and the underlying storage is limited to what a single connection can provide. To leverage more available bandwidth, an ESXi host would require multiple connections from the ESXi host to the storage targets. You might need to configure multiple datastores, each using separate connections between the ESXi host and the storage.

https://t.me/learningnets

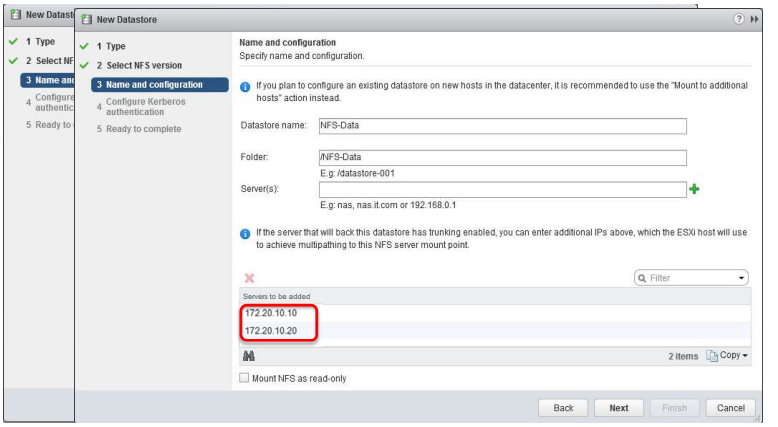
The table shows the recommended configuration for NFS multipathing.

External Switches Support Cross-Stack EtherChannel	External Switches Do Not Support Cross-Stack EtherChannel
Configure one VMkernel port.	Configure two or more VMkernel ports on different virtual switches on different subnets.
Configure NIC teaming by using adapters attached to separate physical switches.	Configure NIC teaming with adapters attached to the same physical switch.
Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.	Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.
To use multiple links, configure NIC teams with the IP hash load-balancing policy.	To use multiple links, allow the VMkernel routing table to make decisions on which link to send packets (requires multiple datastores).

https://t.me/learningnets

6-66 Enabling Session Trunking for NFS 4.1

Multiple IP addresses are configured for each NFS v4.1 datastore.



NFS v4.1 supports native multipathing and session trunking. To enable multipathing, the administrator enters multiple server IP addresses when configuring the datastore.

https://t.me/learningnets

6-67 Lab 7: Accessing NFS Storage

Configure access to an NFS datastore

1. Configure Access to NFS Datastores
2. View NFS Storage Information

https://t.me/learningnets

6-68 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe NFS components
- Describe the differences between NFS v3 and NFS v4.1
- Configure and manage NFS datastores

https://t.me/learningnets



https://t.me/learningnets

6-70 Learner Objectives

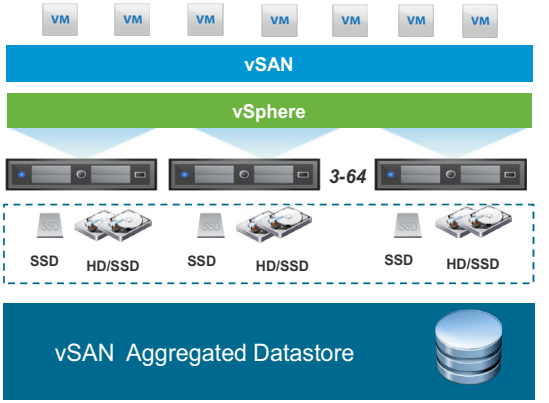
By the end of this lesson, you should be able to meet the following objectives:

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Describe the steps for configuring vSAN
- Explain how to create and use vSAN storage policies

https://t.me/learningnets

6-71 About vSAN

A single vSAN datastore is created using storage from multiple hosts and disks in the cluster.



Starting with vSphere 5.5, VMware introduced software-defined storage with vSAN datastores. vSAN aggregates local storage devices (flash devices and magnetic disks) to provide a clustered datastore that can be used by virtual machines.

vSAN 6.x provides two different configurations: hybrid and all-flash. The hybrid configuration uses server-based flash devices to provide a cache layer for optimal performance while using magnetic disks to provide capacity and persistent data storage. The all-flash configuration uses flash devices for both the cache layer and capacity layer.

In a vSAN environment, a number of ESXi hosts are configured to form a vSAN cluster. All of the ESXi hosts communicate together through a dedicated vSAN network. Hosts without local storage can share their compute resources and take advantage of the clustered storage resources. Local storage is combined on each host to form up to five local disk groups. A disk group includes one flash cache device and up to seven capacity devices.

The disk groups of all of the ESXi hosts in the vSAN cluster are combined to create a vSAN datastore. Only a single datastore exists per vSAN cluster. Object Store File System (OSFS) enables the VMFS volumes from each host to be combined and mounted as a single datastore. This datastore contains all the virtual machine files, including the VMDK files. Each of the VMDK files can have a different virtual machine storage policy created that defines how data is stored on the disks of the datastore. These virtual machine storage policies are configured to take advantage of the capabilities of vSAN.

https://t.me/learningnets

vSAN datastores help administrators to use software-defined storage in the following ways:

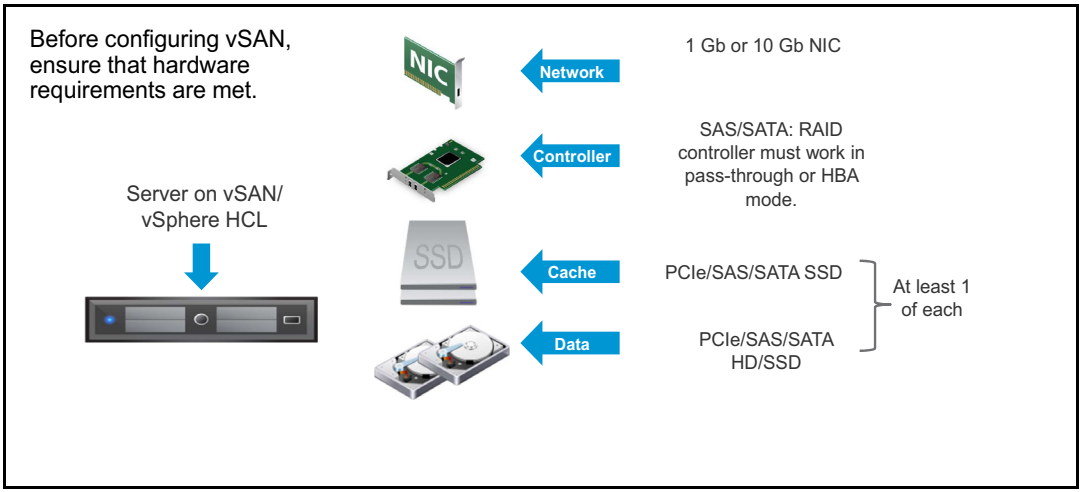
- Storage policy per virtual machine architecture: Multiple policies per datastore allow for each virtual machine to have different storage.
- vSphere and vCenter Server integration: vSAN capability is built in and requires no appliance. You create a vSAN cluster, just like vSphere HA or vSphere DRS.
- Scale-out storage: Up to 64 ESXi hosts can be in a cluster. Scale out by populating new nodes in the cluster or by setting vSAN to scan for empty disks and add them automatically.
- Built-in resiliency: A default policy mirrors all objects for virtual machines that are not configured for the vSAN.

vSAN 6.x introduces several features and enhancements:

- New on-disk format
- New snapshot methodology providing higher performance
- Fault domains designed to withstand rack failure
- Default per-VM storage policy assignable at the datastore level
- Visualization of vSAN datastore usage
- Projection of usage based on possible storage policy configuration
- Resync status displayed in the vSphere Web Client interface
- New disk serviceability functions

https://t.me/learningnets

6-72 vSAN Requirements



The requirements for vSAN start with a machine running vCenter Server. This requirement is necessary for management because vSAN is fully integrated into vSphere.

A vSAN cluster must have a minimum of three hosts that contribute capacity to the cluster. A host must not participate in other clusters beside the vSAN cluster.

Each vSAN cluster requires a dedicated network that each host can communicate with:

- Dedicated 1 Gbps network (10 Gbps recommended) for hybrid disk groups
- Dedicated 10 Gbps network for all-flash disk groups

VMware recommends a 10 Gb network, preferably with a NIC team of 2 x 10 Gb NICs for fault tolerance purposes. Testing has shown that a 1 Gb network can work, but this network size is not supported. The vSAN network must be configured for IPv4 or IPv6 and support multicast.

All hosts in the cluster do not need to have local storage. A host without storage is used for its compute resources while leveraging the vSAN datastore. All hosts that have local storage must have the following devices:

- One Serial Attached SCSI (SAS) or SATA solid-state disk (SSD) or one PCIe flash device and one or more magnetic drives for each hybrid disk group
- One SAS or SATA SSD or PCIe flash device and one or more flash disks with flash capacity enabled for all-flash disk groups

In addition, each host should have 32 GB or more of memory to accommodate a maximum number of five disk groups and a maximum number of seven capacity devices per disk group.

For more information, see *Administering VMware Virtual SAN* at https://www.vmware.com/support/pubs/virtual-san-pubs.html.

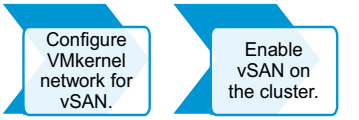
https://t.me/learningnets

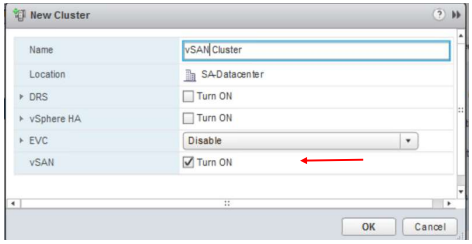
6-73 Configuring vSAN

To use vSAN, you create a host cluster and enable vSAN on the cluster.

To enable the exchange of data in the vSAN cluster, you provide a VMkernel network adapter for vSAN traffic on each ESXi host.

You can enable vSAN when you create a cluster. Or you can edit cluster properties to enable vSAN for an existing cluster.





vSAN datastores are fully integrated with vSphere 6.x.

For information about enabling or disabling a vSAN cluster, see VMware knowledge base article 2058322 at http://kb.vmware.com/kb/2058322.

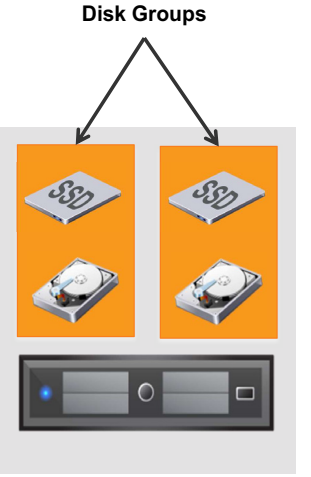
You can configure vSAN only through vSphere Web Client.

https://t.me/learningnets

6-74 Disk Groups

vSAN disk groups composed of flash-based devices and magnetic disks require:

- One flash device:
  - Maximum of one flash device per disk group
- One HDD/SSD:
  - Supports up to seven devices per disk group
- Maximum of five disk groups per host



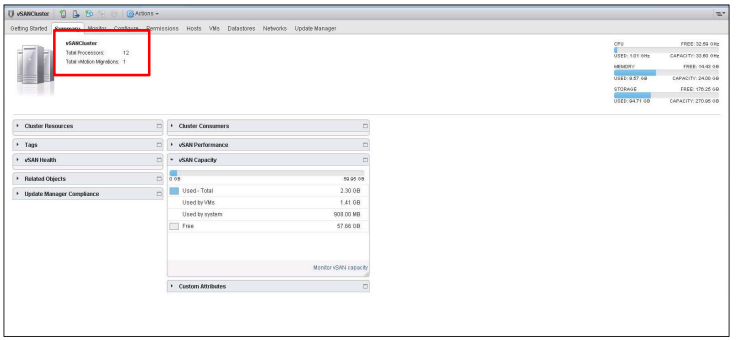
On each ESXi host that contributes its local devices to a vSAN cluster, devices are organized into disk groups. A disk group is a unit of physical storage capacity on a host and a group of physical devices that provide performance and capacity to the vSAN cluster.

The disk groups of all ESXi hosts in a vSAN cluster are combined to create a vSAN datastore. The size of a vSAN datastore is governed by the number of capacity devices per ESXi host and the number of ESXi hosts in the cluster. The vSAN datastore is used to store virtual machine files, including virtual machine disks (VMDKs).

https://t.me/learningnets

6-75 Viewing the vSAN Cluster Summary

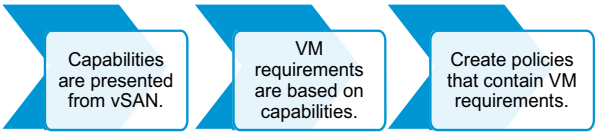
In vSphere Web Client, the **Summary** tab of the vSAN cluster shows the general vSAN configuration information.



Capabilities define the capacity, performance, and availability characteristics of the underlying physical storage. The vSAN cluster presents these capabilities to vCenter Server, where they can be consumed by virtual machines.

Requirements outline the needs of a virtual machine.

Virtual machine storage policies specify the virtual machine requirements so that the virtual machine can be placed appropriately on the vSAN datastore.



After you enable vSAN on a cluster, a single vSAN datastore is created. This datastore appears as another type of datastore on the list of datastores that might be available, including vSphere Virtual Volumes, VMFS, and NFS.

A single vSAN datastore can provide different service levels for each virtual machine or each virtual disk. In vCenter Server, storage characteristics of the vSAN datastore appear as a set of capabilities.

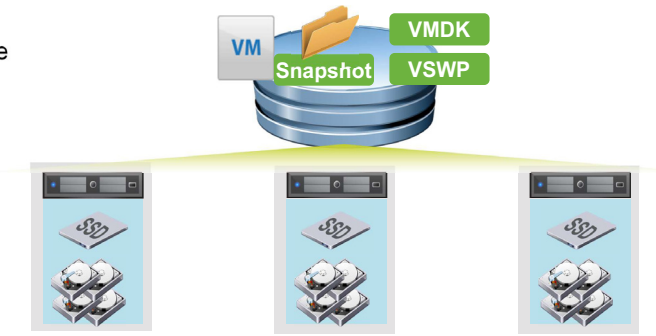
You can reference these capabilities when defining a storage policy for virtual machines. When you later deploy virtual machines, vSAN uses this policy to place virtual machines in the optimal manner based on the requirements of your virtual machine.

https://t.me/learningnets

6-77 Objects in vSAN Datastores

In a vSAN datastore, files are grouped into the following types of objects:

- VM home namespaces
- Virtual disks
- Snapshots
- Swap files
- Memory state



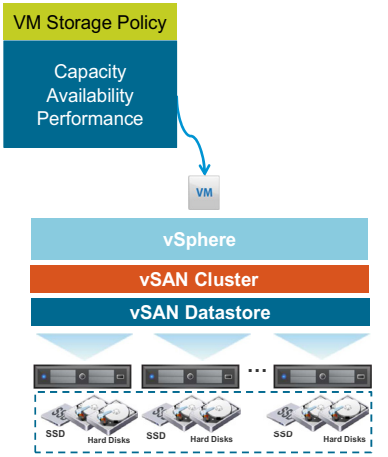
A vSAN cluster stores and manages data as flexible data containers called objects. When you provision a virtual machine on a vSAN datastore, a set of objects is created. The following types of objects are created:

- VM home namespace: Stores the virtual machine metadata (configuration files)
- VMDK: Virtual machine disk
- Snapshot deltas: Created when snapshots of the virtual machine are taken
- VM swap: Virtual machine swap file, which is created when the virtual machine is powered on
- VM memory: Virtual machine's memory state when a virtual machine is suspended or when a snapshot is taken of a virtual machine and its memory state is preserved

https://t.me/learningnets

6-78 Virtual Machine Storage Policies

- Virtual machine storage policies are built before VM deployment to reflect the requirements of the application running in the virtual machine.
- These policies are based on the vSAN capabilities.
- The storage policy for the virtual machine is selected based on the virtual machine's requirements and storage objects for the virtual machine are then created that meet the policy requirements.



Virtual machine storage policies are a set of storage properties that are configured by an administrator to be used by virtual machines. Each of these storage policies is created to reflect a set of capabilities that meets the availability, performance, and storage requirements of some of the virtual machines in an environment. These storage policies are created in advance of the deployment of virtual machines.

When a virtual machines is deployed, the administrator chooses the VM storage policy that meets its requirements for the creation of its VMDK file.

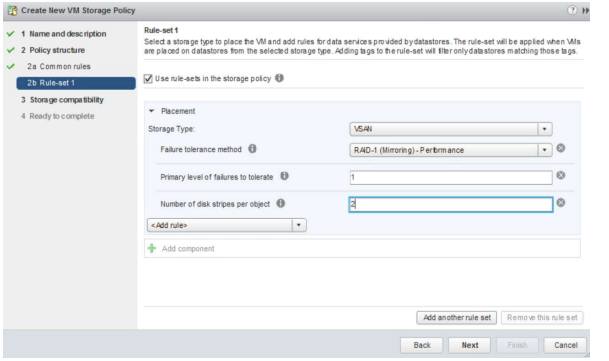
https://t.me/learningnets

6-79 Configuring Virtual Machine Storage Policies

Virtual machine storage policies are essential to virtual machine provisioning.

The policies control:

- Which type of storage is provided for the virtual machine
- How the virtual machine is placed within the storage
- Which data services are offered for the virtual machine



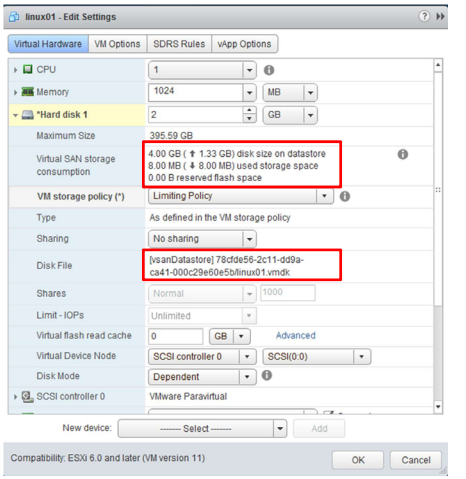
https://t.me/learningnets

6-80 Viewing Virtual Machine vSAN Datastores

The consumption of vSAN storage is based on the virtual machine's storage policy.

The virtual machine's hard disk view provides the following information:

- A summary of the total storage size and the used storage space
- A display of the virtual machine storage policy
- The location of disk files on a vSAN datastore



https://t.me/learningnets

6-81 Review of Learner Objectives

You should be able to meet the following objectives:

- Explain the purpose of a vSAN datastore
- Describe the architecture and requirements of vSAN configuration
- Describe the steps for configuring vSAN
- Explain how to create and use vSAN storage policies

https://t.me/learningnets

6-82 Key Points

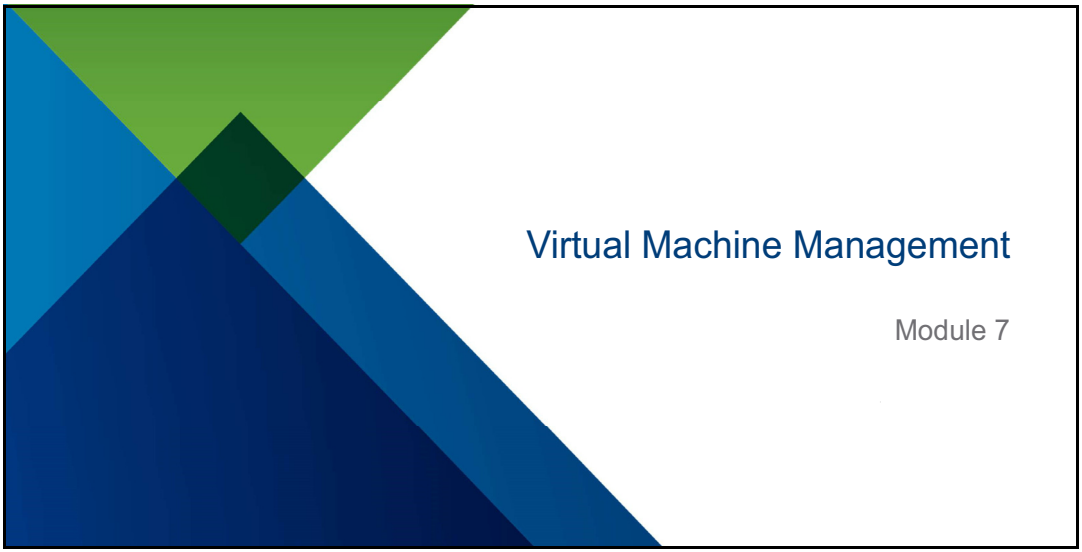
- You use datastores to hold virtual machine files.
- Shared storage is integral to vSphere features such as vSphere vMotion, vSphere HA, and vSphere DRS.
- vSAN clusters direct-attached server disks to create shared storage designed for virtual machines.
- vSAN enables the use of vSphere HA, vSphere vMotion, and vSphere DRS without requiring external shared storage.

Questions?

https://t.me/learningnets

MODULE 7

Virtual Machine Management



https://t.me/learningnets

7-2 You Are Here

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

https://t.me/learningnets

7-3 Importance

Because virtual machines are the basis of all IT operations in a virtual infrastructure, you should know how to properly deploy, configure, and manage them.

Failure to properly populate your virtual machine inventory might have unsatisfactory consequences.

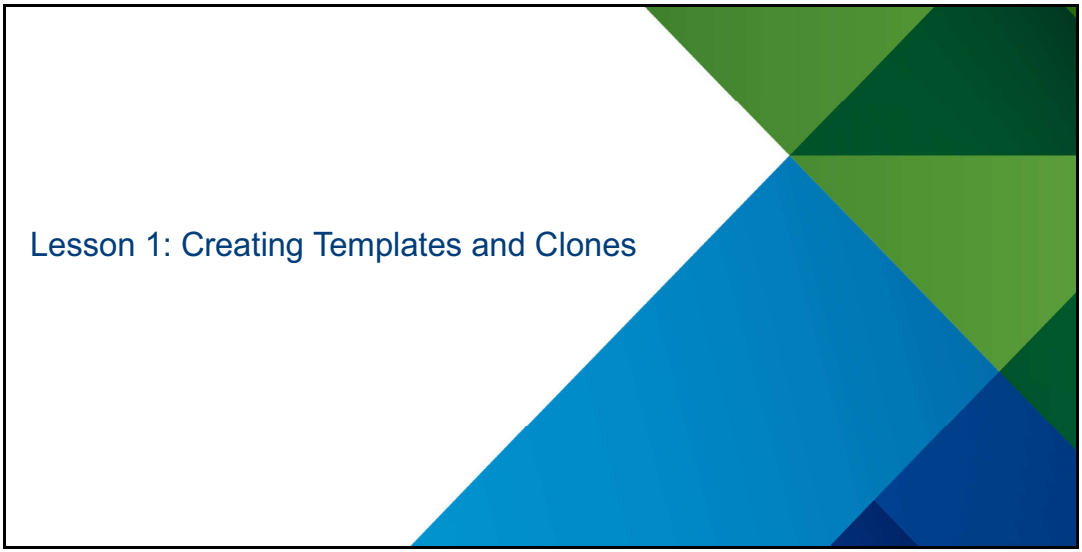
https://t.me/learningnets

7-4 Module Lessons

Lesson 1:	Creating Templates and Clones
Lesson 2:	Working with Content Libraries
Lesson 3:	Modifying Virtual Machines
Lesson 4:	Migrating Virtual Machines
Lesson 5:	Creating Virtual Machine Snapshots

https://t.me/learningnets

7-5 Lesson 1: Creating Templates and Clones



https://t.me/learningnets

7-6 Learner Objectives

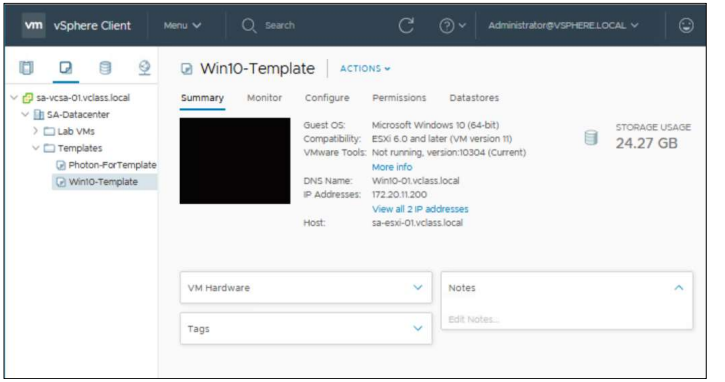
By the end of this lesson, you should be able to meet the following objectives:

- Create a template
- Deploy a virtual machine from a template
- Clone a virtual machine
- Enable guest operating system customization by vCenter Server
- Upgrade a virtual machine's hardware
- Create an instant clone of a virtual machine

https://t.me/learningnets

7-7 Using Templates

A template is a master copy of a virtual machine. A template is used to create and provision new virtual machines.



A template is an image that typically includes:

- A guest operating system
- A set of applications
- A specific virtual machine configuration that provides virtual counterparts to hardware components

Creating templates makes the provisioning of virtual machines much faster and less error-prone than provisioning physical machines and creating a new virtual machine by using the Create New Virtual machine wizard.

Templates coexist with virtual machines in the inventory. You can organize collections of virtual machines and templates into arbitrary folders and apply permissions to virtual machines and templates. Virtual machines can be changed into templates without the need to make a full copy of the virtual machine files and the creation of a new object.

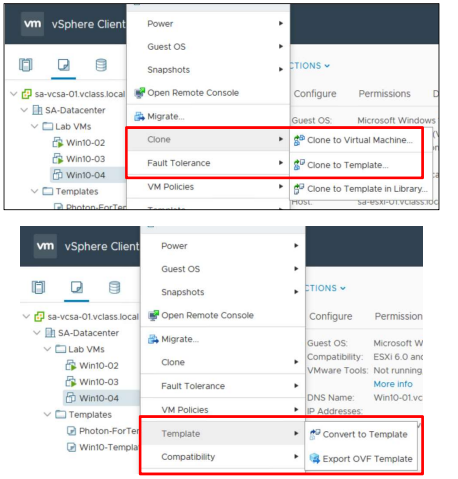
You can deploy a virtual machine from a template. The deployed virtual machine is added to the folder that the user selected when the template was created.

https://t.me/learningnets

7-8 Creating Templates

You can create templates in the following ways:

- Clone the virtual machine to a template: The virtual machine can be powered on or powered off.
- Convert the virtual machine to a template: The virtual machine must be powered off.
- Clone a template: You use this action to create a template that is based on an existing template.



Clone to Template offers you a choice of format in which to store the virtual machine's virtual disks:

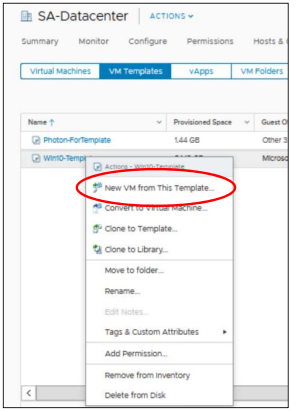
- Same format as source
- Thin-provisioned format
- Thick-provisioned lazy-zeroed format
- Thick-provisioned eager-zeroed format

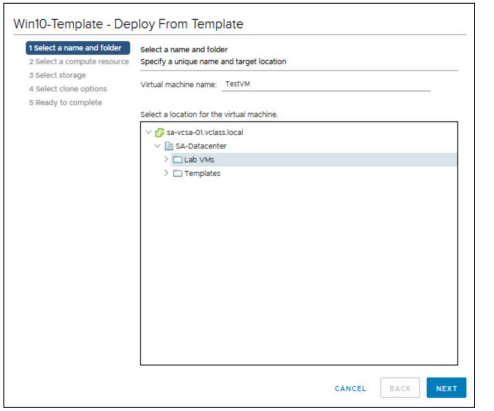
Convert to Template does not offer a choice of format and leaves the virtual machine's disk file intact.

https://t.me/learningnets

7-9 Deploying Virtual Machines from a Template

To deploy a virtual machine, you must provide information such as the virtual machine name, inventory location, host, datastore, and guest operating system customization data.



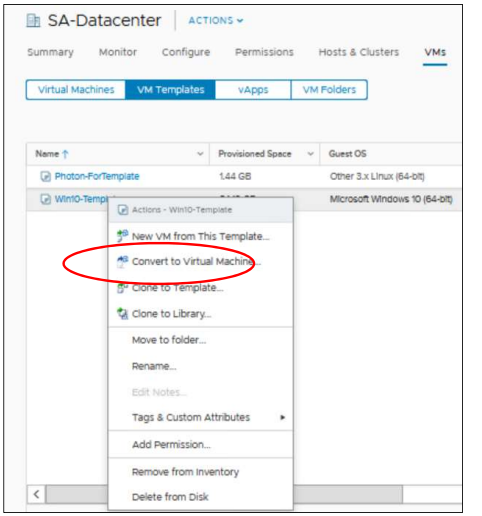


https://t.me/learningnets

7-10 Updating Templates

Update a template to include new patches, make system changes, and install new applications.

- Convert the template to a virtual machine.
- Place the virtual machine on an isolated network to prevent user access.
- Make appropriate changes to the virtual machine.
- Convert the virtual machine to a template.



To update your template to include new patches or software, you do not have to create a new template. Instead, first convert the template to a virtual machine. This conversion enables you to power on the virtual machine.

For added security, prevent users from accessing the virtual machine while you are updating it. Either disconnect the virtual machine from the network or place it on an isolated network.

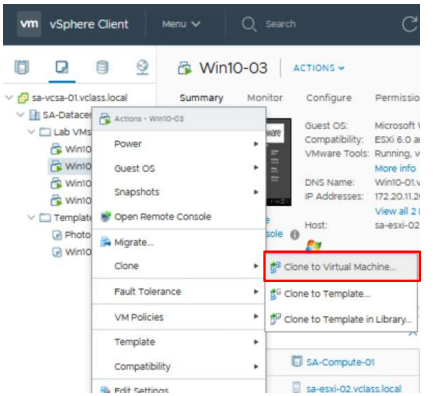
Log in to the virtual machine's guest operating system and apply the patch or install the software. When you have finished, power off the virtual machine and convert it to a template again.

https://t.me/learningnets

7-11 Cloning Virtual Machines

Cloning a virtual machine creates a virtual machine that is an exact copy of the original.

- Cloning is an alternative to deploying a virtual machine.
- The virtual machine being cloned can be powered on or powered off.



Cloning a virtual machine is an alternative to deploying a virtual machine from a template. Cloning a virtual machine creates a duplicate of the virtual machine with the same configuration and installed software as the original.

You can customize the guest operating system of the clone to change the virtual machine name, network settings, and other properties. Customizing the guest operating system prevents conflicts that might occur when a virtual machine and a clone with identical guest operating system settings are deployed simultaneously.

To clone a virtual machine, you must be connected to vCenter Server. You cannot clone virtual machines if you use VMware Host Client to manage a host directly.

The virtual machine that you clone can be powered on or powered off. When you clone a virtual machine that is powered on, services and applications are not automatically quiesced when the virtual machine is cloned.

https://t.me/learningnets

7-12 Customizing the Guest Operating System

Use the Customize guest OS option to make virtual machines, created from the same template or clone, unique.

Customizing a guest operating system enables you to change:

- Computer name
- Network settings
- License settings
- Windows Security Identifier

When cloning or deploying virtual machines from a template:

- You can create a specification to prepare the guest operating systems.
- Specifications are stored in the vCenter Server database.
- You can edit specifications in the Customization Specification Manager.
- Windows and Linux operating systems are supported.

When you clone a virtual machine or deploy it from a template, you can customize its guest operating system to change properties, including:

- Computer name
- Network settings
- License settings

Customizing guest operating systems can help prevent conflicts that occur when virtual machines with identical settings are deployed, such as conflicts because of duplicate computer names or IP addresses.

You can specify the customization settings by using the Guest Customization wizard during cloning or deployment. Or, you can create customization specifications, which are customization settings stored in the VMware vCenter Server® database. During cloning or deployment, you can select a customization specification to apply to the new virtual machine. Use the Customization Specification Manager to manage customization specifications that you create with the Guest Customization wizard.

The guest operating system that is being customized must have VMware Tools installed. The guest operating system must also be installed on a disk attached to SCSI node 0:0 in the virtual machine configuration.

To enable guest operating system customization, vCenter Server must first be configured for this task. To customize virtual machines running operating systems that predate Windows 2008 and

https://t.me/learningnets

Windows Vista, you must install Microsoft Sysprep tools on the vCenter Server system. Sysprep tools are built into the Windows 2008, Windows Vista, and later Windows operating system.

Customization of Linux guest operating systems requires that Perl be installed in the Linux guest operating system.

Guest operating system customization is supported on multiple Linux distributions. To verify customization support for Linux distributions and compatible ESXi hosts, see *VMware Compatibility Guide* at http://www.vmware.com/resources/compatibility.

For more about guest operating system customization, see *vSphere Virtual Machine Administration Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

7-13 About Instant Clones

Instant clones are generated rapidly from a source virtual machine, and they resume from the exact same state as the source.

The processor state, virtual device state, memory state, and disk state of the destination virtual machine are identical to those of the source virtual machine.

Snapshot-based disk sharing similar to linked clones is used to to provide storage efficiency.

Use cases for instant clones:

- 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: Able to replicate VMs from staging to production, and the converse, with the identical running state.

In vSphere 6.7, you can create instant clones of virtual machines only through API calls.

Through instant cloning, users can create powered-on virtual machines from the running state of another powered-on virtual machine. The source virtual machine does not lose its state due to the cloning process. User can move to just-in-time provisioning, given the speed and state-persisting nature of this operation.

Instant cloning is very convenient for large-scale application deployments because it ensures memory efficiency and allows the creation of many virtual machines on a single host.

To avoid network conflicts, instant cloning allows customization of the virtual hardware of the destination virtual machine during an cloning operation. For example, you can customize the MAC addresses of the virtual NICs or the serial and parallel port configurations of the destination virtual machine.

During an instant clone operation, the source virtual machine is stunned for a short period of time, less than 1 second. While the source virtual machine is stunned, a new writable delta disk is generated for each virtual disk, and a checkpoint is taken and transferred to the destination virtual machine. The destination virtual machine then powers on by using the source's checkpoint. After the destination virtual machine is fully powered on, the source virtual machine resumes running.

Instant clone virtual machines are fully independent vCenter Server inventory objects. You can manage instant clone virtual machines like regular virtual machines without any restrictions.

https://t.me/learningnets

7-14 Removing a Virtual Machine

You can remove a virtual machine in the following ways:

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

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

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

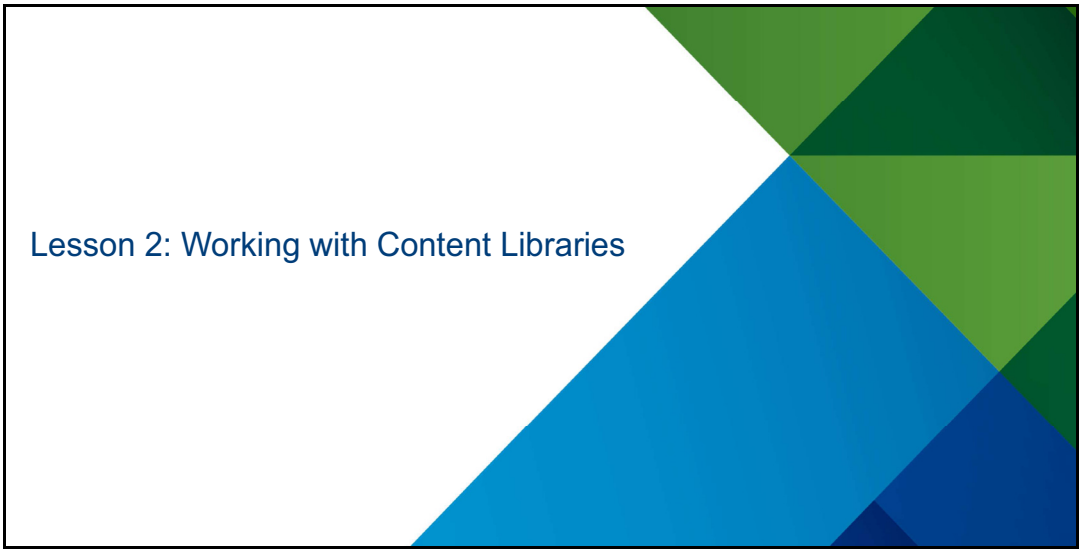
https://t.me/learningnets

7-15 Review of Learner Objectives

You should be able to meet the following objectives:

- Create a template
- Deploy a virtual machine from a template
- Clone a virtual machine
- Enable guest operating system customization by vCenter Server
- Upgrade a virtual machine's hardware
- Create an instant clone of a virtual machine

https://t.me/learningnets



https://t.me/learningnets

7-17 Learner Objectives

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

- Identify benefits of a content library
- Describe types of content libraries
- Deploy a virtual machine from the content library

https://t.me/learningnets

7-18 About Content Libraries

Content libraries are repositories of OVF templates and other files that can be shared and synchronized across vCenter Server systems.



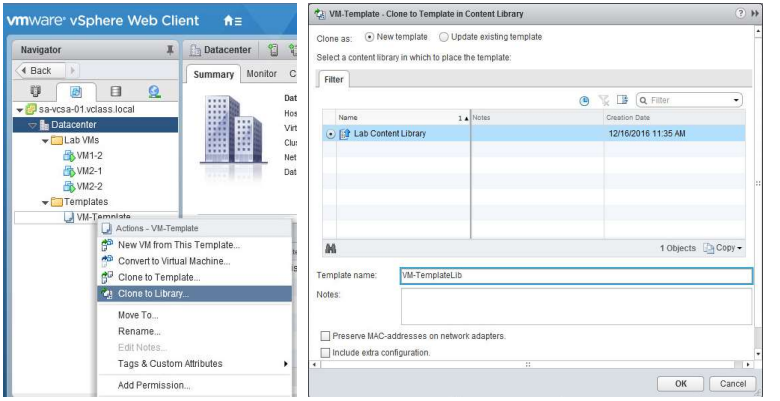
The content library helps manage templates and other file types used in vCenter Server instances globally. Organizations might have multiple vCenter Server instances in data centers around the globe. On these vCenter Server instances, organizations might have a collection of templates, ISO images, and so on. The challenge is that all of these items are independent of one another, with different versions of these files and templates on various vCenter Server instances.

The content library is the solution. IT can store OVF templates, ISO images, or any other file types in a central location. The templates, images, and files can be published, enabling other content libraries to subscribe to and download content. The content library keeps content up to date by periodically synchronizing with the publisher, ensuring that the latest version is available.

https://t.me/learningnets

7-19 Adding Templates to a Content Library

Library items include VM templates, VMware vSphere® vApp™ templates, or other VMware objects that can be contained in a content library.



Library items are VM templates, VMware vSphere® vApp™ templates, or other VMware objects that can be contained in a content library. VMs and vApps have several files, such as log files, disk files, memory files, and snapshot files that are part of a single library item. You can create library items in a specific local library or remove items from a local library. You can also upload files to an item in a local library so that the libraries subscribed to it can download the files to their NFS or SMB server, or datastore.

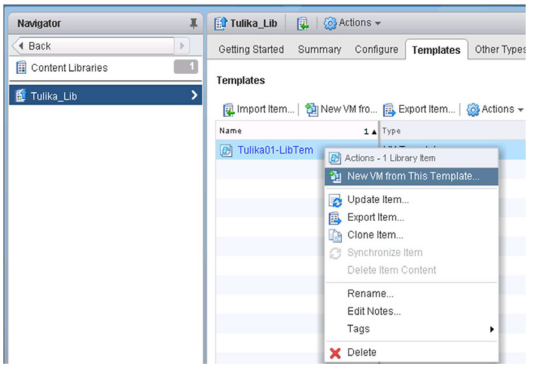
https://t.me/learningnets

7-20 Deploying VMs from Templates in a Content Library

The templates in the library can be used to deploy virtual machines and vApps.

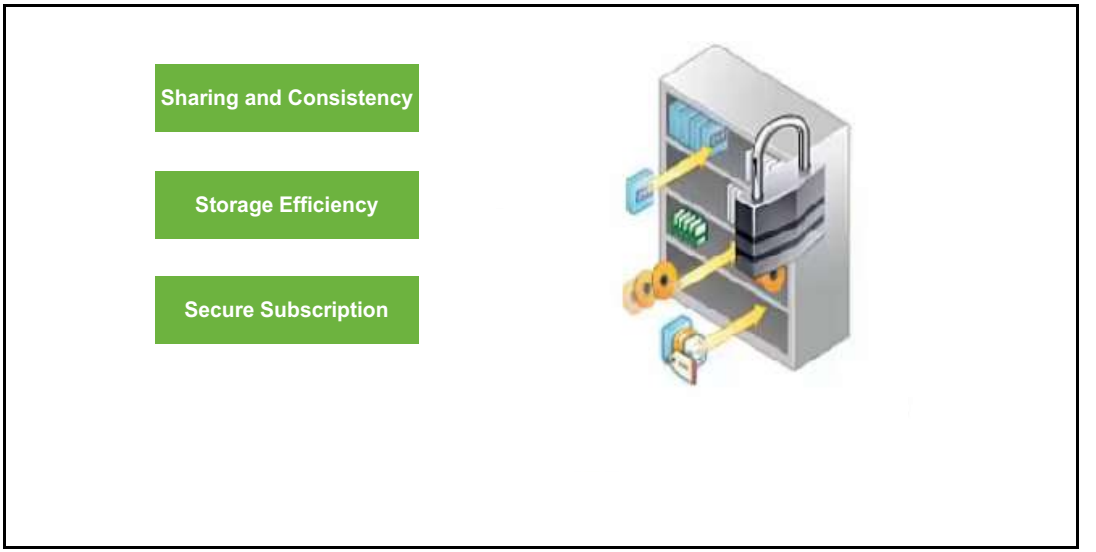
Each VM template, vApp template, or other type of file in a library is a library item.

You can mount an ISO file directly from a content library.



https://t.me/learningnets

7-21 Benefits of Content Libraries



The Content Library API provides services that allow you to create and manage content libraries programmatically. You can create a local library and publish it for the entire virtual environment.

Using content libraries, administrators can perform the following functions:

- Store, version, and share content: Storage and consistency are key reasons to install and use a content library. Sharing the content and ensuring that the content is kept up to date is a major task. For example, assume that you have a main vCenter Server instance and that you create a central content library to store the master copies of OVF templates, ISO images, and other file types. You can publish this content library to allow other libraries, which can be located anywhere in the world, to subscribe and download an exact copy of the data. When an OVF template is added, modified, or deleted from the published catalog, the subscriber synchronizes with the publisher and the libraries are updated with the latest content.
- Perform distributed file management.
- Publish content to public or subscribed content.
- Synchronize content across sites and vCenter Server systems.
- Back up content libraries using vSphere datastores, vCenter directory, or NFS.
- Mount an ISO file directly from a content library.

Subscribers might not always want to download and store all the content from the published library. They might not have the need or the space to store all of the content available to them. Content libraries can address such issues and implement storage efficiencies.

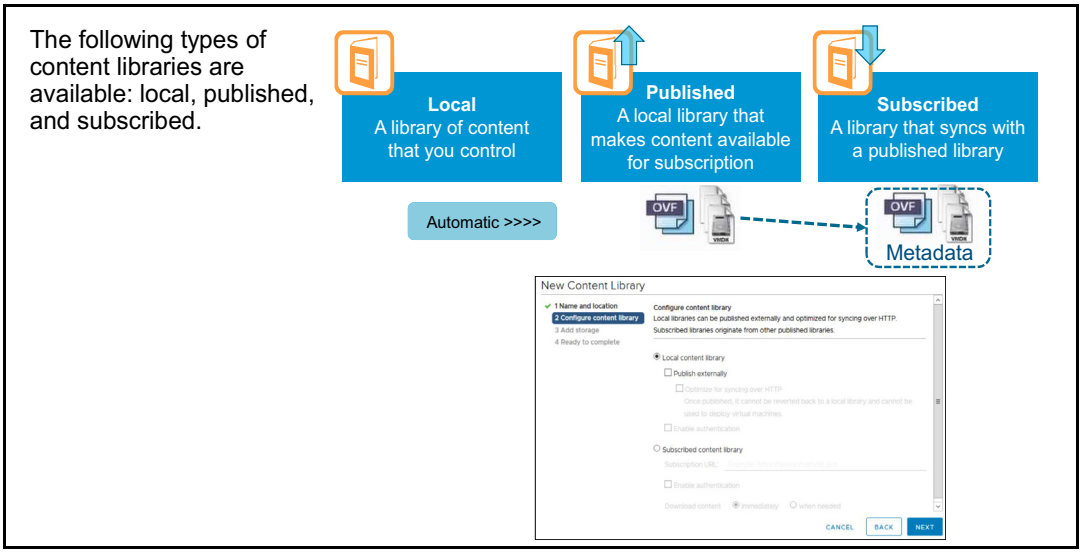
https://t.me/learningnets

When a publisher updates its content, subscribers are notified of the changes. A subscriber can choose to receive either the full data or only the metadata. When the subscriber chooses to receive the full data, the subscriber can choose to download either the entire library or individual items.

Because content libraries are made globally available for subscription, concerns about security are expected. As a security measure, content libraries can be password-protected during publication.

https://t.me/learningnets

7-22 Types of Content Libraries



You can create local libraries and subscribed libraries:

- You can create a local library as the source for content you want to save or share. Create the local library on a single vCenter Server instance. You can add items to a local library or remove them.
- You can publish a local library and this content library service endpoint can be accessed by other vCenter Server instances in your virtual environment. When you publish a library, you can configure the authentication method, which a subscribed library must use to authenticate to it.
- You can create a subscribed library and populate its content by synchronizing to a 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 located 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. After synchronization, both libraries contain the same items or the subscribed library contains the metadata for the items.

https://t.me/learningnets

7-23 Labs

Lab 8: Using Templates and Clones

Lab 9: Using Content Libraries

https://t.me/learningnets

7-24 Lab 8: Using Templates and Clones

Deploy a new virtual machine from a template and clone a virtual machine

1. Create a Virtual Machine Template
2. Create Customization Specifications
3. Deploy a Virtual Machine from a Template

https://t.me/learningnets

7-25 Lab 9: 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 Virtual Machine from a VM Template in the Content Library
4. Delete a Content Library

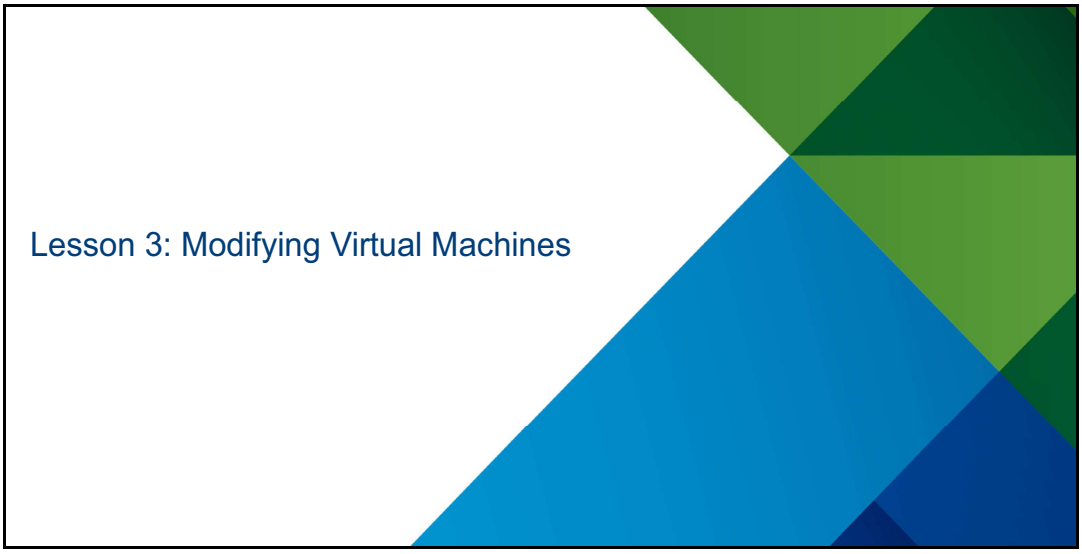
https://t.me/learningnets

7-26 Review of Learner Objectives

You should be able to meet the following objectives:

- Identify benefits of a content library
- Describe types of content libraries
- Deploy a virtual machine from the content library

https://t.me/learningnets



https://t.me/learningnets

7-28 Learner Objectives

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

- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk
- Add a raw device mapping disk to a virtual machine

You might have to modify a virtual machine's configuration, for example, to add a network adapter or a virtual disk. All virtual machine changes can be made while the virtual machine is powered off. Some hardware changes can be made to the virtual machine while it is powered on.

Besides adding virtual hardware, you can also remove virtual hardware and set various virtual machine options.

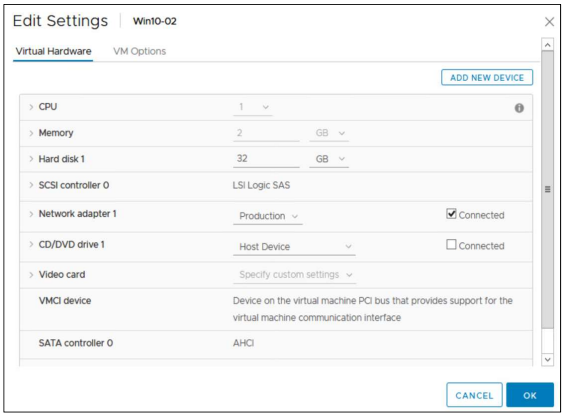
All virtual machine configuration is done in the virtual machine Edit Settings dialog box.

https://t.me/learningnets

7-29 Modifying Virtual Machine Settings

You can modify a virtual machine's configuration in its Edit Settings dialog box:

- Add virtual hardware:
  - Some hardware can be added while the virtual machine is powered on.
- Remove virtual hardware:
  - Some hardware can be removed only when the virtual machine is powered off
- Set virtual machine options.
- Control a virtual machine's CPU and memory resources.



Adding devices to a physical server or removing devices from a physical server requires you to physically interact with the server in the data center. When you use virtual machines, resources can be added dynamically without a disruption in service. You must shut down a virtual machine to remove hardware, but you can reconfigure the virtual machine without entering the data center.

CPU and memory can be added while the virtual machine is powered on. These features are called the CPU hot-plug and memory hot-add options, which are supported only on guest operating systems that support this feature. These options are disabled by default. To use these hot-plug features, the following requirements must be satisfied:

- You must install VMware Tools.
- The virtual machine must use hardware version 7 or later.
- The guest operating system in the virtual machine must support CPU and memory hot-plug features.
- The hot-plug options must be enabled in the **CPU** or **Memory** options on the **Virtual Hardware** tab.

See *vSphere Virtual Machine Administration Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

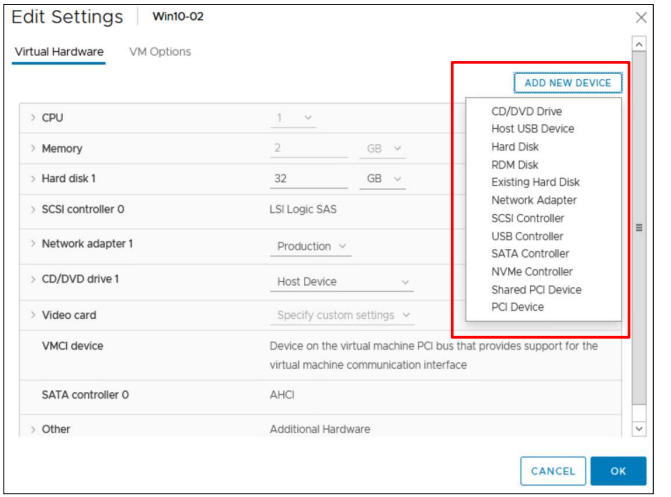
https://t.me/learningnets

7-30 Hot-Pluggable Devices

The hot-plug option enables you to add resources to a running virtual machine.

Examples of hot-pluggable devices: USB controllers, Ethernet adapters, and hard disk devices.

With supported guest operating systems, you can also add CPU and memory while the virtual machine is powered on.



An RDM is a file that has a `.vmdk` extension, but the file contains only disk information describing the mapping to the LUN on the ESXi host. The data is stored on the LUN.

An RDM supports the following compatibility modes:

- Physical compatibility (pass-through) mode: Allows the guest operating system to access the hardware directly. Physical compatibility is useful if you are using SAN-aware applications in the virtual machine. But a LUN configured for physical compatibility cannot be cloned or made into a template. When a snapshot of the virtual machine with a physical RDM is taken, a snapshot of the physical RDM is not taken. LUNs configured for pass-through mode can be a maximum of 64 TB in size. For physical compatibility mode RDM, the file name is `-rdmp.vmdk`. If a virtual machine with a physical compatibility mode RDM is powered on, you can use vSphere Storage vMotion to migrate it to a new datastore, but the contents of the LUN remain intact. The `-rdmp.vmdk` file migrates to the new datastore. The same is true if the powered-on virtual machine is migrated using shared-nothing vSphere vMotion. If a virtual machine with a physical compatibility mode RDM is powered off, its files can be migrated to a new datastore. The user can choose whether the contents of the LUN move to a `-flat.vmdk` file on the new datastore (converting it to a virtual disk) or remain in the LUN (and only the `-rdmp.vmdk` file migrates).
- Virtual compatibility mode: Allows the virtual machine to use VMware snapshots and other advanced features. Virtual compatibility enables the LUN to behave as if it were a virtual disk. When you clone the disk, make a template out of it, or migrate it (if the migration involves copying the disk), the contents of the LUN can optionally be copied to a virtual disk (`-flat.vmdk`) file. Regardless of the power state of the virtual machine, if it has a virtual

https://t.me/learningnets

compatibility mode RDM attached, the user can choose whether the contents of the LUN move to a `-flat.vmdk` file on the destination datastore or remain as a raw LUN (and only the `-rdm.vmdk` file moves to the new datastore).

https://t.me/learningnets

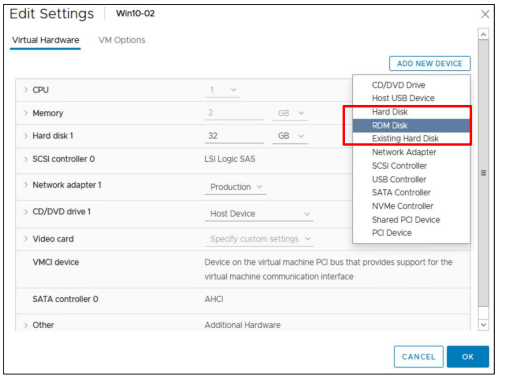
7-31 Creating Raw Device Mapping Files

An RDM file (`-rdm.vmdk`) enables a virtual machine to gain direct access to a physical LUN.

Encapsulating disk information in the RDM file enables the VMkernel to lock the LUN so that only one virtual machine can write to the LUN.

You must define the following items when creating an RDM file:

- Target LUN: LUN that the RDM file will map to
- Mapped datastore: Stores the RDM file with the virtual machine or on a different datastore
- Compatibility mode:
  - Physical
  - Virtual
- Virtual device node



You can increase the size of a virtual disk that belongs to a powered-on virtual machine if the virtual machine lacks snapshots.

After you increase the size of a virtual disk, you might need to increase the size of the file system on this disk. Use the appropriate tool in the guest operating system to enable the file system to use the newly allocated disk space.

You can use RDMs in virtual compatibility or physical compatibility modes.

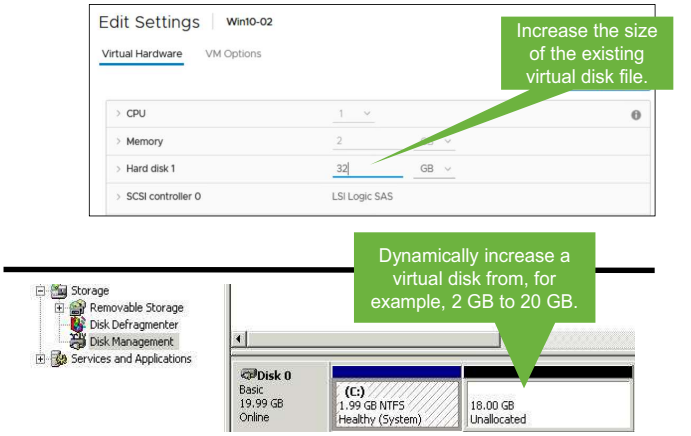
- Virtual mode specifies full virtualization of the mapped device. In virtual mode, the VMkernel sends only READ and WRITE to the mapped device. The mapped device appears to the guest operating system exactly the same as a virtual disk file in a VMFS volume. The real hardware characteristics are hidden. If you are using a raw disk in virtual mode, you can realize the benefits of VMFS such as advanced file locking for data protection and snapshots for streamlining development processes. Virtual mode is also more portable across storage hardware than physical mode, presenting the same behavior as a virtual disk file.
- Physical mode specifies minimal SCSI virtualization of the mapped device, allowing the greatest flexibility for SAN management software.In physical mode, the VMkernel passes all SCSI commands to the device, with one exception: the REPORT LUNs command is virtualized so that the VMkernel can isolate the LUN to the owning virtual machine. Otherwise, all physical characteristics of the underlying hardware are exposed. Physical mode is useful to run SAN management agents or other SCSI target-based software in the virtual machine.

https://t.me/learningnets

7-32 Dynamically Increasing Virtual Disk Size

You can increase the size of a virtual disk that belongs to a powered-on virtual machine:

- It must not have snapshots attached.
- It might require system tools to make the new space usable.



When you inflate a thin-provisioned disk, the inflated virtual disk occupies the entire datastore space originally provisioned to it. Inflating a thin-provisioned disk converts a thin disk to a virtual disk in thick-provision format.

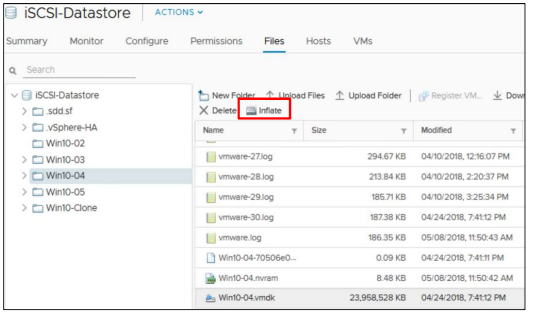
https://t.me/learningnets

7-33 Inflating Thin-Provisioned Disks

Thin-provisioned virtual disks can be converted to a thick, eager-zeroed format.

To inflate a thin-provisioned disk:

- The virtual machine must be powered off.
- Right-click the virtual machine's .vmdk file and select Inflate.



Or you can use vSphere Storage vMotion and select a thick-provisioned disk as the destination.

You can use the **VM Options** tab to modify things like the display name used for the virtual machine and the type of guest operating system installed. The location and name of the configuration file (`.vmx` file) and the location of the virtual machine's directory are also shown.

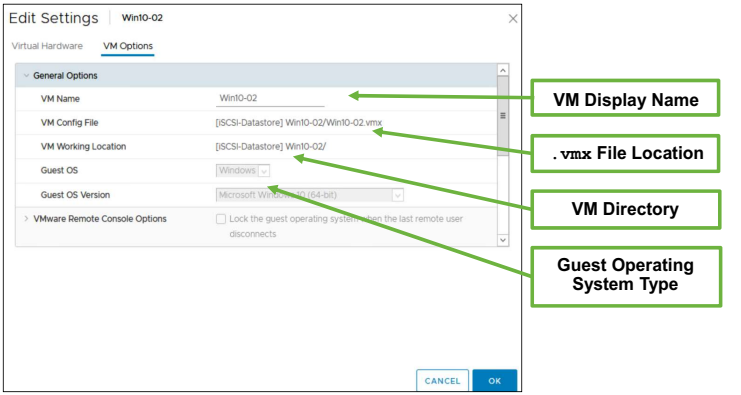
You can select the text for the configuration file and the working location to copy and paste them into a document. However, only the display name and the guest operating system type can be modified.

Changing the display name does not change the names of all of the virtual machine files or the directory that the virtual machine is stored in. When a virtual machine is created, the filenames and the directory name associated with the virtual machine are based on its display name. But changing the display name later does not modify the filename and the directory name.

https://t.me/learningnets

7-34 Virtual Machine Options

On the **VM Options** tab, you can set or change virtual machine options to run VMware Tools scripts, control user access to the remote console, configure startup behavior, and more.



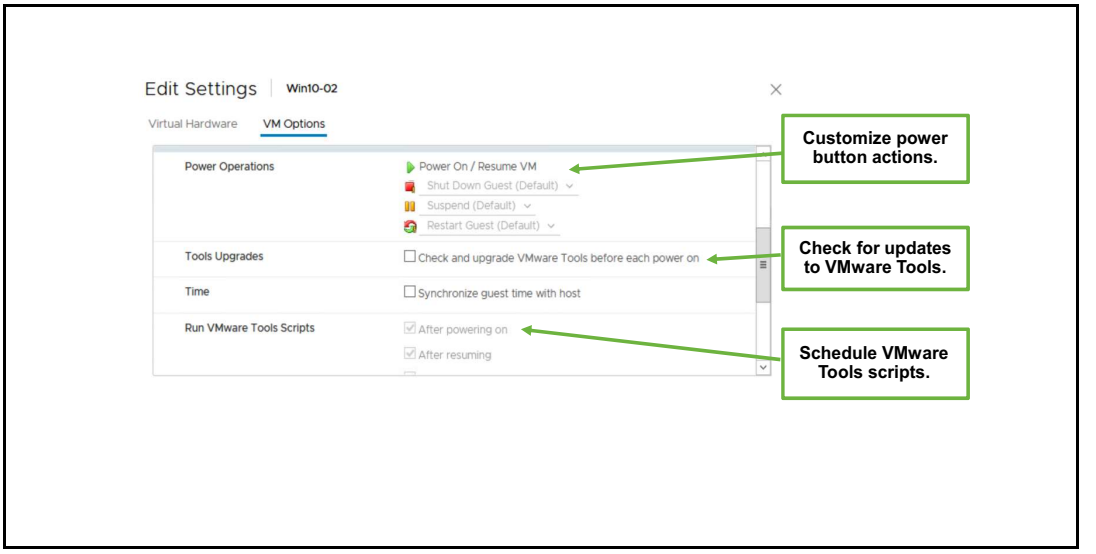
The VMware Tools panel controls how VMware Tools in the virtual machine responds to certain external events. You can use these controls to customize the power buttons on the virtual machine. The virtual machine must be powered off to change these settings.

The **Check and upgrade VMware Tools before each power on** check box can be configured to check whether a newer version of VMware Tools exists. If a newer version exists, VMware Tools is upgraded when the virtual machine is power cycled. The **Synchronize guest time with host** check box enables the guest operating system's clock to synchronize with the host.

For information about time keeping best practices for the guest operating systems that you are using, see VMware knowledge base articles 1318 at http://kb.vmware.com/kb/1318 and 1006427 at http://kb.vmware.com/kb/1006427.

https://t.me/learningnets

7-35 VMware Tools Options

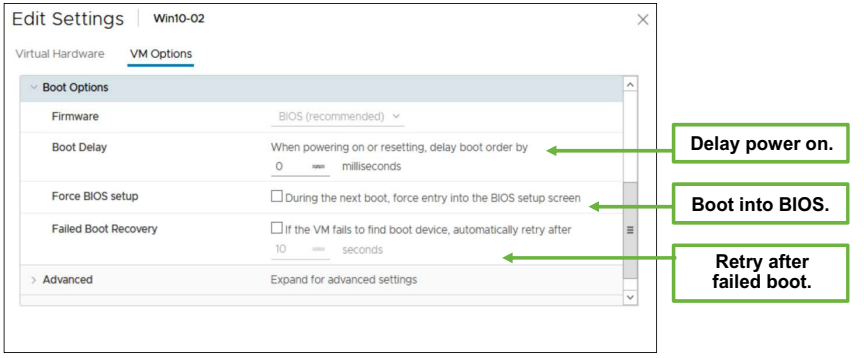


Occasionally, you might need to set the boot options:

- When you build a virtual machine and select a guest operating system, **BIOS** or **EFI** is selected 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 necessary. However, you must do this before installing the guest OS.
- The Boot Delay panel enables you to set a delay between the time when a virtual machine is turned on and the guest operating system starts to boot. A delayed boot can help stagger virtual machine startups when several virtual machines are being powered on.
- You can use the Force BIOS Setup panel to change the BIOS settings, like when forcing a virtual machine to boot from a CD/DVD. The next time the virtual machine powers on, it goes straight into the BIOS. A forced entry into the BIOS is much easier than powering on the virtual machine, opening a console, and quickly trying to press the F2 key to go into the BIOS.
- You can use the Failed Boot Recovery panel to have the virtual machine retry booting after 10 seconds (the default) if the virtual machine fails to find a boot device.

https://t.me/learningnets

7-36 Boot Options



https://t.me/learningnets

7-37 Lab 10: Modifying Virtual Machines

Modify a virtual machine's hardware and add a raw LUN to a virtual machine

1. Clone a Powered-On Virtual Machine
2. Increase the Size of a VMDK File
3. Adjust Memory Allocation on a Virtual Machine
4. Rename a Virtual Machine in the vCenter Server Appliance Inventory
5. Add and Remove a Raw LUN on a Virtual Machine

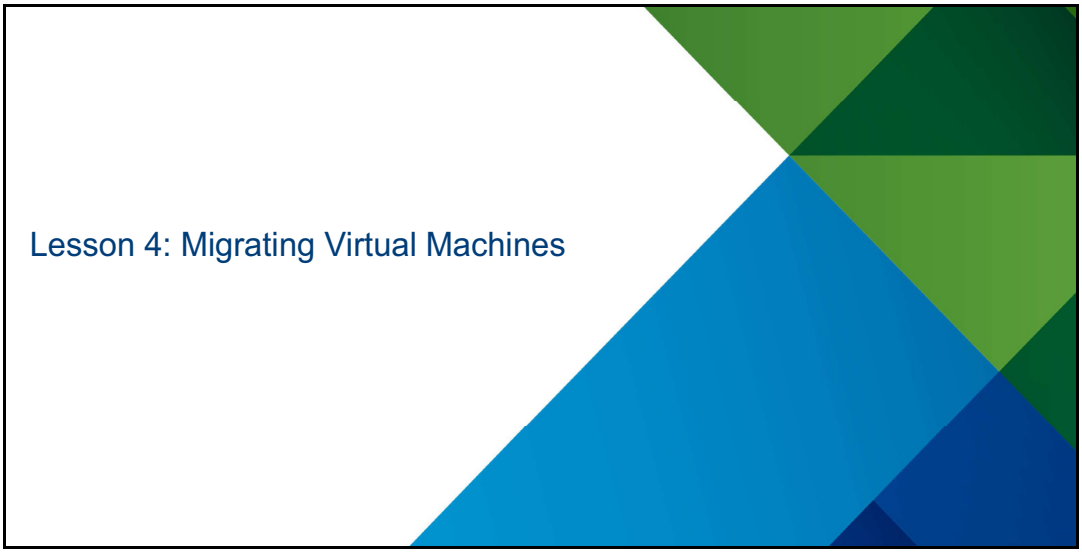
https://t.me/learningnets

7-38 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk
- Add a raw device mapping disk to a virtual machine

https://t.me/learningnets



https://t.me/learningnets

7-40 Learner Objectives

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

- Verify vSphere vMotion requirements, including CPU constraints and guidelines
- Perform a vSphere vMotion migration
- Perform a vSphere Storage vMotion migration
- Perform a shared-nothing vSphere vMotion migration
- Explain how cluster Enhanced vMotion Compatibility and per-VM EVC works with vSphere vMotion

You can move virtual machines from one host or storage location to another location using hot or cold migration. For example, with vSphere vMotion you can move powered-on virtual machines away from a host to perform maintenance, to balance loads, to colocate virtual machines that communicate with each other, to move virtual machines apart to minimize fault domain, to migrate to new server hardware, and so on.

You can perform several types of migration according to the virtual machine resource type:

- Change compute resource only: Moving a virtual machine but not its storage to another compute resource, such as a host, cluster, resource pool, or vApp.
- Change storage only: Moving a virtual machine's storage, including virtual disks, configuration files, or a combination of these, to a new datastore available to the same host.
- Change both compute resource and storage: Moving a virtual machine to another host and moving its disk or virtual machine folder to another datastore. When you move a virtual machine network between distributed switches, the network configuration and policies that are associated with the network adapters of the virtual machine are transferred to the target switch.

https://t.me/learningnets

7-41 Migrating Virtual Machines

Migration means moving a virtual machine from one host, datastore, or vCenter Server system to another host, datastore, or vCenter Server system.

Types of migrations:

- Cold: Migrate a powered-off virtual machine to a new host or datastore.
- Suspended: Migrate a suspended virtual machine to a new host or datastore.
- vSphere vMotion: Migrate a powered-on virtual machine to a new host.
- vSphere Storage vMotion: Migrate a powered-on virtual machine's files to a new datastore.
- Shared-nothing vSphere vMotion: Migrate a powered-on virtual machine to a new host and a new datastore simultaneously.

A deciding factor behind using a particular migration technique is the purpose of performing the migration. For example, you might need to stop a host for maintenance but keep the virtual machines running. Use vSphere vMotion to migrate the virtual machines instead of performing a cold or suspended virtual machine migration. If you must move a virtual machine's files to another datastore to better balance the disk load or transition to another storage array, use vSphere Storage vMotion.

Some migration techniques, such as vSphere vMotion migration, have special hardware requirements that must be met to function properly. Other techniques, such as a cold migration or a suspended virtual machine migration, do not have special hardware requirements to function properly.

https://t.me/learningnets

7-42 Comparing Migration Types

Migration Type	Virtual Machine Power State	Change Host or Datastore?	Across vCenter Server Instances?	Shared Storage Required?	CPU Compatibility
Cold	Off	Host or datastore or both	Yes	No	Different CPU families allowed
Suspended	Suspended	Host or datastore or both	Yes	No	Must meet CPU compatibility requirements
vSphere vMotion	On	Host	Yes	Yes	Must meet CPU compatibility requirements
vSphere Storage vMotion	On	Datastore	No	No	N/A
Shared-nothing vSphere vMotion	On	Both	Yes	No	Must meet CPU compatibility requirements

vSphere vMotion migrates running virtual machines from one ESXi host to another ESXi host with no disruption or downtime. vSphere vMotion enables vSphere DRS to migrate running virtual machines from one host to another to balance the load.

With vSphere vMotion, the entire state of the virtual machine is moved from one host to another while the data storage remains in the same datastore.

The state information includes the current memory content and all the information that defines and identifies the virtual machine. The memory content includes transaction data and whatever bits of the operating system and applications are in memory. The definition and identification information stored in the state includes all the data that maps to the virtual machine hardware elements, including:

- BIOS
- Devices
- CPU
- MAC addresses for the Ethernet cards

https://t.me/learningnets

7-43 vSphere vMotion Migration

A vSphere vMotion migration moves a powered-on virtual machine from one host to another.

vSphere vMotion provides the following capabilities:

- Improves overall hardware use
- Enables continuous virtual machine operation while accommodating scheduled hardware downtime
- Allows vSphere DRS to balance virtual machines across hosts

A vSphere vMotion migration consists of the following steps:

1. Shadow VM is created on the destination host.
2. The virtual machine's memory state is copied over the vSphere vMotion network from the source host to the target host through the vSphere vMotion network. Users continue to access the virtual machine and, potentially, update pages in memory. A list of modified pages in memory is kept in a memory bitmap on the source host.
3. After completing the first pass of memory state copy, another pass of memory copy is performed to copy any pages that changed during the last iteration. Continue this iterative memory copying until no changed pages remain.
4. After most of the virtual machine's memory is copied from the source host to the target host, the virtual machine is quiesced. No additional activity occurs on the virtual machine. In the quiesce period, vSphere vMotion transfers the virtual machine device state and memory bitmap to the destination host.

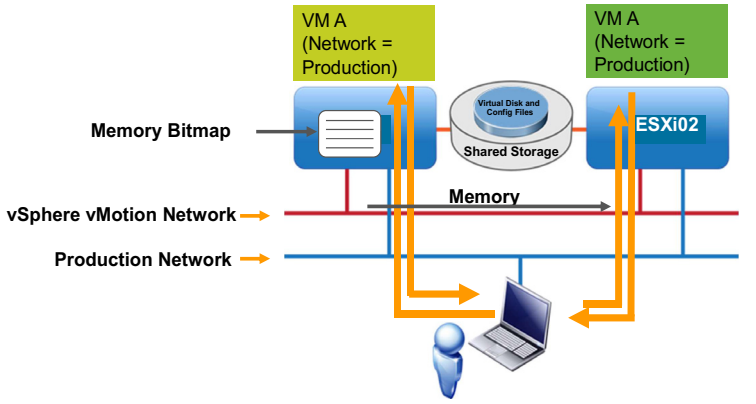
https://t.me/learningnets

5. Immediately after the virtual machine is quiesced on the source host, the virtual machine is initialized and starts running on the target host. A Gratuitous Address Resolution Protocol (GARP) request notifies the subnet that virtual machine A's MAC address is now on a new switch port.
6. Users access the virtual machine on the target host instead of the source host.
7. The memory pages that the virtual machine was using on the source host are marked as free.

https://t.me/learningnets

7-44 vSphere vMotion Migration Workflow

The source host (ESXi01) and the destination host (ESXi02) have access to the shared datastore that holds the virtual machine's files.



The vSphere vMotion migration produces an error in certain conditions. When an error is encountered, the migration does not proceed until you fix the error.

By default, migrating a virtual machine with vSphere vMotion produces an error if the virtual machine has an active connection to an internal virtual switch (that is, a switch that is not directly connected to a physical NIC). The intent of this restriction was to ensure server availability through a live migration, and to avoid problems where the target ESXi host might be misconfigured without an active connection to the physical network. However, this restriction does not take into account when a bridged-mode virtual machine is being used to filter traffic between the workload and the physical network. You may enable vSphere vMotion when using network security or service virtual machines in bridged mode. For information, see VMware knowledge base article 1006701 at http://kb.vmware.com/kb/1006701.

You cannot use vSphere vMotion to migrate a virtual machine that uses a virtual device that is backed by a device on the ESXi host. Disconnect these devices before migrating the virtual machine. vSphere vMotion also produces warnings in certain conditions, for example, when a virtual machine is configured to access a local CD/DVD drive or floppy image but is not connected to it. The vSphere vMotion migration still proceeds even if warnings have not been addressed.

For the complete list of vSphere vMotion migration requirements, see *vCenter Server and Host Management Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

7-45 vSphere vMotion Migration Requirements

A virtual machine must meet the following requirements:

- It should not have an active connection to an internal virtual switch, because migrating such a virtual machine produces an error.
- It must not have CPU affinity configured.
- It must not have a connection to a virtual device, such as a CD/DVD or floppy drive, with a local image mounted.
- If its swap file is not accessible to the destination host, vSphere vMotion must be able to create a swap file that is accessible to the destination host before migration can begin.
- If it uses an RDM disk, the RDM file and the LUN to which it maps must be accessible by the destination host.

Each operation, such as a migration with vSphere vMotion or cloning a virtual machine, is assigned a resource cost. Each host, datastore, or network resource, has a maximum cost that it can support at any one time. Any new migration or provisioning operation that causes a resource to exceed its maximum cost does not proceed immediately, but is queued until other operations complete and release resources. Each of the network, datastore, and host limits must be satisfied for the operation to proceed.

For the complete list of limits on simultaneous migrations, see *vCenter Server and Host Management Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

The source and destination host must meet certain requirements for a vSphere vMotion migration to be successful:

- SAN visibility of virtual disks
- Gigabit Ethernet (or greater) interconnection
- Consistent network configuration, both physical and virtual
- Source and destination server CPUs from the same compatibility group

For better performance, you can configure multiple NIC vMotion. To have multi-NIC vMotion, you can dedicate two or more port groups to the vSphere vMotion traffic, respectively every port group must have a vMotion VMkernel adapter associated with it. Then you can connect one or more physical NICs to every port group. In this way, multiple physical NICs are used for vMotion, which

https://t.me/learningnets

results in greater bandwidth. For the detailed steps for configuring multiple NIC vMotion, see VMware knowledge base article 2007467 at https://kb.vmware.com/kb/2007467.

https://t.me/learningnets

7-46 Host Requirements for vSphere vMotion Migration

Source and destination hosts must have the following characteristics:

- Accessibility to all of the virtual machine's storage (Fibre Channel, iSCSI, or NAS):
  - 128 concurrent vSphere vMotion migrations are possible per VMFS or NFS datastore.
- At least a 1 Gigabit Ethernet (1GigE) network:
  - Each active vSphere vMotion process requires a minimum throughput of 250Mbit/sec on the vSphere vMotion network.
  - Concurrent vSphere vMotion migrations are limited to four on a 1 Gbps network.
  - Concurrent vSphere vMotion 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 EVC or compatibility masks.

CPU compatibility between the source host and the target host is a vSphere vMotion requirement that must be met. For example, if hyperthreading is enabled on the source host and disabled on the destination host, the vSphere vMotion migration continues because the VMkernel handles this difference in characteristics. But if the source host processor supports SSE4.1 instructions and the destination host processor does not support them, the hosts are considered incompatible and the vSphere vMotion migration fails. SSE4.1 instructions are application-level instructions that bypass the virtualization layer and might cause application instability if mismatched after a migration with vSphere vMotion.

https://t.me/learningnets

7-47 CPU Constraints on vSphere vMotion Migration

CPU Characteristics	Exact Match Required	Reason
Clock speeds, cache sizes, hyperthreading, and number of cores	N/A	These are virtualized away by the VMkernel.
Manufacturer (Intel or AMD) family and generation (Opteron4, Intel Westmere)	Applicable	Instruction sets contain many small differences.
Presence or absence of SSE3, SSSE3, or SSE4.1 instructions	Applicable	Multimedia instructions are usable directly by applications.
Virtualization hardware assist	For 32-bit VMs: N/A	This is virtualized away by the VMkernel.
	For 64-bit VMs on Intel: Applicable	Intel 64-bit with VMware implementation uses Intel VT.

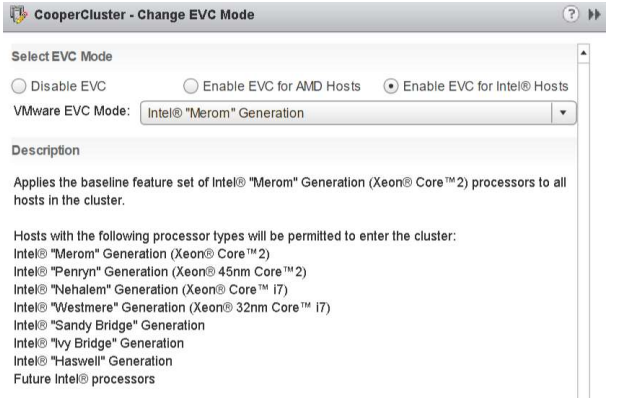
(Enhanced vMotion Compatibility (EVC) ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. Presenting the same CPU feature set prevents migrations with vSphere vMotion from failing because of incompatible CPUs.

For information about EVC processor support, see VMware knowledge base article 1003212 at http://kb.vmware.com/kb/1003212.

https://t.me/learningnets

7-48 Other Cluster Settings: EVC for vSphere vMotion Migration

EVC is a cluster feature that prevents vSphere vMotion migrations from failing because of incompatible CPUs.



EVC facilitates safe vSphere vMotion migration across a range of CPU generations. With EVC, you can use vSphere vMotion to migrate virtual machines among CPUs that would otherwise be considered incompatible.

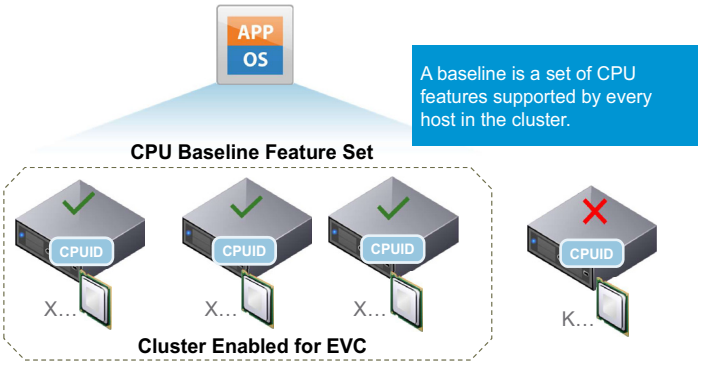
EVC allows vCenter Server to enforce vSphere vMotion compatibility among all hosts in a cluster by forcing hosts to expose a common set of CPU features (baseline) to virtual machines. A baseline is a set of CPU features that are supported by every host in the cluster. When you configure EVC, you set all host processors in the cluster to present the features of a baseline processor. After they have been enabled for a cluster, hosts that are added to the cluster are automatically configured to the CPU baseline.

Hosts that cannot be configured to the baseline are not permitted to join the cluster. Virtual machines in the cluster always see an identical CPU feature set, no matter which host they happen to run on. Because this process is automatic, EVC is easy to use and requires no specialized knowledge of CPU features and masks.

https://t.me/learningnets

7-49 CPU Baselines for an EVC Cluster

EVC works at the cluster level, using CPU baselines to configure all processors included in the cluster that are enabled for EVC.



Before you create an EVC cluster, ensure that the hosts you intend to add to the cluster meet the requirements.

EVC automatically configures hosts whose CPUs feature Intel FlexMigration and AMD-V Extended Migration technologies to be compatible with vSphere vMotion with hosts that use older CPUs. EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the real CPUs on the hosts differ.

You can use one of the following methods to create an EVC cluster:

- Create an empty cluster with EVC enabled and then move hosts into the cluster.
- Enable EVC on an existing cluster.

For EVC to function properly, the applications on the virtual machines must be written to use the CPUID machine instruction to discover CPU features as recommended by the CPU vendors. vSphere cannot support EVC with applications that do not follow the CPU vendor recommendations to discover CPU features.

To determine which EVC modes are compatible with your CPU, search the VMware Compatibility Guide at http://www.vmware.com/resources/compatibility/search.php?deviceCategory=server. Search for the server model or CPU family, and click the entry in the CPU Series column to display the compatible EVC modes.

https://t.me/learningnets

7-50 EVC Cluster Requirements

All hosts in the cluster must meet the following requirements:

- Use CPUs from a single vendor, either Intel or AMD:
  - Use Intel CPUs with Merom microarchitecture and later.
  - Use AMD first-generation Opteron CPUs and later.
- Be enabled for hardware virtualization: AMD-V or Intel VT
- Be enabled for execution-disable technology: AMD No eXecute (NX) or Intel eXecute Disable (XD)
- Be configured for vSphere vMotion migration

Applications in virtual machines must be CPU ID compatible.

The server hardware's CPU specifications usually indicate whether or not the CPUs contain the features that affect vSphere vMotion compatibility. If the specifications of a server or its CPU features are unknown, you can display these features with the VMware CPU identification utility. You use this utility to boot a server and determine whether or not its CPUs contain features like SSE3, SSSE3, and NX/XD.

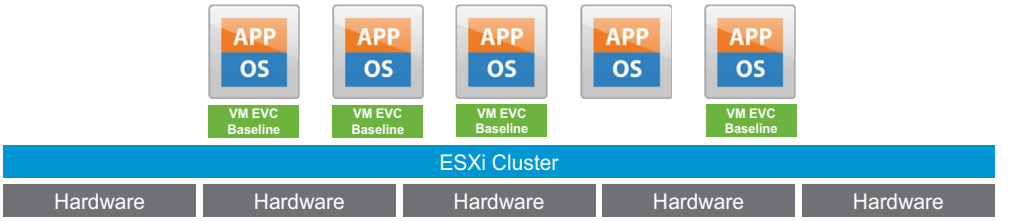
You can download the VMware CPU Identification Utility for ESXi hosts from https://www.vmware.com/support/shared_utilities.html.

https://t.me/learningnets

7-51 Virtual Machine EVC

Virtual machine EVC can be applied to some or all virtual machines in a cluster:

- It facilitates the migration of virtual machines beyond the cluster and across vCenter Server systems and data centers.
- It allows more granular definitions of EVC for specific VMs.
- It is independent of the EVC mode defined at the cluster level.
- Requirements:
  - ESXi 6.7
  - vCenter Server 6.7



Per-VM Enhanced vMotion Compatibility (EVC) enables the EVC mode to become an attribute of the VM rather than the specific processor generation it happens to be booted on in the cluster. This allows for seamless migration between two data centers that sport different processors. Further, the feature is persisted per VM and does not lose the EVC mode during migrations across clusters nor during power cycles.

In this diagram, EVC is not enabled on the cluster. The cluster consists of differing CPU models with different feature sets. The VMs with per-VM EVC can run on any ESXi host that can satisfy the defined EVC mode.

https://t.me/learningnets

7-52 Change CPU Identification Mask Settings

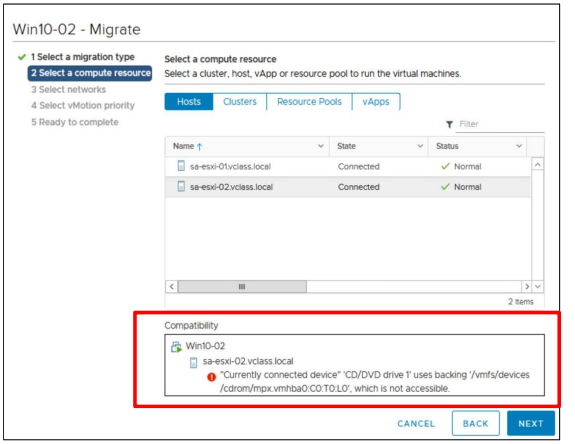
Masking or hiding CPU features can make a virtual machine widely available to ESXi hosts for migration:

- vCenter Server compares the CPU features available to a virtual machine with the CPU features of the destination host to determine whether to allow or disallow migration with vSphere vMotion.
- Masking the NX/XD bits prevents the virtual machine from using these features, but masking provides compatibility that enables you to migrate virtual machines to ESXi hosts that do not include this capability.
- When the NX/XD bit is visible to the guest operating system, the virtual machine can use this feature, but you can migrate the virtual machine only to hosts on which the feature is enabled.

https://t.me/learningnets

7-53 Checking vSphere vMotion Errors

When you select the host and cluster, a validation check is performed to verify that most vSphere vMotion requirements were met.



If validation succeeds, then you can continue in the wizard. If validation does not succeed, a list of vSphere vMotion errors and warnings displays in the Compatibility pane.

Warnings have yellow triangles. Errors have red diamonds. Warnings allow you to perform a vSphere vMotion migration. Errors do not allow you to continue. You must exit the wizard and fix all errors before retrying the migration.

If a failure occurs during the vSphere vMotion migration, the virtual machine is not migrated and continues to run on the source host.

https://t.me/learningnets

7-54 Migration with vSphere Storage vMotion

With vSphere Storage vMotion, you can migrate a virtual machine and its disk files from one datastore to another while the virtual machine is running.

Using vSphere Storage vMotion you can perform the following tasks:

- Move virtual machines off arrays for maintenance or to upgrade.
- Change the disk provisioning type.
- Change virtual machine files on the destination datastore to match the inventory name of the virtual machine. The migration renames all virtual disk, configuration, snapshot, and `.nvram` files.
- Move virtual machines off a storage device to allow maintenance or reconfiguration of the storage device without virtual machine downtime.
- Redistribute virtual machines or virtual disks to different storage volumes to balance capacity or improve performance.

With vSphere Storage vMotion, you can migrate a virtual machine and its disk files from one datastore to another while the virtual machine is running. With vSphere Storage vMotion, you can move virtual machines off of arrays for maintenance or to upgrade. You also have the flexibility to optimize disks for performance, or to transform disk types, which you can use to reclaim space.

You can choose to place the virtual machine and all its disks in a single location, or you can select separate locations for the virtual machine configuration file and each virtual disk. The virtual machine does not change execution host during a migration with vSphere Storage vMotion.

During a migration with vSphere Storage vMotion, you can change the disk provisioning type.

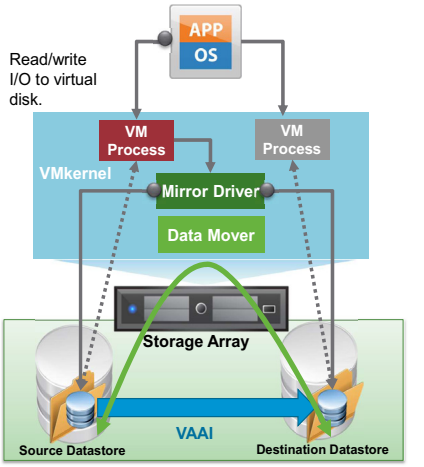
Migration with vSphere Storage vMotion changes virtual machine files on the destination datastore to match the inventory name of the virtual machine. The migration renames all virtual disk, configuration, snapshot, and `.nvram` files. If the new names exceed the maximum filename length, the migration does not succeed.

https://t.me/learningnets

7-55 vSphere Storage vMotion in Action

vSphere Storage vMotion uses an I/O mirroring architecture to copy disk blocks between source and destination:

1. Initiate storage migration.
2. Use the VMkernel data mover or vSphere Storage APIs - Array Integration to copy data.
3. Start a new virtual machine process.
4. Mirror I/O calls to file blocks that are already copied to virtual disk on the destination datastore.
5. Cut over to the destination virtual machine process to begin accessing the virtual disk copy.



The storage migration process does a single pass of the disk, copying all the blocks to the destination disk. If blocks are changed after they are copied, the blocks are synchronized from the source to the destination through the mirror driver, with no need for recursive passes.

This 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 VMware vSphere® Storage APIs - Array Integration, also called hardware acceleration.

https://t.me/learningnets

7-56 vSphere Storage vMotion Guidelines and Limitations

Guidelines:

- Plan the migration and coordinate with administrators.
- Perform migrations during off-peak hours.
- Ensure that the host has access to source datastores and target datastores.

Limitations:

- Independent virtual machine disks must be in persistent mode.

A virtual machine and its host must meet certain resource and configuration requirements for the virtual machine disks (VMDKs) to be migrated with vSphere Storage vMotion. One of the requirements is that the host on which the virtual machine is running must have access both to the source datastore and to the target datastore.

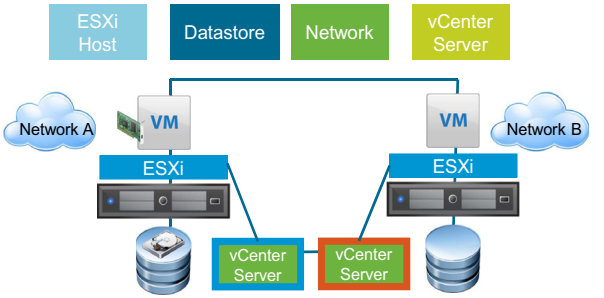
During a migration with vSphere Storage vMotion, you can change the disk provisioning type. Migration with vSphere Storage vMotion changes virtual machine files on the destination datastore to match the inventory name of the virtual machine. The migration renames all virtual disk, configuration, snapshot, and `.nvram` files. If the new names exceed the maximum filename length, the migration does not succeed.

https://t.me/learningnets

7-57 Shared-Nothing vSphere vMotion Migration

Shared-nothing vSphere vMotion migration enables a virtual machine to change its host, datastores, networks, and vCenter Server instances simultaneously, even if the two hosts do not have a shared storage.

This technique combines vSphere vMotion and vSphere Storage vMotion into a single operation.



You can migrate virtual machines beyond storage accessibility boundaries and between hosts, within and across clusters, data centers, and vCenter Server instances.

The shared-nothing vSphere vMotion is useful for performing cross-cluster migrations, when the target cluster machines might not have access to the source cluster's storage. Processes on the virtual machine continue to run during the migration with vSphere vMotion.

Shared-nothing vSphere vMotion migration is useful for virtual infrastructure administration tasks, including:

- Host maintenance: You can move virtual machines from a host to allow host maintenance.
- Storage maintenance and reconfiguration: You can move virtual machines from a storage device to allow maintenance or reconfiguration of the storage device without virtual machine downtime.
- Storage load redistribution: You can manually redistribute virtual machines or virtual disks to different storage volumes to balance capacity or to improve performance.

Shared-nothing vSphere vMotion migration counts against the concurrent limitations of both vSphere vMotion and vSphere Storage vMotion. No more than two concurrent shared-nothing vSphere vMotion migrations are allowed. Running two concurrent shared-nothing vSphere vMotion migrations causes all attempted vSphere Storage vMotion instances to remain queued until one of the active shared-nothing vSphere vMotion migrations is complete. Similarly, shared-nothing vSphere vMotion migration instances also count against vSphere vMotion limits, at most eight concurrent vSphere vMotion instances per host. If two shared-nothing vSphere vMotion migration instances are active, then at most, only six concurrent vSphere vMotion instances are allowed at the same time. If eight vSphere vMotion instances are active, new shared-nothing vSphere vMotion migration attempts are queued until one of the active vSphere vMotion instances is complete.

https://t.me/learningnets

Shared-nothing vSphere vMotion migration behaves exactly like vSphere vMotion with respect to supporting multi-NICs. Likewise, shared-nothing vSphere vMotion migration supports either shared swap or unshared swap migrations similar to vSphere vMotion, with the virtual machine home directory movement becoming an unshared swap migration. Shared-nothing vSphere vMotion migration instances are more expensive, which must be factored in when making migration decisions. Neither vSphere DRS nor VMware vSphere® Storage DRS™ uses shared-nothing vSphere vMotion migration technology. Although neither vSphere DRS nor vSphere Storage DRS recommends shared-nothing vSphere vMotion migrations, users can perform manual shared-nothing vSphere vMotion migrations in or across vSphere Storage DRS or vSphere DRS clusters.

For information about requirements and limitations for shared-nothing vSphere vMotion migration, see *vCenter Server and Host Management Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

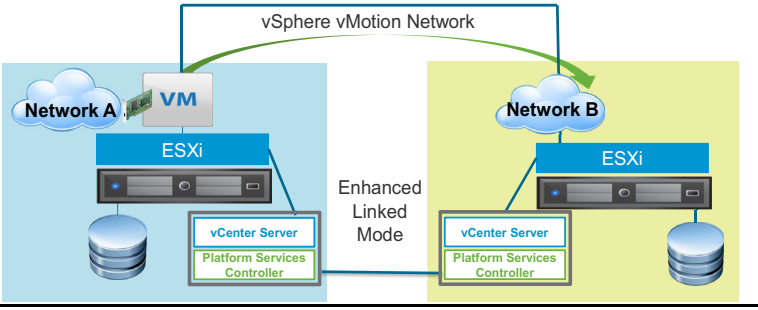
https://t.me/learningnets

7-58 Cross-vCenter migrations

vSphere vMotion can migrate virtual machines between linked vCenter Server systems.

Requirements:

- ESXi hosts and vCenter Server systems must be upgraded to vSphere 6.x.
- vCenter Server instances must be in Enhanced Linked Mode.
- Hosts must be time-synchronized.
- Both of the vCenter Server systems should be the same vCenter Server version.



Migration of virtual machines across vCenter Server systems is helpful in the following VM provisioning cases:

- Balance workloads across clusters and vCenter Server instances.
- Elastically expand or shrink capacity across resources in different vCenter Server instances in the same site or in another geographical area.
- Move virtual machines between environments that have different purposes, for example, from a development environment to production environment.
- Move virtual machines to meet different Service Level Agreements (SLAs) regarding storage space, performance, and so on.

When using vSphere Web Client, both vCenter Server instances must be in Enhanced Linked Mode and must be in the same vCenter Single Sign-On domain. This lets the source vCenter Server to authenticate to the destination vCenter Server.

https://t.me/learningnets

7-59 VMkernel Networking Layer and TCP/IP Stacks

The VMkernel networking layer provides connectivity to hosts and handles the standard system traffic of vSphere vMotion, IP storage, vSphere Fault Tolerance, vSAN, and others.

You can also create VMkernel adapters on the source and target vSphere Replication hosts to isolate the replication data traffic.

TCP/IP stacks at the VMkernel level:

- Default TCP/IP stack
- vSphere vMotion TCP/IP stack
- Provisioning TCP/IP stack
- Custom TCP/IP stacks

Consider the following key points about TCP/IP stacks at the VMkernel level:

- Default TCP/IP stack: Provides networking support for the management traffic between vCenter Server and ESXi hosts and for system traffic such as vSphere vMotion, IP storage, and vSphere Fault Tolerance.
- vSphere vMotion TCP/IP stack: Supports the traffic for live migration of virtual machines. Use the vSphere vMotion TCP/IP stack to provide better isolation for the vSphere vMotion traffic. After you create a VMkernel adapter on the vSphere vMotion TCP/IP stack, you can use only this stack for vSphere vMotion migration on this host. The VMkernel adapters on the default TCP/IP stack are disabled for the vSphere vMotion service. If a live migration uses the default TCP/IP stack while you configure VMkernel adapters with the vSphere vMotion TCP/IP stack, the migration completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are disabled for future vSphere vMotion sessions.
- Provisioning TCP/IP stack: Supports the traffic for virtual machine cold migration, cloning, and snapshot creation. You can use the provisioning 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 virtual machines in long-distance vSphere vMotion. By using the provisioning TCP/IP stack, you can isolate the traffic from the cloning operations on a separate gateway. After you configure a VMkernel adapter with the provisioning TCP/IP stack, all adapters on the default TCP/IP stack are disabled for the provisioning traffic.

https://t.me/learningnets

- 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:

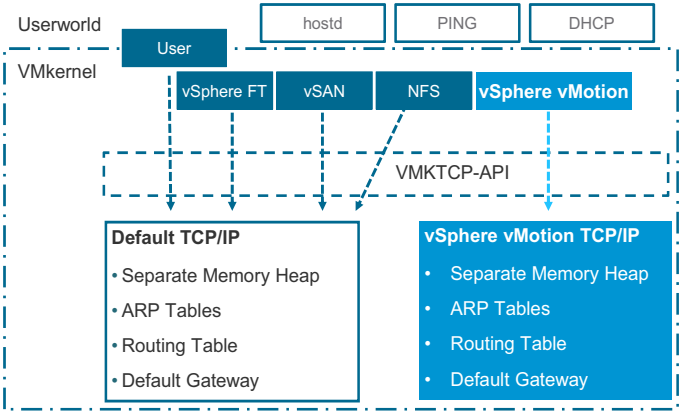
```
esxcli network ip netstack add -N="stack_name"
```

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

7-60 vSphere vMotion TCP/IP Stacks

In vSphere 6.x, each host has a second TCP/IP stack dedicated to vSphere vMotion migration.



vSphere vMotion TCP/IP stacks support the traffic for live migration of virtual machines. Use the vSphere vMotion TCP/IP stack to provide better isolation for the vSphere vMotion traffic. After you create a VMkernel adapter on the vSphere vMotion TCP/IP stack, you can use only this stack for vSphere vMotion migration on this host.

The VMkernel adapters on the default TCP/IP stack are disabled for the vSphere vMotion service. If a live migration uses the default TCP/IP stack while you configure VMkernel adapters with the vMotion TCP/IP stack, the migration completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are disabled for future vSphere vMotion sessions.

https://t.me/learningnets

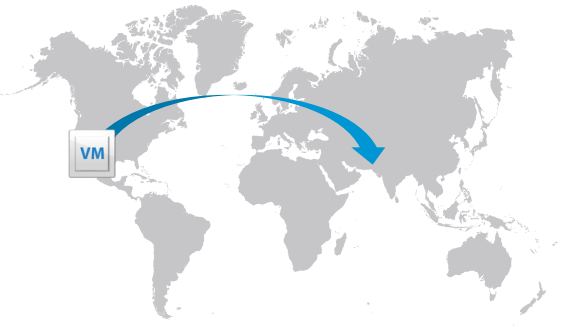
7-61 Long-Distance vSphere vMotion Migration

Long-distance vSphere vMotion migration is an extension of vSphere vMotion migration across vCenter Server instances.

This migration is targeted at environments where vCenter Server systems are spread across large geographic distances and where the latency across sites is high.

Use cases for long-distance vSphere vMotion migration:

- Permanent migrations
- Disaster avoidance
- VMware Site Recovery Manager™ and disaster avoidance testing
- Multisite load balancing
- Follow-the-Sun scenario support



https://t.me/learningnets

7-62 Networking Prerequisite for Long-Distance vSphere vMotion Migration

vSphere vMotion migrations between vCenter Server instances must connect over layer 3 connections:

- Virtual machine network:
  - L2 connection.
  - Same virtual machine IP address available at destination.
- vSphere vMotion network:
  - L3 connection.
  - Secure (recommended if 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

7-63 Network Checks for Migrations Between vCenter Server Instances

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
- vSphere vMotion migration to an internal network, for example, a network without a physical NIC

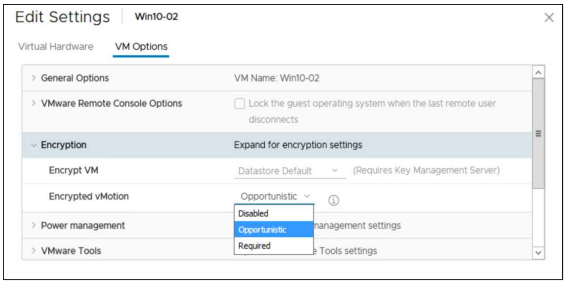
https://t.me/learningnets

7-64 Encrypted vSphere vMotion

vSphere vMotion always uses encryption when migrating encrypted virtual machines.

For virtual machines that are not encrypted, select one of the following encrypted vSphere vMotion options:

- **Disabled**.
- **Opportunistic:** Encrypted vSphere vMotion is used if the source and destination hosts support it.
- **Required:** If the source or destination host does not support encrypted vSphere vMotion, migration with vSphere vMotion fails.



Encrypted vSphere vMotion secures confidentiality, integrity, and authenticity of data that is transferred with vSphere vMotion. Encrypted vSphere vMotion supports all variants of vSphere vMotion, including migration across vCenter Server systems. Encrypted vSphere Storage vMotion is not supported.

You cannot turn off encrypted vSphere vMotion for encrypted virtual machines.

The Opportunistic state is the default option for virtual machines that are not encrypted.

https://t.me/learningnets

7-65 Lab 11: Migrating Virtual Machines

Use vSphere vMotion and vSphere Storage vMotion to migrate virtual machines

1. Migrate Virtual Machine Files from Local Storage to Shared Storage
2. Create a Virtual Switch and a VMkernel Port Group for vSphere vMotion Migration
3. Prepare Virtual Machines to Demonstrate vSphere vMotion Migration
4. Perform vSphere vMotion Migrations of Virtual Machines
5. Perform Compute Resource and Storage Migrations

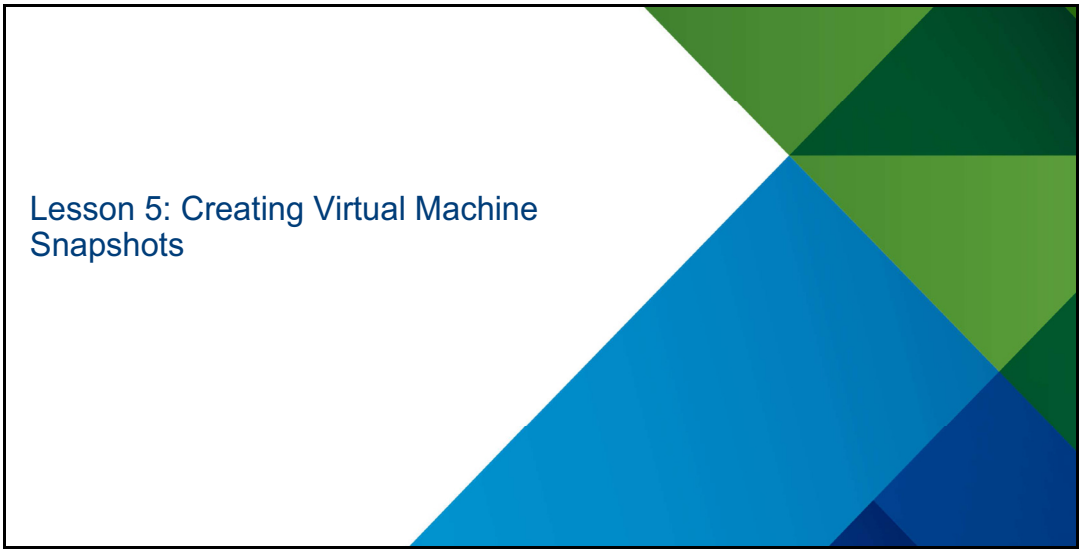
https://t.me/learningnets

7-66 Review of Learner Objectives

You should be able to meet the following objectives:

- Verify vSphere vMotion requirements, including CPU constraints and guidelines
- Perform a vSphere vMotion migration
- Perform a vSphere Storage vMotion migration
- Perform a shared-nothing vSphere vMotion migration
- Explain how cluster Enhanced vMotion Compatibility and per-VM EVC works with vSphere vMotion

https://t.me/learningnets



https://t.me/learningnets

7-68 Learner Objectives

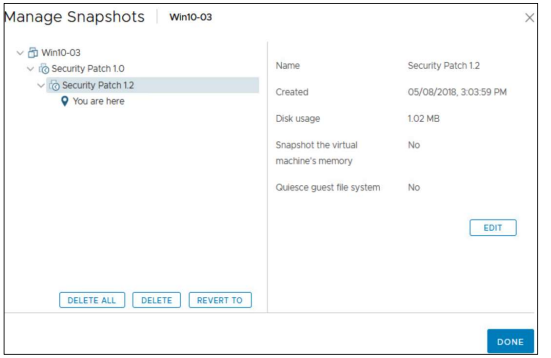
By the end of this lesson, you should be able to meet the following objectives:

- Take a snapshot of a virtual machine
- Manage multiple snapshots
- Delete virtual machine snapshots
- Consolidate snapshots
- Configure the maximum number of snapshots allowed for a virtual machine

https://t.me/learningnets

7-69 Virtual Machine Snapshots

Snapshots enable you to preserve the state of the virtual machine so that you can repeatedly return to the same state.



Snapshots are useful when you want to revert repeatedly to the same state but do not want to create multiple virtual machines. Examples include patching or upgrading the guest operating system in a virtual machine. Snapshots give you the ability to back out of the patch or upgrade process if problems occur during patching or upgrading.

The relationship between snapshots is like the relationship between a parent and a child. Snapshots are organized in a snapshot tree. In a snapshot tree, each snapshot has one parent and one or more children, except for the last snapshot, which has no children.

On the slide, the snapshot named Security patch 1.0 has two child snapshots, each named Security Patch 1.1 and 1.2.

A virtual machine snapshot mainly includes the following states:

- Settings state: Virtual machine's settings (`.nvram` and `.vmx`) and power state
- Disk state: State of the virtual machine's associated disks
- Memory state: Contents of the virtual machine's memory (optional)

https://t.me/learningnets

7-70 Taking Snapshots

You can take a snapshot while a virtual machine is powered on, powered off, or suspended.

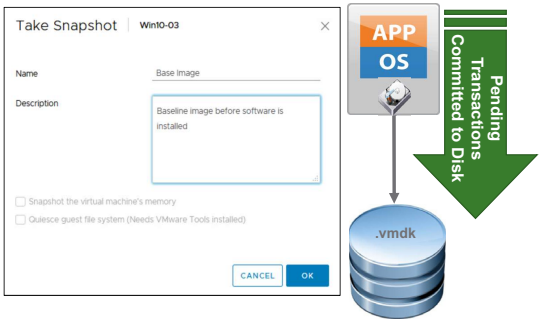
A snapshot captures the following items:

- Virtual machine configuration
- Virtual machine memory state (optional)
- Virtual disks
- Virtual compatibility mode RDMs

A snapshot capture does not include:

- Independent virtual disks (persistent and nonpersistent)
- Physical compatibility mode RDMs

Virtual machine snapshots are not recommended as a virtual machine backup strategy.



A snapshot captures the entire state of the virtual machine at the time that you take the snapshot, including the following states:

- Memory state: The contents of the virtual machine's memory. The memory state is captured only if the virtual machine is powered on and if you select the **Snapshot the virtual machine's memory** check box (selected by default).
- Settings state: The virtual machine settings.
- Disk state: The state of all the virtual machine's virtual disks.

At the time that you take the snapshot, you can also quiesce the guest operating system. This action quiesces the file system of the guest operating system.

Snapshots of physical compatibility mode RDM disks are not supported.

CAUTION

Virtual machine snapshots are not recommended as a virtual machine backup strategy.

https://t.me/learningnets

7-71 Types of Snapshots

A delta or child disk is created when a snapshot is created:

- On the VMFS datastore, the delta disk is a sparse disk.
- Depending on the type of datastore, delta disks use different sparse formats.

Snapshot Type	Datastore Type	Filename	Block Size
VMFSsparse	VMFS5 with virtual disk smaller than 2 TB	#-delta.vmdk	512 bytes
SEsparse	• VMFS6 • VMFS5 with virtual disk larger than 2 TB • Space-efficient (thin provisioned) • Supports disk reclamation (unmap)	#-sesparse.vmdk	4 KB
vsanSparse	vSAN	Delta object	4 MB

Depending on the type of your datastore, delta disks use different sparse formats.

- 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/Os 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. Upon 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 similar to 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 blocks. The unmapping helps to reclaim space allocated by SEsparse once the guest operating system has deleted that data.

https://t.me/learningnets

7-72 Virtual Machine Snapshot Files (1)

A snapshot consists of a set of files:

`-Snapshot#.vmsn`: Configuration state

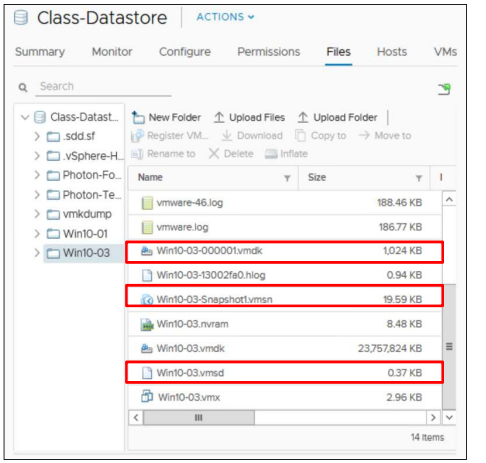
`-Snapshot#.vmem`: Memory state (optional)

`-00000#.vmdk`: Disk descriptor

`-00000#-delta.vmdk`: VMFS5 delta

`-00000#-sesparse.vmdk`: VMFS6 delta

The `.vmsd` file stores the names, descriptions, and relationships for all of the VM's snapshots.



A virtual machine can have one or more snapshots. Each snapshot results in the following files being created:

- Delta disk file: When you take a snapshot of a virtual machine, the state of each virtual disk is preserved, excluding physical RDMs and virtual disks marked as independent. The virtual machine stops writing to its *VM_name*`-flat.vmdk` file, and writes are redirected to *VM_name*`-######-delta.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 virtual machine is powered off.
- Disk descriptor file: *VM_name*`-00000#.vmdk`. This file is a small text file that contains information about the snapshot.
- Configuration state file: *VM_name*`-Snapshot#.vmsn`. # is the next number in the sequence, starting with 1 This file holds the active memory state of the virtual machine when the snapshot was taken, including virtual hardware, power state, and hardware version.
- Memory state file: *VM_name*`-Snapshot#.vmem`. This file is created if the option to include memory state was selected when creating the snapshot. It contains the entire contents of the virtual machine's memory at the time the snapshot of virtual machine was taken.
- Snapshot delta file: *VM_name*`-00000#-delta.vmdk`. This file contains the changes to the virtual disk's data since the snapshot was taken.

https://t.me/learningnets

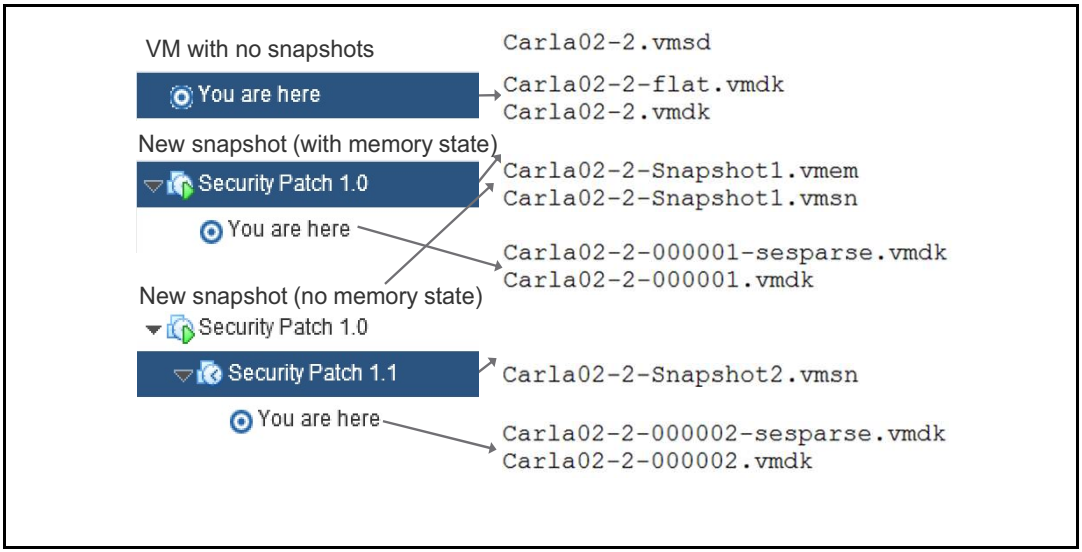
- Snapshot active memory file: *VM_name*-Snapshot#.vmem. This file contains the contents of the virtual machine memory if the option to include memory is selected when creating the snapshot.
- *VM_name*.vmsd is the snapshot list file and is created at the time that the virtual machine is created. It maintains snapshot information for a virtual machine so that it can create a snapshot list in vSphere Web Client. This information includes the name of the snapshot .vmsn file and the name of the virtual disk file.
- The snapshot state file has a .vmsn extension and is used to store the state of a virtual machine when a snapshot is taken. A new .vmsn file is created for every snapshot that is created on a virtual machine and is deleted when the snapshot is deleted. The size of this file varies, based on the options selected when the snapshot is created. For example, including the memory state of the virtual machine in the snapshot increases the size of the .vmsn file.

You can exclude one or more of the VMDKs from a snapshot by designating a virtual disk in the virtual machine as an independent disk. Placing a virtual disk in independent mode is typically done when the virtual disk is created. If the virtual disk was created without enabling independent mode, you must power off the virtual machine to enable it.

Other files might also exist, depending on the virtual machine hardware version. For example, each snapshot of a virtual machine that is powered on has an associated *snapshot_name_number*.vmem file, which contains the guest operating system main memory, saved as part of the snapshot.

https://t.me/learningnets

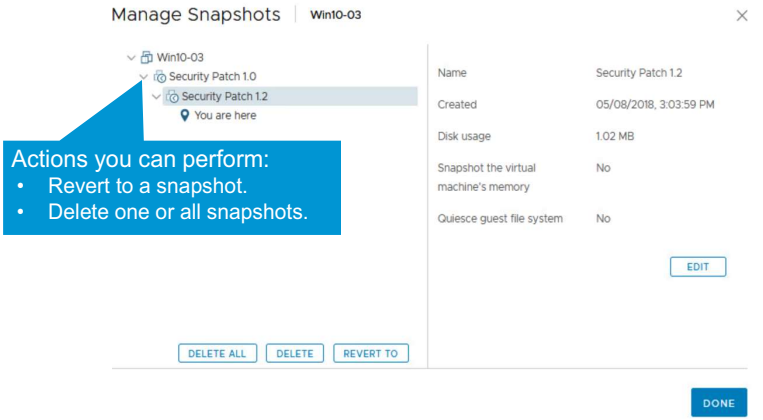
7-73 Virtual Machine Snapshot Files (2)



https://t.me/learningnets

7-74 Managing Snapshots

You can use the **Snapshots** tab to view snapshots for the active virtual machine and to act directly on them.



You can perform the following actions from the **Snapshot** tab:

- **Edit Snapshot**: Edit the snapshot name and description.
- **Delete Snapshot**: Removes the snapshot from the Snapshot Manager and consolidates the snapshot files to the parent snapshot disk and merge with the virtual machine base disk.
- **Delete All Snapshots**: Commits all the intermediate snapshots before the current-state icon (You are here) to the Tulika01-03 VMDK file and removes all snapshots for that virtual machine.
- **Revert to**: Enables you to restore, or revert to, a particular snapshot. The snapshot that you restore becomes the current snapshot.

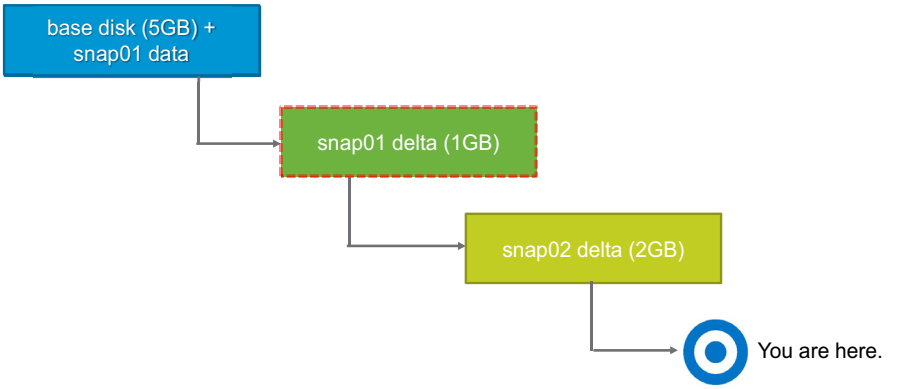
When you revert to a snapshot, you return all these items to the state that they were in at the time that you took the snapshot. If you want the virtual machine to be suspended, powered on, or powered off when you start it, ensure that the virtual machine is in the correct state when you take the snapshot.

Deleting a snapshot (with **Delete Snapshot** or **Delete All Snapshots**) 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

7-75 Deleting Virtual Machine Snapshots (1)

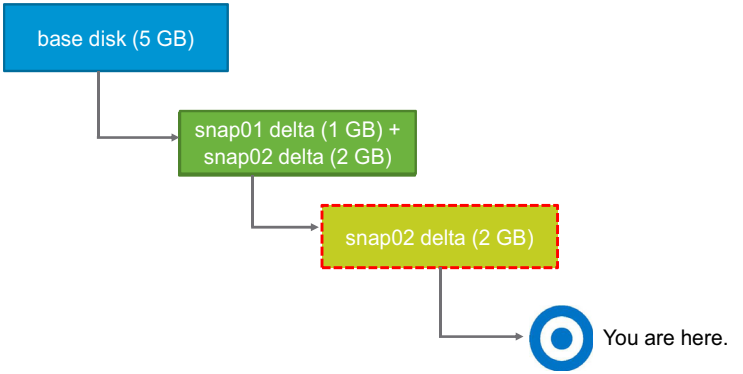
If you delete a snapshot one or more levels above the You are here level, the snapshot state is deleted. The snap01 data is committed into the parent (base disk), and the foundation for snap02 is retained.



https://t.me/learningnets

7-76 Deleting Virtual Machine Snapshots (2)

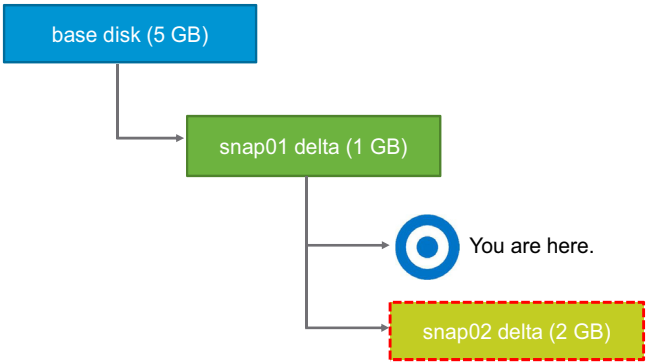
If you delete the latest snapshot, the changes are committed to its parent. The snap02 data is committed into snap01 data, and the snap02 -delta.vmdk file is deleted.



https://t.me/learningnets

7-77 Deleting Virtual Machine Snapshots (3)

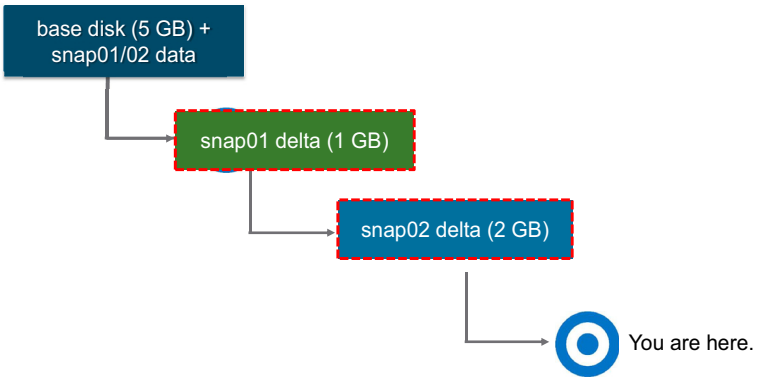
If you delete a snapshot one or more levels below the You are here level, subsequent snapshots are deleted, and you can no longer return to those states. The snap02 data is deleted.



https://t.me/learningnets

7-78 Deleting All Virtual Machine Snapshots

The delete-all-snapshots mechanism uses storage space efficiently. The size of the base disk does not increase. Like a single snapshot deletion, changed blocks in the snapshot overwrite their counterparts in the base disk.



In the slide, snap01 is committed to the base disk before snap02 is committed.

All snapshots before You are here are committed all the way up to the base disk. All snapshots after You are here are discarded.

https://t.me/learningnets

7-79 About Snapshot Consolidation

Snapshot consolidation is a method to commit a chain of snapshots to the base disks when the Snapshot Manager shows that no snapshots exist but the delta files still remain on the datastore.

Snapshot consolidation is intended to resolve problems that might occur with snapshots:

- The snapshot descriptor file is committed correctly, thus the **Snapshot** window shows that all the snapshots are deleted.
- The snapshot files (`-delta.vmdk`) are still part of the virtual machine.
- Snapshot files continue to expand until the virtual machine runs out of datastore space.

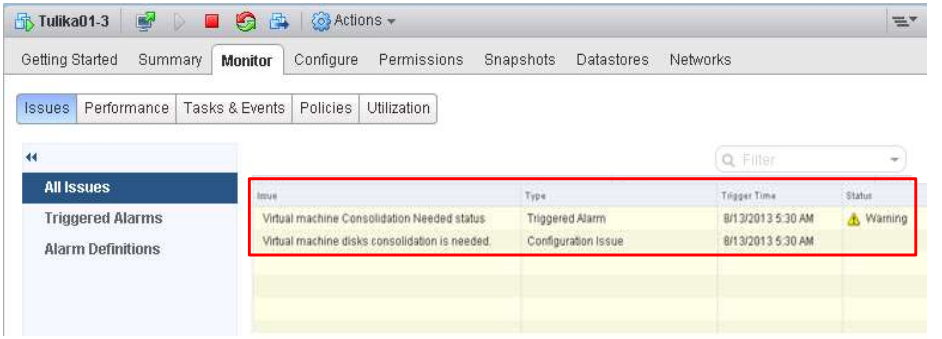
Snapshot consolidation is a way to clean unneeded snapshot delta files from a datastore. If no snapshots are registered for a virtual machine, but snapshot delta files exist, snapshot consolidation commits the chain of the snapshots indicated by the delta files and removes them.

If consolidation is not performed, the snapshot files might expand to the point of consuming all the remaining space on the virtual machine's datastore.

https://t.me/learningnets

7-80 Discovering When to Consolidate Snapshots

A warning on the **Monitor > Issues** tab of the virtual machine notifies users 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 is displayed, the user can use vSphere Web Client to commit the snapshots, rather than committing snapshots from a command-line session.

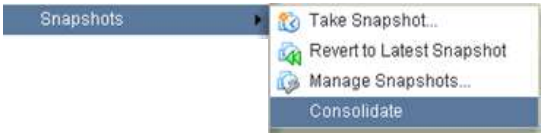
In the example, the **Monitor** tab displays a consolidation issue.

https://t.me/learningnets

7-81 Performing Snapshot Consolidation

After the snapshot consolidation warning appears, users can use vSphere Web Client to consolidate the snapshots:

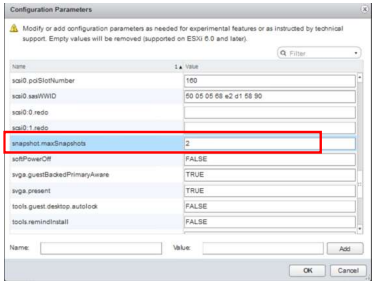
- Select **Snapshots > Consolidate** to reconcile snapshots.
- All snapshot delta disks are committed to the base disks.



https://t.me/learningnets

7-82 Setting the Maximum Number of Snapshots Allowed for a VM

You can set the maximum number of snapshots for a virtual machine by editing the snapshot.maxSnapshots = *n* variable of the configuration file.



https://t.me/learningnets

7-83 Lab 12: Managing Virtual Machines

Perform virtual machine management tasks

1. Unregister a Virtual Machine from the vCenter Server Appliance Inventory
2. Register a Virtual Machine in the vCenter Server Appliance Inventory
3. Unregister and Delete Virtual Machines from the Datastore
4. Take Snapshots of a Virtual Machine
5. Add Files and Take Another Snapshot of a Virtual Machine
6. Revert the Virtual Machine to a Snapshot
7. Delete an Individual Snapshot
8. Delete All Snapshots

https://t.me/learningnets

7-84 Review of Learner Objectives

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
- Configure the maximum number of snapshots allowed for a virtual machine

https://t.me/learningnets

7-85 Key Points

- vCenter Server provides features for provisioning virtual machines, such as templates and cloning.
- By deploying virtual machines from a template, you can create many virtual machines easily and quickly.
- You can use vSphere vMotion to move virtual machines while they are powered on.
- You can use vSphere Storage vMotion to move virtual machines from one datastore to another datastore.
- You can use virtual machine snapshots to preserve the state of the virtual machine so that you can return repeatedly to the same state.

Questions?

https://t.me/learningnets

https://t.me/learningnets

MODULE 8

Resource Management and Monitoring



https://t.me/learningnets

8-2 You Are Here

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

https://t.me/learningnets

8-3 Importance

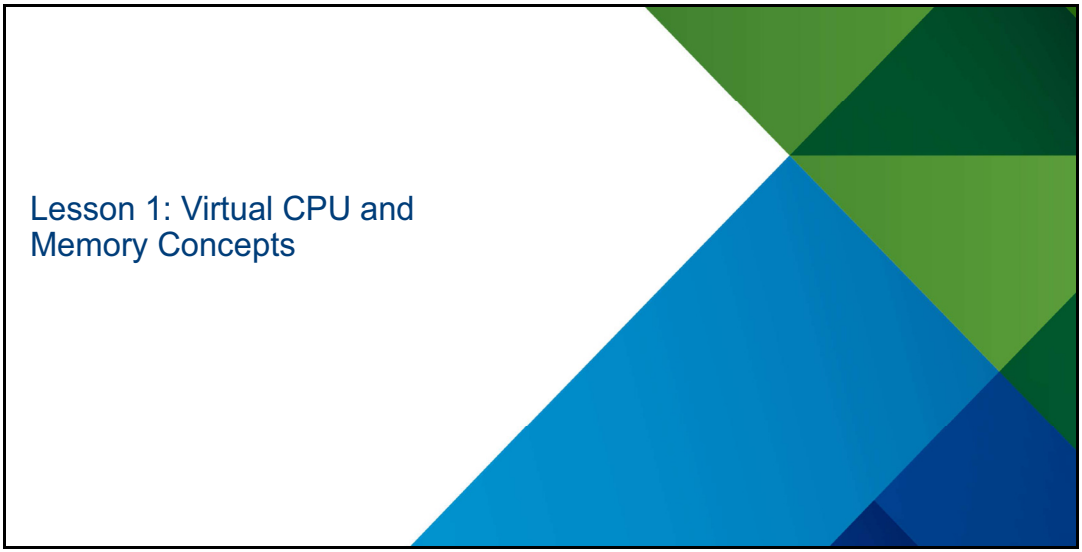
Although the VMkernel works proactively to avoid resource contention, maximizing performance requires both analysis and ongoing monitoring.

https://t.me/learningnets

8-4 Module Lessons

Lesson 1:	Virtual CPU and Memory Concepts
Lesson 2:	Resource Controls and Resource Pools
Lesson 3:	Monitoring Resource Use
Lesson 4:	Using Alarms

https://t.me/learningnets



https://t.me/learningnets

8-6 Learner Objectives

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

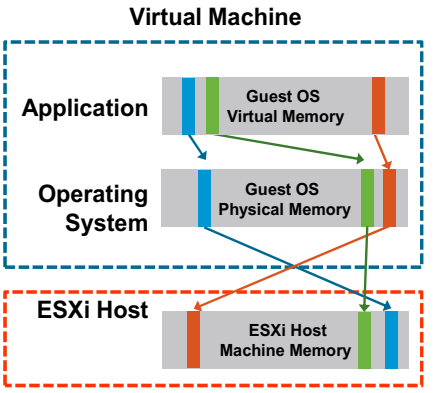
- Discuss CPU and memory concepts in a virtualized environment
- Describe what overcommitment of a resource means
- Identify additional technologies that improve memory usage
- Describe how VMware vSphere® Virtual Symmetric Multiprocessing works and how hyperthreading is used by the VMkernel

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 virtual machine.



When running a virtual machine, the VMkernel creates a contiguous addressable memory space for the virtual machine. This memory space has the same properties as the virtual memory address space presented to applications by the guest operating system. This memory space allows the VMkernel to run multiple virtual machines simultaneously while protecting the memory of each virtual machine from being accessed by others. From the perspective of an application running in the virtual machine, the VMkernel adds an extra level of address translation that maps the guest physical address to the host physical address.

https://t.me/learningnets

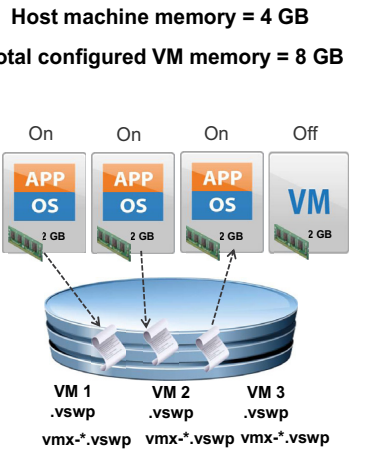
8-8 Virtual Machine Memory Overcommitment

Memory is overcommitted when the combined configured memory footprint of all powered-on virtual machines exceeds that of the host memory sizes.

Virtual machines do not always use their full allocated memory.

An ESXi host transfers memory from idle virtual machines to virtual machines that need more memory to improve memory usage.

Memory overhead is stored in a swap file (`.vswp`).



The total configured memory sizes of all virtual machines might exceed the amount of available physical memory on the host. However, this condition does not necessarily mean memory is overcommitted. Memory is overcommitted when the working memory size of all virtual machines exceeds that of the ESXi host's physical memory size.

Because of the memory management techniques used by the ESXi host, your virtual machines can use more virtual RAM than the available physical RAM on the host. For example, you can have a host with 4 GB of memory and run four virtual machines with 2 GB of memory each. In that case, the memory is overcommitted. For example, if all four virtual machines are idle, the combined consumed memory might be below 4 GB. However, if all virtual machines are actively consuming memory, then their memory footprint might exceed 4 GB and the ESXi host might become overcommitted. An ESXi host can run out of memory if virtual machines consume all reservable memory in an overcommitted-memory environment. Although the powered-on virtual machines are not affected, a new virtual machine might fail to power on due to lack of memory. Overcommitment makes sense because, typically, some virtual machines are lightly loaded while others are more heavily loaded, and relative activity levels vary over time.

Extra memory from a virtual machine is gathered into a `.vswp` swap file. The memory overcommitment process on the host uses the `vmx-*.vswp` swap file to gather and track memory overhead. Memory from this file is swapped out to disk when host machine memory is overcommitted.

https://t.me/learningnets

8-9 Memory Reclamation Techniques

Memory reclamation occurs when the free memory value of the host is 6 percent or less, indicating that the host cannot meet the memory requirements.

The ESXi host can use the following methods to reclaim virtual machine memory:

- Economize use of physical memory pages:
  - Transparent page sharing allows pages with identical contents to be stored only once.
- Deallocate memory from one virtual machine to another:
  - The ballooning mechanism becomes active when memory is scarce, then forces virtual machines to use their own paging areas.
- Memory compression tries to reclaim some memory performance when memory contention is high.
- Host-level SSD swapping:
  - Use of a solid-state drive on the host for a host cache swap file might increase performance.
- Page virtual machine memory out to disk:
  - Using VMkernel swap space is the last resort, because of poor performance.

The VMkernel uses various techniques when reclaiming memory on an ESXi host. Each technique is described in the order that the VMkernel uses it:

- The first technique economizes the use of physical memory pages and is called transparent page sharing (TPS). TPS allows pages with identical contents to be stored only once. Workloads consisting of multiple virtual machines often consume less memory than they would when running on physical machines. The effect of TPS is that it allows hosts to efficiently support higher levels of memory overcommitment.
- During times of memory contention, the VMkernel looks for opportunities to reclaim either idle or allocated but unused memory from virtual machines. The VMkernel transfers memory from idle virtual machines to virtual machines that need more memory, using the `vmmemctl` driver that is typically installed with VMware Tools. The `vmmemctl` driver is also called the memory balloon driver.
- Memory compression is another technique that the VMkernel uses to reclaim host physical memory. This technique attempts to reclaim some memory by compressing pages when contention is high to avoid swapping out to a virtual machine swap file.
- Another option for reclaiming memory is to create a host cache on a solid-state drive (SSD) on the host machine. This host cache will be used after TPS, ballooning, and memory compression have been tried, and before using a swap file on disk. Having the host cache file on an SSD drive can improve the performance of this technique.

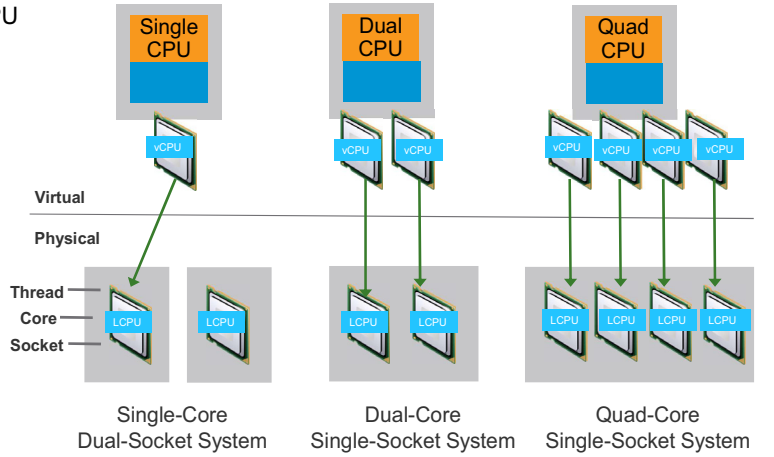
https://t.me/learningnets

- The last and least desirable technique employed is to page virtual machine memory out to disk. Swapping to disk is used when other techniques are temporarily unable to reclaim memory quickly enough to satisfy current system demand. A significant performance penalty is observed when this technique is used.

https://t.me/learningnets

8-10 Configuring Multicore Virtual CPUs

With multicore virtual CPU support, you can control the number of cores per virtual socket in a virtual machine.



You can configure a virtual machine with up to 128 virtual CPUs (vCPUs). The VMkernel includes a CPU scheduler that dynamically schedules vCPUs on the physical CPUs of the host system.

The VMkernel scheduler considers socket-core-thread topology when making scheduling decisions. Intel and AMD have developed processors that combine multiple processor cores into a single integrated circuit, called a socket in this discussion. A socket is a single package that can have one or more physical CPUs, with each core having one or more logical CPUs (LCPUs) or threads. LCPUs provide the core with the ability to schedule one thread of execution. Each LCPU provides the core with the ability to schedule one thread of execution. On the slide, the first system is a single-core, dual-socket system that has two cores and therefore two LCPUs.

When a vCPU of a single-vCPU or multi-vCPU virtual machine must be scheduled, the VMkernel maps the vCPU to an available logical processor.

https://t.me/learningnets

8-11 About Hyperthreading

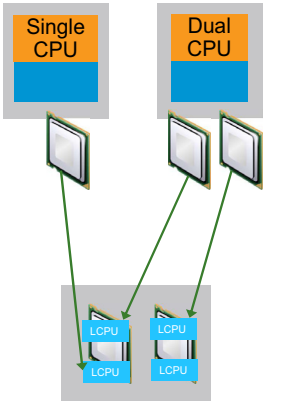
Hyperthreading enables a core to execute two threads or sets of instructions at the same time.

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 on the host system, ESXi can execute two threads at the same time on each processor core (physical CPU). Hyperthreading provides more scheduler throughput. That is, hyperthreading provides more logical CPUs on which vCPUs can be scheduled. The drawback of hyperthreading is that it does not double the power of a core. So, if both threads of execution need the same on-chip resources at the same time, one thread has to wait. Still, ESXi uses recent advances in hyperthreading technology and, on systems that use these new technologies, performance is improved.

An ESXi host enabled for hyperthreading should behave almost exactly like a standard system. Logical processors on the same core have adjacent CPU numbers. Logical processors 0 and 1 are on the first core, logical processors 2 and 3 are on the second core, and so on.

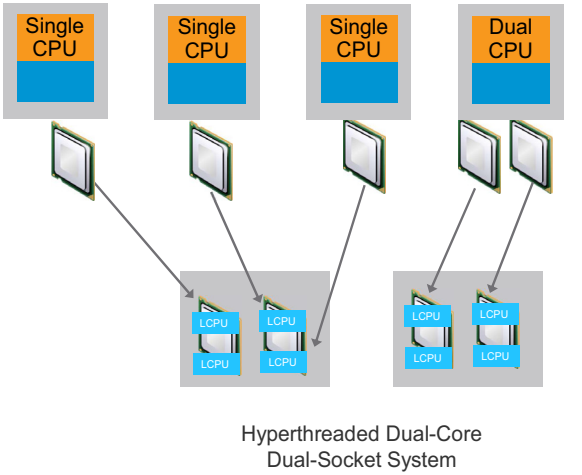
To ensure that hyperthreading is functioning, consult the hardware documentation to see whether or not the BIOS includes support for hyperthreading. Then, enable hyperthreading in the system BIOS. Some manufacturers call this option Logical Processor and others call it Enable Hyperthreading.

Use vSphere Web Client to ensure that hyperthreading for your host is turned on. To access the hyperthreading option, go to the host's **Summary** tab and select **CPUs** under Hardware.

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 virtual machines, providing capabilities similar to traditional symmetric multiprocessing (SMP) systems. The VMkernel intelligently manages processor time to guarantee that the load is spread smoothly across processor cores in the system. Every 2 to 40 milliseconds (depending on the socket-core-thread topology), the VMkernel looks to migrate vCPUs from one logical processor to another to keep the load balanced. The VMkernel does its best to schedule virtual machines with multiple vCPUs on two different cores rather than on two logical processors on the same core. But if necessary, the VMkernel can map two vCPUs from the same virtual machine to threads on the same core.

If a logical processor has no work, it is put into a halted state. This action frees its execution resources and allows the virtual machine running on the other logical processor on the same core to use the full execution resources of the core. Because the VMkernel scheduler properly accounts for this halt time, a virtual machine running with the full resources of a core is charged more than a virtual machine running on a half core. This approach to processor management ensures that the server does not violate the ESXi resource allocation rules.

https://t.me/learningnets

8-13 Review of Learner Objectives

You should be able to meet the following objectives:

- Discuss CPU and memory concepts in a virtualized environment
- Describe what overcommitment of a resource means
- Identify additional technologies that improve memory usage
- Describe how VMware vSphere® Virtual Symmetric Multiprocessing works and how hyperthreading is used by the VMkernel

https://t.me/learningnets



https://t.me/learningnets

8-15 Learner Objectives

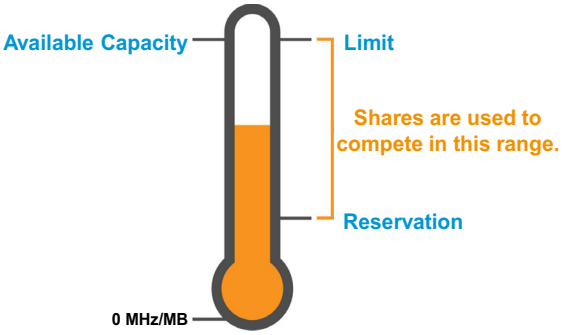
By the end of this lesson, you should be able to meet the following objectives:

- Assign share values for CPU and memory resources
- Describe how virtual machines compete for resources
- Create a resource pool
- Set resource pool attributes
- Establish CPU and memory reservations and limits
- Describe expandable reservations
- Use a vApp to manage virtual machines

https://t.me/learningnets

8-16 Shares, Limits, and Reservations

A virtual machine powers on only if its reservation can be guaranteed.



Because virtual machines simultaneously use the resources of an ESXi host, resource contention can occur.

For proper resource management, vSphere has mechanisms to enable less, more, or an equal amount of access to a defined resource. vSphere also prevents a virtual machine from consuming large amounts of a resource and grants a guaranteed amount of a resource to a virtual machine whose performance is not adequate or requires a certain amount of a resources to run properly.

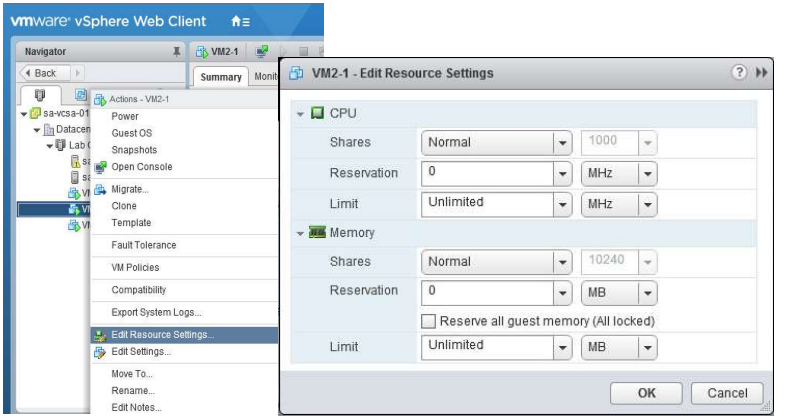
When host memory or CPU is overcommitted, a virtual machine's allocation target is somewhere between its specified reservation and specified limit, depending on the virtual machine's shares and the system load. vSphere uses a share-based allocation algorithm to achieve efficient resource use for all virtual machines and to guarantee a given resource to the virtual machines that need it most. Three configurable parameters control a virtual machine's access to a given resource:

- Share: A share is a value that specifies the relative priority or importance of a virtual machine's access to a given resource.
- Limit: Limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a virtual machine.
- Reservation: A reservation specifies the guaranteed minimum allocation for a virtual machine.

https://t.me/learningnets

8-17 Defining Resource Settings for Individual VMs

When available resource capacity does not meet the demands of the resource consumers, you might need to customize the amounts of resources that are allocated to VMs.



You must understand how shares, limits, and reservations work. If you overprovision your ESXi host on memory and CPU, you need a tool to ensure that the right machines get the correct amount of resources. Shares, limits, and reservations serve this purpose.

A reservation is a guarantee of either memory or CPU for a virtual machine. You define the reservation in MB or MHz.

Shares specify the relative importance of a VM or resource pool. If a VM has twice as many shares of a resource as another VM, it is entitled to consume twice that resource when these two VMs are competing.

A limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a VM.

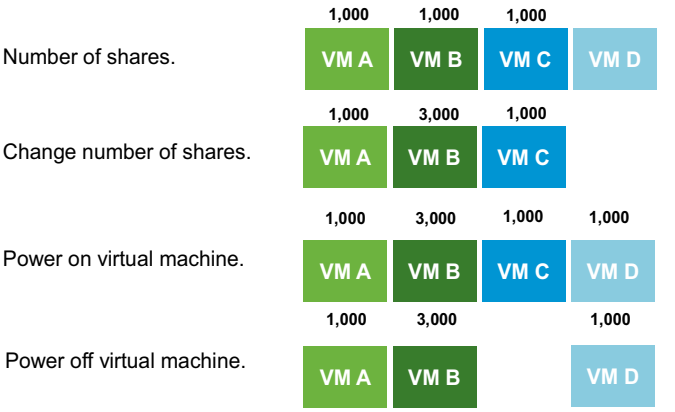
In most cases, it is not necessary to specify a limit. Specifying limits has the following benefits and drawbacks:

- Benefits: Assigning a limit is useful if you start with a small number of virtual machines and want to manage user expectations. The performance deteriorates as you add more virtual machines. You can simulate having fewer resources available by specifying a limit.
- Drawbacks: You might waste idle resources if you specify a limit. The system does not allow virtual machines to use more resources than the limit, even when the system is underused and idle resources are available. Specify the limit only if you have good reasons for doing so.

https://t.me/learningnets

8-18 Resource Sharing by VMs

Virtual machines are resource consumers. The default resource settings that you assign during creation work well for most machines.



The proportional share mechanism applies to CPU, memory, storage I/O, and network I/O allocation. The mechanism operates only when virtual machines are contending for the same resource.

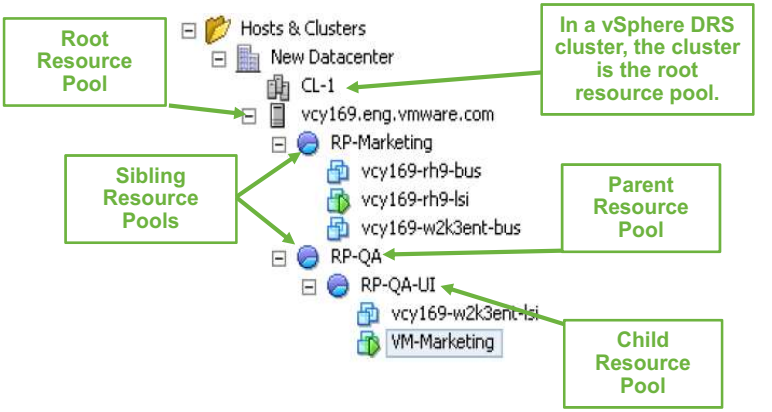
Shares guarantee that a virtual machine is given a certain amount of a resource (CPU, RAM, storage I/O, or network I/O). For example, consider the third row of virtual machines on the slide, where virtual machine D has been powered on with 1,000 shares. Before it was powered on, a total of 5,000 shares were available, but virtual machine D's addition increases the total shares to 6,000. The result is that the other virtual machines decline in value. But each virtual machine's share value still represents a minimum guarantee. Virtual machine A is still guaranteed one-sixth of the resource because it owns one-sixth of the shares.

You can add shares to a virtual machine while it is running, and it will get more access to that resource (assuming competition for the resource). When you add a virtual machine, it gets shares too. The virtual machine's share amount factors into the total number of shares, but the existing virtual machines are guaranteed not to be starved for the resource. When you delete or power off a virtual machine, fewer total shares remain, so the surviving virtual machines get more access.

https://t.me/learningnets

8-19 About Resource Pools

A resource pool is a logical abstraction of hierarchically managed CPU and memory resources.



A resource pool is a logical abstraction for managing resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources. In the on-slide example, RP-QA is the parent resource pool for RP-QA-UI. RP-Marketing and RP-QA are siblings.

A resource pool allows you as the administrator to divide and allocate resources to virtual machines and other resource pools. A resource pool allows you to control the aggregate CPU and memory resources of the compute resource. The compute resource can be either a standalone host or a vSphere DRS cluster. Resource pools are also used to delegate privileges to other users and groups.

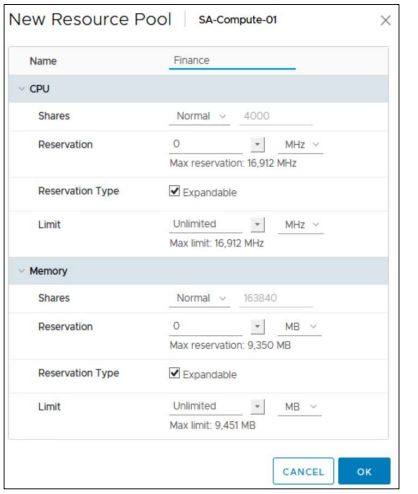
The topmost resource pool is called the root resource pool. Each standalone host and each vSphere DRS cluster has an (invisible) root resource pool that groups the resources of that host or cluster. The root resource pool does not appear, because the resources of the host (or cluster) and the root resource pool are always the same.

https://t.me/learningnets

8-20 Resource Pool Attributes

You can create a child resource pool of any ESXi host, resource pool, or vSphere DRS cluster:

- Shares: Low, Normal, High, Custom
- Reservations: In MHz or GHz, MB or GB
- Limit:
  - In MHz or GHz, MB or GB.
  - Unlimited access, by default, up to maximum amount of resource accessible.
- Reservation type:
  - Expandable selected: Virtual machines and subpools can draw from this pool's parent.
  - Expandable deselected: Virtual machines and subpools can draw only from this pool, even if its parent has free resources.



Like virtual machines, a resource pool has reservation, limit, and share values for CPU and memory resources.

Resource pools also have a setting that enables reservations to be expanded. The Expandable Reservation attribute allows a resource pool that cannot satisfy a reservation request to search through its hierarchy to find unreserved capacity to satisfy the reservation request.

https://t.me/learningnets

8-21 Reasons to Use Resource Pools

Resource pools can provide these benefits:

- Flexible hierarchical organization
- Isolation between pools and sharing in pools
- Access control and delegation
- Separation of resources from hardware
- Management of sets of virtual machines running a multitier service
- Ability to prioritize virtual machine workloads

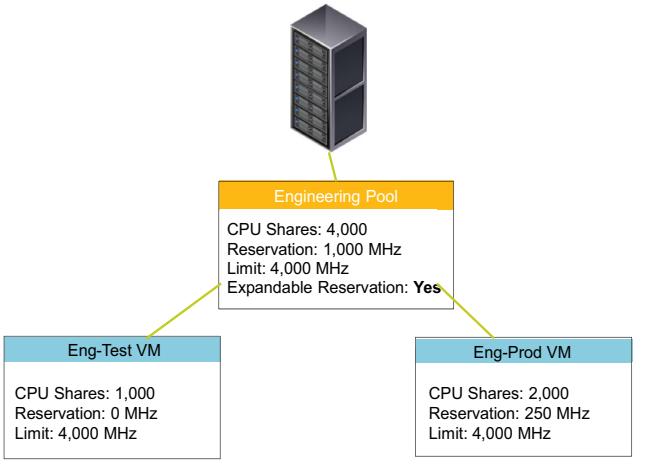
With resource pools, you can delegate control over resources of a standalone host or a vSphere DRS cluster. Using resource pools can result in the following benefits:

- Flexible hierarchical organization: Add, remove, or reorganize resource pools or change resource allocations as needed.
- Isolation between pools, sharing in pools: Top-level administrators can make a pool of resources available to a department-level administrator.
- Access control and delegation: Virtual machine creation and management are performed within the boundaries of the resources to which the resource pool is entitled. Delegation is usually done with permission settings.
- Separation of resources from hardware: If you are using vSphere DRS clusters, the resources of all hosts are always assigned to the cluster.
- Management of sets of virtual machines running a multitier service: Group virtual machines for a multitier service in a resource pool.

https://t.me/learningnets

8-22 Resource Pool Example

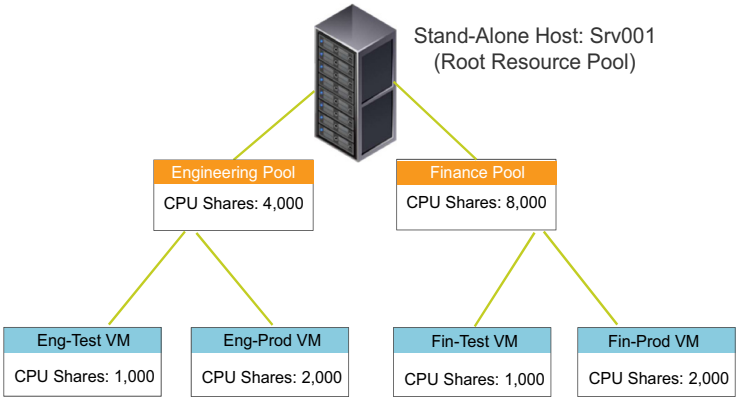
This example shows where resource attributes are set on a resource pool.



https://t.me/learningnets

8-23 Resource Pool Example: CPU Shares

In this example, the Finance resource pool has twice as many CPU shares as the Engineering resource pool. The Finance resource pool is entitled to twice as many CPU resources as the Engineering resource pool.



Resource pools can be organized hierarchically. The root resource pool is the topmost resource pool. The root resource pool includes the sum for all CPUs (in megahertz) and the sum of all the installed RAM (in megabytes) available in the compute environment (standalone host or cluster).

On the slide, the root resource pool is a standalone host named Svr001. It has 12,000 MHz of CPU and 4 GB of RAM, available for use by other resource pools or virtual machines.

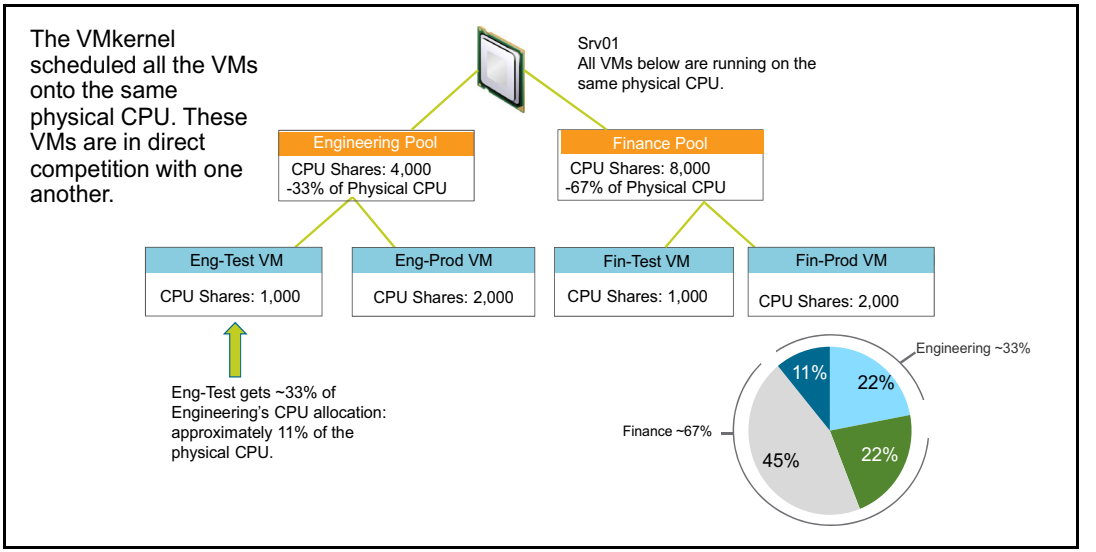
Except for the root resource pool, every resource pool has a parent resource pool. A resource pool might contain child resource pools or only virtual machines that are powered on in it.

A child resource pool is used to allocate resources from the parent resource pool for the child's consumers. Administrative control can also be delegated to individuals or organizations. A child resource pool cannot exceed the capacity of the parent resource pool. Creating a child pool reserves resources from the parent pool, whether or not virtual machines in the child pool are powered on.

Shares specify the relative priority or importance of either a resource pool or virtual machine. If a resource pool has twice as many shares of a resource as another resource pool, it is entitled to consume twice as much of that resource. The same thing can be applied to virtual machines.

https://t.me/learningnets

8-24 Resource Pool Example: CPU Contention



The Engineering pool gets 33 percent of that CPU and splits its allotment between virtual machines Eng-Test and Eng-Prod. Likewise, the Finance pool gets 67 percent of that CPU and splits its allotment between virtual machines Fin-Test and Fin-Prod. A virtual machine's resource settings are constrained by the resources of the resource pool to which the virtual machine belongs.

The virtual machine Eng-Test gets approximately 33 percent of the CPU allocation of the Engineering resource pool [1,000/(1,000+2,000)]. This figure is equal to about 11 percent of the physical CPU (33 percent of 33 percent equals approximately 11 percent). Each of the virtual machines gets a percentage of the physical CPU allocated to its resource pool that is based on its individual share allocation.

The example on the slide uses approximations to explain how the number of shares affects the amount of CPU allocated to a virtual machine.

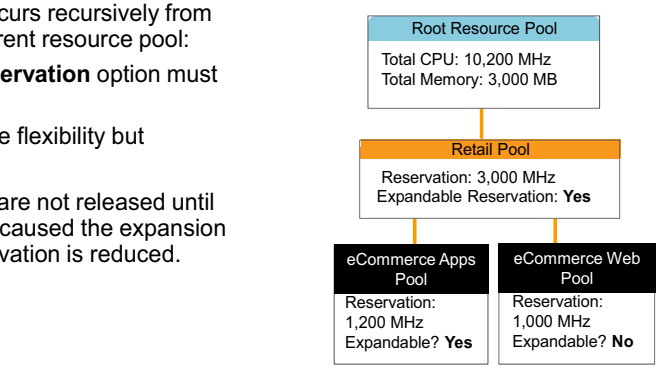
https://t.me/learningnets

8-25 About Expandable Reservations

Borrowing resources occurs recursively from the ancestors of the current resource pool:

- The **Expandable Reservation** option must be enabled.
- This option offers more flexibility but less protection.

Expanded reservations are not released until the virtual machine that caused the expansion is shut down or its reservation is reduced.



A mismanaged or badly sized expandable reservation might claim all unreserved capacity.

Expandable reservation allows a resource pool that cannot satisfy a reservation request to search through its hierarchy to find unreserved capacity to satisfy the reservation request.

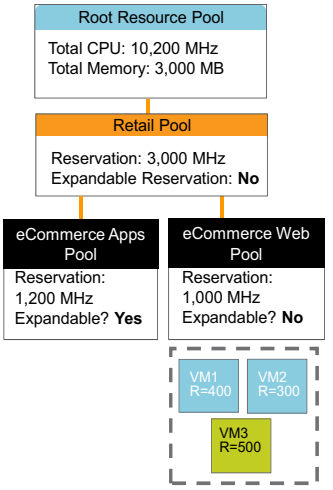
On the slide, the child resource pool eCommerce Apps has the **Expandable Reservation** option enabled. The reservation of a child resource pool cannot exceed that of its parent. The search for unused resources goes through the ancestry of the root resource pool or to the first resource pool where the **Expandable Reservation** option is not enabled.

Use expandable reservation carefully. A single child resource pool can use all of its parent's available resources, leaving nothing directly available for other child resource pools. This can happen if virtual machines are improperly moved around to get them to start based upon resource reservations.

You might want to disable the **Expandable Reservation** option when you are giving a fixed amount of resources to a group. For example, an IT administrator whose customers are different organizations in the company that have paid for a fixed amount of CPU and memory resources might want to disable the **Expandable Reservation** option.

https://t.me/learningnets

8-26 Example of Expandable Reservation (1)



eCommerce resource pools reserve 2,200 MHz of the 3,000 MHz that the Retail pool has reserved.

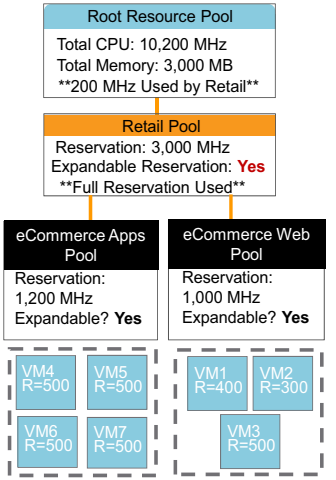
Power on virtual machines in the eCommerce Web pool.

With the **Expandable Reservation** option disabled on the eCommerce Web pool, VM3 cannot be started with a reservation of 500 MHz:

- Lower the virtual machine reservation.
- Enable the **Expandable Reservation** option.
- Increase the eCommerce Web pool's reservation.

https://t.me/learningnets

8-27 Example of Expandable Reservation (2)



Enable the **Expandable Reservation** option on the eCommerce Web pool.

The system considers the resources available in the child resource pool and its direct parent resource pool.

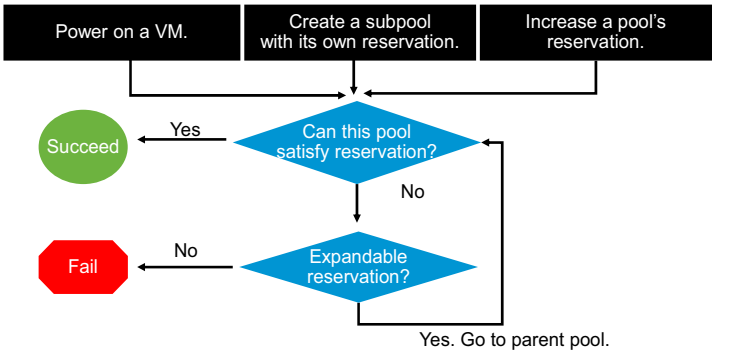
The virtual machine's reservation is charged against the reservation for eCommerce Web pool.

eCommerce Web pool's reservation is charged against the reservation for Retail.

https://t.me/learningnets

8-28 Admission Control for CPU and Memory Reservations

Admission control is used to ensure that you do not allocate resources that are not available.



Certain operations must satisfy admission control:

- Powering on a virtual machine
- Creating a resource pool with its own reservation
- Increasing a resource pool's reservation

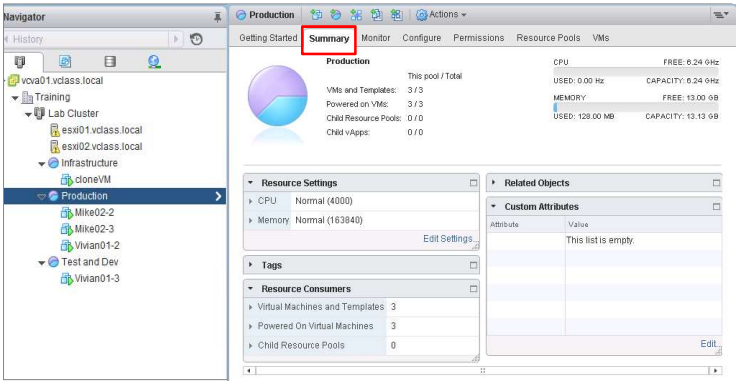
If the object (virtual machine or resource pool) resides in a resource pool with an expandable reservation, the parent of the current resource pool is consulted, if necessary, to satisfy the reservation.

Reservations cannot be overcommitted.

https://t.me/learningnets

8-29 Resource Pool Summary Tab

The resource pool's **Summary** tab displays information that applies to the host machine and its resources.



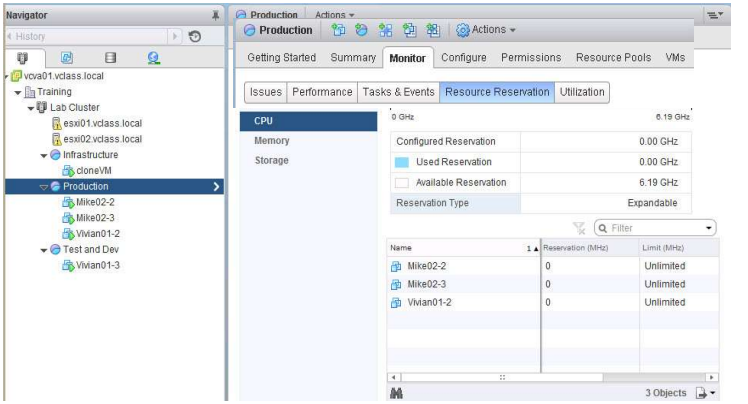
Consider these important points about the **Summary** tab:

- The Resource Settings pane displays CPU and Memory share settings.
- The Resource Consumers pane displays the number of virtual machines, number of powered-on virtual machines, and child resource pools that are in the selected resource pool.
- The Tags pane displays tags assigned to objects that reside in the resource pool.

https://t.me/learningnets

8-30 Resource Reservation Tab

On the **Resource Reservation** tab, you can view information about a resource pool's CPU, memory, and storage resources.



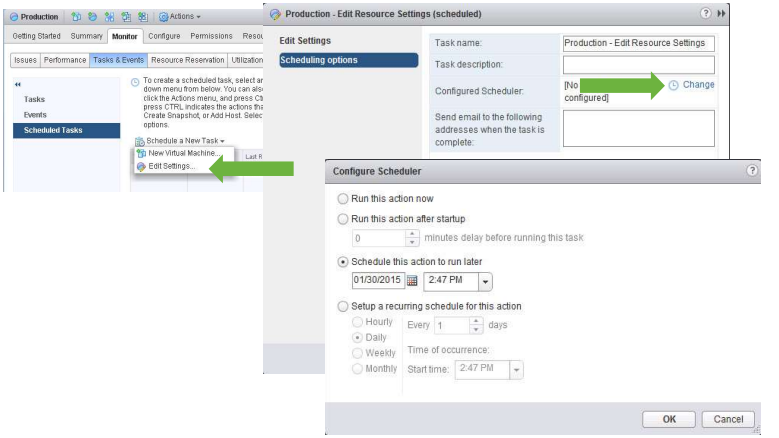
The following information is displayed for the pool's CPU and memory resources:

- Amount of CPU and memory reservation that is configured
- Type of reservation used (expandable or fixed)
- Amount of reservation in use by virtual machines and child pools
- Amount of CPU and memory that is available to be reserved

https://t.me/learningnets

8-31 Scheduling Changes to Resource Settings

You can schedule a task to change the resource settings of a resource pool or virtual machine.



You can configure the task to change the shares, reservation, and limit for CPU, memory, or for both so that you can accommodate changing business priorities.

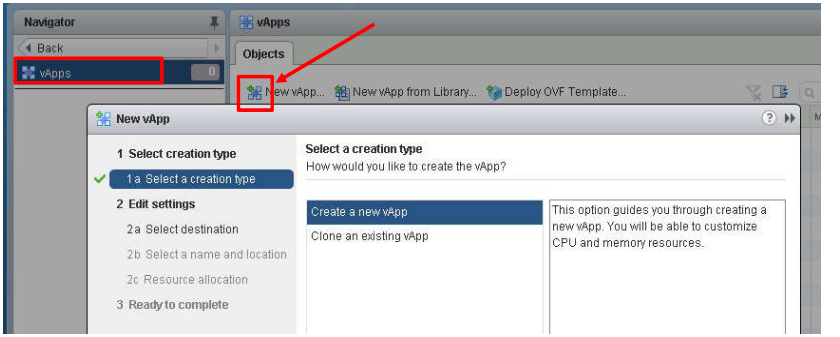
For example, at the end of each quarter, you can give financial applications more CPU and memory resources than internal applications. In a retail organization, you can double the CPU and memory resource reservations for the virtual machines running the online store applications during the month of December.

https://t.me/learningnets

8-32 Managing Virtual Machines with vApps

A vApp is an object in the vCenter Server inventory:

- A vApp is a container for one or more virtual machines.
- A vApp can be used to package and manage multitiered applications.



You can use vSphere as a platform for running applications, such as multitiered applications. The applications can be packaged to run directly on top of vSphere.

In vSphere Web Client, a vApp is represented in the Hosts and Clusters view and in the VM and Templates view.

A vApp is a container for one or more virtual machines or vApps. A vApp shares functionality with virtual machines. A vApp can power on and power off and it can be cloned. The distribution format for a vApp can be either Open Virtualization Format (OVF) or Open Virtualization Appliance (OVA). The differences between these formats are as follows:

- An OVF file is a collection of virtual machine files. The OVF file is an XML file that has information about the virtual disk files in the directory. When you export a virtual machine as an OVF file, a directory is created that has an OVF file and the VMDK files.
- OVA is the portable virtual machine format from XenSource, a third-party product. The OVA file is a single file that can be considered an archive, like a ZIP file, of all the files that belong to the OVF directory.

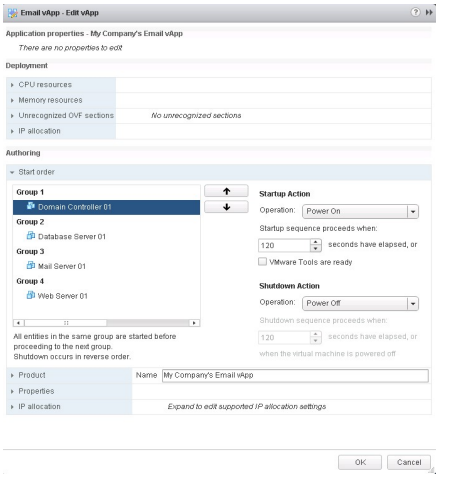
https://t.me/learningnets

8-33 vApp Characteristics

You can configure several vApp settings by right-clicking the vApp:

- CPU and memory allocation
- IP allocation policy

You can also configure the virtual machine startup and shutdown order.



After creating the vApp, you can modify certain vApp settings:

- Resource allocation: Determines how CPU and memory should be allocated for the vApp
- IP allocation policy: Determines how IP addresses are allocated for the vApp
- Advanced settings: Includes product and vendor information, and custom properties

You can change the order in which virtual machines (and nested vApps) in a vApp start up and shut down. You change this order by assigning virtual machines to groups. All entities in the same group are started before those in the next group. Shutdown is done in the reverse order. You can also specify delays and actions performed at startup and shutdown.

On the slide, virtual machine Domain Controller 01 is started first. Virtual machine Database Server 01 starts after 120 seconds.

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

https://t.me/learningnets

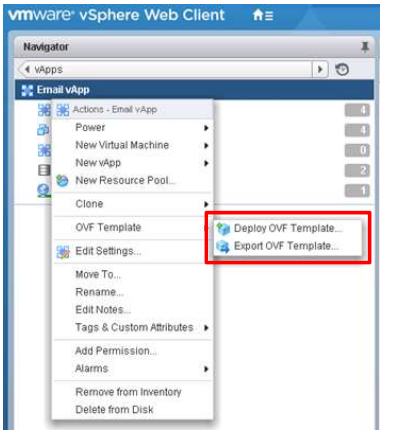
8-34 Exporting and Deploying vApps

When exporting a vApp as an OVF or OVA template:

- You can share the vApp with others.
- You can use the vApp for archive purposes.

When deploying an OVF template:

- You can deploy multitier vApps.
- You can deploy OVF from VMware Virtual Appliance Marketplace.



https://t.me/learningnets

8-35 Lab 13: Managing Resource Pools

In vCenter Server, create and use resource pools on an ESXi host

1. Create CPU Contention
2. Create Resource Pools
3. Verify Resource Pool Functionality

https://t.me/learningnets

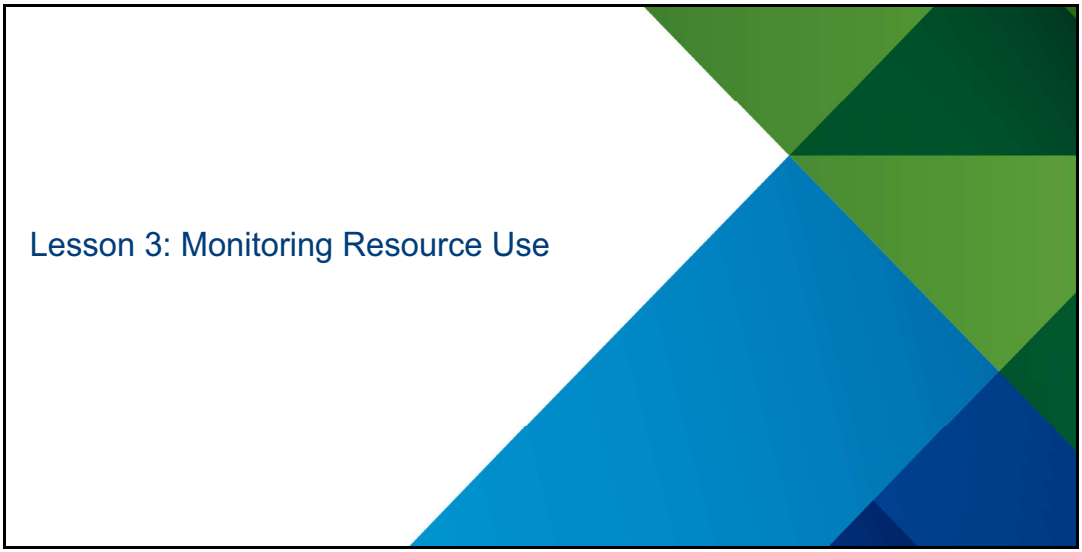
8-36 Review of Learner Objectives

You should be able to meet the following objectives:

- Assign share values for CPU and memory resources
- Describe how virtual machines compete for resources
- Create a resource pool
- Set resource pool attributes
- Establish CPU and memory reservations and limits
- Describe expandable reservations
- Use a vApp to manage virtual machines

•

https://t.me/learningnets



https://t.me/learningnets

8-38 Learner Objectives

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

- Use the performance-tuning methodology and resource monitoring tools
- Use performance charts to view and improve performance
- Monitor the key factors that can affect the virtual machine's performance: CPU, memory, disk, and network bandwidth use

https://t.me/learningnets

8-39 Performance-Tuning Methodology

Follow these best practices for performance-tuning your vSphere infrastructure:

- Assess performance:
  - Use appropriate monitoring tools.
  - Record a numerical benchmark before changes.
- Identify the limiting resource.
- Make more resources available:
  - Allocate more resources.
  - Reduce competition.
  - Log your changes.
- Benchmark again.

Do not make casual changes to production systems.

The best practice for performance tuning is to take a logical step-by-step approach:

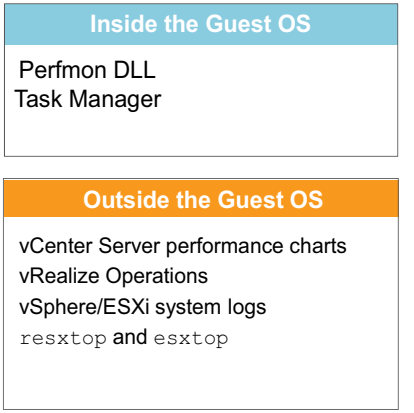
1. For a complete view of the performance situation of a virtual machine, use monitoring tools in the guest operating system and in vCenter Server. Record benchmarks before you make changes.
2. Identify the resource that the virtual machine relies on the most. That resource is most likely to affect the virtual machine's performance if the virtual machine is constrained by it.
3. Give a virtual machine more resources or decrease the resources of other virtual machines.
4. After making more of the limiting resource available to the virtual machine, take another benchmark and record changes.

Be extra cautious when making changes to production systems because a change might have a negative effect on the performance of the virtual machines.

https://t.me/learningnets

8-40 Resource-Monitoring Tools

Many resource-monitoring and performance-monitoring tools are available to administrators to use with vSphere.



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.

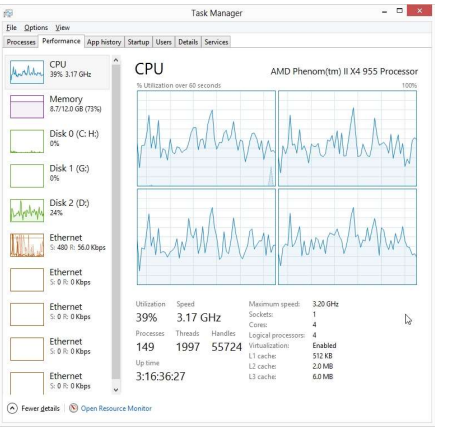
A partial list of these resource-monitoring tools is shown on the slide.

https://t.me/learningnets

8-41 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



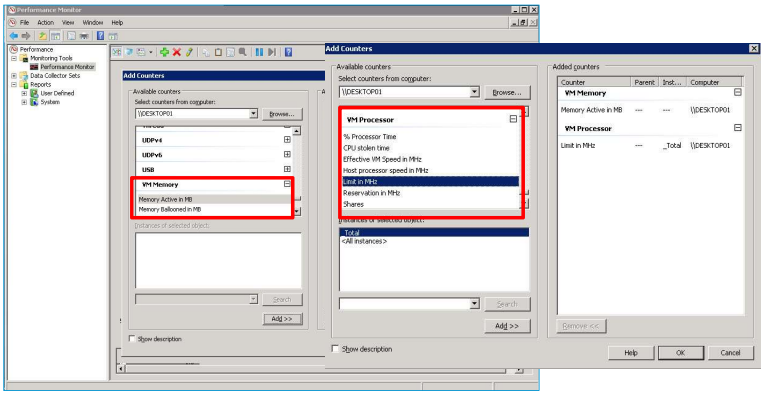
Windows Task Manager helps you measure CPU and memory use in the guest operating system.

The measurements that you take with tools in the guest operating system reflect resource usage of the guest operating system, not necessarily of the virtual machine itself.

https://t.me/learningnets

8-42 Using Perfmon to Monitor VM Resources

The Perfmon DLL in VMware Tools provides virtual machine processor and memory objects to access host statistics inside a virtual machine.



VMware Tools includes a library of functions called the Perfmon DLL. Perfmon allows you to access key host statistics in a guest virtual machine. The Perfmon performance objects (VM Processor and VM Memory) allow you to view actual CPU and memory usage alongside 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 virtual machine'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.

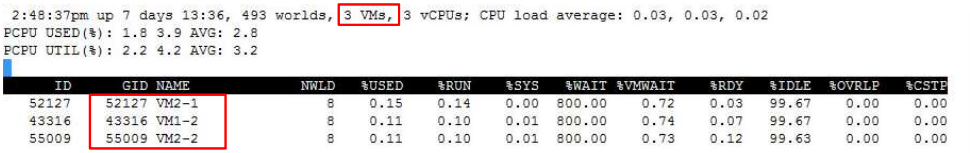
https://t.me/learningnets

8-43 Using esxtop to Monitor VM Resources

`esxtop` is the primary real-time performance monitoring tool for vSphere:

- This tool can be run from the host's local ESXi Shell as `esxtop`.
- This tool can also be run remotely from vSphere CLI as `resxtop`.
- `esxtop` works like the `top` performance utility found in Linux operating systems.

In this example, you enter lowercase `c` and uppercase `V` to view CPU metrics for virtual machines.



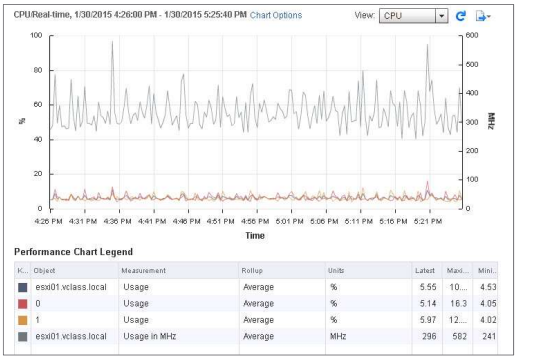
You can run the `esxtop` utility by using vSphere ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

https://t.me/learningnets

8-44 Monitoring Inventory Objects with Performance Charts

The vSphere statistics subsystem collects data on the resource usage of inventory objects:

- Counters and metric groups
- Collection levels and collection intervals
- Data availability



Data on a wide range of metrics is collected at frequent intervals, processed, and archived in the vCenter Server database.

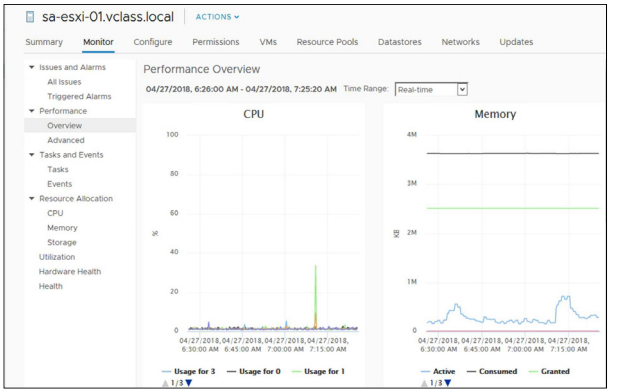
You can access statistical information through command-line monitoring utilities or by viewing performance charts in vSphere Web Client and vSphere Client.

- Counters and metric groups: vCenter Server systems and hosts use data counters to query for statistics. A data counter is a unit of information relevant to a given inventory object or device. Each counter collects data for a different statistic in a metric group. For example, the disk metric group includes separate data counters to collect data for disk read rate, disk write rate, and disk usage. Statistics for each counter are rolled up after a specified collection interval. Each data counter consists of several attributes that are used to determine the statistical value collected.
- Collection levels and collection intervals: Collection levels determine the number of counters for which data is gathered during each collection interval. Collection intervals determine the time period during which statistics are aggregated, calculated, rolled up, and archived in the vCenter Server database. Together, the collection interval and collection level determine how much statistical data is collected and stored in your vCenter Server database.
- Data availability: Real-time data appears in the performance charts only for hosts and virtual machines that are powered on. Historical data appears for all supported inventory objects, but might be unavailable during certain circumstances.

https://t.me/learningnets

8-45 Working with Overview Performance Charts

The overview performance charts display the most common metrics for an object in the inventory.



Host's Performance Charts Partial Overview Panel

Two types of VMware performance charts are available in vSphere Web Client: overview charts and advanced charts.

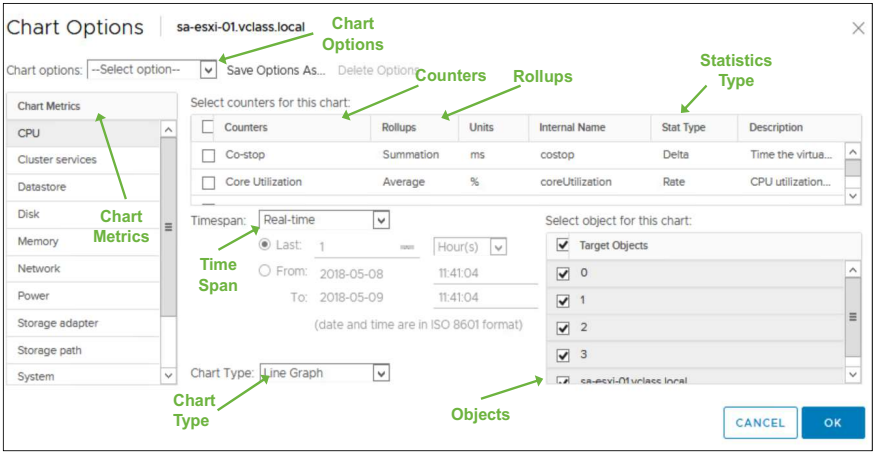
The overview performance charts show the performance statistics that VMware considers most useful for monitoring performance and diagnosing problems.

Depending on the object that you select in the inventory, the performance charts in the Overview panel provide you with a quick visual representation of how your host or virtual machine is doing. The example shows a partial view of the overview performance charts for an ESXi host.

https://t.me/learningnets

8-46 Working with Advanced Performance Charts

Advanced charts support data counters that are not supported in other performance charts.



In vSphere Web Client, you can customize the appearance of the advanced performance charts.

Advanced charts have the following features:

- More information: 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.
- Export data to a spreadsheet.
- Save data to an image file or a spreadsheet.

https://t.me/learningnets

8-47 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 generated for the past hour at a 20-second specificity. Historical information is information generated for the past day, week, month, or year, at varying specificities.

By default, vCenter Server has four collection intervals: day, week, month, and year. Each interval specifies a length of time that statistics are archived in the vCenter Server database. You can configure which intervals are enabled and for what period of time. You can also configure the number of data counters used during a collection interval by setting the collection level. Together, the collection interval and the collection level determine how much statistical data is collected and stored in your vCenter Server database. For example, using the table, past-day statistics show one data point every 5 minutes, for a total of 288 samples. Past-year statistics show 1 data point per day, or 365 samples.

Real-time statistics are not stored in the database. They are stored in a flat file on ESXi hosts and in memory on vCenter Server systems. ESXi hosts collect real-time statistics only for the host or the virtual machines available on the host. Real-time statistics are collected directly on an ESXi host every 20 seconds. If you query for real-time statistics, vCenter Server queries each host directly for the data. vCenter Server does not process the data at this point. vCenter Server only passes the data to vSphere Web Client and vSphere Client.

On ESXi hosts, the statistics are kept for 30 minutes, after which 90 data points are collected. The data points are aggregated, processed, and returned to vCenter Server. vCenter Server then archives the data in the database as a data point for the day collection interval.

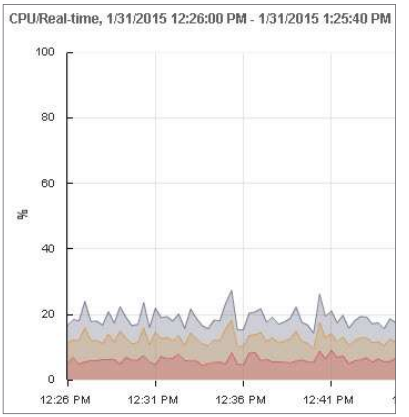
To ensure that performance is not impaired when collecting and writing the data to the database, cyclical queries are used to collect data counter statistics. The queries occur for a specified collection interval. At the end of each interval, the data calculation occurs.

https://t.me/learningnets

8-48 About Chart Types

Depending on the metric type and object, performance metrics are displayed in different types of charts:

- Line chart: Each instance is shown separately.
- Bar chart: Each instance is a bar in the chart.
- Pie chart: Each instance is a slice in a circular pie.
- Stacked chart: Graphs are stacked on top of one another.



You can use different chart types to change how metrics appear:

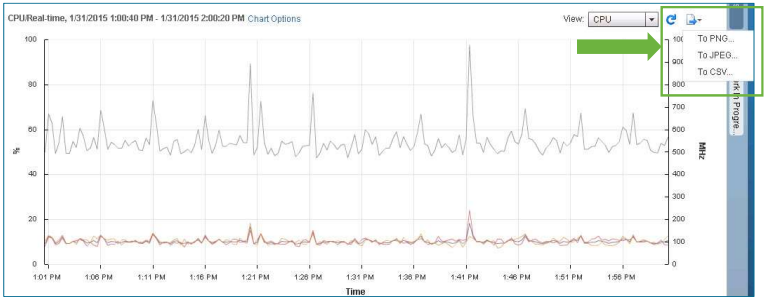
- Line chart: Displays metrics for a single inventory object. The data for each performance counter is plotted on a separate line in the chart. For example, a network chart for a host can contain two lines: one showing the number of packets received and one showing the number of packets transmitted.
- Bar chart: Displays storage metrics for datastores in a selected data center. Each datastore is represented as a bar in the chart. Each bar displays metrics based on the file type: virtual disks, snapshots, swap files, and other files.
- Pie chart: Displays storage metrics for a single object, based on the file types or virtual machines. For example, a pie chart for a datastore can display the amount of storage space occupied by the virtual machines taking up the largest space.
- Stacked chart: Displays metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term Other. For example, a host's stacked CPU usage chart displays CPU usage metrics for the five virtual machines on the host that are consuming the most CPU resources. The Other amount contains the total CPU usage of the remaining virtual machines. The metrics for the host itself are displayed in separate line charts. Stacked charts are useful in comparing resource allocation and usage across multiple hosts or virtual machines. By default, the 10 child objects with the highest data counter values appear.

https://t.me/learningnets

8-49 Saving Charts

You click the **Save Chart** icon above the graph to save performance chart information.

You can save information in PNG, JPEG, and CSV formats.



In vSphere Web Client, you can save data from the advanced performance charts to a file in various graphics formats or in Microsoft Excel format. When you save a chart, you select the file type and save the chart to the location of your choice.

https://t.me/learningnets

8-50 About Objects and Counters

The performance charts graphically display CPU, memory, disk, network, and storage metrics for devices and entities managed by vCenter Server.

Objects are instances or aggregations of devices:

- Examples: vCPU0, vCPU1, vmhba1:1:2, and aggregate over all NICs

Counters identify which statistics to collect.

- Examples:
  - CPU: Used time, ready time, usage (%)
  - NIC: Network packets received
  - Memory: Memory swapped

vCenter Server allows the user to determine how much or how little information about a specific device type is displayed. You can control the amount of information a chart displays by selecting one or more objects and counters.

An object refers to an instance for which a statistic is collected. For example, you might collect statistics for an individual CPU, 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

8-51 About Statistics Types

The statistics type is the unit of measurement used during the statistics interval.

Statistics Type	Description	Example
Rate	Value over the current interval	CPU use (MHz)
Delta	Change from previous interval	CPU ready time
Absolute	Absolute value, independent of interval	Memory active

The statistics type refers to the measurement used during the statistics interval and is related to the unit of measurement.

The statistics type is one of the following:

- Rate: Value over the current statistics interval
- Delta: Change from 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

8-52 About Rollup

Rollup is the conversion function between statistics intervals:

- 5 minutes of past-hour statistics are converted to 1 past-day value:
  - Fifteen 20-second statistics are rolled up into a single value.
- 30 minutes of past-day statistics are converted to 1 past-week value:
  - Six 5-minute statistics are rolled up into a single value.
- Other rollup types: Minimum, Maximum

Rollup Type	Conversion Function	Sample Statistic
Average	Average of data points	CPU use (average)
Summation	Sum of data points	CPU ready time (milliseconds)
Latest	Last data point	Uptime (days)

When looking at different historical intervals, data is displayed at different specificities. Past-hour statistics are shown at a 20-second specificity, and past-day statistics are shown at a 5-minute specificity. The averaging that is done to convert from one time interval to another is called rollup.

Different rollup types are available. The rollup type determines the type of statistical values returned for the counter:

- Average: The data collected during the interval is aggregated and averaged.
- Minimum: The minimum value is rolled up.
- Maximum: The maximum value is rolled up.

  The minimum and maximum values are collected and displayed only in collection level 4. Minimum and maximum rollup types are used to capture peaks in data during the interval. For real-time data, the value is the current minimum or current maximum. For historical data, the value is the average minimum or average maximum.

  For example, the following information for the CPU usage chart shows that the average is collected at collection level 1 and 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)

https://t.me/learningnets

Statistics levels include:

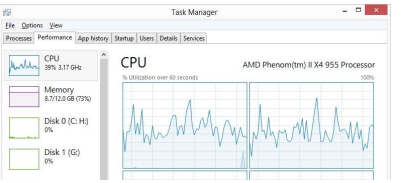
- Summation: The data collected is summed. The measurement displayed in the performance chart represents the sum of data collected during the interval.
- Latest: The data collected during the interval is a set value. The value displayed in the performance chart represents the current value.

For example, if you look at the CPU Used counter in a CPU performance chart, the rollup type is summation. This means that for a given 5-minute interval, the sum of all of the 20-second samples in that interval is represented.

https://t.me/learningnets

8-53 Interpreting Data from Tools

vCenter Server monitoring tools and guest OS monitoring tools provide different points of view.



Task Manager in Guest Operating System



CPU Usage Chart for Host

The key to interpreting performance data is to observe the range of data from the guest operating system, the virtual machine, and the host's perspective.

The CPU usage statistics in Task Manager, for example, do not give you the complete picture. You should also view CPU usage for the virtual machine and the host on which the virtual machine is located. Use the performance charts in vCenter Server to view this data.

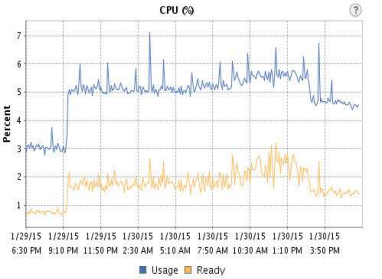
https://t.me/learningnets

8-54 CPU-Constrained VMs

If CPU usage is continuously high, the virtual machine is constrained by the CPU. However, the host might have enough CPU for other virtual machines to run.

Multiple virtual machines are constrained by the CPU if the following conditions are present:

- High CPU usage in the guest operating system
- Relatively high CPU ready values for the virtual machines



Single Virtual Machine CPU Usage

To check if a virtual machine is being constrained by CPU resources, check CPU usage in the guest operating system using, for example, Task Manager.

If CPU usage is high, check the virtual machine's CPU use. In vSphere Web Client or vSphere Client, select the virtual machine in the inventory, click the **Monitor** tab, and click **Performance**. Use either the overview charts or the advanced charts to view CPU usage. The slide displays an advanced chart tracking a virtual machine's CPU usage.

If a virtual machine's CPU usage remains high over a period of time, the virtual machine is constrained by CPU. Other virtual machines on the host might have enough CPU resources to satisfy their needs.

If more than one virtual machine is constrained by CPU, the key indicator is CPU ready time. Ready time refers to the interval when a virtual machine is ready to execute instructions but cannot because it cannot get scheduled onto a CPU. Several factors affect the amount of ready time:

- Overall CPU use: You are more likely to see ready time when usage is high because the CPU is more likely to be busy when another virtual machine becomes ready to run.
- Number of resource consumers (in this case, guest operating systems): When a host is running a larger number of virtual machines, the scheduler is more likely to need to queue a virtual machine behind virtual machines that are already running or queued.

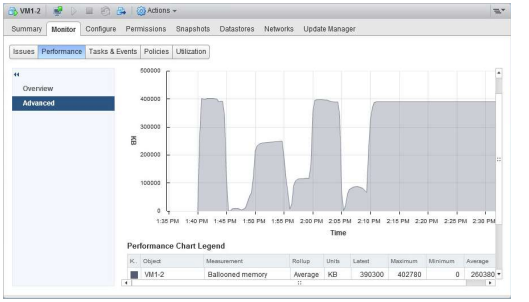
A good ready time value varies from workload to workload. To find a good ready time value for your workload, collect ready time data over time for each virtual machine. When you have this ready time data for each virtual machine, estimate how much of the observed response time is ready time. If the shortfalls in meeting response-time targets for the applications appear largely because of the ready time, take steps to address the excessive ready time.

https://t.me/learningnets

8-55 Memory-Constrained VMs

Check the virtual machine's ballooning activity to determine if the virtual machine is constrained for memory:

- If ballooning activity is high, this state might not be a problem if all virtual machines have sufficient memory.
- If ballooning activity is high and the guest OS is swapping, then the virtual machine is constrained for memory.

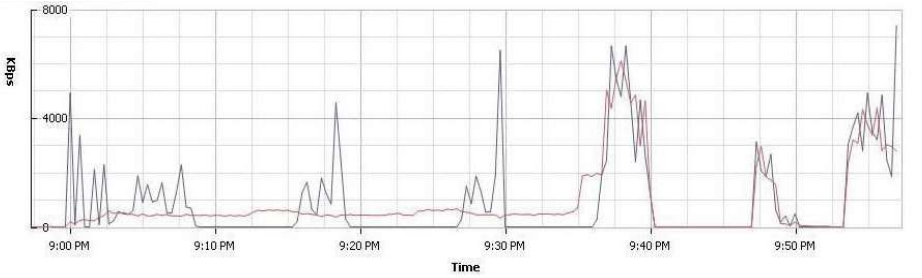


When a virtual machine experiences ballooning activity, some of the guest operating system's physical memory is being reclaimed from the virtual machine by the balloon driver. If a virtual machine experiences high ballooning values, this might not be a problem if the virtual machine continues to have the memory that it needs. But if a virtual machine experiences high ballooning activity over time and its guest operating system starts to page, the virtual machine might be constrained for memory.

https://t.me/learningnets

8-56 Memory-Constrained Hosts

If active host-level swapping is occurring, then the host memory is overcommitted.



Performance Chart Legend

K..	Object	Measurement	Rollup 1 ▲	Units
	sa-esxi-01.vclass.local	Swap in rate	Average	KBps
	sa-esxi-01.vclass.local	Swap out rate	Average	KBps

If multiple virtual machines are being constrained for memory, you see high ballooning activity and the guest operating systems paging. You also see the virtual machine itself being swapped in and out by the VMkernel. This serious situation indicates that the host memory is overcommitted and needs to be increased.

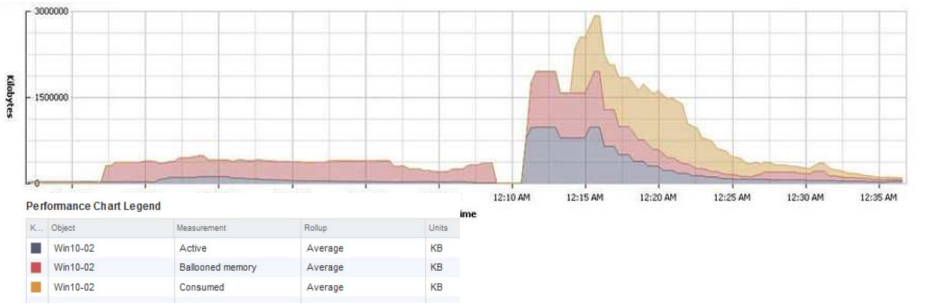
https://t.me/learningnets

8-57 Monitoring VM Active Memory

You can monitor for increases in active memory on the host.

Host active memory is active physical memory used by virtual machines and the VMkernel.

If the amount of active memory is high, virtual machines might become memory-constrained.



A general memory counter to monitor over time is a host's active memory counter. Host active memory refers to the amount of physical memory that is actively being used by virtual machines and the VMkernel. If the active memory of certain virtual machines is continuously high, this state might lead to those virtual machines being constrained by memory.

https://t.me/learningnets

8-58 Disk-Constrained VMs

Disk-intensive applications can saturate the storage or the path.

If you suspect that a virtual machine is constrained by disk access:

- Measure the throughput and latency between the virtual machine and storage.
- Use the advanced performance charts to monitor:
  - Read rate and write rate
  - Read latency and write latency

Disk performance problems are commonly caused by saturating the underlying physical storage hardware. You can use the vCenter Server advanced performance charts to measure storage performance at different levels. These charts can provide you with insight into a virtual machine's performance by allowing you to monitor everything, from the virtual machine's datastore to a specific storage path.

If you select a host object, you can view throughput and latency for a datastore, a storage adapter, or a storage path. The storage adapter charts are available only for Fibre Channel storage. The storage path charts are available for Fibre Channel and iSCSI storage, not for NFS.

If you select a virtual machine object, you can view throughput and latency for the virtual machine's datastore or specific virtual disk.

To monitor throughput, view the Read rate and Write rate counters. To monitor latency, view the Read latency and Write latency counters.

https://t.me/learningnets

8-59 Monitoring Disk Latency

To determine disk performance problems, monitor two disk latency data counters:

- Kernel command latency:
  - This counter is the average time spent in the VMkernel per SCSI command.
  - High numbers (greater than 2 or 3 ms) represent either an overworked array or an overworked host.
- Physical device command latency:
  - This counter is the average time the physical device takes to complete a SCSI command.
  - High numbers (greater than 15 or 20 ms) represent a slow or overworked array.

A reliable way to determine whether or not your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. Use the advanced performance charts to view these statistics. In particular, monitor the following counters:

- Kernel command latency: This data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0 through 1 milliseconds. If the value is greater than 4 milliseconds, the virtual machines on the ESXi host are trying to send more throughput to the storage system than the configuration supports.
- Physical device command latency: This data counter measures the average amount of time, in milliseconds, for the physical device to complete a SCSI command. Depending on your hardware, a number greater than 15 milliseconds indicates that the storage array might be slow or overworked.

https://t.me/learningnets

8-60 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 virtual machine is constrained by the network:

- Verify that VMware Tools is installed. The enhanced network drivers are available only if VMware Tools is installed.
- Measure the effective bandwidth between the virtual machine and its peer system.
- Check for dropped receive packets and dropped transmit packets.

Like disk performance problems, network performance problems are commonly caused by saturating a network link between client and server. Use a tool like Iometer, or a large file transfer, to measure the effective bandwidth.

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine if packets are being dropped, use the advanced performance charts to examine the droppedTx and droppedRx network counter values of a virtual machine.

In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. In some instances, large packets can result in high network latency. When network packets are small, more packets are transferred, but the network speed is slower because more CPU is required to process the data.

https://t.me/learningnets

8-61 Lab 14: Monitoring Virtual Machine Performance

Use the system monitoring tools to reflect the CPU workload

1. Create the CPU Workload
2. Use Performance Charts to Monitor CPU Utilization
3. Undo Changes Made to the Virtual Machines

https://t.me/learningnets

8-62 Review of Learner Objectives

You should be able to meet the following objectives:

- Use the performance-tuning methodology and resource monitoring tools
- Use performance charts to view and improve performance
- Monitor the key factors that can affect the virtual machine's performance: CPU, memory, disk, and network bandwidth use

https://t.me/learningnets



https://t.me/learningnets

8-64 Learner Objectives

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

- Create alarms with condition-based triggers
- Create alarms with event-based triggers
- View and acknowledge triggered alarms

https://t.me/learningnets

8-65 About Alarms

An alarm is a notification that occurs in response to selected events or conditions that occur with an object in the inventory.

Default alarms exist for various inventory objects. Many default alarms are defined for hosts and virtual machines.

You can create custom alarms for a wide range of inventory objects:

- Virtual machines
- Hosts, clusters, and data centers
- Datastores and datastore clusters
- Networks, distributed switches, and distributed port groups

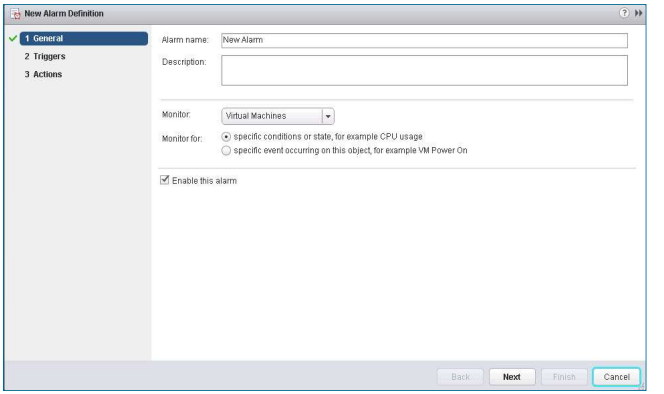
Alarms are notifications in response to selected events or conditions that occur with an object in the inventory. VMware provides a set of predefined alarms for most objects in the vCenter Server inventory. For example, alarms exist for host, virtual machine, and resource pool memory and CPU usage. You can also define custom alarms for virtual machines, hosts, clusters, data centers, datastores, datastore clusters, networks, distributed switches, and distributed port groups.

The predefined alarms are configurable. If the predefined alarms do not address the condition, state, or event that you want to monitor, define custom alarms instead of modifying predefined alarms.

https://t.me/learningnets

8-66 Alarm Settings

To monitor your environment, you can create and modify alarm definitions in vSphere Web Client.



On the General page, you name the alarm, give it a description, and give it an alarm type.

You also select what to monitor:

- Specific conditions or state: A condition-based alarm. You can create condition-based alarms for virtual machines, hosts, and datastores.
- Specific events occurring on this object: An event-based alarm. You can create event-based alarms for virtual machines, hosts, clusters, data centers, datastores, datastore clusters, networks, distributed virtual switches, and distributed virtual port groups.

You can also use the General page to enable or disable the alarm by selecting or deselecting the **Enable this alarm** check box.

https://t.me/learningnets

8-67 Alarm Triggers

An alarm requires a trigger:

- A condition or state trigger monitors the current condition or state:
  - Examples:
    - A virtual machine's current snapshot is above 2 GB in size.
    - A host is using 90 percent of its total memory.
    - A datastore is disconnected from all hosts.
- An event trigger monitors events:
  - Examples:
    - The health of a host's hardware has changed.
    - A license has expired in the data center.
    - A host has left the distributed switch.

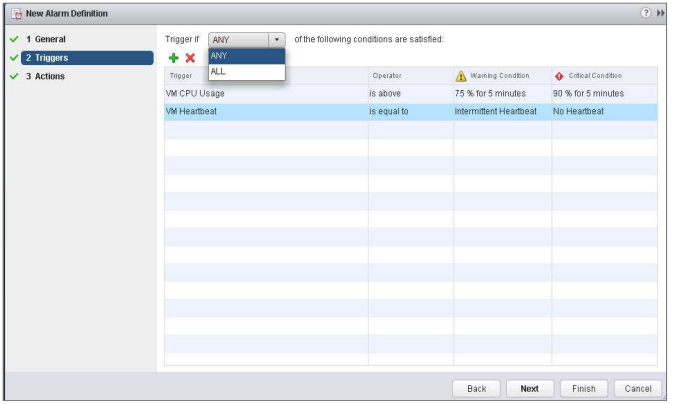
You configure alarm triggers to generate warnings and alerts when the specified criteria are met. Alarms have condition or state triggers and event triggers.

- Condition or state triggers: Monitor the current condition or state of virtual machines, hosts, and datastores. Conditions or states include power states, connection states, and performance metrics such as CPU and disk usage.
- Event triggers: Monitor events that occur in response to operations occurring with a managed object in the inventory or the vCenter Server system. For example, an event is recorded each time a virtual machine (which is a managed object) is cloned, created, deleted, deployed, and migrated.

https://t.me/learningnets

8-68 Configuring Condition or State Triggers

Condition or state triggers monitor metrics for a host, virtual machine, or datastore.



On the slide, you can configure a condition trigger so that the following cases are true:

- A virtual machine's CPU usage must be above 75 percent for more than 5 minutes to generate a warning.
- A virtual machine's CPU usage must be above 90 percent for more than 5 minutes to generate an alert.

Time periods are used to ensure that the metric conditions are valid and not caused by incidental spikes.

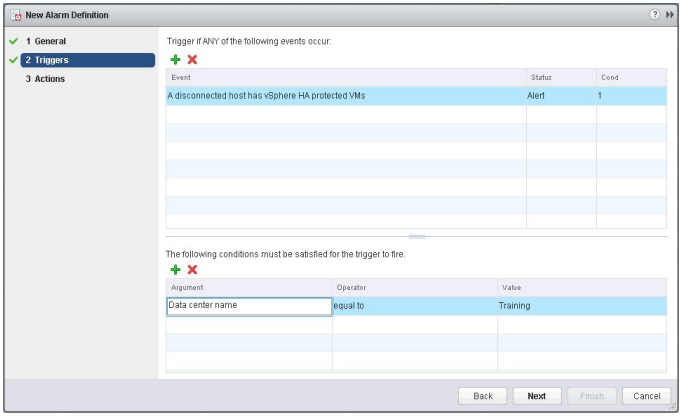
Also on the slide, you can configure a state trigger to generate an alert if a virtual machine has no heartbeat. When the triggering conditions are no longer true, a triggered alarm resets itself and no longer triggers.

If you add multiple triggers, you can choose to trigger the alarm if any one of the conditions is satisfied or if all of the conditions are satisfied.

https://t.me/learningnets

8-69 Configuring Event Triggers

Event triggers monitor the current state of a host, virtual machine, or datastore.



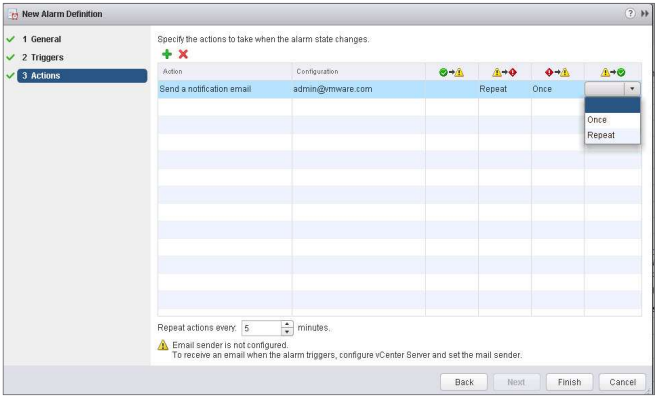
Event triggers do not rely on thresholds or durations. They use arguments, operators, and values to identify the triggering condition. In the example, the event trigger monitors the A disconnected host has vSphere HA protected VMs event.

Whenever host that is housing virtual machines to be protected by vSphere HA is disconnected, it triggers the alert. A condition has also been configured to trigger the alert only if the guest is in a data center named Training.

https://t.me/learningnets

8-70 Configuring Actions

You can define actions that the system performs when the alarm is triggered or changes status.



Alarms are composed of a trigger and an action. An action is the operation that occurs in response to the trigger, for example, sending an email notification to one or more administrators.

In the New Alarm Definition dialog box, use the **Actions** tab to specify actions to take when the alarm is triggered. Colors and shapes are used to denote alarm severity: a green circle is normal, a yellow triangle is a warning, and a red diamond is an alert.

You can set alarms to trigger when the state changes:

- From a green circle to a yellow triangle
- From a yellow triangle to a red diamond
- From a red diamond to a yellow triangle
- From a yellow triangle to a green circle

For every action, you can specify an option for each color transition:

- **Empty** indicates no interest in the transition.
- **Once** tells vCenter Server to perform the action only one time.
- **Repeat** tells vCenter Server to repeat the action until another color change occurs. The default is 5 minutes. The maximum is 2 days. The mentioned default time value is related to **Repeat action every** drop-down menu at the bottom of the dialog box.

https://t.me/learningnets

Every alarm type has the following actions:

- Send a notification email
- Send a notification trap
- Run a command

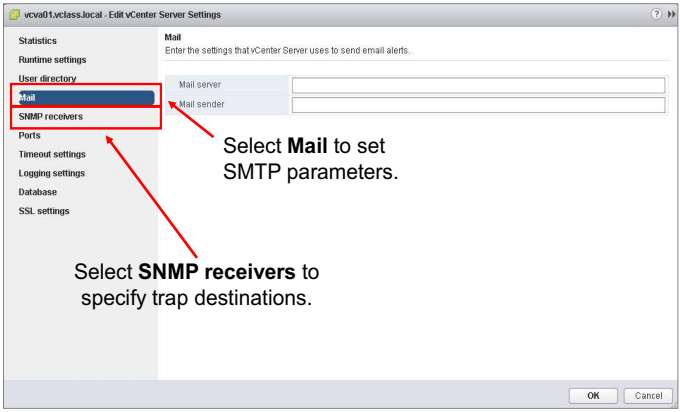
Virtual machine alarms and host alarms have more actions, such as:

- Migrate virtual machine
- Power on a virtual machine
- Power off a virtual machine
- Reboot guest on virtual machine
- Suspend a virtual machine
- Reboot host
- Shut down host

https://t.me/learningnets

8-71 Configuring vCenter Server Notifications

You must configure the email address of the sender account to enable vCenter Server operations, such as sending email notifications as alarm actions.

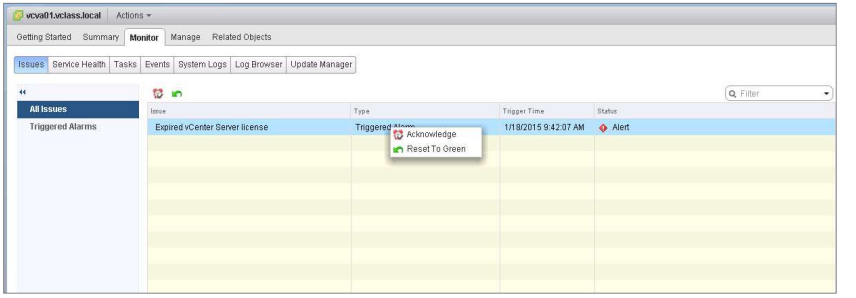


You can configure up to four receivers of SNMP traps. They must be configured in numerical order. Each SNMP trap requires a corresponding host name, port, and community.

https://t.me/learningnets

8-72 Viewing and Acknowledging Triggered Alarms

You use the Acknowledge alarm feature to track when triggered alarms are addressed.



After you acknowledge an alarm in vSphere Web Client or vSphere Client, its actions are discontinued. Alarms are neither cleared nor reset when acknowledged.

Acknowledging an alarm lets other users know that you are taking ownership of the issue. For example, a host has an alarm set on it that monitors CPU usage and sends an email to an administrator when the alarm is triggered. The host CPU usage spikes, triggering the alarm which sends an email to the host's administrator. The administrator acknowledges the triggered alarm to let other administrators know that the problem is being addressed, and to prevent the alarm from sending more email messages. However, the alarm is still visible in the system.

https://t.me/learningnets

8-73 Lab 15: Using Alarms

Use the vCenter Server Appliance alarm feature

1. Create a Virtual Machine Alarm to Monitor a Condition
2. Create a Virtual Machine Alarm to Monitor an Event
3. Trigger Virtual Machine Alarms and Acknowledge the Alarms
4. Disable Virtual Machine Alarms

https://t.me/learningnets

8-74 Review of Learner Objectives

You should be able to meet the following objectives:

- Create alarms with condition-based triggers
- Create alarms with event-based triggers
- View and acknowledge triggered alarms

https://t.me/learningnets

8-75 Key Points

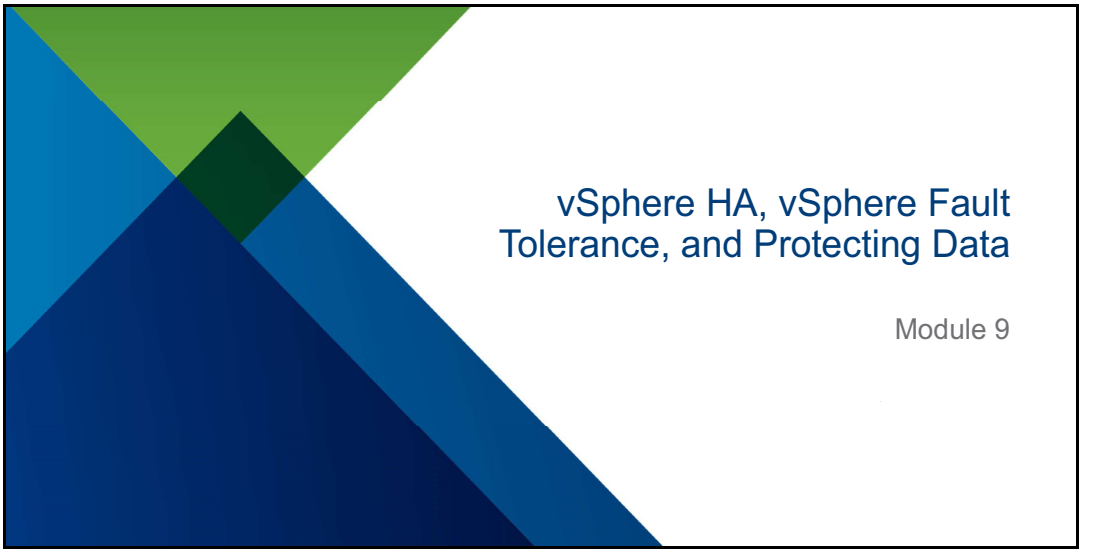
- For proper resource management, vSphere has mechanisms to enable less, more, or an equal amount of access to a defined resource.
- vSphere prevents a virtual machine from consuming large amounts of a resource. It grants a guaranteed amount of a resource to a virtual machine whose performance is not adequate or requires a certain amount of a resource to run properly.
- You can use a resource pool to divide and allocate resources to virtual machines and other resource pools.
- A vApp is a container for one or more virtual machines. The vApp can be used to package and manage related applications.
- You can use the **Performance** tab to monitor the performance of a host or a virtual machine in real time or over a period of time.
- You use alarms to monitor the vCenter Server inventory and send notifications when selected events or conditions occur.

Questions?

https://t.me/learningnets

MODULE 9

vSphere HA, vSphere Fault Tolerance, and Protecting Data



https://t.me/learningnets

9-2 You Are Here

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

https://t.me/learningnets

9-3 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.

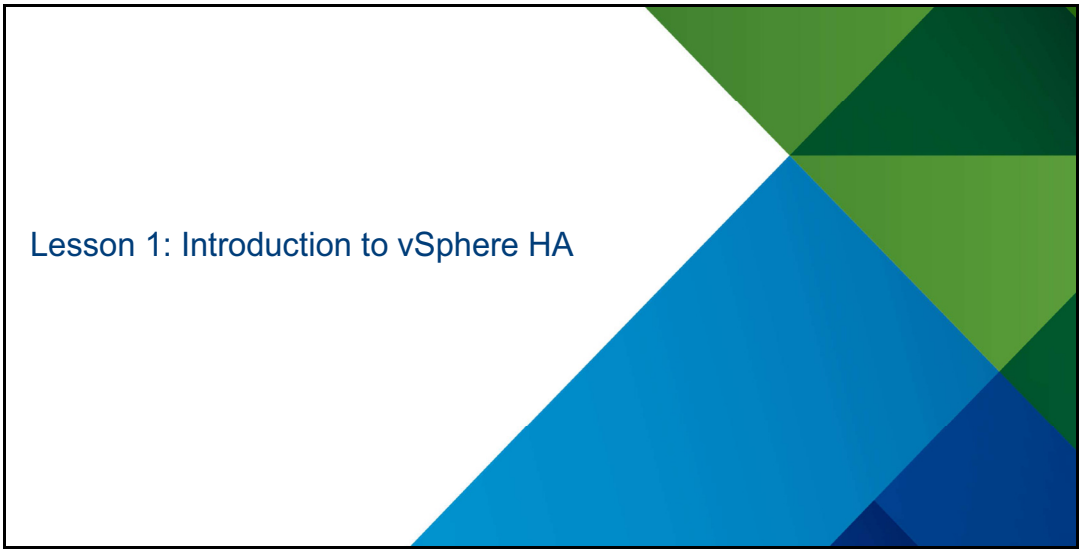
Configuring highly available, computer-based services is extremely important for an organization to remain competitive in contemporary business environments.

https://t.me/learningnets

9-4 Module Lessons

Lesson 1:	Introduction to vSphere HA
Lesson 2:	vSphere HA Architecture
Lesson 3:	Configuring vSphere HA
Lesson 4:	Introduction to vSphere Fault Tolerance
Lesson 5:	vSphere Replication

https://t.me/learningnets



https://t.me/learningnets

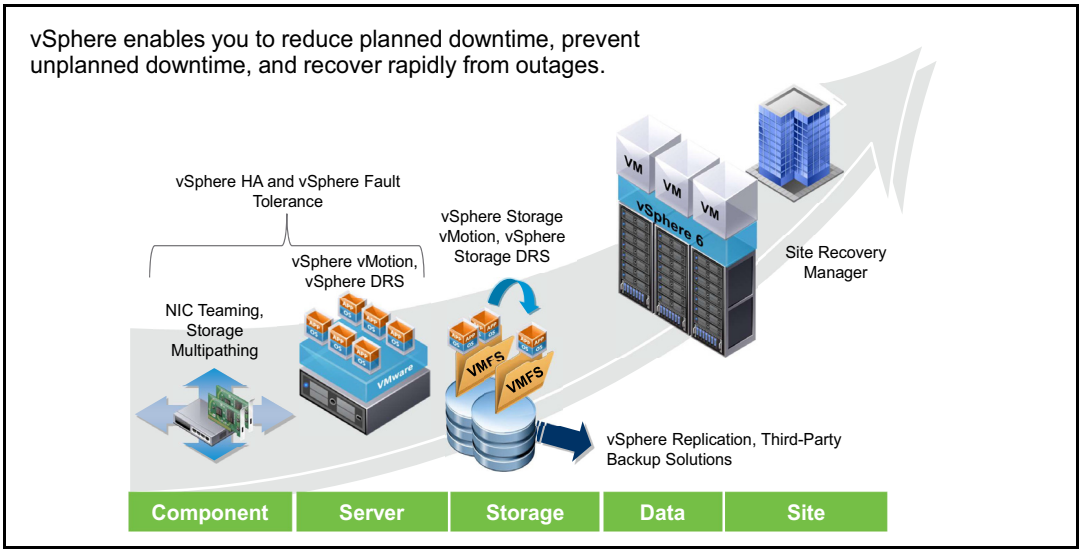
9-6 Learner Objectives

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

- Describe options that you can configure to make your vSphere environment highly available
- Discuss the response of vSphere HA when an ESXi host, a virtual machine, or an application fails

https://t.me/learningnets

9-7 Protection at Every Level



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

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

- Provide higher availability independent of hardware, operating system, and applications.
- Reduce planned downtime for common maintenance operations.
- Provide automatic recovery in cases of failure.

Many methods ensure high availability in a virtualized environment. vSphere uses the following technologies to ensure that virtual machines running in the environment remain available:

- Virtual machine migration
- Multiple I/O adapter paths
- Virtual machine load balancing
- Fault tolerance
- Disaster recovery tools

https://t.me/learningnets

vSphere HA provides a base level of protection for your virtual machines by restarting virtual machines in the event of a host failure. vSphere Fault Tolerance provides a higher level of availability, allowing users to protect any virtual machine from a host failure with no loss of data, transactions, or connections. vSphere Fault Tolerance provides continuous availability by ensuring that the states of the primary and secondary VMs are identical at any point in the instruction execution of the virtual machine.

vSphere vMotion and vSphere Storage vMotion keep virtual machines available during a planned outage, for example, when hosts or storage must be taken offline for maintenance. System recovery from unexpected storage failures is simple, quick, and reliable with the encapsulation property of virtual machines. vSphere Storage vMotion can be used to support planned storage outages resulting from upgrades to storage arrays to newer firmware or technology and VMFS upgrades.

VMware vSphere® Replication™ enables a vSphere platform to protect virtual machines natively by copying their disk files to another location where they are ready to be recovered.

Virtual machine encapsulation is leveraged by third-party backup applications that support file and image-level backups using vSphere Data Protection to protect data. Backup solutions play prominent roles in recovering from deleted files or disks and corrupt or infected guest operating systems or file systems.

Site Recovery Manager enables you to quickly restore your organization's IT infrastructure, shortening the time that you experience a business outage. Site Recovery Manager automates setup, failover, and testing of disaster recovery plans. Site Recovery Manager requires that vCenter Server be installed at the protected site and at the recovery site. Site Recovery Manager also requires either host-based replication through vSphere Replication or preconfigured array-based replication between the protected site and the recovery site.

https://t.me/learningnets

9-8 vCenter Server Availability: Recommendations

Make vCenter Server and the components that it relies on highly available.

vCenter Server relies on these major components:

- vCenter Server database:
  - Create a cluster for the database.
  - Set up native backup and restore.
- Authentication identity source:
  - For example, vCenter Single Sign-On (a service within Platform Services Controller) and Active Directory.
  - Set up with multiple redundant servers.

Methods for making vCenter Server available:

- Use vSphere HA and vCenter Server High Availability to protect the vCenter Server virtual machine.

To provide high availability to vCenter Server, provide high availability for the components that it uses:

- vCenter Server database
- Authentication identity source, such as vCenter Single Sign-On or Active Directory (AD). vCenter Single Sign-On is one of the common infrastructure services within Platform Services Controller. If your environment uses AD, ensure that it is set up with multiple redundant servers.

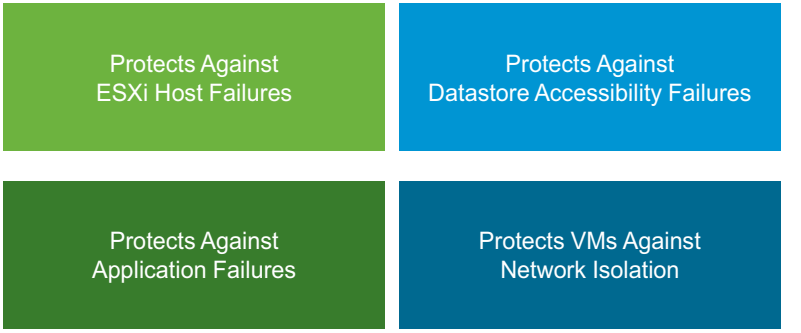
High availability for vCenter Server can be implemented by using vSphere HA. vSphere HA protects against hardware and operating system failures.

To protect Center Server Appliance, you can use vCenter Server High Availability.

https://t.me/learningnets

9-9 About vSphere HA

vSphere HA uses multiple ESXi hosts that are configured as a cluster to provide rapid recovery from outages and cost-effective high availability for applications running in virtual machines.



vSphere HA protects application availability in the following ways:

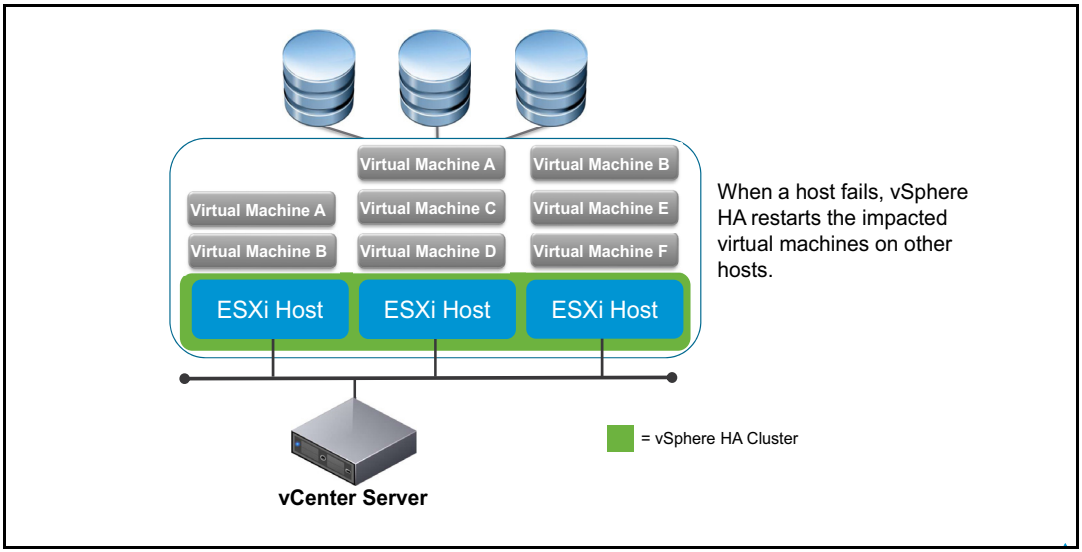
- It protects against a server failure by restarting the virtual machines on other hosts within the cluster.
- It protects against application failure by continuously monitoring a virtual machine and resetting it in the event that a failure is detected.
- It protects against datastore accessibility failures and provides automated recovery for affected virtual machines. With VM Component Protection (VMCP), the affected virtual machines are restarted on other hosts that still have access to the datastores.
- It protects virtual machines against network isolation by restarting them if their host becomes isolated on the management or vSAN network. This protection is provided even if the network has become partitioned.

Unlike other clustering solutions, vSphere HA protects all workloads using the infrastructure itself:

- You do not need to install special software within the application or virtual machine. All workloads are protected by vSphere HA. After vSphere HA is configured, no actions are required to protect new virtual machines. They are automatically protected.
- You can combine vSphere HA with vSphere DRS to protect against failures and to provide load balancing across the hosts within a cluster.

https://t.me/learningnets

9-10 vSphere HA Scenario: ESXi Host Failure

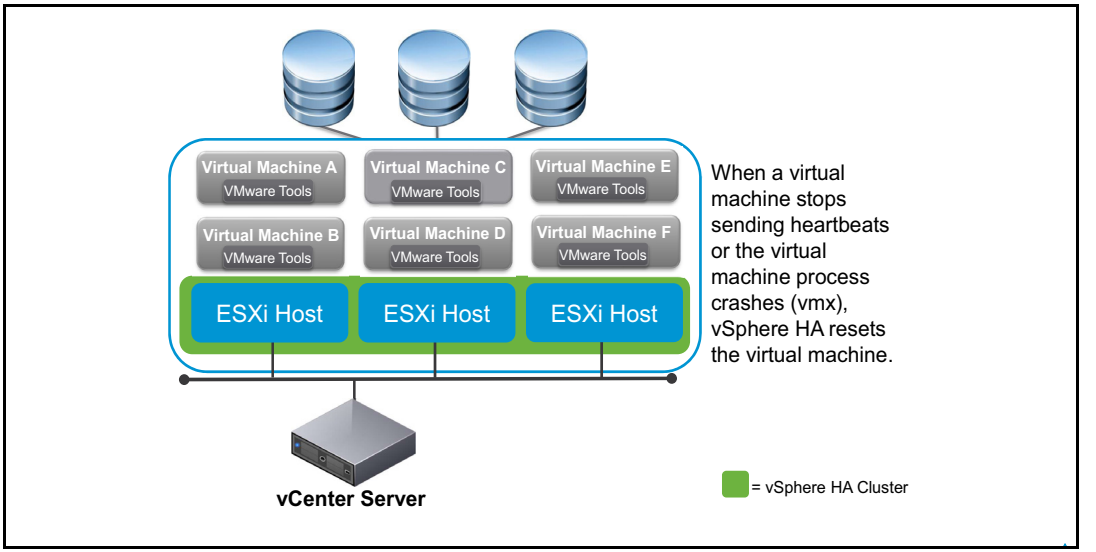


vSphere HA is also able to determine whether a ESXi host is isolated or has crashed. If an ESXi host crashes, vSphere HA has the responsibility of restarting virtual machines that were running on the failed host on the remaining hosts in the cluster.

In every cluster, downtime depends on how long it takes your guest operating systems and applications to restart when the virtual machine is failed over.

https://t.me/learningnets

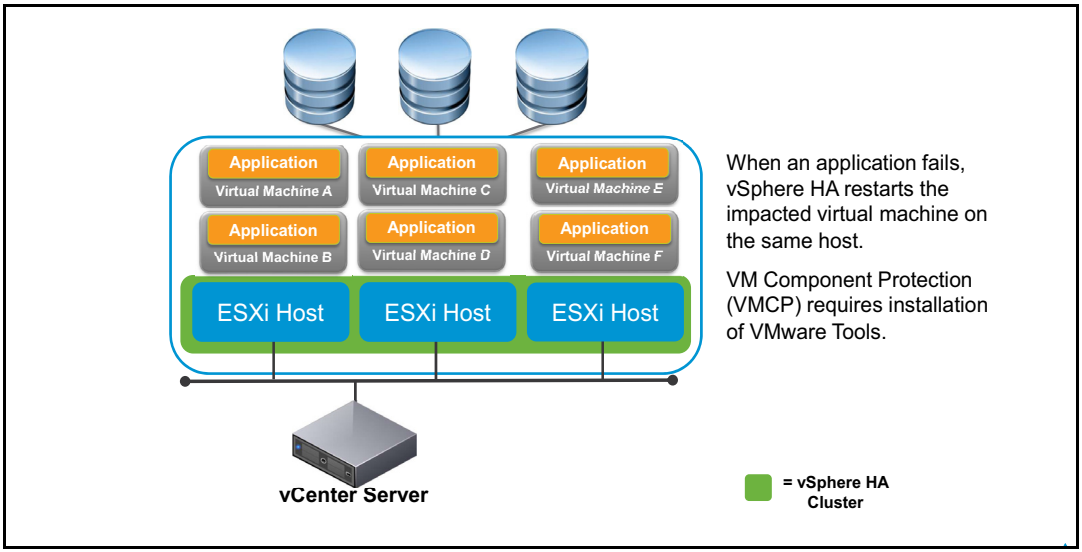
9-11 vSphere HA Scenario: Guest Operating System Failure



If virtual machine monitoring is enabled, the vSphere HA agent on each individual host monitors VMware Tools in each virtual machine running on the host. When a virtual machine stops sending heartbeats, the guest operating system is reset. The virtual machine stays on the same host.

https://t.me/learningnets

9-12 vSphere HA Scenario: Application Failure



The agent on each host can optionally monitor heartbeats of applications running in each virtual machine. When an application fails, the virtual machine on which the application was running is restarted on the same host. Application monitoring requires a third-party application monitoring agent designed to work with virtual machine application monitoring.

https://t.me/learningnets

9-13 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 virtual machines.

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:

- Permanent device loss (PDL):
  - Unrecoverable loss of accessibility.
  - Occurs when a storage device reports the datastore is no longer accessible by the host.
  - Response can be either **Issue events** or **Power off and restart VMs**.
- 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**.

https://t.me/learningnets

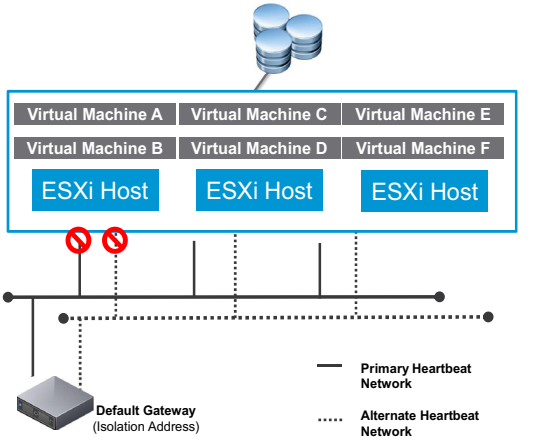
9-14 vSphere HA Scenario: Protecting VMs Against Network Isolation

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:

- vSphere HA tries to ping the cluster isolation addresses. A network isolation address is an IP address that is pinged to determine whether a host is isolated from the network.
- If pinging fails, the host declares that it is isolated from the network.

This protection is provided even if the network has become partitioned.



If you ensure that the network infrastructure is sufficiently redundant and that at least one network path is available at all times, host network isolation is less likely to occur.

https://t.me/learningnets

9-15 Importance of Redundant Heartbeat Networks

In a vSphere HA cluster, heartbeats have the following characteristics:

- They are sent between the master host and the slave hosts.
- They are used to determine whether a master host or a slave host has failed.
- They are sent over a heartbeat network.

Redundant heartbeat networks ensure reliable failure detection and minimize the chance of host-isolation scenarios.

Heartbeat network implementation:

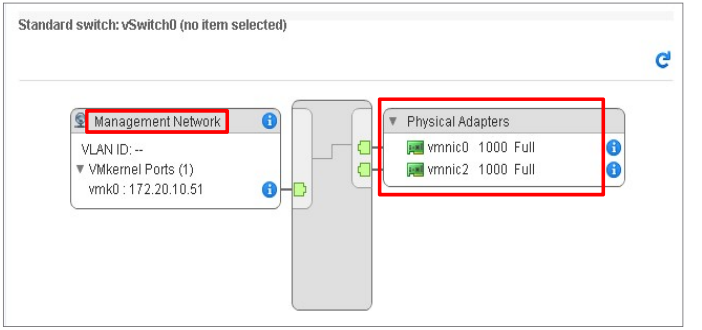
- Implemented by using a VMkernel port that is marked for management.
- Implemented by using a VMkernel port that is marked for vSAN traffic when vSAN is in use.

VMware recommends redundant heartbeat networking for your vSphere HA cluster. If you do not provide redundancy, your failover setup has a single point of failure. When a master host's connection fails, a second connection is still available to send heartbeats to other hosts.

https://t.me/learningnets

9-16 Redundancy Using NIC Teaming

You can use NIC teaming to create a redundant heartbeat network on ESXi hosts.



NIC Teaming on an ESXi Host

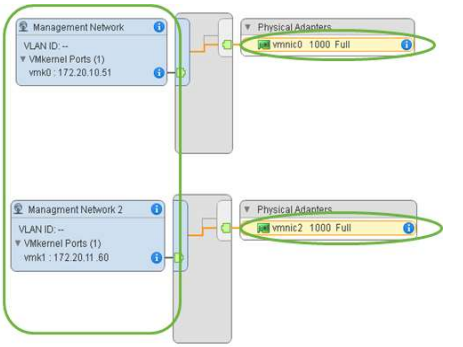
https://t.me/learningnets

9-17 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 enables 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, vSphere HA sends heartbeats over both management network connections. If one path fails, vSphere HA still sends and receives heartbeats over the other path.

https://t.me/learningnets

9-18 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe options that you can configure to make your vSphere environment highly available
- Discuss the response of vSphere HA when an ESXi host, a virtual machine, or an application fails

https://t.me/learningnets



https://t.me/learningnets

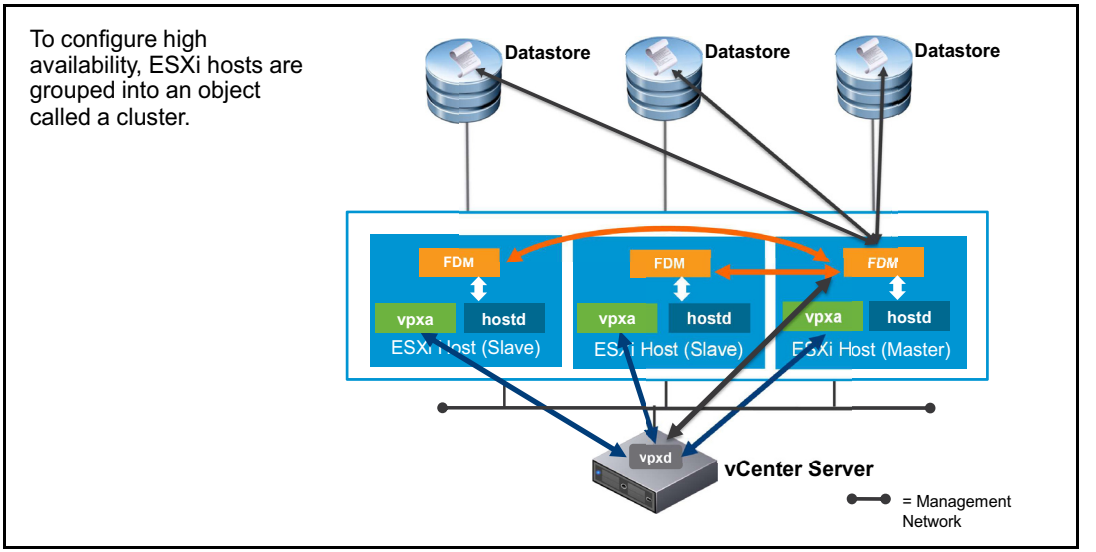
9-20 Learner Objectives

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

- Describe the heartbeat mechanisms used by vSphere HA
- Identify and discuss other failure scenarios
- Recognize vSphere HA design considerations

https://t.me/learningnets

9-21 vSphere HA Architecture: Agent Communication



When vSphere HA is enabled, the Fault Domain Manager (FDM) service starts on the member hosts. After the FDM agents have started, the cluster hosts are said to be in a fault domain. Hosts cannot participate in a fault domain if they are in maintenance mode, in standby mode, or disconnected from vCenter Server. A host can be in only one fault domain at a time.

The fault domain is managed by a master host. All other hosts are called slave hosts. FDM agents on slave hosts all communicate with FDM on the master host.

To determine which host will be the master, an election process takes place. The system that can access the greatest number of datastores is elected the master. If more than one cluster hosts see the same number of datastores, the election process determines the master host by using the host managed object ID (MOID) assigned by vCenter Server.

The election process for a new master completes in approximately 15 seconds and occurs under these circumstances:

- vSphere HA is enabled.
- The master encounters a system failure because of one of the following factors:
  - The master host is placed in maintenance mode.
  - The master host is placed in standby mode.
  - vSphere HA is reconfigured.
- The slave hosts cannot communicate with the master host due to a network problem.

https://t.me/learningnets

During the election process, the candidate vSphere HA agents communicate with each other over the management network or the vSAN network in a vSAN cluster by using User Datagram Protocol (UDP). All network connections are point-to-point. After the master agent has been determined, master and slave hosts communicate using secure TCP. When vSphere HA is started, vCenter Server contacts the master host agent and sends a list of hosts with membership in the cluster along with the cluster configuration. That information is saved to local storage on the master host and then pushed out to the slave hosts in the cluster. If additional hosts are added to the cluster during normal operation, the master agent sends an update to all hosts in the cluster.

The master host provides an interface for vCenter Server to query the state of and report on the health of the fault domain and virtual machine availability. vCenter Server tells the vSphere HA agent which virtual machines to protect along with their virtual machine-to-host compatibility list. The agent learns about state changes through hostd and vCenter Server learns of these through vpxa. The master host monitors the health of the slave hosts and takes responsibility for virtual machines that were running on a failed slave host.

A slave host monitors the health of virtual machines running locally and sends state changes to the master host. A slave host also monitors the health of the master host.

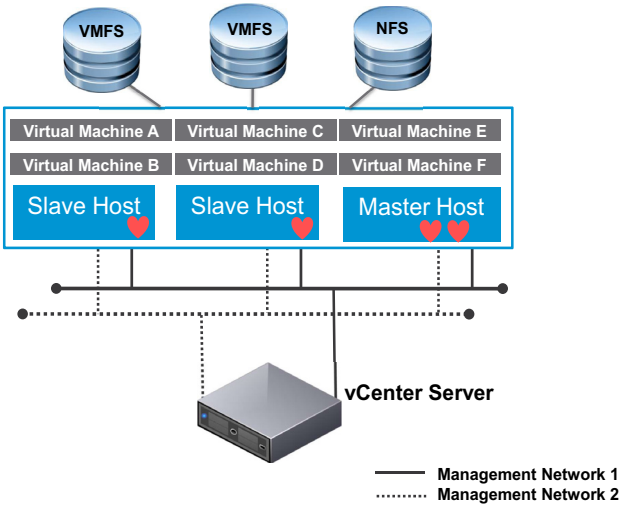
vSphere HA is configured, managed, and monitored through vCenter Server. The cluster configuration data is maintained by the vpxd process which runs on the vCenter Server system. Cluster configuration changes are reported by the vpxd process to the master agent. The master agent advertises a new copy of the cluster configuration information and each slave fetches an updated copy. Each slave writes the updated configuration information to local storage. A list of protected virtual machines is stored on each datastore. The virtual machine list is updated after each user-initiated power-on (protected) and power off (unprotected) operation. The virtual machine list is updated after vCenter Server observes these operations.

A virtual machine becomes protected when an operation results in a power on. Reverting a virtual machine to a snapshot with memory state causes the virtual machine to power on and become protected. Similarly, a user action that causes the virtual machine to power off, for example, reverting to a snapshot without memory state or a standby operation performed in the guest, causes the virtual machine to become unprotected.

https://t.me/learningnets

9-22 vSphere HA Architecture: Network Heartbeats

The master host sends periodic heartbeats to the slave hosts so that the slave hosts know that the master host is alive.



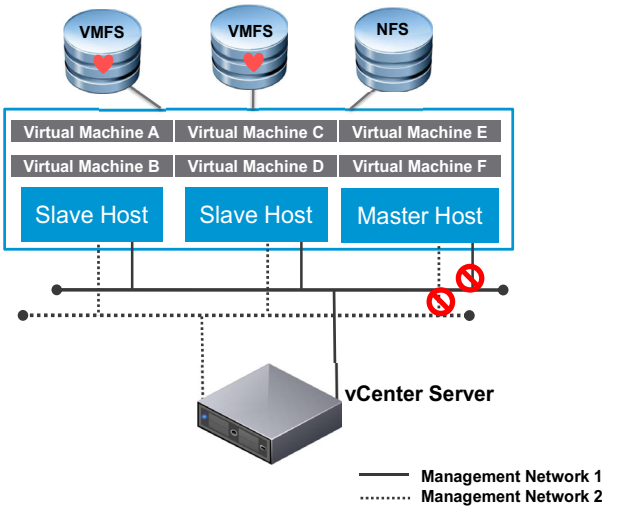
Heartbeats are sent to each slave host from the master host over all configured management networks. However, slave hosts use only one management network to communicate with the master. If the management network used to communicate with the master fails, the slave switches to another management interface to communicate with the master host.

If the slave host does not respond within predefined timeout period, the master host declares the slave host as agent unreachable. When a slave host is not responding, the master host attempts to determine the cause of the slave host's inability to respond. The master host must determine whether the slave host crashed, is not responding because of a network failure, or the vSphere HA agent is in an unreachable state.

https://t.me/learningnets

9-23 vSphere HA Architecture: Datastore Heartbeats

Datastores are used as a backup communication channel to detect virtual machine and host heartbeats.



The datastore heartbeats are used to make the distinction between a failed and isolated or partitioned host. vSphere HA tries to restart virtual machines only in one of these situations:

- A host has failed (no network heartbeats, no ping, no datastore heartbeats).
- A host becomes isolated and the cluster's configured host isolation response is Power off or Shut down.

https://t.me/learningnets

9-24 vSphere HA Failure Scenarios

vSphere HA can identify and respond to various types of failures:

- Slave host failure
- Master host failure
- Network failure (host isolation)

VMCP enables vSphere HA to detect and respond to storage-related failures:

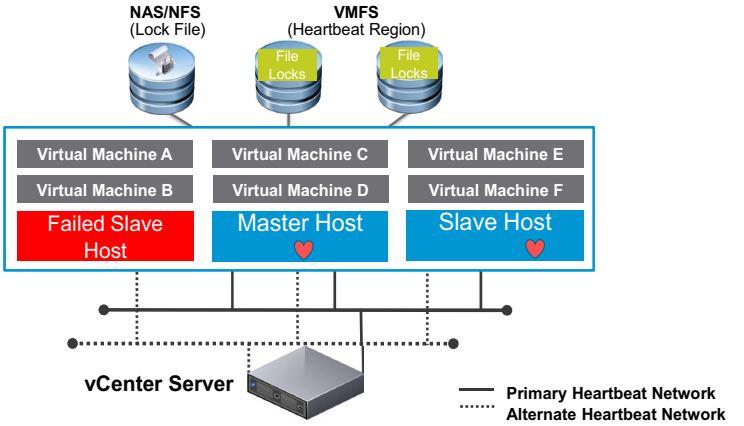
- Datastore access failure
  - APD
  - PDL

vSphere HA can also determine whether an ESXi host is isolated or has crashed. Isolation refers to when an ESXi host cannot see traffic coming from the other hosts in the cluster and cannot ping the default gateway. If an ESXi host crashes, vSphere HA must restart the virtual machines that were running on the failed host on the remaining hosts in the cluster. If the ESXi host is isolated because it cannot ping the default gateway (or isolation address) and sees no management network traffic, the host executes the Host Isolation Response.

https://t.me/learningnets

9-25 Failed Slave Hosts

When a slave host does not respond to the network heartbeat issued by the master host, the master vSphere HA agent tries to identify the cause.



The master host must determine whether the slave host has crashed or is not responding because of a network problem, for example, a misconfigured firewall rule or component failure. The type of failure dictates how vSphere HA responds.

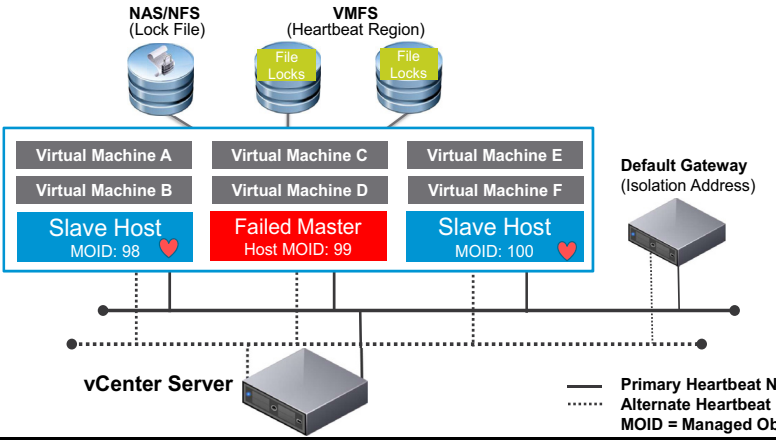
When heartbeats cannot be obtained using the network, the master host must determine whether the slave host has had a network failure or if the system has crashed. The master host checks for both responses to pings and datastore heartbeats. If pings and datastore heartbeats are not present, the host is declared dead. The absence of both a network and datastore heartbeat indicates full host failure.

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-<X>-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 slave host suffers from network failure or host failure. 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

9-26 Failed Master Hosts

When the master host is placed in maintenance mode or crashes, the slave hosts detect that the master host is no longer issuing heartbeats.



In this case, an election must take place to determine a new master host. The host with access to the greatest number of datastores is elected the master. If all slave hosts have equal datastore access, the election process selects a new master host using the highest lexically ordered MOID assigned by vCenter Server when the host was added to the vCenter Server inventory.

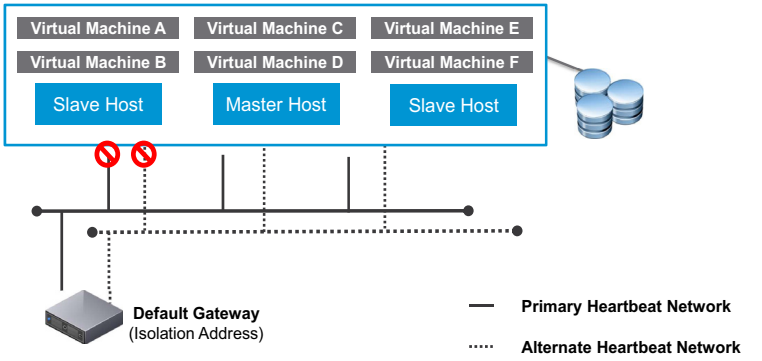
If the master host fails, the slave participates in a new master election. When a new master is selected, it reads MAC and IP addresses of the hosts and virtual machines from a host list that is stored on a datastore. The host list is used to determine whether or not the master should accept a connection from a slave.

The new master reads a file on each datastore that contains the state of all virtual machines and determines which virtual machines are protected by vSphere HA and might have to be restarted. The new master first identifies which virtual machines are running on the slave hosts. Through the process of elimination, it determines which virtual machines require restarting. This same process is used to restart virtual machines after a total cluster failure.

https://t.me/learningnets

9-27 Isolated Hosts

If the host does not receive any heartbeats or observe any election traffic on the management and cannot ping its default gateway, the host is isolated.



A host is declared isolated when the following two conditions occur:

- When the host is not receiving network heartbeats
- When the host cannot ping its isolation addresses

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 vSphere HA agent must determine if that host is still alive, and merely isolated, by checking for datastore heartbeats. Datastore heartbeats are used by vSphere HA only when a host becomes isolated or partitioned.

https://t.me/learningnets

9-28 Design Considerations

Host isolation events can be minimized through good design:

- Implement redundant heartbeat networks.
- Implement redundant isolation addresses.

If host isolation events do occur, good design enables vSphere HA to determine whether the isolated host is still alive.

The VM network should be physically separate from the heartbeat network.

Implement datastores so that they are separated from the management network by using one or both of the following approaches:

- Fibre Channel over fiber optic.
- Physically separating your IP storage network from the management network.

If a datastore is based on Fibre Channel, the datastore access is not disrupted by the network failure. When using datastores based on IP storage (for example, NFS, iSCSI, or Fibre Channel over Ethernet), you must physically separate (or logically separate, if physical separation is impossible) the IP storage network and the management network (the heartbeat network).

https://t.me/learningnets

9-29 Virtual Machine Storage Failures

With an increasing number of virtual machines and datastores on each host, storage connectivity problems have high costs but are infrequent.

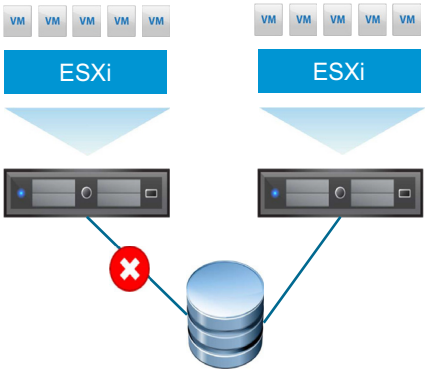
Connectivity problems are due to:

- Network or switch failure
- Array misconfiguration
- Power outage

Storage connectivity problems affect virtual machine availability:

- Virtual machines on affected hosts are difficult to manage.
- Applications with attached disks crash.

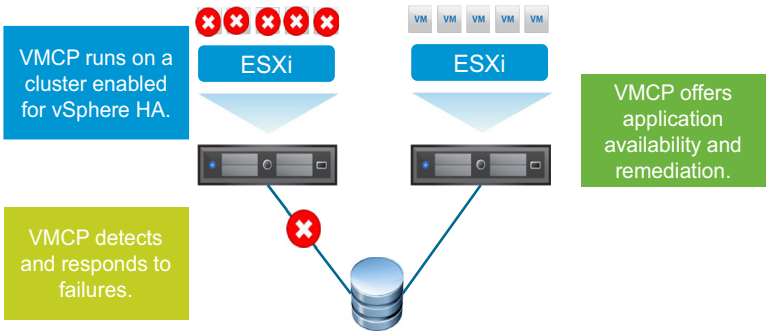
VMCP is not supported with vSAN.



https://t.me/learningnets

9-30 About VMCP

VMCP protects against storage failures on a virtual machine.



Only vSphere HA clusters that contain ESXi 6.x hosts can be used to enable VMCP.

If VMCP is enabled, vSphere HA can detect datastore accessibility failures and provide automated recovery for affected virtual machines.

VMCP provides protection against datastore accessibility failures that can affect a virtual machine running on a host in a vSphere HA cluster. When a datastore accessibility failure occurs, the affected host can no longer access the storage path for a specific datastore. You can determine the response that vSphere HA will give to such a failure, ranging from the creation of event alarms to virtual machine restarts on other hosts.

Only vSphere HA clusters that contain ESXi 6.x hosts can be used to enable VMCP. Clusters containing hosts from an earlier release cannot enable VMCP. Such hosts cannot be added to a cluster enabled for VMCP.

https://t.me/learningnets

9-31 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe the heartbeat mechanisms used by vSphere HA
- Identify and discuss other failure scenarios
- Recognize vSphere HA design considerations

https://t.me/learningnets



https://t.me/learningnets

9-33 Learner Objectives

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

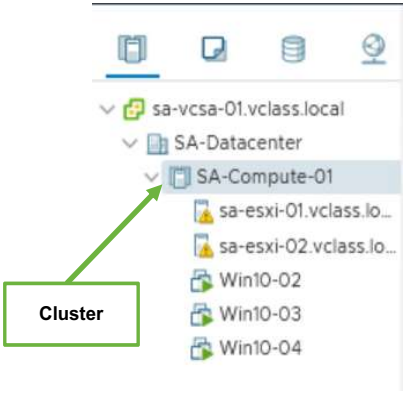
- Recognize the requirements for creating and using a vSphere HA cluster
- Configure a vSphere HA cluster

https://t.me/learningnets

9-34 About 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.

Features such as vSphere HA, vSphere DRS, and vSAN can be enabled in a cluster.



vSphere HA uses multiple ESXi hosts that are configured as a cluster to provide rapid recovery from outages and high availability for applications running in virtual machines. The resources are managed by vCenter Server. Features such as vSphere HA, vSphere DRS, and vSAN can be enabled in a cluster.

https://t.me/learningnets

9-35 vSphere HA Prerequisites

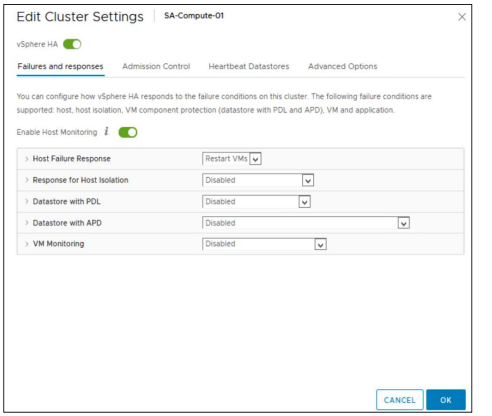
You must follow the guidelines to create a vSphere HA cluster:

- Use either VMware vSphere® Essentials Plus Kit or VMware vSphere® Standard Edition™.
- A cluster can contain between two and 64 ESXi hosts.
- All hosts must be configured with static IP addresses. If you are using DHCP, you must ensure that the address for each host persists across reboots.
- All hosts must have at least one management network in common.
- For virtual machine monitoring to work, VMware Tools must be installed in every virtual machine.
- Only vSphere HA clusters that contain ESXi 6.x hosts can be used to enable VMCP.

https://t.me/learningnets

9-36 Configuring vSphere HA Settings

When you create or configure a vSphere HA cluster, you must configure settings that determine how the feature works.



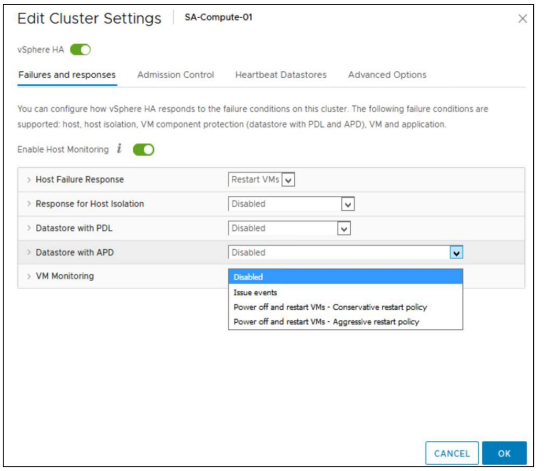
In both vSphere Web Client and vSphere Client, you can configure the following vSphere HA settings:

- Availability failure conditions and responses: Provide settings for host failure responses, host isolation, virtual machine monitoring, and VMCP.
- Admission control: Enable or disable admission control for the vSphere HA cluster and choose a policy for how it is enforced.
- Heartbeat datastores: Specify preferences for the datastores that vSphere HA uses for datastore heartbeating.
- Advanced options: Customize vSphere HA behavior by setting advanced options.

https://t.me/learningnets

9-37 vSphere HA Settings: Failure and Responses

You use the Failures and responses pane to configure a cluster's response if a failure occurs.

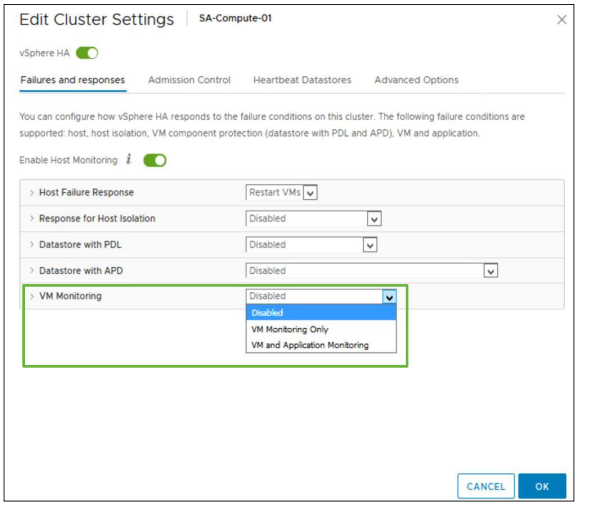


The Failure and Responses pane enables you to configure how your cluster should function when problems are encountered. You can specify the vSphere HA cluster's response for host failures and isolation. You can also configure VMCP actions when permanent device loss and all paths down situations occur and enable VM monitoring.

https://t.me/learningnets

9-38 vSphere HA Settings: Virtual Machine Monitoring

You use VM Monitoring settings to control the monitoring of virtual machines.



By default, **VM Monitoring** is set to **Disabled**.

Virtual machine monitoring restarts individual virtual machines if their VMware Tools heartbeats are not received or if the guest operating system has not issued an I/O for the last 2 minutes (by default). If neither criteria is met, it is most likely because the guest operating system has failed. In such a case, the virtual machine monitoring service determines that the virtual machine has failed and the virtual machine resets to restore service.

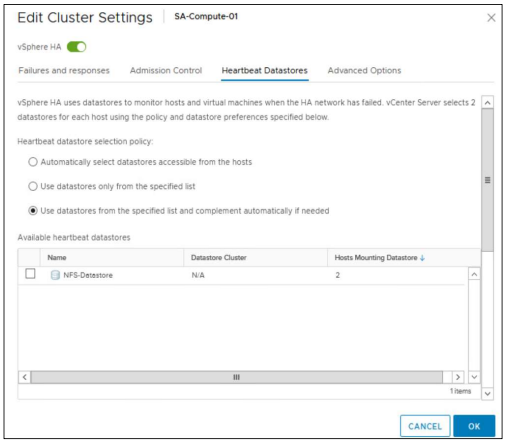
You can configure the level of monitoring sensitivity. Highly sensitive monitoring results in a more rapid conclusion that a failure has occurred. Although unlikely, highly sensitive monitoring might lead to falsely identifying failures when the virtual machine or application is still working but heartbeats have not been received due to factors like resource constraints. Low-sensitivity monitoring results in longer interruptions in service between actual failures and virtual machines being reset. Select an option that is an effective compromise for your needs.

Select **VM Monitoring Only** to reset individual virtual machines if their heartbeats are not received within a set time. You can also select **VM and Application Monitoring** to enable application monitoring.

https://t.me/learningnets

9-39 vSphere HA Settings: Heartbeat Datastores

A heartbeat file is created on the selected datastores and is used if the management network fails.



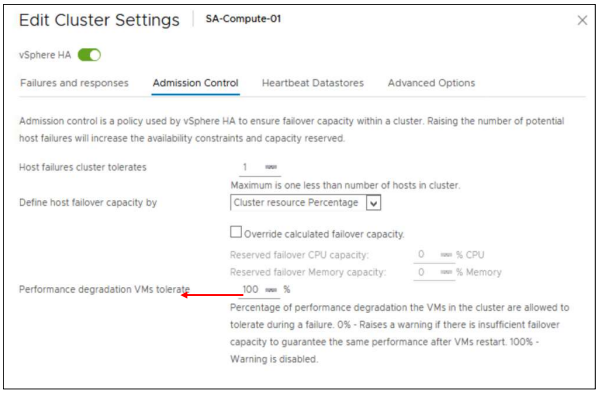
Using datastore heartbeating, the master host determines whether a host has failed or a network isolation has occurred. If datastore heartbeating from the host stops, then the host is considered failed. In this case, the failed host's virtual machines are started on another host in the vSphere HA cluster.

Datastore heartbeating takes checking the health of a host to another level by checking more than the management network to determine a host's health. You can configure a list of datastores to monitor for a particular host. Or you can let vSphere HA make the decision. You can also combine both methods.

https://t.me/learningnets

9-40 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 virtual machine resource reservations are respected.



After you create a cluster, admission control allows you to specify whether or not virtual machines can be started if they violate availability constraints. The cluster reserves resources to allow failover for all running virtual machines on the specified number of hosts.

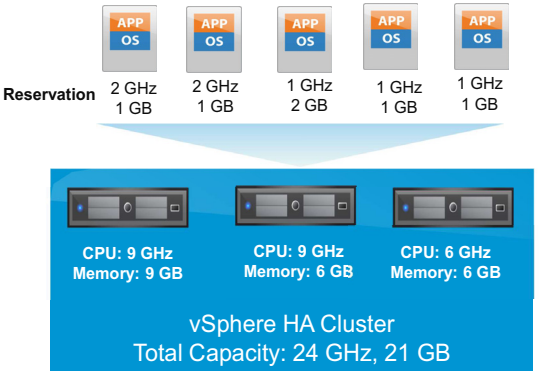
The admission control settings include:

- **Disabled**: (Not recommended) This option disables admission control, allowing the virtual machines violating availability constraints to power on.
- **Slot Policy**: A slot is a logical representation of memory and CPU resources. With the slot policy option, vSphere HA calculates the slot size, determines how many slots each host in the cluster can hold, and consequently determines the current failover capacity of the cluster.
- **Cluster resource Percentage**: (Default) This 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

9-41 Example: Admission Control Using Cluster Resources Percentage

- Total VM reservations: 7 GHz and 6 GB.
- Current CPU failover capacity is 70%:
  - ((24 GHz - 7 GHz)/24 GHz) = 70.8%
- Current memory failover capacity is 71%:
  - ((21 GB - 6 GB)/21 GB) = 71.4%

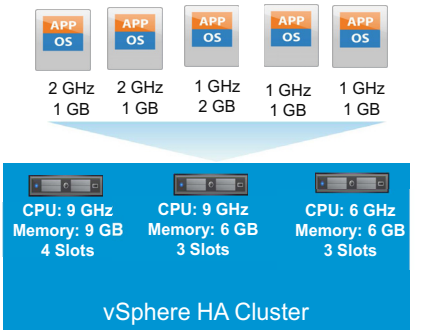


https://t.me/learningnets

9-42 Example: Admission Control Using Slots

A slot is the memory and CPU reservation required for any powered-on VM in the cluster:

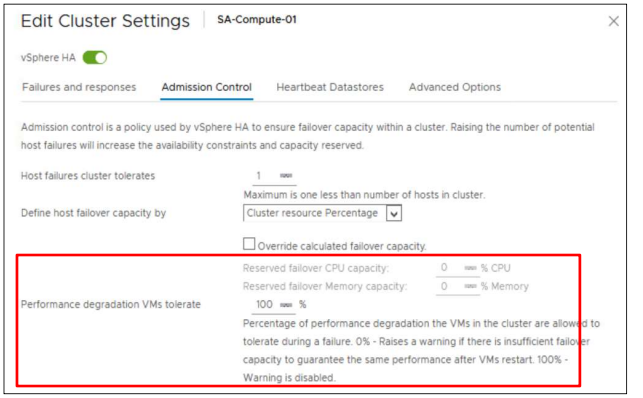
- Calculate the slot size: 2 GHz CPU and 2 GB memory.
- Calculate the slots per host. In the example, the cluster has 10 slots.
- The current failover capacity is one host. In the example, the failover capacity is four (the largest host has four slots):
  - If Host1 fails, six slots remain in the cluster, which is sufficient for all five of the powered-on virtual machines.
  - If both Host1 and Host2 fail, only three slots remain, which is insufficient.
- If the current failover capacity is less than the configured failover capacity, disallow more VMs from powering on.



https://t.me/learningnets

9-43 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 usage exceeds the failover capacity percentage. The resource reduction calculation takes into account a virtual machine's reserved memory and memory overload to decide whether to permit it to power on, migrate, or have reservation changes.

Setting the Performance degradation VMs tolerate threshold allows you to specify when the occurrence of a configuration issue should arise. For example:

- The default value is 100 percent, which produces no warnings.
- If you reduce the threshold to 0 percent, a warning is generated as soon as cluster usage 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 usage x 20 percent. When the current usage minus the performance reduction exceeds the available capacity, a configuration notice is issued.

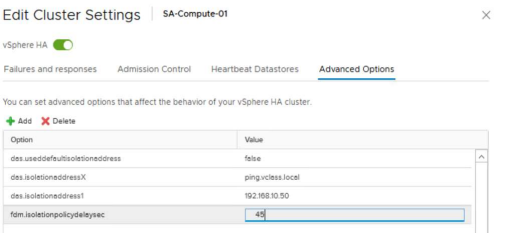
The Performance degradation VMs tolerate threshold is not available unless vSphere DRS is enabled.

https://t.me/learningnets

9-44 vSphere HA Settings: Advanced Options

You set advanced vSphere HA options to customize vSphere HA behavior:

- To force a cluster not to use the default isolation address (default gateway):
  - das.usedefaultisolationaddress = false
- To force a cluster to ping alternate isolation addresses:
  - das.isolationaddressX = ping reachable address
- To force a cluster to wait beyond the default 30-second isolation action window:
  - fdm.isolationpolicydelaysec = > 30 sec
- To force maximum bound on the memory and CPU slot size:
  - das.slotmeminmb = 100
  - das.slotcpuinmhz = 32



You can set advanced options that affect the behavior of your vSphere HA cluster. For more details, see *vSphere Availability Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

9-45 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.

Optionally, you can configure a delay when a certain condition is met.

Priority tier:

- VMs can be grouped into tiers indicating their startup priority. All VMs in the priority 1 tier receive their resources first and are powered on.
- Only when all the VMs in tier 1 have met their defined restart condition does vSphere HA continue to the VMs in the priority 2 tier, and so on.

After a host failure, virtual machines are assigned to other hosts with unreserved capacity, with the highest priority virtual machines placed first. The process continues to those virtual machines with lower priority until all have been placed or no more cluster capacity is available to meet the reservations or memory overhead of the virtual machines. A host then restarts the virtual machines assigned to it in priority order.

If there are insufficient resources, vSphere HA waits for more unreserved capacity to become available, for example, due to a host coming back online, and then retries the placement of these virtual machines. To reduce the chance of this situation occurring, configure vSphere HA admission control to reserve more resources for failures. Admission control allows you to control the amount of cluster capacity that is reserved by virtual machines, which is unavailable to meet the reservations and memory overhead of other virtual machines if a failure occurs.

https://t.me/learningnets

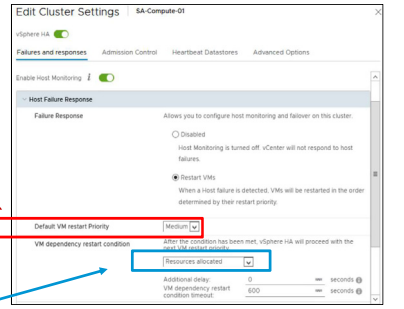
9-46 Configuring vSphere HA Orchestrated Restart

Regular VMs are put in the Medium restart priority by default, unless the restart priority is explicitly set using VM overrides.

Exceptions:

- Agent VMs always start first, and the restart priority is nonconfigurable.
- vSphere Fault Tolerance secondary VMs fail over before regular VMs. Primary VMs follow normal restart priority.

Priorities:
Lowest (5), Low, Medium (default), High, Highest (1)

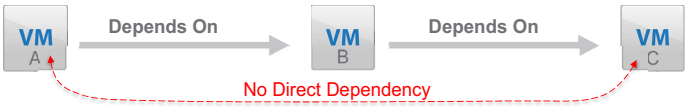


Conditions:
- Resources allocated
- Powered on
- Guest heartbeats detected
- App heartbeat detected

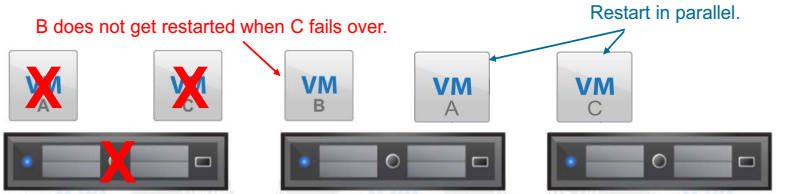
https://t.me/learningnets

9-47 VM Dependencies in Orchestrated Restart

VMs can depend only on other VMs of the same or higher priority. Only direct dependencies are supported. VM-to-VM dependency is a hard rule. Creating cyclical dependencies causes VM restart to fail.

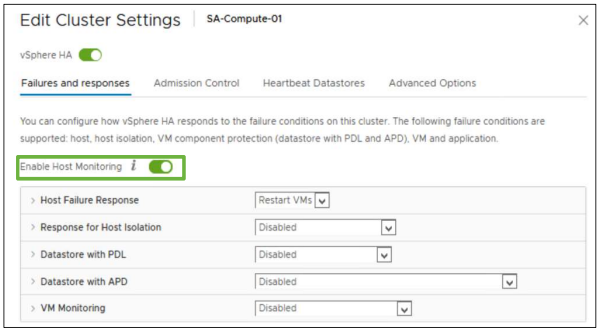


In vSphere 6.5 and later, vSphere HA restarts VMs only from a failed host. Configure affinity rules to keep VMs on the same host if necessary.



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 virtual machines.



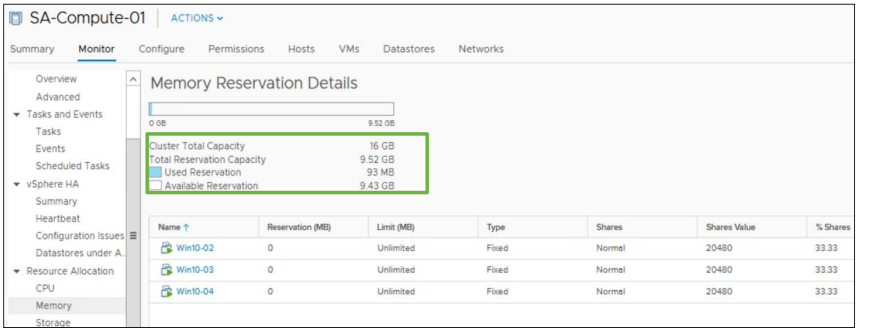
The following network maintenance suggestions can help you avoid the false detection of host failure and network isolation due to dropped vSphere HA heartbeat:

- Changing your network hardware or networking settings can interrupt the heartbeats used by vSphere HA to detect host failures, and might result in unwanted attempts to fail over virtual machines. When making changes to the networks where your clustered ESXi hosts reside, suspend host monitoring and place the host in maintenance mode.
- 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

9-49 Cluster Resource Reservation

The Resource Allocation page of the **Monitor** tab reports total cluster CPU, memory, memory overhead, storage capacity, the capacity reserved by VMs, and how much capacity remains available.

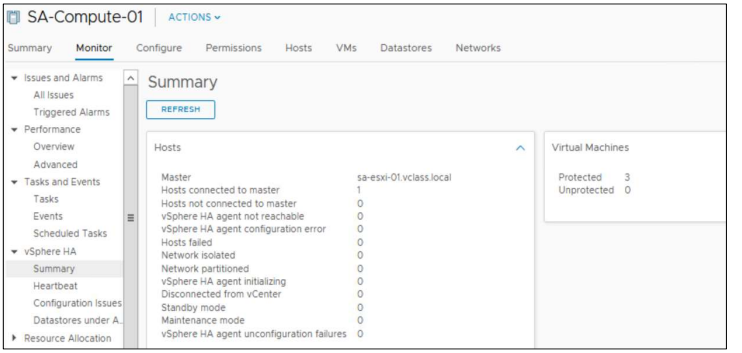


vCenter Server uses vSphere HA admission control to ensure that sufficient resources are available in a cluster to provide failover protection and to ensure that virtual machine resource reservations are respected.

https://t.me/learningnets

9-50 Monitoring Cluster Status

You can monitor the status of a vSphere HA cluster on the Summary page of the Monitor tab.



You can monitor the status of a vSphere HA cluster on the Summary page of the vSphere HA panel, on the **Monitor** tab.

Configuration issues and other errors can occur for your cluster or its hosts that adversely affect proper vSphere HA operation. You can monitor these errors on the Configuration Issues page, which is accessible from the same area.

https://t.me/learningnets

9-51 Lab 16: Using vSphere HA

Use vSphere HA functionality

1. Create a Cluster Enabled for vSphere HA
2. Add Your ESXi Hosts to the Cluster
3. Test the vSphere HA Functionality
4. View the vSphere HA Cluster Resource Usage
5. Manage vSphere HA Slot Size
6. Configure a vSphere HA Cluster with Strict Admission Control

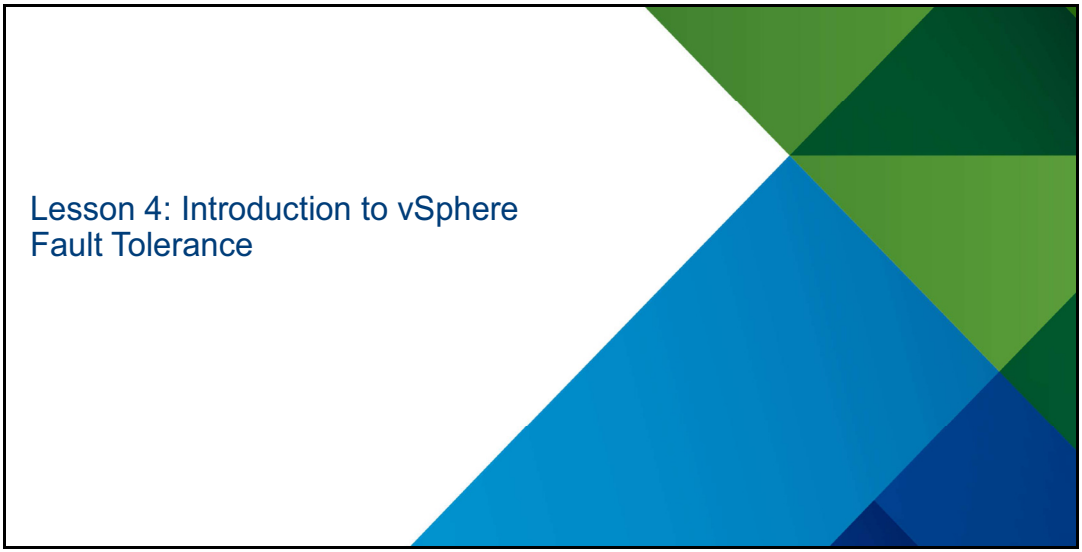
https://t.me/learningnets

9-52 Review of Learner Objectives

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

9-54 Learner Objectives

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

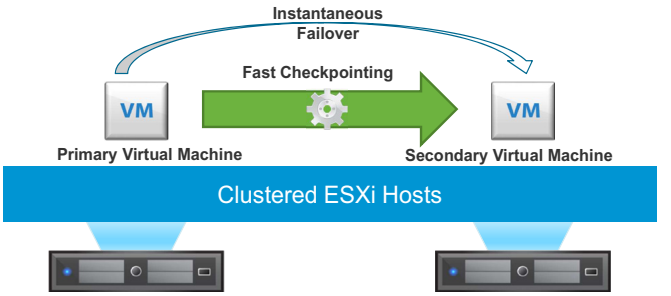
- Describe the benefits of using vSphere Fault Tolerance
- Explain how vSphere Fault Tolerance works
- Describe main vSphere Fault Tolerance features
- Describe how vSphere Fault Tolerance works with vSphere HA and vSphere DRS
- Enable vSphere Fault Tolerance from vSphere Web Client

https://t.me/learningnets

9-55 About vSphere Fault Tolerance

vSphere Fault Tolerance provides instantaneous failover and continuous availability:

- Zero downtime
- Zero data loss
- No loss of TCP connections



You can use vSphere Fault Tolerance for your virtual machines to ensure business continuity with higher levels of availability and data protection than is offered by vSphere HA.

vSphere Fault Tolerance is built on the ESXi host platform. vSphere Fault Tolerance provides continuous availability by having identical virtual machines run on separate hosts.

You can use vSphere Fault Tolerance for most mission-critical virtual machines. vSphere Fault Tolerance provides continuous availability for such a virtual machine by creating and maintaining another virtual machine that is identical and continuously available to replace it in the event of a failover situation.

The protected virtual machine is called the primary virtual machine. The duplicate virtual machine is called the secondary virtual machine. The secondary virtual machine is created and runs on a different host to the primary virtual machine. The secondary virtual machine's execution is identical to that of the primary virtual machine. The secondary virtual machine can take over at any point without interruption, thus providing fault-tolerant protection.

The primary virtual machine and the secondary virtual machine continuously monitor the status of each other to ensure that fault tolerance is maintained. A transparent failover occurs if the host running the primary virtual machine fails, in which case the secondary virtual machine is immediately activated to replace the primary virtual machine. A new secondary virtual machine is started, and fault tolerance redundancy is reestablished automatically. If the host running the secondary virtual machine fails, it is also immediately replaced. In either case, users experience no interruption in service and no loss of data.

https://t.me/learningnets

9-56 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 eight 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 CPUs 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 (EVC) feature is enabled. This process enables fault-tolerant virtual machines to benefit from better initial placement.

When a cluster has EVC enabled, vSphere DRS makes the initial placement recommendations for fault-tolerant virtual machines and allows you to assign a vSphere DRS automation level to primary VMs. The secondary VM always assumes the same setting as its associated primary VM.

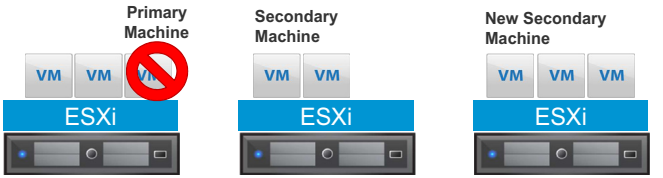
When vSphere Fault Tolerance is used for virtual machines in a cluster that has EVC disabled, the fault-tolerant virtual machines are given the disabled vSphere DRS automation level. In such a cluster, each primary VM is powered on only on its registered host, and its secondary VM is automatically placed.

https://t.me/learningnets

9-57 vSphere Fault Tolerance with vSphere HA and vSphere DRS

vSphere Fault Tolerance works with vSphere HA and vSphere DRS:

- vSphere HA:
  - Is required for vSphere Fault Tolerance
  - Restarts failed VMs
  - Is vSphere Fault Tolerance aware
- vSphere DRS:
  - Selects which hosts run the primary and secondary VM, when a VM is powered on
  - Does not automatically migrate fault-tolerant VMs



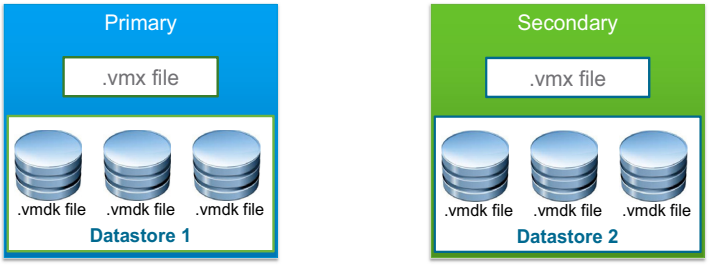
A fault-tolerant virtual machine and its secondary copy are not allowed to run on the same host. This restriction ensures that a host failure cannot result in the loss of both virtual machines.

https://t.me/learningnets

9-58 Redundant VMDK Files

vSphere Fault Tolerance creates two complete virtual machines.

Each virtual machine has its own `.vmx` configuration file and `.vmdk` files. Each virtual machine can be on a different datastore.

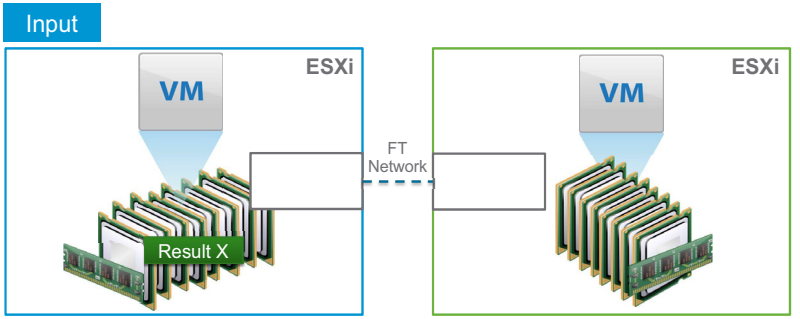


vSphere Fault Tolerance provides failover redundancy by creating two full virtual machine copies. The virtual machine files can be placed on the same datastore. However, VMware recommends that you place these files on separate datastores to provide recovery from datastore failures.

https://t.me/learningnets

9-59 vSphere Fault Tolerance Checkpoint

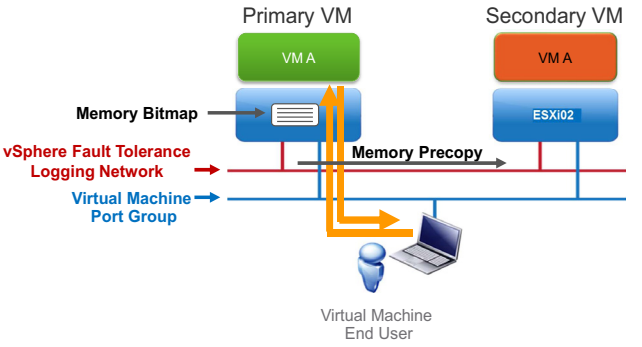
Changes on the primary machine are not processed on the secondary machine. The memory is updated on the secondary machine.



https://t.me/learningnets

9-60 vSphere Fault Tolerance: Precopy

Using vSphere Fault Tolerance, a second VM is created on the secondary host. The memory of the source virtual machine is then copied to the secondary host.

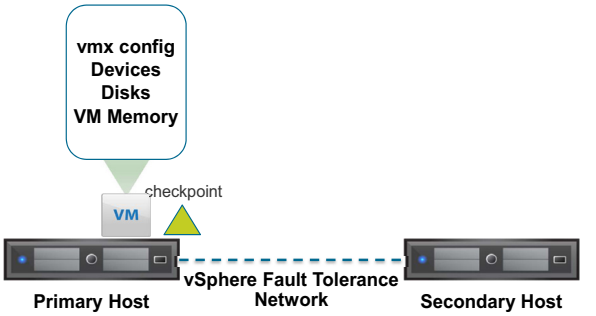


https://t.me/learningnets

9-61 vSphere Fault Tolerance Fast Checkpointing

The symmetric multiprocessor vSphere Fault Tolerance checkpoint interval is dynamic.

It adapts to maximize the workload performance.



To enable vSphere Fault Tolerance for SMP virtual machines, you must do very fast, continuous copying (checkpointing) of the primary host virtual machine. The primary host virtual machine is copied (checkpointed) periodically, and the copies are sent to a secondary host. If the primary host crashes, the virtual machine continues on the secondary host at the point of its last network send.

The goal is to take checkpoints of virtual machines at least every 10 milliseconds with small CPU overhead in the critical path. You can move noncritical path processing to another processor core. The primary virtual machine is continuously copied (checkpointed), and these copies (checkpoints) are sent to a secondary host.

The initial complete copy (checkpoint) is created using a modified form of vSphere vMotion migration to the secondary host. The primary virtual machine holds each outgoing network packet until the following copy (checkpoint) has been sent to the secondary host.

If the primary host crashes, the virtual machine can be resumed on the backup from the last complete copy (checkpoint).

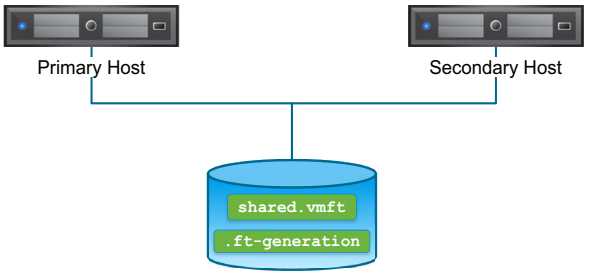
In vSphere Fault Tolerance, checkpoint data makes up the last changed pages of memory. The source virtual machine is paused to access this memory. This pause is typically under one second.

https://t.me/learningnets

9-62 vSphere Fault Tolerance Shared Files

vSphere Fault Tolerance has shared files:

- `shared.vmft` prevents UUID change.
- `.ft-generation` is for the split-brain condition.



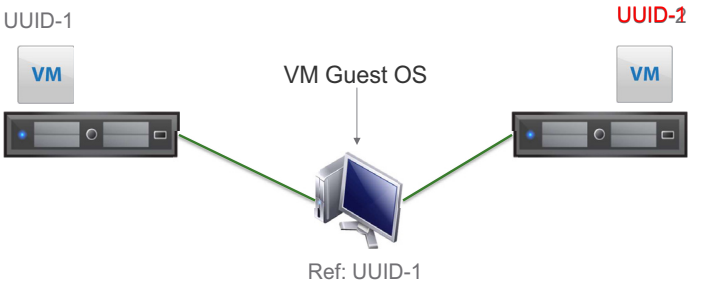
A virtual machine split-brain condition can occur when a host becomes isolated or partitioned from a master host, and the master host cannot communicate with it using heartbeat datastores. In this situation, the master host cannot determine that the host is alive and so declares the host as dead.

The master host tries to restart the virtual machines that are running on the isolated or partitioned host. This attempt succeeds if the virtual machines remained running on the isolated or partitioned host and the isolated or partitioned host lost access to the virtual machines' datastores when it became isolated or partitioned. A split-brain condition exists because two instances of the virtual machine exist. The `.ftgeneration` file ensures that only one instance of the virtual machine can read from or write to the virtual machine's virtual disks.

https://t.me/learningnets

9-63 About the shared.vmft File

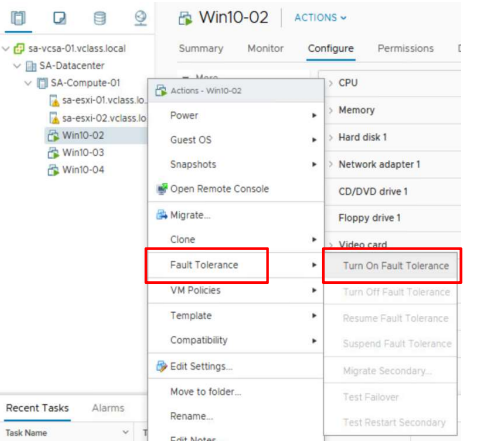
The `shared.vmft` file, which is found on a shared datastore, is the vSphere Fault Tolerance metadata file. This file contains the primary and secondary instance UUIDs and the primary and secondary vmx paths.



https://t.me/learningnets

9-64 Enabling vSphere Fault Tolerance on a Virtual Machine

You can turn on vSphere Fault Tolerance for a virtual machine through vSphere Web Client.



After you have taken all of the required steps for enabling vSphere Fault Tolerance for your cluster, you can use the feature by turning it on for individual virtual machines.

Before vSphere Fault Tolerance can be turned on, validation checks are performed on a virtual machine.

After these checks are passed and you turn on vSphere Fault Tolerance for a virtual machine, new options are added to the Fault Tolerance section of its context menu. These options include turning off or disabling vSphere Fault Tolerance, migrating the secondary virtual machine, testing failover, and testing restart of the secondary virtual machine.

When vSphere Fault Tolerance is turned on, vCenter Server resets the virtual machine's memory limit and sets the memory reservation to the memory size of the virtual machine. While vSphere Fault Tolerance remains turned on, you cannot change the memory reservation, size, limit, number of virtual CPUs, or shares. You also cannot add or remove disks for the virtual machine. When vSphere Fault Tolerance is turned off, any parameters that were changed are not reverted to their original values.

https://t.me/learningnets

9-65 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe the benefits of using vSphere Fault Tolerance
- Explain how vSphere Fault Tolerance works
- Describe main vSphere Fault Tolerance features
- Describe how vSphere Fault Tolerance works with vSphere HA and vSphere DRS
- Enable vSphere Fault Tolerance from vSphere Web Client

https://t.me/learningnets



https://t.me/learningnets

9-67 Learner Objectives

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

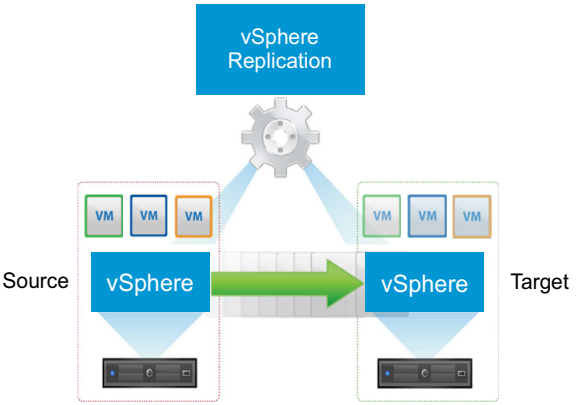
- Describe vSphere Replication
- Configure vSphere Replication
- Recover replicated virtual machines.
- Identify the benefits of offloaded backup processing and changed-block tracking

https://t.me/learningnets

9-68 About vSphere Replication

vSphere Replication is an extension for vCenter Server.

It provides hypervisor-based virtual machine replication and recovery.



vSphere Replication is an alternative to storage-based replication. vSphere Replication protects virtual machines from partial or complete site failures by replicating the virtual machines between the following sites:

From a source site to a target site:

- Within a single site from one cluster to another
- From multiple source sites to a shared remote target site

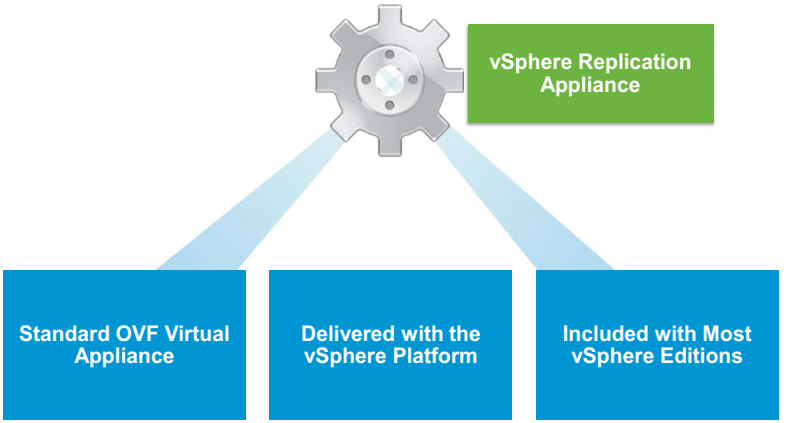
vSphere Replication provides several benefits as compared to storage-based replication:

- Data protection at lower cost per virtual machine than storage array-based replication
- A replication solution that enables flexibility in storage vendor selection at the source and target sites
- Overall lower cost per replication than storage array-based replication

https://t.me/learningnets

9-69 About the vSphere Replication Appliance

The vSphere Replication appliance provides all the components required to perform vSphere replication.



The contents of the vSphere Replication appliance include:

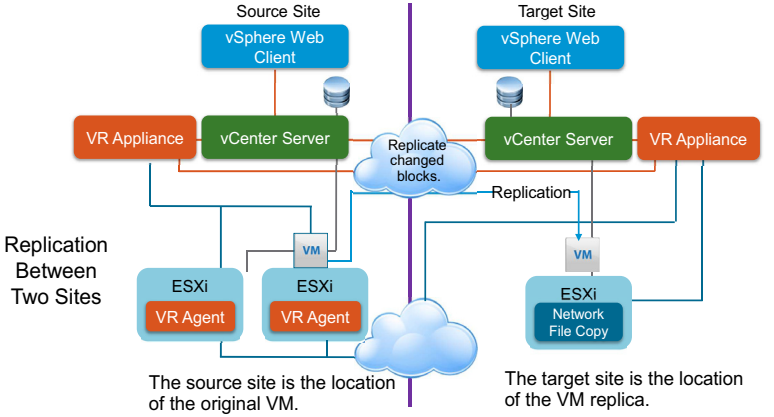
- A vSphere Replication server that provides the core of the vSphere Replication infrastructure
- An embedded database that stores replication configuration and management information
- A vSphere Replication management server:
  - Configures the vSphere Replication server
  - Enables, manages, and monitors replications
  - Authenticates users and checks their permissions to perform vSphere Replication operations
- A plug-in to vSphere Web Client that provides a user interface for vSphere Replication

You can use vSphere Replication immediately after you deploy the appliance. The vSphere Replication appliance provides the virtual appliance management interface (VAMI) that is used to 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

9-70 Replication Functions

vSphere Replication enables replication of a VM from a source site to a target site, monitoring and managing the replication status, and recovering the VM at the target site.

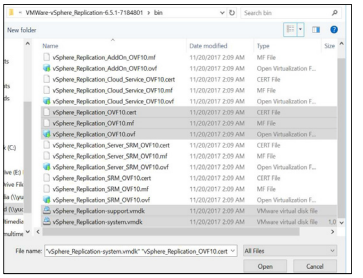


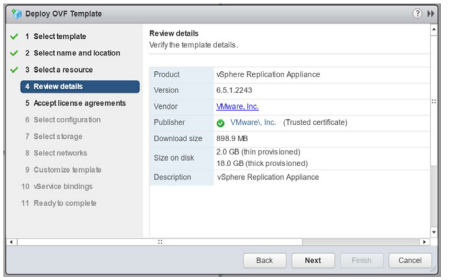
https://t.me/learningnets

9-71 Deploying the vSphere Replication Virtual Appliance

To use vSphere Replication, you must vSphere Web Client to deploy the vSphere Replication appliance on an ESXi host:

- Download the installation package to a local directory, or obtain its URL.
- Use the standard vSphere OVF deployment wizard to deploy the appliance:
  - Select **vCenter > Hosts and Clusters**.
  - Right-click a host and select **Deploy OVF template**.





The vSphere Replication appliance registers as an extension with the corresponding vCenter Server instance.

The vSphere Replication 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. Additional vSphere Replication servers that you deploy are themselves virtual appliances. You must register any additional vSphere Replication server with the vSphere Replication appliance on the corresponding site.

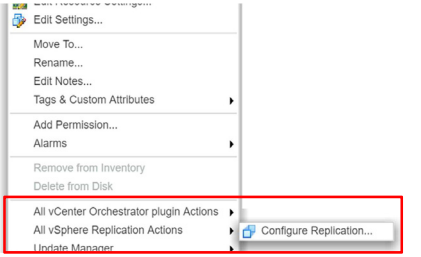
The vSphere Replication appliance provides the VAMI. You can use this interface to reconfigure the vSphere Replication database, network settings, public-key certificates, and passwords for the appliances.vSphere Replication can be deployed with either IPv4 or IPv6 address. Mixing IP addresses, for example having a single appliance with an IPv4 and an IPv6 address, is not supported. To register as an extension, vSphere Replication relies on the VirtualCenter.FQDN property of the vCenter Server. When an IPv6 address is used for vSphere Replication, the VirtualCenter.FQDN property must be set to a fully qualified domain name that can be resolved to an IPv6 address or to a literal address. When operating with an IPv6 address, vSphere Replication requires that all components in the environment, such as vCenter Server and ESXi hosts are accessible using the IPv6 address.

https://t.me/learningnets

9-72 Configuring vSphere Replication for a Single VM

To configure vSphere Replication for a virtual machine, perform one of the following actions in vSphere Web Client:

- Select the virtual machine in the inventory and select **Actions > All vSphere Replication Actions > Configure**.
- Right-click the virtual machine in the inventory and select **All vSphere Replication Actions > Configure**.



vSphere Replication can protect individual virtual machines and their virtual disks by replicating them to another location.

When you configure replication, you set a recovery point objective (RPO) to determine the period of time between replications. For example, an RPO of 1 hour seeks to ensure that a virtual machine loses no more than 1 hour of data during the recovery. For smaller RPOs, less data is lost in a recovery, but more network bandwidth is consumed keeping the replica up to date.

Every time that a virtual machine reaches its RPO target, vSphere Replication records approximately 3800 bytes of data in the vCenter Server events database. If you set a low RPO period, this can quickly create a large volume of data in the database. To avoid creating large volumes of data in the vCenter Server events database, limit the number of days that vCenter Server retains event data. See the section on configuring the database retention policy in *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-vcenter-server-67-host-management-guide.pdf. Alternatively, set a higher RPO value.

vSphere Replication guarantees crash consistency amongst all the disks that belong to a virtual machine. If you use Microsoft Volume Shadow Copy Service (VSS) quiescing, you might obtain a higher level of consistency. The available quiescing types are determined by the virtual machine's operating system. For VSS quiescing support for Windows virtual machines, see Compatibility Matrixes for vSphere Replication 5.8 at https://www.vmware.com/support/vsphere-replication/doc/vsphere-replication-compat-matrix-5-8.html.

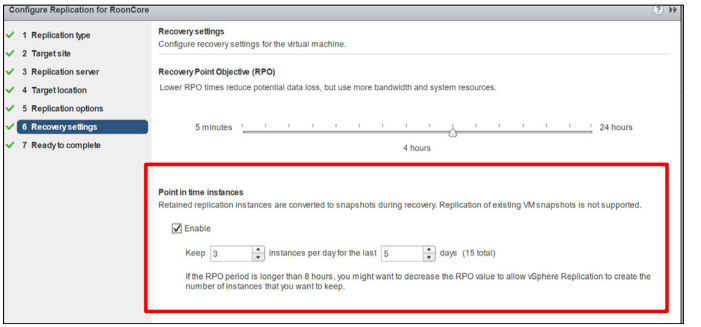
You can use vSphere Replication with a vSAN datastore on the source and target sites. See *Using vSphere Replication with Virtual SAN Storage* at http://docs.vmware.com for the limitations when using vSphere Replication with vSAN.

https://t.me/learningnets

9-73 Configuring Multiple Points in Time and the Recovery Point Objective

You use vSphere Web Client to configure multiple points in time during vSphere Replication configuration.

The value that you set with the **Recovery Point Objective (RPO)** slider affects replication scheduling.



When you configure a virtual machine for replication, the vSphere Replication agent sends changed blocks in the virtual machine disks from the source site to the target site, where they are applied to the copy of the virtual machine. This process occurs independently of the storage layer. vSphere Replication performs an initial full synchronization of the source virtual machine and its replica copy. You can use replication seeds to reduce the amount of time and bandwidth required for the initial replication.

During replication configuration, you can set an RPO and enable retention of instances from multiple points in time.

As administrator, you can monitor and manage the status of the replication. You can view information for incoming and outgoing replications, source and target site status, replication issues, warnings, and errors.

When you manually recover a virtual machine, vSphere Replication creates a copy of the virtual machine connected to the replica disk, but does not connect any of the virtual network cards to port groups. You can review the recovery and status of the replica virtual machine and attach it to the networks. You can recover virtual machines at different points in time, such as the last known consistent state. vSphere Replication presents the retained instances as ordinary virtual machine snapshots to which you can revert the virtual machine.

vSphere Replication stores replication configuration data in its embedded database. You can also configure vSphere Replication to use an external database.

https://t.me/learningnets

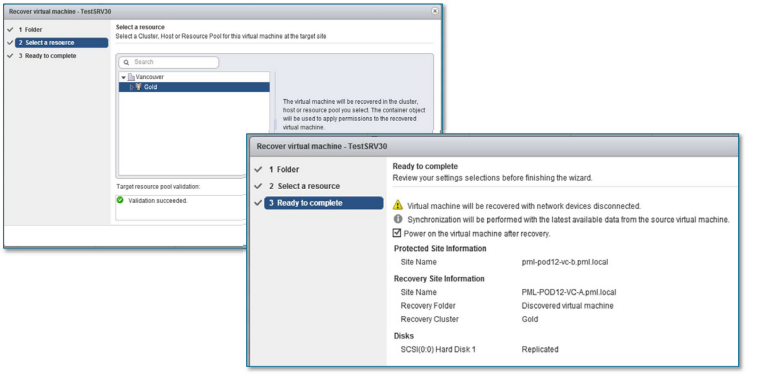
You can replicate a virtual machine between two sites. vSphere Replication is installed on both source and target sites. Only one vSphere Replication appliance is deployed on each vCenter Server system. You can deploy additional vSphere Replication servers.

https://t.me/learningnets

9-74 Recovering Virtual Machines

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.



You use the Recover virtual machine wizard in vSphere Web Client at the target site to perform the recovery.

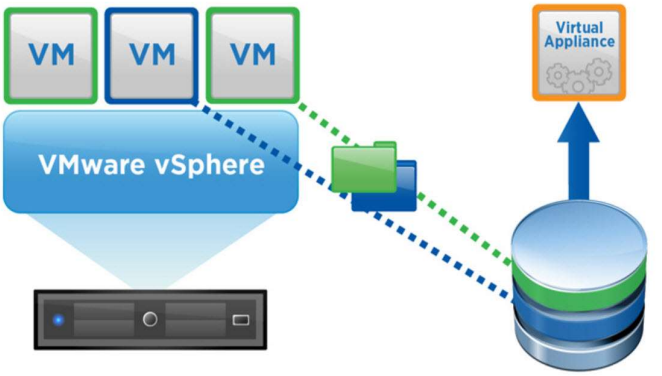
You are asked to select either to recover the virtual machine with all the latest data or to recover the virtual machine with the most recent data available on the target site. If you select **Recover with recent changes** to avoid data loss, vSphere Replication performs a full synchronization of the virtual machine from the source site to the target site before recovering the virtual machine. This option requires that the data of the source virtual machine be accessible. You can only select this option if the virtual machine is powered off. If you select **Recover with latest available data**, vSphere Replication recovers the virtual machine by using the data from the most recent replication on the target site, without performing synchronization. Selecting this option results in the loss of any data that has changed since the most recent replication. Select this option if the source virtual machine is inaccessible or if its disks are corrupted.

vSphere Replication validates the input that you provided and recovers the virtual machine. If successful, the virtual machine status changes to Recovered. The virtual machine appears in the inventory of the target site.

https://t.me/learningnets

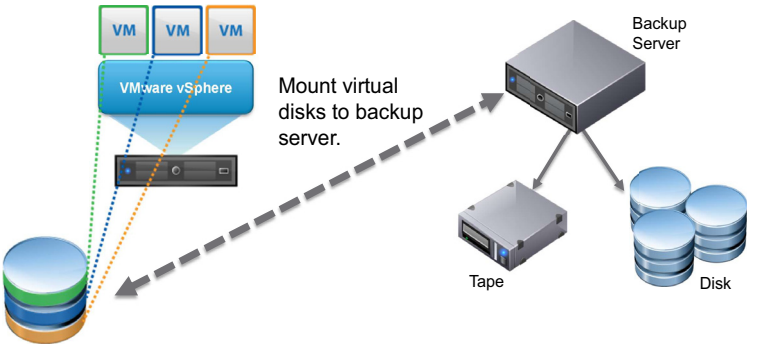
9-75 vSphere Storage APIs - Data Protection

vSphere Storage APIs - Data Protection enables backup products to perform centralized, efficient, off-host, LAN-free backups of vSphere virtual machines.



https://t.me/learningnets

9-76 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.

Perhaps one of the biggest bottlenecks that limits backup performance can be the backup server that is handling all the backup coordination tasks. One of these backup tasks is copying data from point A to point B. Other backup tasks do a lot of CPU processing. For example, tasks are performed to determine what data to back up and what not to back up. Other tasks are performed to deduplicate data and compress data that is written to the target.

A server with insufficient CPU resources can greatly reduce backup performance. It is important not to skimp on resources for your backup server. A physical server or virtual machine with an ample amount of memory and CPU capacity is necessary for the best backup performance possible.

The motivation to use LAN-free backups is to reduce the stress on the physical resources of the ESXi host when virtual machines are backed up. LAN-free backups reduce the stress by offloading backup processing from the ESXi host to a backup proxy server.

You can configure your environment for LAN-free backups to the backup server, also called the backup proxy server. For LAN-free backups, the backup server must be able to access the storage managed by the ESXi hosts on which the virtual machines to back up are running.

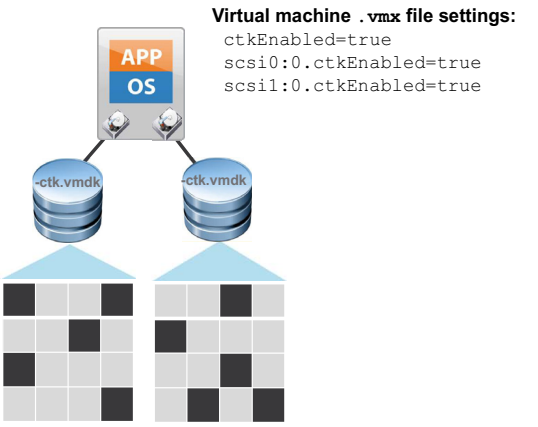
If you are using NAS or direct-attached storage, ensure that the backup proxy server is accessing the volumes with a network-based transport. If you will be running a direct SAN backup, zone the SAN and configure the disk subsystem host mappings. The host mappings must be configured so that all ESXi hosts and the backup proxy server access the same disk volumes.

https://t.me/learningnets

9-77 About Changed-Block Tracking

With changed-block tracking, you copy only file blocks that have changed since the last backup.

Changed-block tracking enables faster incremental backups.



Changed-block tracking (CBT) is a VMkernel feature that keeps track of the storage blocks of virtual machines as they change over time. The VMkernel keeps track of block changes on virtual machines, which enhances the backup process for applications that have been developed to exploit vSphere Storage APIs - Data Protection.

To increase the efficiency of image-level backups, vSphere Data Protection uses the CBT feature, which greatly reduces the backup time of a given virtual machine image and provides the ability to process a large number of virtual machines within a particular backup window.

By leveraging CBT during restores, vSphere Data Protection offers fast and efficient recoveries of virtual machines to their original location. During a restore process, vSphere Data Protection uses CBT to determine which blocks have changed since the last backup. The use of CBT reduces data transfer within the vSphere environment during a recovery operation and, more importantly, reduces the recovery time.

Additionally, Data Protection evaluates the workload between both restore methods (full image restore or a recovery leveraging CBT) and performs the method resulting in the fastest restore time. This is useful in scenarios where the change rate since the last backup in a virtual machine being restored is very high and the overhead of a CBT analysis operation would be more costly than a direct full-image recovery. Data Protection determines which method results in the fastest image recovery times for virtual machines in the environment.

https://t.me/learningnets

9-78 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe vSphere Replication
- Configure vSphere Replication
- Recover replicated virtual machines.
- Identify the benefits of offloaded backup processing and changed-block tracking

https://t.me/learningnets

9-79 Key Points

- vSphere HA restarts virtual machines on the remaining hosts in the cluster.
- Hosts in vSphere HA clusters have a master-slave relationship.
- You implement redundant heartbeat networks either with NIC teaming or by creating additional heartbeat networks.
- vSphere Fault Tolerance provides zero downtime for applications that must be available at all times.
- vSphere Replication can be used to protect virtual machines as part of a disaster recovery strategy.
- vSphere Replication is the only hypervisor-based replication solution that operates at the individual VMDK level, enabling replication between datastores hosted on any storage.

Questions?

https://t.me/learningnets

MODULE 10

vSphere DRS



https://t.me/learningnets

10-2 You Are Here

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

https://t.me/learningnets

10-3 Importance

As you scale your vSphere environment, you must be aware of the vSphere features and functions that will help you manage the hosts in your environment.

With vSphere DRS, you can improve service levels by guaranteeing appropriate resources to virtual machines. And you can monitor and manage more infrastructure per system administrator.

https://t.me/learningnets

10-4 Learner Objectives

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

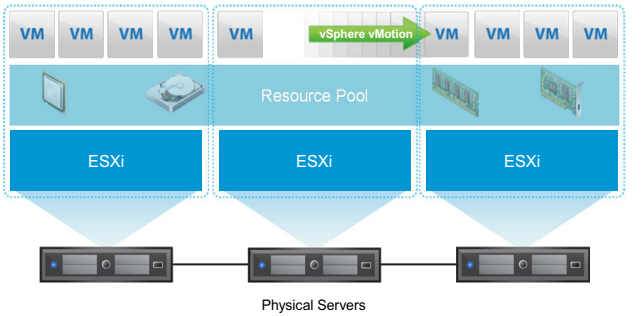
- Describe the functions of a vSphere DRS cluster
- Create a vSphere DRS cluster
- Review information about a vSphere DRS cluster
- Remove a host from a vSphere DRS cluster

https://t.me/learningnets

10-5 About vSphere DRS

vSphere DRS aggregates computing capacity across a collection of servers into logical resource pools.

vSphere DRS allocates available resources among the virtual machines based on predefined rules that reflect business needs and changing priorities.



vSphere DRS helps improve resource allocation across all hosts and resource pools. vSphere DRS collects resource use information for all hosts and virtual machines in the cluster and gives recommendations (or migrates virtual machines) in one of two situations:

- Initial placement: When you power on a virtual machine in the cluster for the first time, vSphere DRS either places the virtual machine or makes a recommendation.
- Load balancing: vSphere DRS attempts to improve resource use across the cluster by performing automatic migrations of virtual machines (vSphere vMotion) or by providing a recommendation for virtual machine migrations.

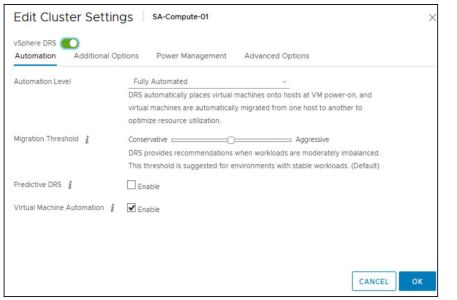
https://t.me/learningnets

10-6 vSphere DRS Cluster Prerequisites

vSphere DRS works best when the VMs meet vSphere vMotion migration requirements.

To use vSphere DRS for load balancing, the hosts in the cluster must be part of a vSphere vMotion migration network. If they are not, vSphere DRS can still make initial placement recommendations.

To use shared storage, configure all hosts in the cluster. All hosts in the cluster must have access to the same datastores.



An VMware ESXi™ host that is added to a vSphere DRS cluster must meet certain prerequisites to use cluster features:

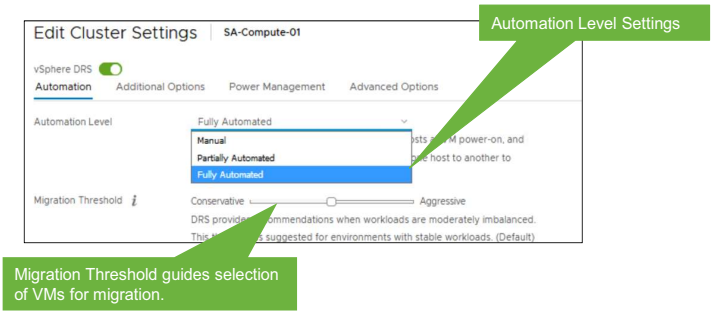
- vSphere DRS works best if the virtual machines meet vSphere vMotion requirements.
- To use vSphere DRS for load balancing, the hosts in your cluster must be a part of a vSphere vMotion network.
- Configure all managed hosts to use shared storage: VMFS, vSAN, vSphere Virtual Volumes, or NFS datastores.
- Place the disks of all virtual machines on shared storage that is accessible by source and destination hosts.

vSphere DRS clusters can be created, or vSphere DRS can be enabled for existing vSphere HA or vSAN clusters.

https://t.me/learningnets

10-7 vSphere DRS Cluster Settings: Automation Level

Configure the automation level for the initial placement of virtual machines and dynamic balancing while virtual machines are running.



To create a vSphere DRS cluster, right-click your data center in the inventory and select **New Cluster**. The New Cluster dialog box appears. Enter a descriptive name for your cluster, and select the **Turn On VMware DRS** check box.

In the vSphere DRS window that appears (shown on the slide), you define the automation level. The automation level determines whether vSphere DRS makes migration recommendations or automatically places virtual machines on hosts. vSphere DRS makes placement decisions when a virtual machine powers on and when virtual machines must be rebalanced across hosts in the cluster. The following automation levels are available:

- **Manual**: When you power on a virtual machine, vSphere DRS displays a list of recommended hosts on which to place the virtual machine. When the cluster becomes imbalanced, vSphere DRS displays recommendations for virtual machine migration.
- **Partially automated**: When you power on a virtual machine, vSphere DRS places it on the best-suited host. When the cluster becomes imbalanced, vSphere DRS displays recommendations for manual virtual machine migration.
- **Fully automated**: When you power on a virtual machine, vSphere DRS places it on the best-suited host. When the cluster becomes imbalanced, vSphere DRS migrates virtual machines from overused hosts to underused hosts to ensure balanced use of cluster resources.

The migration threshold determines how aggressively vSphere DRS selects to migrate virtual machines:

- Level 1 (Conservative): Applies only priority 1 recommendations. vCenter Server applies only recommendations that must be taken to satisfy cluster constraints, such as affinity rules and host maintenance.

https://t.me/learningnets

- Level 2: Apply priority 1 and priority 2 recommendations. vCenter Server applies recommendations that promise a significant improvement to the cluster's load balance.
- Level 3 (default): Apply priority 1, priority 2, and priority 3 recommendations. vCenter Server applies recommendations that promise at least good improvement to the cluster's load balance.
- Level 4: Apply priority 1, priority 2, priority 3, and priority 4 recommendations. vCenter Server applies recommendations that promise even a moderate improvement to the cluster's load balance.
- Level 5 (Aggressive): Apply all recommendations. vCenter Server applies recommendations that promise even a slight improvement to the cluster's load balance.

https://t.me/learningnets

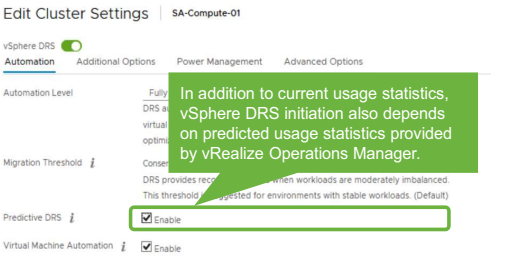
10-8 vSphere DRS Cluster Settings: Predictive DRS

vSphere DRS and VMware vRealize® Operations Manager™ combine data to predict future demand and determine when and where hot spots will occur.

The vSphere DRS data collector retrieves statistics from the following sources:

- Resource usage statistics from hosts
- Predictive statistics from the vRealize Operations Manager server

The predicted usage statistics always take precedence over current usage statistics.



When the demands for resources on a virtual machine changes, vSphere DRS itself does not have the ability to predict these resource demands. However, by using statistics from VMware vRealize® Operations Manager™, vSphere DRS can analyze the expected workload on virtual machines in the near future. vSphere DRS can then use the predicted workload values with the current usage values to make placement decisions in advance of the resources being requested.

If the current usage is higher than the predicted usage, the current usage statistics take precedence and these values are used to generate recommendations.

If the predicted usage is higher than the current usage, the predicted usage value is used as current demand of the VM. Thus, the resources are available in advance of the VM needing them, avoiding any performance impact that might occur.

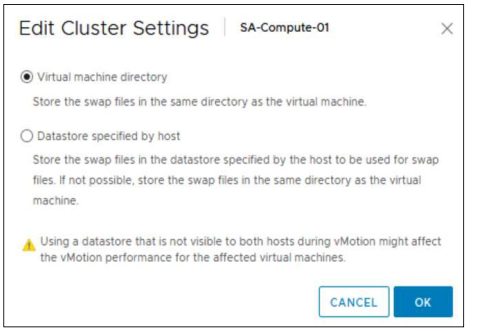
By default, Proactive DRS uses values for one hour in advance to ensure timely migration.

https://t.me/learningnets

10-9 Other Cluster Settings: Swap File Location

By default, swap files for a virtual machine are located on a datastore in the folder containing other virtual machine files.

You can place virtual machine swap files on an alternative datastore.



By default, the swap file is created in the same location as the virtual machine's configuration file, which could either be on a VMFS datastore, 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 virtual machine is powered on. If this file cannot be created, the virtual machine cannot power on. Instead of accepting the default, you can also use the following options:

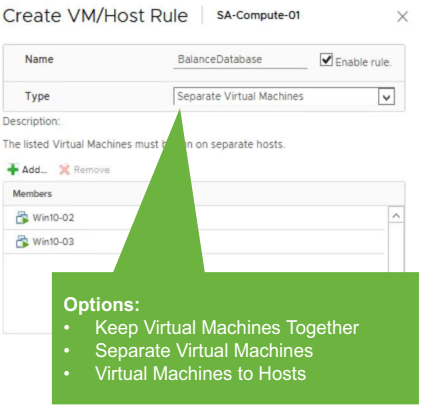
- Use per-virtual machine configuration options to change the datastore to another shared storage location.
- Use host-local swap, which allows you to specify a datastore stored locally on the host. This allows you to swap at a per-host level, saving space on the SAN. However, it can lead to a slight degradation in performance for vSphere vMotion because pages swapped to a local swap file on the source host must be transferred across the network to the destination host. Currently, vSAN and vSphere Virtual Volumes datastores cannot be specified for host-local swap.

https://t.me/learningnets

10-10 vSphere DRS Cluster Settings: Virtual Machine Affinity

vSphere DRS affinity rules specify that selected virtual machines be placed either on the same host (affinity) or on separate hosts (anti-affinity).

- Affinity rules: Use for multi-virtual machine systems where virtual machines communicate heavily with one another.
- Anti-affinity rules: Use for multi-virtual machine systems where load balancing or high availability is desired.



After a vSphere DRS cluster is created, you can edit its properties to create rules that specify affinity. The following types of rules exist:

- Affinity rules: vSphere DRS keeps certain virtual machines together on the same host (for example, for performance reasons).
- Anti-affinity rules: vSphere DRS ensures that certain virtual machines are not together (for example, for availability reasons).

The slide shows an anti-affinity rule that requires two virtual machines (Greg01 and Greg01-2) to be placed on different hosts, possibly due to availability or performance reasons. Conversely, affinity rules can be used to keep certain virtual machines on the same host because of increased locality or performance benefits, for example, if virtual machines are communicating heavily with one another.

If two rules conflict, you are prevented from enabling both.

When you add or edit a rule, and the cluster is immediately in violation of that rule, the system continues to operate and tries to correct the violation.

For vSphere DRS clusters that have a default automation level of manual or partially automated, migration recommendations are based on both rule fulfillment and load balancing.

https://t.me/learningnets

10-11 vSphere DRS Cluster Settings: DRS Groups

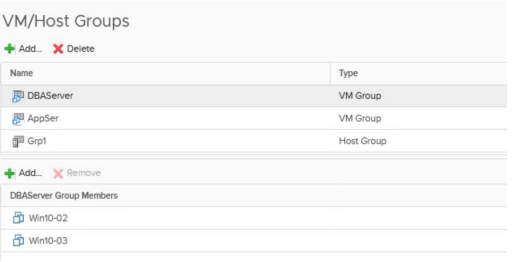
VM or host groups are used in defining VM-Host affinity rules.

Types of VM or host groups:

- VM group: A virtual machine can belong to multiple virtual machine DRS groups.
- Host group: A host can belong to multiple host DRS groups.

A virtual machine can belong to multiple VM groups.

A host can belong to multiple host groups.



The third vSphere DRS affinity rule option is to set a VM-Host affinity rule. This type of rule specifies whether virtual machines can or cannot be run on a host. For ease of administration, virtual machines can be placed in VM or host groups. You can create one or more VM or host groups in a vSphere DRS cluster, each consisting of one or more virtual machines. On the slide, DBAServer is a virtual machine DRS group. A virtual machine can belong to more than one virtual machine DRS group.

Likewise, a host DRS group consists of one or more ESXi hosts. A host can belong to more than one host DRS group. On the slide, Grp1 is a host DRS group.

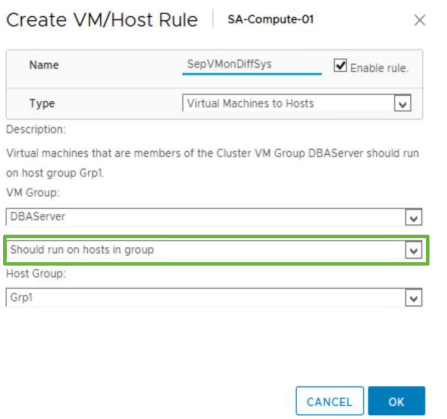
The main use for DRS groups is to assist in defining DRS rules known as the VM-Host affinity rules.

https://t.me/learningnets

10-12 vSphere DRS Cluster Settings: VM-Host Affinity Rules

A VM-Host affinity rule:

- Specifies an affinity (or anti-affinity) relationship between a virtual machine DRS group and a host DRS group
- Is either a required rule or a preferential rule



Options:
- Must run on hosts in group
- Must not run on hosts in group
- Should 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 virtual machine DRS group can run on the members of a specific host DRS group.

Unlike an affinity rule for virtual machines which specifies affinity (or anti-affinity) between individual virtual machines, a VM-Host affinity rule specifies an affinity relationship between a group of virtual machines and a group of hosts. Rules are either required or preferential.

A VM-Host affinity rule includes the following components:

- Virtual machine DRS group
- Host DRS group
- Designation of whether the rule is a requirement (must) or a preference (should) and whether it is affinity (run on) or anti-affinity (not run on)

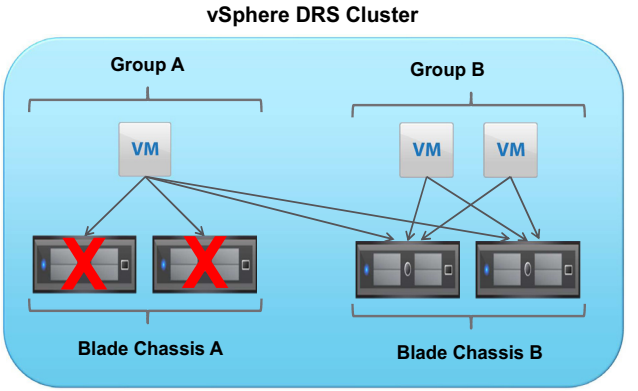
Because VM-Host affinity rules are cluster-based, the virtual machines and hosts that are included in a rule must all reside in the same cluster. If a virtual machine is removed from the cluster, it loses its membership from all virtual machine groups, including the DRS group affiliation, even if it is later returned to the cluster.

https://t.me/learningnets

10-13 VM-Host Affinity Rule: Preferential

A preferential rule is softly enforced and can be violated if necessary.

Example: Separate virtual machines on different blade systems.



A preferential rule is softly enforced. Preferential rules can be violated to allow the proper functioning of vSphere DRS, vSphere HA, and VMware vSphere® Distributed Power Management™.

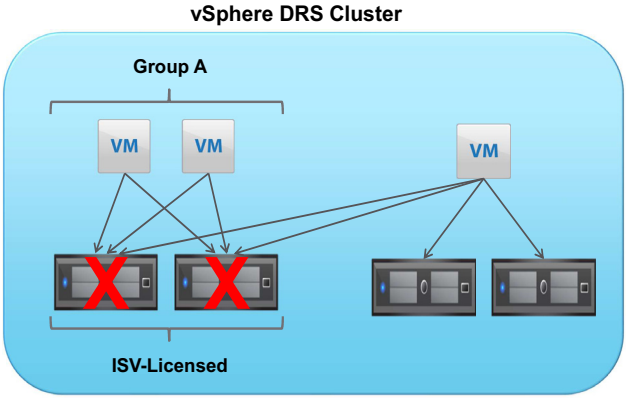
A preferential rule might be used for separating virtual machines onto different blade systems for better performance. On the slide, Group A and Group B are virtual machine DRS groups. Blade Chassis A and Blade Chassis B are host DRS groups. The goal is to force the virtual machines in Group A to run on the hosts in Blade Chassis A and to force the virtual machines in Group B to run on the hosts in Blade Chassis B. If the hosts fail, vSphere HA restarts the virtual machines on the other hosts in the cluster. If the hosts are put into maintenance mode or become overused, vSphere DRS moves the virtual machines on the other hosts in the cluster.

https://t.me/learningnets

10-14 VM-Host Affinity Rule: Required

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 virtual machines has licensing restrictions. You can enforce this rule when the software running in your virtual machines has licensing restrictions. You can place such virtual machines in a DRS group. Then create a rule that requires the virtual machines to run on a host DRS group, which contains hosts with the required licenses.

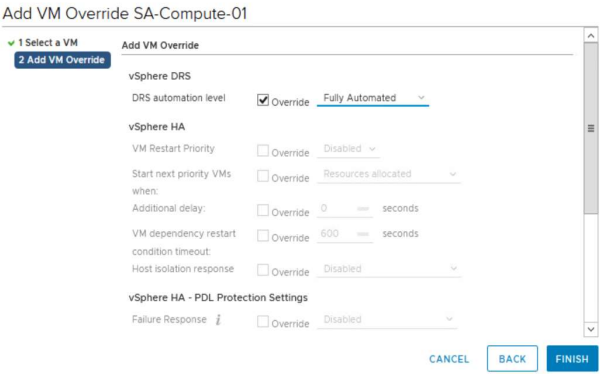
When you create a VM-Host affinity rule that is based on the licensing or hardware requirements of the software running in your virtual machines, you are responsible for ensuring that the groups are properly set up. The rule does not monitor the software running in the virtual machines nor does it know which third-party licenses are in place on which ESXi hosts.

On the slide, Group A is a virtual machine DRS group. You can force Group A to run on hosts in the group called ISV-Licensed to ensure that the virtual machines in Group A run on hosts that have the required licenses. But if the hosts in group ISV-Licensed fail, vSphere HA cannot restart the virtual machines in Group A on hosts that are not in the group. If the hosts in group ISV-Licensed are put into maintenance mode or become overused, vSphere DRS cannot move the virtual machines in Group A to hosts that are not in the group.

https://t.me/learningnets

10-15 vSphere DRS Cluster Settings: VM-Level Automation

You can customize the automation level for individual virtual machines in a cluster to override the automation level set on the entire cluster.



Setting the automation level for individual virtual machines enables you to fine-tune automation to suit your needs. For example, you might have a virtual machine that is especially critical to your business and you would like more control over its placement. Set its automation level to manual. If a virtual machine's automation level is set to disabled, vCenter Server does not migrate that virtual machine or provide migration recommendations for it.

As a best practice, enable automation. Select the automation level based on your environment and level of comfort. For example, if you are new to vSphere DRS clusters, you might select **Partially Automated** because you want control over the placement of virtual machines. When you are comfortable with what vSphere DRS does and how it works, you might set the automation level to **Fully Automated**. Set the automation level of **Manual** on virtual machines over which you want to exercise more control, such as your business-critical virtual machines.

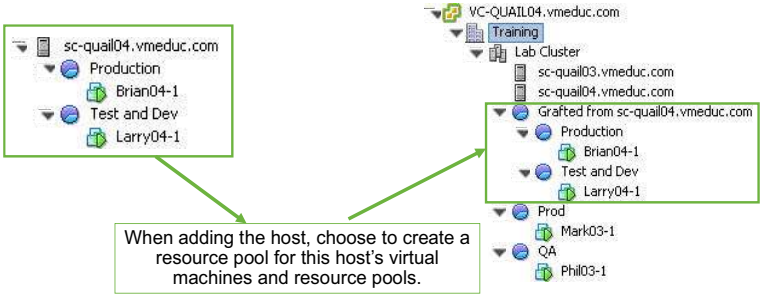
https://t.me/learningnets

10-16 Adding a Host to a Cluster

When adding or moving a host into a vSphere DRS cluster, you can keep the resource pool hierarchy of the existing host.

If vSphere DRS is enabled, host resource pools are lost.

Example: Add sc-quail04 to Lab Cluster.



To add a host to a vSphere DRS cluster, drag an ESXi host onto the cluster object in the inventory. Use the Add Host wizard to complete the process.

When adding a host with resource pools to a vSphere DRS cluster, you must decide on the placement of the resource pools. By default, the resource pool hierarchy is discarded and the host is added at the same level as the virtual machines. However, you can choose to retain the resource pool hierarchy of the existing host by ensuring that vSphere DRS is enabled. You can also choose to graft the host's resource pools onto the cluster's resource pool hierarchy. Grafting refers to adding the branches of the host's tree to the branches of the cluster's tree, as fruit tree branches are grafted onto rootstock.

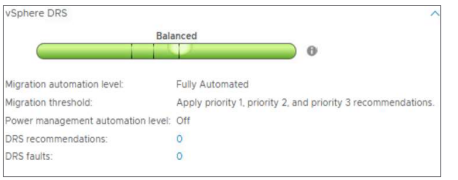
vSphere Web Client does not allow saving the resource pool as Grafted from *host_name*. You must enter a name for the resource pool.

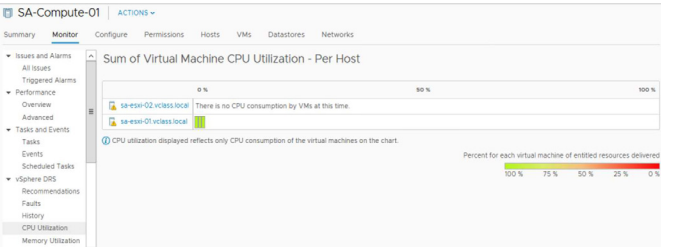
https://t.me/learningnets

10-17 Viewing vSphere DRS Cluster Information

The cluster **Summary** tab provides information specific to vSphere DRS.

Clicking the vSphere DRS page on the **Monitor** tab displays CPU, memory, and network utilization per host.





The vSphere DRS pane on the cluster's **Summary** tab appears only when vSphere DRS is enabled. This section provides the following information related to vSphere DRS:

- Selected automation levels
- Number of vSphere DRS recommendations and faults
- Configured migration threshold

A priority level for each migration recommendation is computed using the load-imbalance metric of the cluster. This metric is displayed as Current host load standard deviation on the cluster's **Summary** tab in vSphere Web Client. A higher load imbalance leads to higher-priority migration recommendations.

On the **Monitor** tab, click the vSphere DRS page. Click the **CPU Utilization**, the **Network Utilization**, or the **Memory Utilization** link to view a chart of all the hosts in the cluster and how their CPU and memory resources are allocated to each virtual machine.

For CPU use, the virtual machine information is represented by a colored box. If you point to the colored box, the virtual machine's CPU usage information appears. If the virtual machine is receiving the resources that it is entitled to, the box is green. Green means that 100 percent of the virtual machine's entitled resources 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).

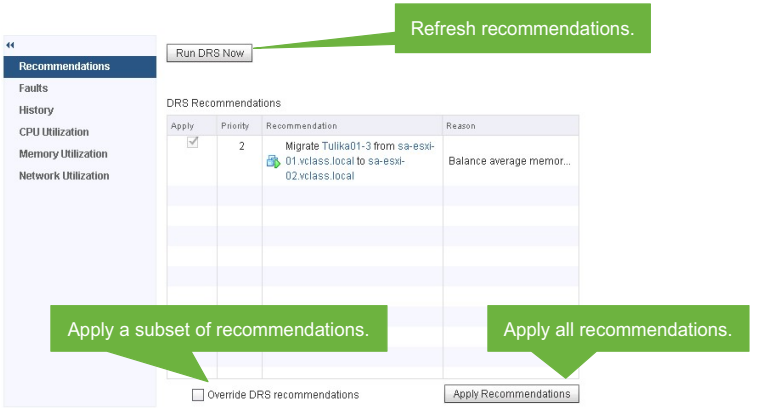
For memory usage, the virtual machine boxes are not color-coded because the relationship between consumed memory and entitlement is often not easily categorized.

For network usage, the displayed network data reflects all traffic across physical network interfaces on the host.

https://t.me/learningnets

10-18 Viewing vSphere DRS Recommendations

The **vSphere DRS** tab displays information about the vSphere DRS recommendations made for the cluster, the faults that occurred when applying such recommendations, and the history of vSphere DRS actions.



When you select the cluster object, click the **Monitor** tab and the **vSphere DRS** tab. The first three views are: Recommendations, Faults, and History.

In the Recommendations view, you can see the current set of recommendations generated for optimizing resource usage in the cluster through either migrations or power management. Only manual recommendations awaiting user confirmation appear on this list:

To refresh the recommendations, click **Run DRS Now**.

To apply all recommendations, click **Apply Recommendations**.

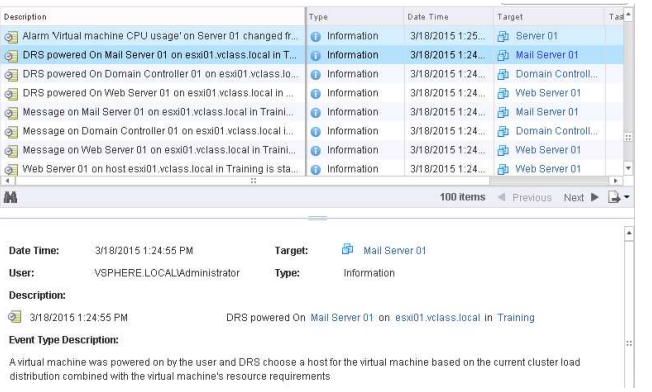
To apply a subset of the recommendations, select the **Override DRS recommendations** check box. Select the check box next to each desired recommendation and click **Apply Recommendations**.

https://t.me/learningnets

10-19 Monitoring Cluster Status

Select the cluster object in the inventory to view the state of the cluster and all its hosts and virtual machines.

You can view the cluster's **Tasks** and **Events** tabs for more information.



vSphere Web Client indicates whether a vSphere DRS cluster is valid, overcommitted (yellow), or invalid (red). vSphere DRS clusters become overcommitted or invalid for several reasons:

- A cluster might become overcommitted if a host fails.
- A cluster becomes invalid if vCenter Server is unavailable and you use VMware Host Client to power on virtual machines.
- A cluster becomes invalid if the user reduces the reservation on a parent resource pool while a virtual machine is in the process of failing over.
- If changes are made to hosts or virtual machines by using vSphere Web Client while vCenter Server is unavailable, those changes take effect. When vCenter Server becomes available again, you might find that clusters have turned red or yellow because cluster requirements are no longer met.

For more information about cluster states, see *vSphere Resource Management Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

10-20 Maintenance Mode and Standby Mode

To service a host in a cluster (for example, to install more memory or remove a host from a cluster), you must place the host in maintenance mode:

- VMs on the host should be migrated to another host or shut down.
- You cannot power on or migrate VMs to a host entering maintenance mode.
- While in maintenance mode, the host does not allow you to deploy or power on a VM.
  - When DRS is in fully automated mode, Powered on VM is automatically migrated from a host that is being placed into maintenance mode.

When a host is placed in standby mode, it is powered off.

Standby mode is used by VMware vSphere® Distributed Power Management™ to optimize power usage.

You place a host in maintenance mode when you need to service it, for example, to install more memory. A host enters or leaves maintenance mode only as the result of a user request.

Virtual machines that are running on a host entering maintenance mode must be shut down or migrated to another host, either manually or automatically, by vSphere DRS. The host continues to run the Enter Maintenance Mode task until all virtual machines are powered down or moved away. You cannot power on virtual machines or migrate virtual machines to a host entering maintenance mode.

When no more running virtual machines are on the host, the host's icon indicates that it has entered maintenance mode. The host's **Summary** tab indicates the new state. While in maintenance mode, the host prevents you from deploying or powering on a virtual machine.

When a host machine is placed in standby mode, it is powered off.

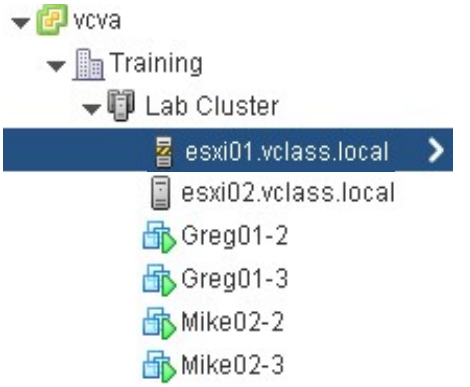
Normally, host machines are placed in standby mode by vSphere DPM to optimize power usage. When a host is placed in standby mode, it is powered off. You can also place a host in standby mode manually. However, the next time that vSphere DRS runs, it might undo your change or recommend that you undo the changes. If you want a host to remain powered off, place it in maintenance mode and turn it off.

https://t.me/learningnets

10-21 Removing a Host from the vSphere DRS Cluster

Before removing a host from a vSphere DRS cluster, consider the following issues:

- The resource pool hierarchy remains with the cluster.
- Because a host must be in maintenance mode, all virtual machines running on the host are powered off.
- The resources available for the cluster decrease.



To remove a host from a cluster, right-click the host in the inventory and select **Enter Maintenance Mode**. After the host is in maintenance mode, drag it to a different inventory location, for example, the data center or another cluster.

Before you remove a host from a vSphere DRS cluster, consider the following issues:

- When you remove a host from a cluster, the host retains only the root resource pool, even if you used a vSphere DRS cluster and grafted the host resource pool when you added the host to the cluster. The hierarchy remains with the cluster. You can create a host-specific resource pool hierarchy.
- When a host is put into maintenance mode, all its running virtual machines must be shut down, suspended, or migrated to other hosts by using vSphere vMotion. Virtual machines with disks on local storage must be powered off, suspended, or migrated to another host and datastore using shared-nothing vSphere vMotion migration. When you remove the host from the cluster, the virtual machines that are currently associated with the host are also removed from the cluster.
- If you remove a host from a cluster, the resources available for the cluster decrease. If the cluster still has enough resources to satisfy the reservations of all virtual machines and resource pools in the cluster, the cluster adjusts resource allocation to reflect the reduced amount of resources. If the cluster lacks the resources to satisfy the reservations of all resource pools but has enough resources to satisfy the reservations for all virtual machines, an alarm is issued. vSphere DRS continues to run.

https://t.me/learningnets

10-22 Disabling vSphere DRS and Restoring a Resource Pool Tree

You may turn off vSphere DRS for a cluster.

If you disable vSphere DRS, the resource pools are removed from the cluster.

When vSphere DRS is disabled, the cluster's resource pool hierarchy and affinity rules are not reestablished when vSphere DRS is turned back on.

To avoid losing the resource pools, save a snapshot of the resource pool tree on your local machine.

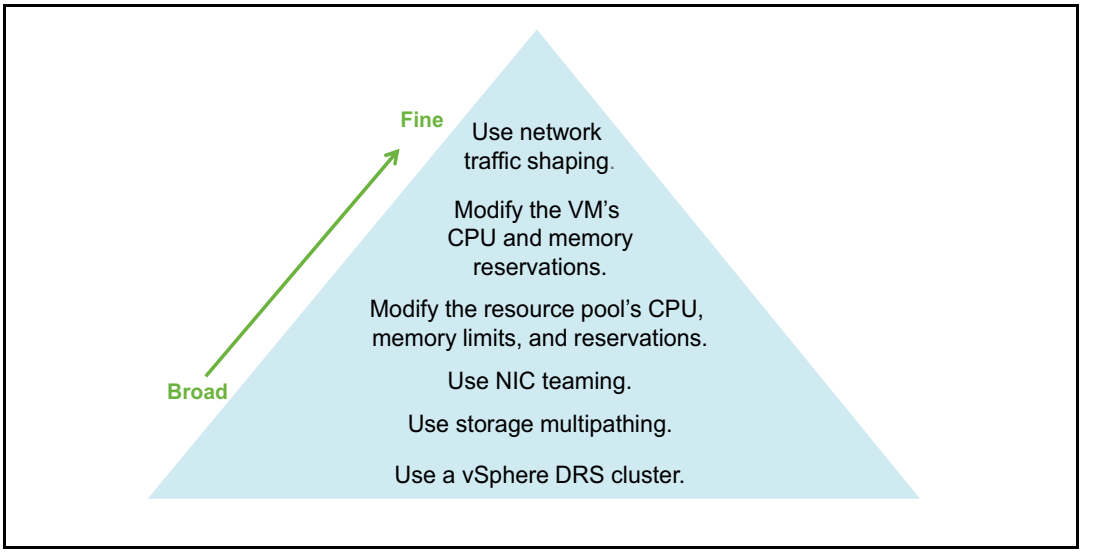
You can restore a previously saved resource pool tree snapshot when you enable vSphere DRS.

You can use the snapshot to restore the resource pool when you enable vSphere DRS. To restore a resource pool tree, you need to meet the following prerequisites:

- vSphere DRS must be turned on.
- You can restore a snapshot only on the same cluster where it was taken.
- No other resource pools are present in the cluster.

https://t.me/learningnets

10-23 Improving Virtual Machine Performance Methods



Review the methods used to improve a virtual machine's performance. The methods are listed from specific to broad. Specific methods affect a particular virtual machine. Broad methods affect several entities. The following methods can be used to improve the performance of a virtual machine:

- If a virtual machine is network-constrained, use network traffic shaping to give it more network bandwidth during its peak hours.
- If a virtual machine is constrained by memory, add memory shares or increase the virtual machine's memory reservation.
- If a virtual machine is constrained by CPU, add CPU shares or increase the virtual machine's CPU reservation.
- If a virtual machine is constrained by CPU or memory, increase the limits or reservations of the resource pool that the virtual machine belongs to.
- If the network load is imbalanced, you can use NIC teaming to balance the network load across multiple physical network adapters.
- If the disk I/O load is imbalanced, you can use storage multipathing to balance the disk I/O load across multiple paths to a datastore.
- If the virtual machine load is imbalanced, you can place hosts in a vSphere DRS cluster to balance the virtual machine load across all hosts in the cluster.

https://t.me/learningnets

10-24 Using vSphere HA with vSphere DRS

vSphere HA might not be able to fail over virtual machines for several reasons:

- vSphere HA admission control is disabled and resources are insufficient in the remaining hosts to power on all of the failed VMs.
- Required VM-Host affinity rules prevent vSphere HA from failing over.
- Sufficient aggregated resources exist, but they are fragmented across hosts. In such cases, vSphere HA uses vSphere DRS to try to adjust the cluster by migrating VMs to defragment the resources.

When vSphere HA performs failover and restarts virtual machines on different hosts, its first priority is the immediate availability of all virtual machines. After the virtual machines are restarted, the hosts in which they were powered on are usually heavily loaded, and other hosts are comparatively lightly loaded.

vSphere HA is closely integrated with vSphere DRS. When a failover occurs, vSphere HA first checks whether resources are available on that host for the failover. If resources are not available, vSphere HA asks vSphere DRS to accommodate for these where possible. Inefficient use of resources might lead to defragmented resources throughout the cluster. vSphere HA requests but cannot guarantee defragmentation of resources to accommodate for this virtual machine's resource requirements.

In addition, vSphere DRS flattens shares and limits on virtual machines before failover. This flattening process ensures that virtual machines get their entitled resources if they would have failed over to the correct resource pool.

https://t.me/learningnets

10-25 Lab 17: Implementing vSphere DRS Clusters

Implement a vSphere DRS cluster

1. Create a Load Imbalance
2. Create a vSphere DRS Cluster
3. Verify Proper vSphere DRS Cluster Functionality
4. Create, Test, and Disable a VM Affinity Rule
5. Create, Test, and Disable an Anti-Affinity Rule
6. Create, Test, and Disable a VM-Host Affinity Rule

https://t.me/learningnets

10-26 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe the functions of a vSphere DRS cluster
- Create a vSphere DRS cluster
- Review information about a vSphere DRS cluster
- Remove a host from a vSphere DRS cluster

https://t.me/learningnets

10-27 Key Points

- vSphere DRS clusters provide automated resource management for multiple ESXi hosts.
- vSphere DRS works best if the virtual machines meet vSphere vMotion migration requirements.

Questions?

https://t.me/learningnets

MODULE 11

vSphere Update Manager



https://t.me/learningnets

11-2 You Are Here

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

https://t.me/learningnets

11-3 Importance

Over time, your vSphere environment might undergo changes in its hardware or software configuration, including software updates and patches.

vSphere Update Manager implements changes to your vSphere environment in an orderly, controlled, and systematic fashion, assisting in the manageability and scalability of your environment.

https://t.me/learningnets

11-4 Learner Objectives

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

- Describe vSphere Update Manager functionality
- Use vSphere Update Manager to create and attach a baseline

https://t.me/learningnets

11-5 About vSphere Update Manager

vSphere Update Manager enables centralized, automated patch and version management for ESXi hosts, virtual machine hardware, VMware Tools, and virtual appliances:

- vSphere Update Manager reduces security risks by:
  - Reducing the number of vulnerabilities
  - Eliminating security breaches that exploit older vulnerabilities
- vSphere Update Manager reduces the diversity of systems in an environment by:
  - Making management easier
  - Reducing security risks
- vSphere Update Manager keeps hosts running more smoothly by:
  - Patching, including bug fixes
  - Making troubleshooting easier

You can use vSphere Update Manager with either vCenter Server on Windows or vCenter Server Appliance.

vSphere Update Manager enables centralized, automated patch and version management for vSphere and supports ESXi hosts, virtual machine hardware, VMware Tools, and virtual appliances. Updates that you specify can be applied to ESXi hosts, virtual machine hardware, and virtual appliances that you scan. With vSphere Update Manager, you can perform the following tasks:

- Scan for compliance and apply updates to virtual machine hardware, appliances, and hosts.
- Directly upgrade hosts, virtual machine hardware, VMware Tools, and virtual appliances.
- Apply third-party software on hosts.

Keeping the patch versions up to date for virtual machine hardware and ESXi hosts helps reduce the number of vulnerabilities in an environment and the range of problems requiring solutions. All systems require ongoing patching and reconfiguration or other solutions. Reducing the diversity of systems in an environment and keeping them in compliance are security best practices. Additionally, because patches include bug fixes, vSphere Update Manager keeps environments operating properly and without service interruption or errors.

vCenter Server Appliance delivers vSphere Update Manager as an optional service. vSphere Update Manager is bundled in vCenter Server Appliance.

VMware maintains its infrastructure through the use of an automated tool called vSphere Update Manager. vSphere Update Manager can be used to apply VMware patches to hosts and guest updates to VMware Tools in an automated manner. To apply patches, vSphere Update Manager uses a user-created baseline to apply to the chosen infrastructure element. These baseline updates can be applied to a single object or many objects simultaneously.

https://t.me/learningnets

11-6 vSphere Update Manager Capabilities

vSphere Update Manager has the following capabilities:

- Automated patch downloading:
  - Automated patch downloading begins with information-only downloading.
  - It is scheduled at regular configurable intervals.
- Baselines and baseline groups creation.
- Ability to attach baselines and baseline groups to inventory objects.
- Scanning: Inventory systems are scanned for baseline compliance.
- Remediation: Inventory systems that are not compliant can be automatically patched.

vSphere Update Manager uses a set of operations to ensure effective patch and upgrade management.

This process begins by downloading information about a set of security patches. One or more of these patches are aggregated to form a baseline. Multiple baselines can be added to a baseline group. You can use baseline groups to combine different types of baselines and scan and remediate an inventory object against all of them as a whole. If a baseline group contains both upgrade and patch baselines, the upgrade runs first.

A collection of virtual appliances and ESXi hosts can be scanned for compliance with a baseline or a baseline group and remediated (updated or upgraded). These processes can be started manually or through scheduled tasks.

https://t.me/learningnets

11-7 vSphere Update Manager Components

vSphere Update Manager requires network connectivity with vCenter Server and includes several components:

- vSphere Update Manager server component can be installed on the same computer as Windows vCenter Server or on a different computer. vSphere Update Manager is integrated with vCenter Server Appliance and is delivered as an optional service.
- Client component: vSphere Update Manager functionality appears on the **Update Manager** tab when an ESXi host is selected in vSphere Web Client.
- Database:
  - The vSphere Update Manager server and Update Manager Download Service that you install on Windows require a database to store and organize server data.
  - Oracle or Microsoft SQL Server database is required (only for vSphere Update Manager on Windows)

https://t.me/learningnets

11-8 Configuring vSphere Update Manager

You can modify the vSphere Update Manager configuration only if you have the correct privileges.

Click the **Update Manager** icon on the Home page of vSphere Web Client to access administrative settings:

- Network connectivity: Modify the IP address or host name for the patch store.
- Download settings: Modify the download location for new patches and virtual appliance upgrades.
- Download schedule: Change the default schedule for checking upgrades, if your environment requires more frequent or less frequent checks.
- Notification check schedule: Configure the time and frequency with which vSphere Update Manager checks for patch recalls or the release of patch fixes.
- Virtual machine settings: Configure whether vSphere Update Manager should take snapshots of virtual machines before applying updates.
- Host and cluster settings: Enable Quick Boot, configure VM power state, and temporarily disable vSphere HA, vSphere Fault Tolerance, and vSphere DPM.

You can modify the vSphere Update Manager settings only if you have the privileges to configure the vSphere Update Manager settings and service. These permissions must be assigned on the vCenter Server system with which vSphere Update Manager is registered.

Connect VMware vSphere® Update Manager Client™ to a vCenter Server system with which vSphere Update Manager is registered. On the Home page of vSphere Web Client, click the **Update Manager** icon to access administrative settings:

- Network Connectivity: The network ports are configured during installation. You can modify the IP address or host name for the patch store in the vSphere Update Manager network connectivity settings.
- Download Source: If your deployment system is connected to the Internet, you can directly download ESXi patches and extensions, as well as virtual appliance upgrades.
- Proxy Settings: You can configure vSphere Update Manager to download updates from the Internet by using a proxy server.
- Checking for Updates (Download Schedule): vSphere Update Manager checks for virtual appliance upgrades, host patches, and extensions at regular intervals. Generally, the default schedule settings are sufficient, but you can change the schedule if your environment requires more or less frequent checks.
- Notification Check Schedule: By default vSphere Update Manager checks for notifications about patch recalls, patch fixes, and alerts at certain time intervals. You can modify this schedule. By default, the task to check for notifications and to send notifications alerts is enabled and is called the VMware vSphere Update Manager Check Notification task. By

https://t.me/learningnets

modifying this task, you can configure the time and frequency with which vSphere Update Manager checks for patch recalls or the release of patch fixes and sends notifications to the email addresses you specify.

- Virtual Machine Settings: By default, vSphere Update Manager is configured to take snapshots of virtual machines before applying updates. If the remediation fails, you can use the snapshot to return the virtual machine to the state before the remediation. vSphere Update Manager does not take snapshots of fault-tolerant virtual machines and virtual machines that are running virtual machine hardware version 3. If you decide to take snapshots of such virtual machines, the remediation might fail. You can choose to keep snapshots indefinitely or for a fixed period of time.
- Host and Cluster Settings: You determine how you want vSphere Update Manager to behave with hosts and clusters.

For more information, see *Installing and Administering VMware vSphere Update Manager* at https://www.vmware.com/support/pubs/vum_pubs.html. See also *vSphere Upgrade Guide* at https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html.

https://t.me/learningnets

11-9 Baselines and Baseline Groups

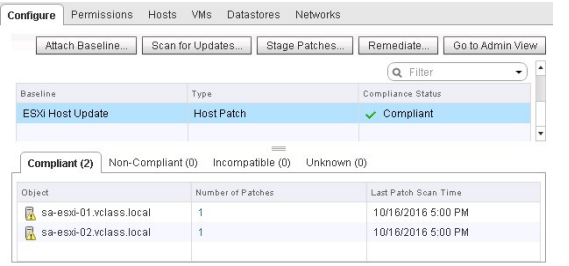
A baseline includes one or more patches, extensions, or upgrades.

vSphere Update Manager includes the following baselines by default:

- Host critical dynamic patch baseline
- Host noncritical dynamic patch baseline
- VMware Tools upgrade baseline
- VMware hardware upgrade baseline
- Virtual appliance upgrade baseline

A baseline group includes multiple baselines.

Baseline groups can contain one upgrade baseline and one or more patch and extension baselines.



Baselines can be upgrade, extension, or patch baselines. Baselines contain a collection of one or more patches, extensions, or upgrades.

Baseline groups are assembled from existing baselines and might contain one upgrade baseline and one or more patch and extension baselines, or they might contain a combination of multiple patch and extension baselines. When you scan hosts, virtual machines, and virtual appliances, you evaluate them against baselines and baseline groups to determine their level of compliance.

To create, edit, or delete baselines and baseline groups, you must have the Manage Baseline privilege. To attach baselines and baseline groups, you must have the Attach Baseline privilege. Privileges must be assigned on the vCenter Server system with which vSphere Update Manager is registered.

vSphere Update Manager includes the following default dynamic patch and upgrade baselines:

- Critical Host Patches (Predefined): Check ESXi hosts for compliance with all critical patches.
- Non-Critical Host Patches (Predefined): Check ESXi hosts for compliance with all optional patches.
- VMware Tools Upgrade to Match Host (Predefined): Checks virtual machines for compliance with the latest VMware Tools version on the host. vSphere Update Manager supports upgrading of VMware Tools for virtual machines on hosts that are running ESXi 5.0 and later.
- VM Hardware Upgrade to Match Host (Predefined): Checks the virtual hardware of a virtual machine for compliance with the latest version supported by the host. vSphere Update Manager supports upgrading to virtual hardware version vmx-11 on hosts that are running ESXi 6.x.
- VA Upgrade to Latest (Predefined): Checks virtual appliance compliance with the latest released virtual appliance version.

https://t.me/learningnets

11-10 Creating and Editing Patch or Extension Baselines

You can create custom patch, extension, and upgrade baselines to meet the needs of your specific deployment by using the New Baseline wizard:

- Create a fixed patch baseline:
  - Fixed baselines consist of a set of patches that do not change as patch availability changes.
- Create a dynamic patch baseline:
  - Dynamic baselines consist of a set of patches that meet certain criteria.
- Create a host extension baseline:
  - Extension baselines contain additional software for ESXi hosts. This additional software might be VMware or third-party software.
- Filter patches or extensions in the New Baseline wizard:
  - When you create a patch or extension baseline, you can filter the patches and extensions available in the vSphere Update Manager repository to find specific patches and extensions to exclude or include in the baseline.

You create and manage baselines in the **Update Manager** tab. You can remediate hosts against baselines that contain patches or extensions. Depending on the patch criteria you select, patch baselines can be either dynamic or fixed.

Dynamic patch baselines contain a set of patches, which updates automatically according to patch availability and the criteria that you specify. Fixed baselines contain only patches that you select, regardless of new patch downloads.

Extension baselines contain additional software modules for ESXi hosts. This additional software might be VMware software or third-party software. You can install additional modules by using extension baselines, and update the installed modules by using patch baselines.

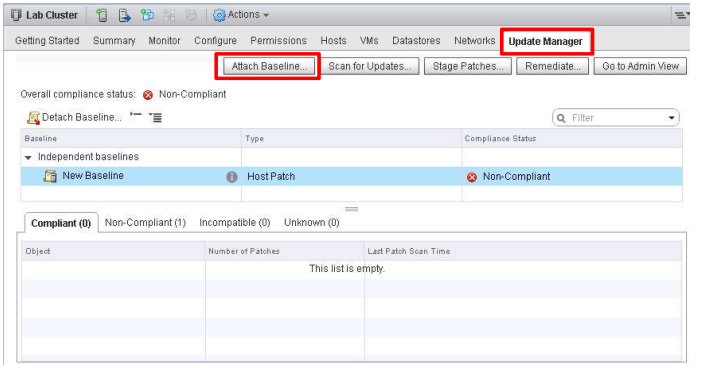
If your vCenter Server system is part of a connected group in vCenter Linked Mode, and you have more than one vSphere Update Manager instance, patch and extension baselines that you create are not applicable to all inventory objects managed by other vCenter Server systems in the group. Baselines are specific for the vSphere Update Manager instance you select.

https://t.me/learningnets

11-11 Attaching Baselines

To view compliance information and scan objects in the inventory against baselines or baseline groups, you must first attach baselines or baseline groups to these objects.

You can attach baselines or baseline groups to objects from the **Update Manager** tab in vSphere Web Client.



To view compliance information and remediate objects in the inventory against specific baselines and baseline groups, you must first attach baselines and baseline groups to these objects.

To use baselines and baseline groups, you must attach baselines and baseline groups to selected inventory objects, such as virtual machines, virtual appliances, hosts, or container objects.

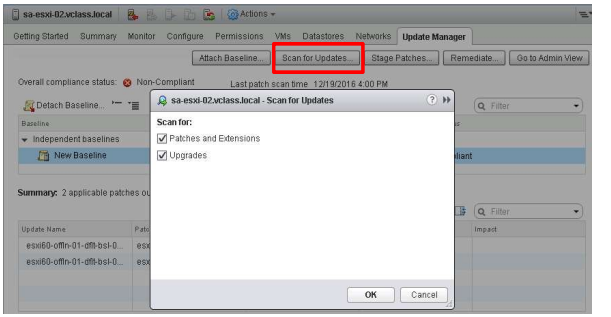
Although you can attach baselines and baseline groups to individual objects, a more efficient method is to attach them to container objects, such as folders, vApps, clusters, and data centers. Individual vSphere objects inherit baselines attached to the parent container object. Removing an object from a container removes the inherited baselines from the object.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, you can attach baselines and baseline groups to objects managed by the vCenter Server system with which vSphere Update Manager is registered. Baselines and baseline groups that you attach are specific to the vSphere Update Manager instance that is registered with the vCenter Server system.

https://t.me/learningnets

11-12 Scanning for Updates

Scanning evaluates the inventory object against the baseline or baseline group.



Scanning is the process in which attributes of a set of hosts, virtual machines, or virtual appliances are evaluated against patches, extensions, and upgrades in the attached baselines and baseline groups. You can configure vSphere Update Manager to scan virtual machines, virtual appliances, and ESXi hosts against baselines and baseline groups by either scheduling or manually initiating scans to generate compliance information.

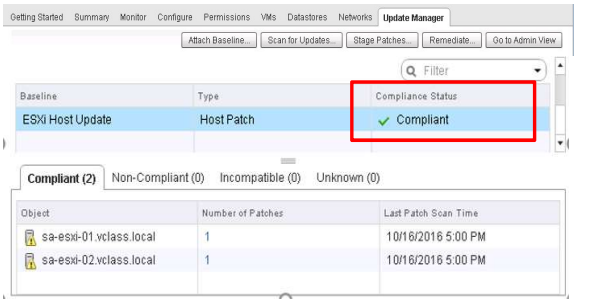
If the object that you select is a container object, all child objects are also scanned. The larger the virtual infrastructure and the higher up in the object hierarchy that you begin the scan, the longer the scan takes.

After you have an inventory object attached to a baseline, perform a scan by selecting the object and selecting **Scan for updates**.

https://t.me/learningnets

11-13 Reviewing vSphere Object Compliance

You can review compliance information for the virtual machines, virtual appliances, and hosts against baselines and baseline groups that you attach.



When you select a container object, you view the overall compliance status of the attached baselines, as well as all the individual compliance statuses. If you select an individual baseline attached to the container object, you see the compliance status of the baseline.

If you select an individual virtual machine, appliance, or host, you see the overall compliance status of the selected object against all attached baselines and the number of updates. If you further select an individual baseline attached to this object, you see the number of updates grouped by the compliance status for that baseline.

The Compliant status indicates that a vSphere object is compliant with all baselines in an attached baseline group or with all patches, extensions, and upgrades in an attached baseline. A compliant state requires no further action. If a baseline contains patches or upgrades that are not relevant to the target object, both the individual updates and baselines or baseline groups that contain them are treated as not applicable, and represented as compliant.

The Non-Compliant status indicates that one or more baselines in a baseline group, or one or more patches, extensions, or upgrades in a baseline are applicable to the target object but are not installed (missing) on the target. You must remediate the target object to make it compliant.

https://t.me/learningnets

11-14 Remediating Objects

You can remediate virtual machines, templates, virtual appliances, and hosts:

- You can perform the remediation immediately or schedule it for a later date.
- Host remediation runs in different ways, depending on the types of baselines that you attach and whether or not the host is in a cluster.
- For ESXi hosts in a cluster, the remediation process is sequential by default.
- The remediation of hosts in a cluster requires you to temporarily disable cluster features such as vSphere DPM and vSphere HA admission control.

You can remediate virtual machines, virtual appliances, and hosts using either user-initiated remediation or scheduled remediation at a time that is convenient for you.

You can remediate virtual machines and appliances together.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, you can remediate only the inventory objects managed by the vCenter Server system with which vSphere Update Manager is registered.

Host remediation runs in different ways depending on the types of baselines you attach and whether the host is in a cluster or not.

For ESXi hosts in a cluster, the remediation process is sequential by default. With vSphere Update Manager 6 and later, you can select to run host remediation in parallel.

When you remediate a cluster of hosts sequentially and one of the hosts fails to enter maintenance mode, vSphere Update Manager reports an error, and the process stops and fails. The hosts in the cluster that are remediated stay at the updated level. The ones that are not remediated after the failed host remediation are not updated. If a host in a cluster enabled for vSphere DRS runs a virtual machine on which vSphere Update Manager or vCenter Server is installed, vSphere DRS first attempts to migrate the virtual machine running vCenter Server or vSphere Update Manager to another host, so that the remediation succeeds. In case the virtual machine cannot be migrated to another host, the remediation fails for the host but the process does not stop. vSphere Update Manager proceeds to remediate the next host in the cluster.

The host upgrade remediation of ESXi hosts in a cluster proceeds only if all hosts in the cluster can be upgraded.

https://t.me/learningnets

Remediation of hosts in a cluster requires that you temporarily disable cluster features, such as vSphere DPM and vSphere HA admission control. You should also turn off vSphere Fault Tolerance if it is enabled on any of the virtual machines on a host, and disconnect the removable devices connected to the virtual machines on a host, so that they can be migrated with vSphere vMotion. Before you start a remediation process, you can generate a report that shows which cluster, host, or virtual machine has the cluster features enabled.

When you remediate a cluster of hosts in parallel, vSphere Update Manager remediates multiple hosts concurrently.

During parallel remediation, if vSphere Update Manager encounters an error when remediating a host, it ignores the host and the remediation process continues for the other hosts in the cluster. vSphere Update Manager continuously evaluates the maximum number of hosts it can remediate concurrently without disrupting vSphere DRS settings.

You can limit the number of concurrently remediated hosts to a specific number.

vSphere Update Manager remediates hosts that are part of a vSAN cluster sequentially, even if you select the option to remediate them in parallel. By design, only one host from a vSAN cluster can be in a maintenance mode at any time.

https://t.me/learningnets

11-15 Patch Recall Notification

At regular intervals, vSphere Update Manager contacts VMware to download notifications about patch recalls, new fixes, and alerts:

- **Notification Check Schedule** is selected by default.

On receiving patch recall notifications, vSphere Update Manager takes the following actions:

- It generates a notification in the notification tab.
- It no longer applies the recalled patch to any host. The patch is flagged as recalled in the database.
- It deletes the patch binaries from its patch repository.

vSphere Update Manager does not uninstall recalled patches from ESXi hosts. It waits for a newer patch and applies that patch to make a host compliant.

At regular intervals, vSphere Update Manager contacts VMware to download information (notifications) about patch recalls, new fixes, and alerts. You can change the schedule by modifying the **Notification Check Schedule** setting in the vSphere Update Manager **Configuration** tab.

When patches with problems or potential problems are released, these patches are recalled in the metadata, and vSphere Update Manager marks them as recalled. If you try to install a recalled patch, vSphere Update Manager notifies you that the patch is recalled and does not install it on the host. If you have already installed such a patch, vSphere Update Manager notifies you that the recalled patch is installed on certain hosts. vSphere Update Manager also deletes all the recalled patches from the vSphere Update Manager patch repository.

When a new patch is released, vSphere Update Manager downloads it and prompts you to install it to fix the problems that the recalled patch might cause. If you try to install the recalled patch, vSphere Update Manager alerts you that the patch is recalled and that you must install a fix.

https://t.me/learningnets

11-16 EAM Integration with vSphere Update Manager

ESX Agent Manager (EAM) is used by VMware NSX® to prepare ESXi hosts with the required NSX data path and host tools extension:

- vSphere Update Manager automatically creates and attaches a system-managed EAM baseline for each ESXi cluster that was prepared by VMware NSX® Manager™.
- When a new version of the NSX data path and host tools extension is released, the EAM baselines automatically update and report that the ESXi cluster is no longer compliant. They also recommend remediation to update the extensions installed on the ESXi hosts.

ESX Agent Manager (EAM) automates the process of deploying and managing vSphere ESX agents, while extending the function of an ESXi host to provide additional services that a vSphere solution requires.

ESX agent virtual machines are similar to services in Windows or Linux. They start when the operating system starts, and they stop when it shuts down. The behavior of ESX agent virtual machines is transparent to the user. A vSphere host reaches the ready state when the ESXi operating system has started and all ESX agent virtual machines have been provisioned and powered on.

EAM enables you to monitor the health of ESX agents and blocks users from performing certain operations on ESX agents that might affect the virtual machines that use them. For example, EAM can prevent an ESX agent virtual machine from being powered off or moved from an ESXi host that contains other virtual machines that use that agent.

https://t.me/learningnets

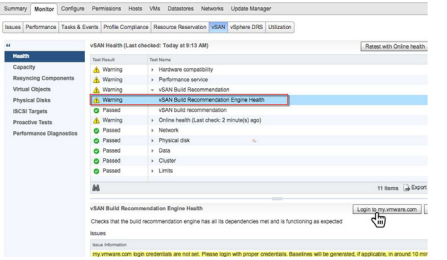
11-17 vSAN Integration with vSphere Update Manager

To ensure that the vSAN cluster is up to date, vSAN 6.6.1 (and later) is integrated with vSphere Update Manager.

vSAN generates system baselines and baseline groups for use with vSphere Update Manager.

The vSAN release catalog maintains information about available releases, preference order for releases, and critical patches needed for each release. The vSAN release catalog is hosted on VMware Cloud™.

vSphere Update Manager determines the correct upgrade path for each vSAN cluster, based on the current vSAN release catalog. vSAN also includes the necessary drivers and patch updates for the recommended release in its system baseline.



vSAN patch and version management is fully integrated in vSAN. vSAN version recommendations are automatically generated using information from the VMware Compatibility Guide, the vSAN release catalog, and awareness of the underlying hardware configuration.

To update VSAN:

1. Add your my.vmware.com credentials to the vSAN Build Recommendation Engine.

   After you log in, vSAN generates a baseline group of recommended updates for each vSAN cluster. vSAN system baselines are listed in the Baselines pane of the **Baselines and Groups** tab.

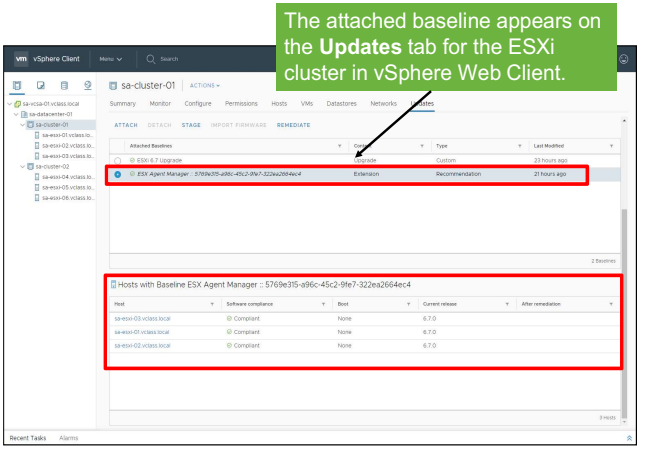
2. To update the vSAN cluster, click the **Update Manager** tab and click **Remediate**.

https://t.me/learningnets

11-18 About System Baselines

System baselines ensure an ESXi cluster's compliance and compatibility with the recommended software, updates, or extensions defined by the baseline:

- System baselines are read-only and are automatically created and maintained by vSphere Update Manager.
- When an ESXi cluster has been prepared by NSX Manager, the EAM baseline automatically attaches to the ESXi cluster and is scanned for compliance.



https://t.me/learningnets

11-19 vSphere Update Manager Integration with vSphere DRS

When performing remediation operations on an ESXi cluster, vSphere Update Manager automatically integrates with vSphere DRS:

- In a vSphere DRS cluster, vSphere Update Manager puts hosts into maintenance mode before vSphere DRS runs. Putting hosts into maintenance mode results in the vSphere DRS cluster hosts being evacuated before being patched.
- When an enter maintenance mode operation is called, certain prechecks are performed to ensure that the ESXi host can enter maintenance mode.
- vSphere reports any configuration issues that might prevent an ESXi host from entering maintenance mode.

If the following checks do not pass, the maintenance mode operation fails and reports why it failed:

- ESXi host level checks:
  - ESXi host cluster participation
  - Powered-on virtual machine capacity
  - vSphere DRS enabled status
  - vSphere DRS fully automated status
- Virtual machine level checks:
  - Virtual machine vSphere DRS manual override
  - Virtual machine vSphere DRS rule conflict
  - vSphere HA admission control violation

https://t.me/learningnets

11-20 Lab 18: Using vSphere Update Manager

Install, configure, and use vSphere Update Manager

1. Modify the Cluster Settings
2. Configure vSphere Update Manager
3. Create a Patch Baseline
4. Attach a Baseline and Scan for Updates
5. Stage the Patches onto the ESXi Host
6. Remediate the ESXi Host

https://t.me/learningnets

11-21 Review of Learner Objectives

You should be able to meet the following objectives:

- Describe vSphere Update Manager functionality
- Use vSphere Update Manager to create and attach a baseline

https://t.me/learningnets

11-22 Key Points

- vSphere Update Manager reduces security vulnerabilities by keeping systems up to date and by reducing the diversity of systems in an environment.

Questions?

https://t.me/learningnets

MODULE 12

vSphere Troubleshooting



https://t.me/learningnets

12-2 You Are Here

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

https://t.me/learningnets

12-3 Importance

You can quickly identify, diagnose, and solve a problem if you use an efficient troubleshooting methodology in a consistent and repeatable manner.

Knowing how to use the right tools to solve various types of problems can save time and maximize your troubleshooting result.

https://t.me/learningnets

12-4 Module Lessons

Lesson 1:	Tools for Troubleshooting
Lesson 2:	Introduction to Troubleshooting

https://t.me/learningnets



https://t.me/learningnets

12-6 Learner Objectives

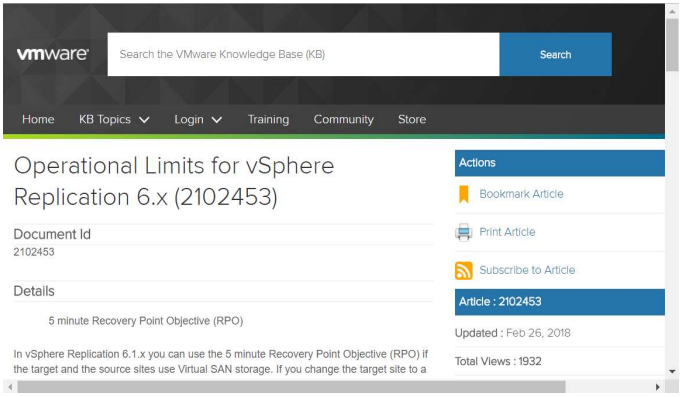
By the end of this lesson, you should be able to meet the following objectives:

- Review troubleshooting tools
- Find important log files
- Use vSphere Syslog Collector

https://t.me/learningnets

12-7 About the VMware Knowledge Base

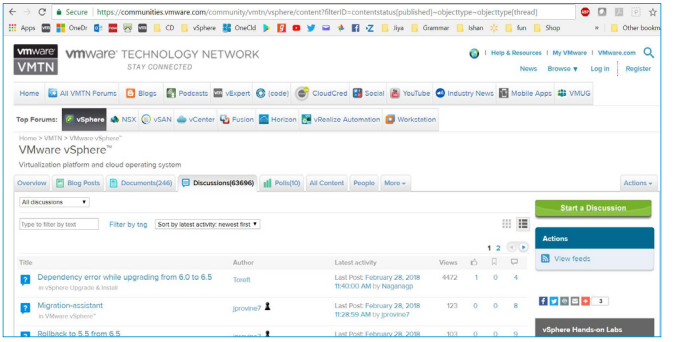
The VMware knowledge base provides support solutions, error messages, and troubleshooting guides.



https://t.me/learningnets

12-8 About the VMware Community

The VMware Community is a family of websites and online communities that includes blogs, discussion forums, documents, and information sharing about products and services, virtualization software and technology, and services provided by VMware and its subsidiaries.



The VMware Community provides a forum for interested users of virtualization and related products and services and for participants in VMware programs such as beta programs, user groups, and customer advisory councils to exchange information.

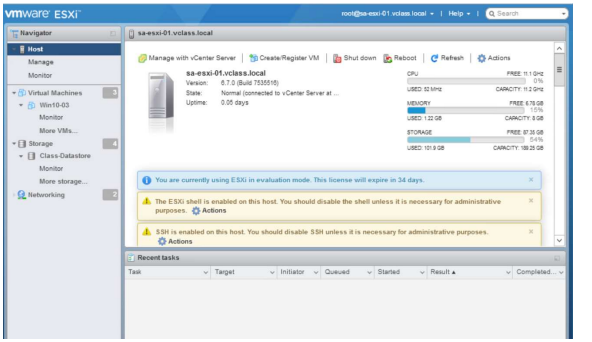
https://t.me/learningnets

12-9 About VMware Host Client

You use VMware Host Client to conduct emergency management when vCenter Server is temporarily unavailable.

VMware Host Client functions include the following operations:

- Basic virtualization operations, such as deploying and configuring virtual machines of various complexity
- Creating and managing networking and datastores
- Advanced tuning of host-level options to improve performance



VMware Host Client is an HTML5-based client that is used to connect to and manage single ESXi hosts.

You can use VMware Host Client to perform administrative and basic troubleshooting tasks, as well as advanced administrative tasks on your target ESXi host. You can also use VMware Host Client to conduct emergency management when vCenter Server is not available.

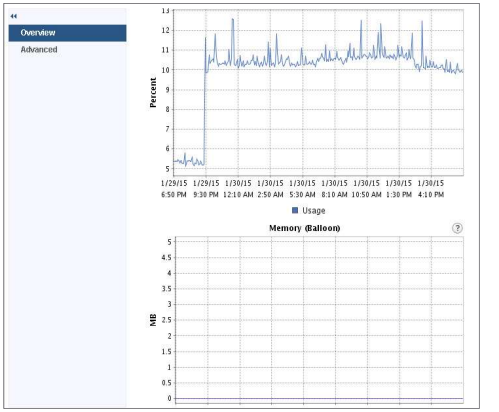
Though they have similar user interfaces, VMware Host Client has different functions from vSphere Web Client. You use vSphere Web Client to connect to vCenter Server and manage multiple ESXi hosts. You use VMware Host Client to manage a single ESXi host.

https://t.me/learningnets

12-10 Monitoring the Virtual Environment with Performance Charts

VMware provides tools to help you monitor your virtual environment and locate the source of potential and current problems:

- Overview charts display multiple data sets in one panel to evaluate resource statistics and display thumbnail charts for child objects. Charts for a parent object and a child object are also available.
- Advanced charts display more information than overview charts. They are configurable, and they can be printed or exported.



Host's Performance Charts
Partial Overview Panel

The vSphere statistics subsystem collects data on the resource usage of inventory objects. Data on a wide range of metrics is collected at frequent intervals and then processed and archived in the vCenter Server database. You can access statistical information through command-line monitoring utilities or by viewing performance charts in vSphere Web Client.

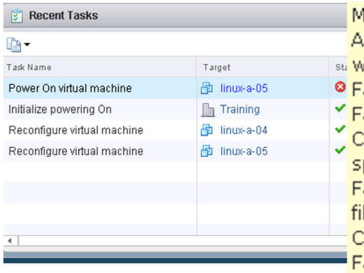
The performance chart options that you can access from the **View** menu depend on the type of inventory object that you select.

https://t.me/learningnets

12-11 Viewing and Interpreting Diagnostic Information

You can view diagnostic messages in the GUI or in log files.

You interpret the diagnostic messages to find the root cause.



Module 'MonitorLoop' power on failed.
An error was received from the ESX host while powering on VM linux-a-05.
Failed to start the virtual machine.
Failed to power on VM.
Could not power on virtual machine: No space left on device.
Failed to extend the virtual machine swap file
Current swap file size is 0 KB.
Failed to extend swap file from 0 KB to 524288 KB.
File system specific implementation of LookupAndOpen[file] failed
File system specific implementation of Lookup[file] failed

https://t.me/learningnets

12-12 Location of vCenter Server Logs

Most of the vCenter Server log files are on vCenter Server Appliance in the `/var/log/vmware/` directory.

Subdirectories exist for vCenter Server components. For example, a `vpxd-svcs` subdirectory exists for vCenter Server services logs.

Some logs, such as `FirstBoot.log`, are in `/var/log`.

```
root@sa-vcsa-01 [ /var/log/vmware ]# ls
applmgmt          rsyslogd        vmdird
applmgmt-audit    rsyslogd-2068   vmdnsd
cis-license       rsyslogd-2078   vmon
cloudvm           sca             vmware-imagebuilder
cm                sso             vmware-sps
content-library   syslog          vmware-updatemgr
eam               vapi            vpostgres
journal           vcha            vpxd
mbcs              vctop           vpxd-svcs
netdumper         vmafd           vsan-health
perfcharts        vmafdd          vsm
psc-client        vmcad           vsphere-client
rbd               vmcam           vsphere-ui
rhttpproxy        vmdir
```

By default, vCenter Server places its log files in a default location. The location is determined based on where the software is installed.

For more information about the location of vCenter Server log files, see VMware knowledge base article 2110014 at http://kb.vmware.com/kb/2110014.

https://t.me/learningnets

12-13 Common Logs

Common logs are common to all vCenter Server systems.

Log Directory	Description
applmgmt	VMware Appliance Management Service
CloudVM	Logs for allotment and distribution of resources between services
CM	VMware Component Manager
rhttpproxy	Reverse Web Proxy
SCA	VMware Service Control Agent
Vapi	VMware vAPI Endpoint
vmaffd	VMware Authentication Framework daemon
vmdird	VMware Directory Service daemon
Vmon	VMware Service Lifecycle Manager
vSphere-Client	VMware vSphere Web Client

https://t.me/learningnets

12-14 Management Node Logs

Management node logs appear on vCenter Server Appliance systems that are combined nodes (single-node deployment) or are third-party nodes.

The table lists some of the more important management node logs.

Log	Description
`AutoDeploy`	VMware vSphere Auto Deploy Waiter
`content-library`	VMware Content Library Service
`EAM`	VMware ESX Agent Manager
`InvSvc`	VMware Inventory Service
`vmcam`	VMware vSphere Authentication Proxy
`vpxd`	VMware VirtualCenter Server
`vPostgres`	vFabric Postgres database service

https://t.me/learningnets

12-15 Platform Services Controller Logs

Platform Services Controller logs appear on Platform Services Controller nodes.

Log	Description
`cis-license`	VMware Licensing Service
`SSO`	VMware Secure Token Service
`Vmcad`	VMware Certificate Authority daemon

https://t.me/learningnets

12-16 Important vCenter Server Log for Troubleshooting

The `vpxd.log` file is the main log file for vCenter Server. Most vCenter Server actions are captured in `vpxd.log`, in `/var/log/vmware/vpxd/`.

```
2017-03-03T17:37:30.809Z info vpxd[7FD7BF66C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-3d] [VpxLRO] -- BEGIN lro-24393 -- ResourceModel -- cis.data.provider.ResourceModel.query -- 5266df8c-2f36-e3c4-1d59-c474
66b43db1(52723af4-e7c1-92c5-02c8-4108140b6ba2)
2017-03-03T17:37:30.810Z info vpxd[7FD7BF66C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-3d] [VpxLRO] -- FINISH lro-24393
2017-03-03T17:37:30.813Z info vpxd[7FD7BF66C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-ef] [VpxLRO] -- BEGIN lro-24394 -- ResourceModel -- cis.data.provider.ResourceModel.query -- 5266df8c-2f36-e3c4-1d59-c474
66b43db1(52723af4-e7c1-92c5-02c8-4108140b6ba2)
2017-03-03T17:37:30.813Z info vpxd[7FD7BF66C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-ef] [VpxLRO] -- FINISH lro-24394
2017-03-03T17:37:30.815Z info vpxd[7FD7BFE7C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-df] [VpxLRO] -- BEGIN lro-24395 -- ResourceModel -- cis.data.provider.ResourceModel.query -- 5266df8c-2f36-e3c4-1d59-c474
66b43db1(52723af4-e7c1-92c5-02c8-4108140b6ba2)
2017-03-03T17:37:30.815Z info vpxd[7FD7BFE7C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-df] [VpxLRO] -- FINISH lro-24395
2017-03-03T17:37:30.816Z info vpxd[7FD7BFE7C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-67] [VpxLRO] -- BEGIN lro-24396 -- ResourceModel -- cis.data.provider.ResourceModel.query -- 5266df8c-2f36-e3c4-1d59-c474
66b43db1(52723af4-e7c1-92c5-02c8-4108140b6ba2)
2017-03-03T17:37:30.816Z info vpxd[7FD7BFE7C700] [Originator@6876 sub=vpxLro opID=urn:vmomi:VirtualMachine:vm-128:596e26b2-df48-4c88-b39c-eb702
ba52e66.properties:01-67] [VpxLRO] -- FINISH lro-24396
```


https://t.me/learningnets

12-17 Location of ESXi Host Logs

Most ESXi host log files are in the `/var/log` directory.

```
[root@sa-esxi-01:/var/log] ls
Xorg.log                 ipmi                    vitd.log
auth.log                 jumpstart-stdout.log    vmauthd.log
boot.gz                  kickstart.log           vmkdevmgr.log
clomd.log                lacp.log                vmkernel.log
configRP.log             nfcd.log                vmkeventd.log
ddecomd.log              osfsd.log               vmksummary.log
dhclient.log             rabbitmqproxy.log       vmkwarning.log
epd.log                  rhttpproxy.log          vmware
esxcli-software.log      sdrsinjector.log        vmware-vmsvc.log
esxupdate.log            shell.log               vobd.log
fdm.log                  smbios.bin              vprobe.log
hbrca.log                storagerm.log           vpxa.log
hostd-probe.log          swapobjd.log            vsanmgmt.log
hostd.log                sysboot.log             vsansystem.log
hostdCgiServer.log       syslog.log              vsanvpd.log
hostprofiletrace.log     tallylog                vvold.log
iofilter-init.log        upitd.log
iofiltervpd.log          usb.log
```

If persistent scratch space is configured, many of these logs are on the scratch volume, `/scratch/log`. The `/var/log` directory contains symbolic links (which are identified as light blue names in the output) to log files in `/scratch/log`. Run the `ls -l` command to list the log files that these symbolic links point to.

https://t.me/learningnets

12-18 Useful ESXi Host Logs for Troubleshooting

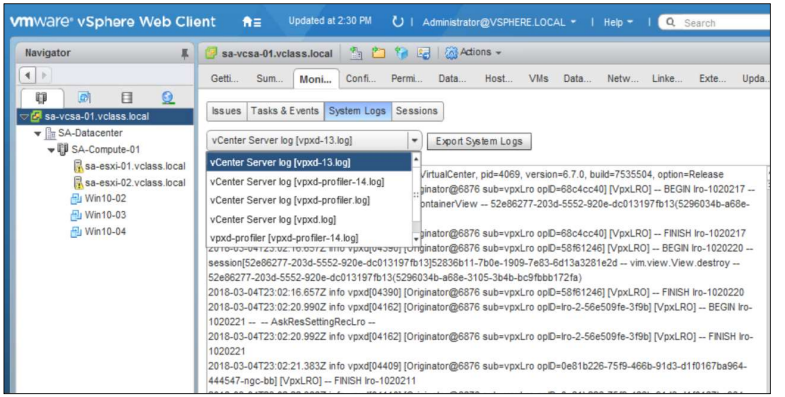
ESXi hosts write to multiple log files, depending on which action is being performed.

Log	Description
`hostd.log`	Host management service logs
`syslog.log`	Management service initialization, watchdogs, scheduled tasks, and DCUI use
`vmkernel.log`	Core VMkernel logs, including device discovery, storage and networking device and driver events, and virtual machine startups
`vmkwarning.log`	A summary of warning and alert log messages excerpted from the VMkernel logs
`vmksummary.log`	A summary of ESXi host startup and shutdown, and an hourly heartbeat with uptime, number of virtual machines running, and service resource consumption

https://t.me/learningnets

12-19 Viewing vCenter Server Log Files in vSphere Web Client

You can use vSphere Web Client to view log files on vCenter Server systems. Use the **Monitor** tab for vCenter Server systems.

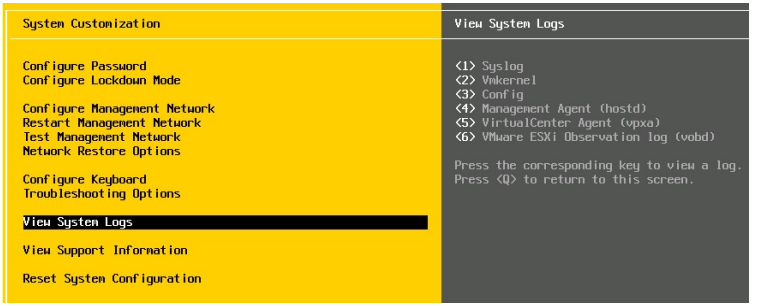


https://t.me/learningnets

12-20 Viewing Log Files in the DCUI

You can use the DCUI to view log files if vCenter Server is not available.

Only the log files for a single ESXi host can be viewed in the DCUI.



In the ESXi host console, press F2 and log in to the DCUI, using the root user name and password. Press the appropriate key to search the log file (/), request help (H), or quit viewing the log file (Q).

https://t.me/learningnets

12-21 About vSphere Syslog Collector

vSphere Syslog Collector provides a single location for all ESXi hosts to write log files.

vSphere Syslog Collector enables logs from multiple hosts to be combined and provides a structure for network logging.

vSphere Syslog Collector is preinstalled on vCenter Server Appliance.

By default, the vSphere Syslog Collector server uses port 514 for TCP and UDP, and port 1514 for SSL.

vSphere Syslog Collector provides a unified architecture for system logging. With vSphere Syslog Collector, you can direct logs from ESXi hosts to a server on the network, rather than save the logs to a local disk.

https://t.me/learningnets

12-22 Review of Learner Objectives

You should be able to meet the following objectives:

- Review troubleshooting tools
- Find important log files
- Use vSphere Syslog Collector

https://t.me/learningnets



https://t.me/learningnets

12-24 Learner Objectives

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

- Define the scope of troubleshooting
- Use a structured approach to solve configuration and operational problems
- Apply a troubleshooting methodology to logically diagnose faults and improve troubleshooting efficiency

https://t.me/learningnets

12-25 Troubleshooting Process

Troubleshooting is a systematic approach to identifying the root cause of a problem and defining steps to resolve the problem.

A typical troubleshooting process involves the following steps:

1. Define the problem.
2. Identify the cause of the problem.
3. Resolve the problem.

The troubleshooting process begins when a user reports a problem. In this context, the user is anyone using the system, from an end user to an administrator. The problem reported by the user might not be the real problem. A user might be reporting symptoms of the problem.

An observed problem might be directly causing the symptoms, but typically the problem has a more fundamental cause.

https://t.me/learningnets

12-26 Definition of a System Problem

A system problem is a fault in a system, or one of its components, that negatively affects the services needed for normal production.

System problems arise from various sources:

- Configuration problems
- Resource contention
- Network attacks
- Software:
  - Bugs
  - Interoperability
- Hardware:
  - Failures
  - Hardware compatibility

Do not assume that you understand the problem after you have identified one symptom. Do a thorough analysis. Verify that nothing else is broken.

A system consists of several components, both software and hardware. For example, an ESXi host consists of components such as CPU, memory, storage, networking, and hypervisor software. A virtual machine consists of various components, such as one or more applications, a guest operating system, and virtual hardware. A problem that occurs in a system can disrupt and negatively affect production services that were functioning normally.

https://t.me/learningnets

12-27 Collecting Symptoms of a Problem

Collecting symptoms is the first step in troubleshooting a problem.

A single root cause often presents itself as several symptoms that users report.

Differentiating between symptoms and the root cause of a problem is imperative.

Symptom	Possible Cause	Root Cause
One or more LUNs on a storage array are not visible to a specific ESXi host.	The LUNs that are not visible are not presented correctly to the ESXi host.	
	Pathing failure.	The network has failed between the ESXi host and the storage array. No redundant path is available.
You cannot connect to vCenter Server with vSphere Web Client.	The VirtualCenter Server service failed to start.	vCenter Server Appliance has a corrupt database.
	Network path between you and vCenter Server Appliance is down.	

Problems can arise in any computing environment. Complex application behaviors, changing demands, and shared infrastructure can lead to problems arising in previously stable environments.

Troubleshooting problems requires an understanding of the interactions between the software and hardware components of a computing environment. Moving to a virtualized computing environment adds new software layers and new types of interactions that must be considered when troubleshooting problems.

https://t.me/learningnets

12-28 Gathering Supplemental Information

Ask questions to gather additional information to define the problem:

- Can the problem be reproduced?
  - Provide a repeatable means to recreate the problem and a way to validate that the problem was resolved.
- What is the scope?
  - Does the problem affect only one object or multiple objects?
- Was the system working before?
  - If so, what changed in the environment or configuration?
- Is the problem a known problem?
  - Consult references, such as release notes, to determine whether the problem is documented.

Proper troubleshooting requires starting with a broad view of the computing environment and systematically narrowing the scope of the investigation as possible sources of problems are eliminated. Troubleshooting efforts that start with a narrowly conceived idea of the source of a problem often get stuck in detailed analysis of one component, when the real source of the problem is elsewhere in the infrastructure.

To quickly isolate the source of a problem, you must adhere to a logical troubleshooting methodology that avoids preconceptions about the source of the problem.

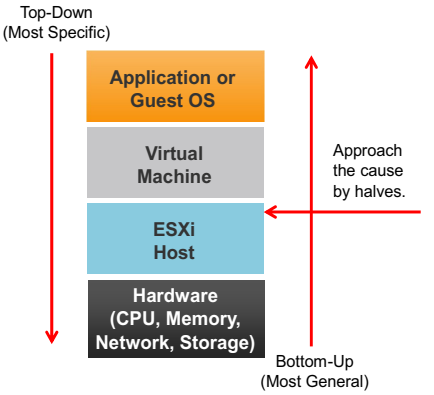
https://t.me/learningnets

12-29 Identifying Possible Causes and Taking the Appropriate Approach

A structured approach to troubleshooting enables you to determine the root cause quickly and effectively.

Based on the problem's characteristics, take one of the following troubleshooting approaches:

- Investigate the cause top-down.
- Investigate the cause bottom-up.
- Approach the cause by halves.



In a VMware virtual environment, the root cause of a problem can occur in any one of the virtual components. Knowing where to start looking for the root cause is often not obvious. Thus, gathering as much information as you can about the problem can help determine which virtual component to check first.

You might take one of the following troubleshooting approaches:

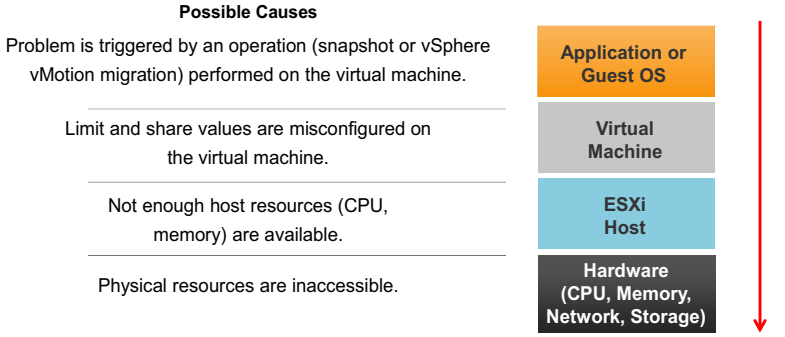
- Top-down: Start troubleshooting in the guest operating system first, then work your way down the stack, then to the virtual machine, then to the ESXi host, and finally to the hardware.
- Bottom-up: Start troubleshooting at the hardware level first, then work your way up the stack to the ESXi host, then to the virtual machine, and finally to the guest operating system.
- Approach cause by halves: Start troubleshooting at the middle of the stack. For example, start with the virtual machine and test possible causes. The test results determine whether you should continue troubleshooting up the stack or down the stack.

https://t.me/learningnets

12-30 Determining the Root Cause

To determine the root cause, test your environment and eliminate possible causes.

Example: The virtual machine stopped responding.



General virtual infrastructure knowledge and knowledge of your specific system configuration are very helpful in identifying possible causes. Prioritize the list of possible causes, ordering them from most probable to least probable. Then test each possible cause to determine the most likely cause of the problem, called the root cause.

In the example, the problem is that a virtual machine has stopped responding. In a nonresponsive system, the operating system seems to be paralyzed and no error messages appear. However, the operating system is still running. Such problems might require guidance from documents, such as VMware knowledge base articles. For example, to troubleshoot a virtual machine that has stopped responding, see VMware knowledge base article 1007819 at http://kb.vmware.com/kb/1007819.

For this problem, you might take a top-down approach. Start with the operations performed on the virtual machine, check the virtual machine configuration, and check for sufficient resources on the host where the virtual machine is located.

https://t.me/learningnets

12-31 Resolving the Problem

After identifying the root cause, assess the impact of the problem on operations:

- High impact: Resolve immediately.
- Medium impact: Resolve when possible.
- Low impact: Resolve during next maintenance window.

Identify possible solutions and their impact on the vSphere environment:

- Short-term solution: Workaround.
- Long-term solution: Reconfiguration.
- Impact analysis: Assess the impact of the solution on operations.

Resolve the problem by implementing the most effective solution.

After identifying the root cause, resolve the problem. To resolve the problem, you identify possible solutions to the problem, then implement a solution.

When determining the best solution, assess the impact that the problem has on normal operations. For example, if the problem causes business-critical applications to be inaccessible, then the impact of the problem is high, and immediate resolution is necessary.

You might decide to first implement a short-term fix so that systems can be brought back online quickly. Before implementing the short-term solution, document all changes that you have made to the system from the time the problem occurred. Also, back up your log files from the time that the problem occurred. Some short-term solutions can be destructive and truncate important log information necessary for additional assistance.

Eventually, you want to implement a more permanent, long-term solution to prevent the problem from recurring.

https://t.me/learningnets

12-32 Example Scenario: Defining the Problem

Problem: vSphere vMotion migration fails with an error.

Scenario: You try to migrate the virtual machine named VM01 from the host named ESXi01 to the host named ESXi02. After a couple of minutes, the vSphere vMotion migration fails with an error.

Is this failure a vSphere vMotion problem or a symptom of an underlying problem?

The error message provides additional information.

In the example, you use the troubleshooting methodology to diagnose a vSphere vMotion migration problem. You use vSphere Web Client to perform a vSphere vMotion migration, but the migration fails with an error.

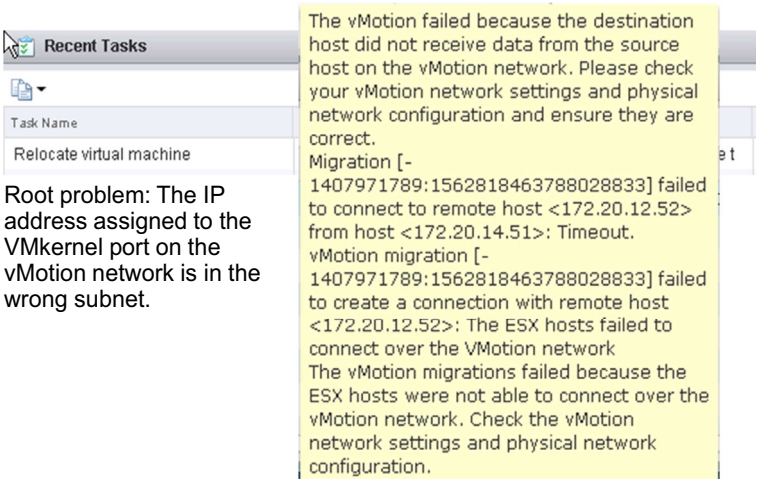
At this point, you cannot tell whether the problem is specific to vSphere vMotion or whether the problem is in the underlying infrastructure, such as storage or networking.

To pinpoint the problem area, gather information about the problem, starting with any diagnostic messages displayed in vSphere Web Client.

https://t.me/learningnets

12-33 Example: Gathering Information

Error messages can help determine the problem.



Root problem: The IP address assigned to the VMkernel port on the vMotion network is in the wrong subnet.

vSphere Web Client displays error messages for the failed vSphere vMotion migration task, with the following information:

- A general system error occurred: The vSphere vMotion migrations failed because the ESXi hosts were not able to connect over the vSphere vMotion network. Check the vSphere vMotion network settings and physical network configuration.
- vSphere vMotion migration failed to create a connection with remote host 172.20.13.52: The ESXi hosts failed to connect over the vSphere vMotion network.
- Migration failed to connect to remote host 172.20.13.52 from host 172.20.12.51: Timeout.

  The IP addresses refer to the vSphere vMotion VMkernel interfaces on the remote host (ESXi02) and the local host (ESXi01).
- The vSphere vMotion migration failed because the destination host did not receive data from the source host on the vSphere vMotion network. Verify that your vSphere vMotion network settings and physical network configuration are correct.

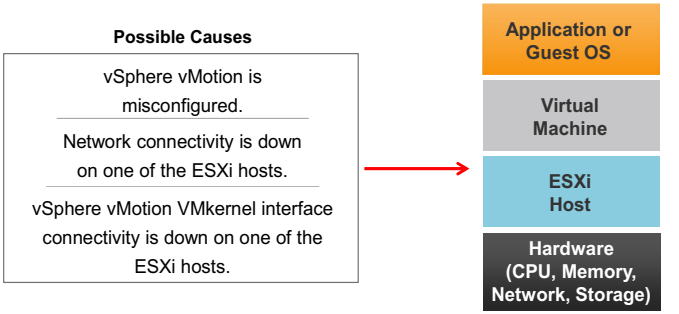
The first error message in the stack is helpful and tells you to check the vSphere vMotion network settings and physical network configuration. All error messages might not be so helpful.

https://t.me/learningnets

12-34 Example: Identifying Possible Causes

Use the information that you gathered to identify possible causes:

- Based on error messages, the vSphere vMotion migration failed because ESXi01 and ESXi02 failed to connect over the network named vMotion.
- This error indicates a possible misconfiguration on the ESXi host.
- Check the connectivity of the vSphere vMotion VMkernel interface.



The error message points to connectivity issues with the network named vMotion, with the following possible causes:

- vSphere vMotion is misconfigured.
- Network connectivity between ESXi01 and ESXi02 is down.
- The vSphere vMotion VMkernel interface connectivity between ESXi01 and ESXi02 is down.

When you initiate vSphere vMotion migration, several compatibility checks are performed before the migration is initiated. Thus, you can eliminate possible causes such as vSphere vMotion not being enabled or incompatible CPUs, because these configuration items are checked before the migration begins.

https://t.me/learningnets

12-35 Example: Resolving the Problem

In this example, assume that the root cause is an incorrect IP address of a VMkernel interface for vSphere vMotion on the ESXi02 host.

Assess the impact of the problem on operations:

- Probably high impact:
  - The problem affects any virtual machine that is migrated to the ESXi02 host.
  - The problem also affects the proper operation of vSphere DRS.

Identify possible solutions to resolve the problem:

- Short-term solution: Do not migrate virtual machines to the ESXi02 host.
- Long-term solution: Fix the IP address of the vSphere vMotion VMkernel interface of the ESXi02 host.

Implementing the solution should not require downtime.

When you have identified the root cause, identify possible solutions to fix the problem. The impact that the problem has on normal operations (high, medium, or low) determines how quickly the solution should be implemented.

Finally, determine the appropriate type of solution for this problem. You might implement a short-term solution so that the system works normally. Document all changes that you made to the system since the problem occurred. Also, back up your log files from the time the problem occurred, because logs rotate and might not be available in the future.

https://t.me/learningnets

12-36 VMware Severity Definitions and Response Targets

The severity of the problem and the service levels of the purchased support program determine the speed and method of response targets.

The table lists the response time targets for providing the initial response. For a full list of VMware's response times and a description of problem severity, go to https://www.vmware.com/support/policies/severity.

Severity	Definition	Mission-Critical Response
Critical 1	Production server or other mission-critical systems are down and no workaround is immediately available.	30 minutes or less, 24x7
Major 2	Major functionality is severely impaired.	2 business hours, 12x7
Minor 3	Partial, noncritical loss of functionality of the software.	4 business hours, 12x5
Cosmetic 4	General usage questions.	8 business hours, 12x5

https://t.me/learningnets

12-37 Review of Learner Objectives

You should be able to meet the following objectives:

- Define the scope of troubleshooting
- Use a structured approach to solve configuration and operational problems
- Apply a troubleshooting methodology to logically diagnose faults and improve troubleshooting efficiency

https://t.me/learningnets

12-38 Key Points

- A structured approach to troubleshooting helps you resolve problems quickly and effectively.
- Differentiating between the symptoms and the problem is an important step in the troubleshooting process.
- Prerequisite knowledge of how the VMware virtual infrastructure works, as well as your knowledge of your system's configuration, is very useful in the troubleshooting process.

Questions?

https://t.me/learningnets