Citrix XenServer 6.0 Administration

Citrix Course CXS-203-1I

Citrix XenServer 6.0 Administration

Citrix Course CXS-203-1I
April 2012
Version 1.1

Citrix authorized
Courseware

Table of Contents

Module 1: Introduction to XenServer 17
Overview 19
XenServer Product Line 20
  Product Simplification 20
  Architectural Changes 21
  Virtual Appliances and Broader Vendor Support 21
  Microsoft System Center Integration 21
  XenDesktop Integration 21
  Enhanced Operating System Support 22
  Other Enhancements and Improvements 22
  Test Your Knowledge: New Features in XenServer 6.0 22
XenServer Architecture Overview 24
  XenServer Architectural Components 24
  Test Your Knowledge: Architectural Components 25
  XenCenter Overview 26
  Resource Pools 27
  XenServer Storage Overview 28
  XenServer Networking Overview 29
  Provisioning Services Overview 29
Test Your Knowledge: XenServer Virtualization Platform 30

Module 2: Installing and Configuring XenServer 31
Overview 33
XenServer Installation 34
  Virtual Machine Storage 34
  Installing Single and Multiple Disk Systems 35
Installing XenServer 36
  Pre-Installation Checklist 36
  XenServer Installation Process 37
  IntelliCache 37
  Configuring NTP 37
  Test Your Knowledge: XenServer Installation 38
XenCenter Installation and Management Consoles 39
  XenCenter Consoles 39
XenServer Licensing Components 40
  License File Management 40
  Citrix License Server Virtual Appliance 40
  Obtaining License Files 41
  License Management Console 41
  Test Your Knowledge: Citrix License Server 42
Managing XenServer Users 43

Key Benefits of Active Directory Integration 43
Configuring Active Directory Integration 44
End-User Authentication Using Active Directory 44
Active Directory Integration 44
External Authentication Process 45
Test Your Knowledge: Integrating XenServer and Active Directory 45
Role-Based Access Control 46
Roles 46
Security Logs 47
Test Your Knowledge: Role-Based Access Control 47

Module 3: XenServer Networking 49
Overview 51
XenServer Networking Overview 52
Network Stacks Supported by XenServer 52
Network Components and Architecture 54
Network Architecture Diagram 54
Network Adapter Drivers 55
Private-Server and External Networks 56
Test Your Knowledge: Network Components 57
VLAN Support and Components 58
Initial Network Setup 60
Test Your Knowledge: VLANs 60
NIC Bond 61
NIC Bonding Architecture 61
Load Balancing 61
Test Your Knowledge: NIC Bond Modes 61
NIC Bond Configuration for Resource Pools 63

Module 4: XenServer Storage Repositories 65
Overview 67
Storage Technologies 68
Virtual Disk Image Formats 68
Storage Technology Comparison 69
Test Your Knowledge: Storage Technologies 70
XenServer Storage Architecture 71
Storage Repository Architecture 72
Multiple Storage Repositories 74
Local Storage Repositories 75
Virtual Disk Size 75
Test Your Knowledge: Storage Repositories 76
Test Your Knowledge: Local Storage 76
Storage Capability Comparison 78
NFS Storage Overview 79
NFS Architecture 79

NFS Virtual Disks 80
NFS Remote Storage Configuration and Recommendations 80
Test Your Knowledge: NFS Storage 80
iSCSI Storage Overview 81
iSCSI Architecture 82
iSCSI Adapters 83
iSCSI Setup Configuration and Recommendations 83
Test Your Knowledge: iSCSI Storage 83
Resizing a Storage Repository 84
Resizing a Storage Repository - iSCSI 84
Resizing a Storage Repository - HBA 84
Fibre Channel SANs Overview 85
Fibre Channel Architecture 85
LUN Device Path 86
Test Your Knowledge: Fibre Channel HBA Management 86
Dedicated NIC Bonds for Remote Storage 87
Dedicated Remote Storage Interface Configuration 87
Storage Multipathing 88
Citrix StorageLink Overview 89
Citrix StorageLink Storage Repository 89
StorageLink Supported Array Types 89
Upgrading to StorageLink with XenServer 6.0 90
Test Your Knowledge: Storage 90

Module 5: Creating and Managing Virtual Machines 91
Overview 93
Windows Virtual Machine Architecture 94
Virtual Memory and Disk Size Limits for Windows Virtual Machines 94
Virtual Device Support for Windows Virtual Machines 95
Installation from an ISO 96
ISO Libraries 96
To Create a Windows Virtual Machine 96
Test Your Knowledge: Creating an ISO Library for a Windows Virtual Machine 96
Linux Distributions 98
Creating Linux Virtual Machines 99
Test Your Knowledge: Using a Linux Demo Template 100
XenServer Tools 101
XenServer Tools for Windows Virtual Machines 101
XenServer Tools for Linux Virtual Machines 101
To Install XenServer Tools 102
Test Your Knowledge: XenServer Tools 102
Life Cycle Operations 103
Suspend and Resume 103
Deleting a Virtual Machine 103
Test Your Knowledge: Life Cycle Operations 104
Virtual Machine Templates 105

- Virtual Machine Snapshots ... 105
  - To Create a Template from a Windows Virtual Machine ... 106
  - Creating a Template from a Linux Virtual Machine ... 106
  - Test Your Knowledge: Determining Which Template to Use ... 107
- Exporting a Virtual Machine ... 108
- Copying a Virtual Machine ... 110
- Assigning Resources to a Virtual Machine ... 112
- Importing a Virtual Machine ... 113
- Modifying Virtual Machine or Template Resources ... 115
  - Virtual Disks ... 115
  - Virtual Disk Sizes ... 115
  - Virtual NICs ... 115
  - Virtual CPUs ... 116
  - Memory ... 116
- Virtual Appliance Packages ... 117
  - To Create Virtual Appliances ... 118
  - Exporting an Appliance Package ... 118
  - To Export an Appliance Package ... 118
  - Importing an Appliance Package ... 119
  - To Import an Appliance Package ... 119
- XenConvert Overview ... 120
  - Physical Machine Conversion ... 120
  - Converting from Physical to Virtual Machine ... 122
  - Performing a P2V Conversion ... 122
- Dynamic Memory Control ... 123
  - Dynamic Memory Control Optimization ... 124
  - Static Memory Range ... 124
  - DMC Behavior when Launching New Virtual Machines ... 125
  - Test Your Knowledge: Dynamic Memory Control ... 125

Module 6: Installing and Configuring Provisioning Services ... 127

- Overview ... 129
- Provisioning Services Technology ... 130
  - Provisioning Services for XenServer ... 130
  - Citrix Provisioning Services Components ... 130
  - Citrix Provisioning Services Infrastructure ... 132
  - Provisioning Services Administration ... 133
  - Administrative Roles ... 134
  - Key Services ... 135
  - DHCP Deployment Options ... 137
  - DHCP Configuration ... 138
  - Test Your Knowledge: Key Services ... 139
- Installation Planning ... 140
  - Provisioning Services Farm Design ... 140
  - Farms ... 140
  - SQL Database Server ... 141

- SQL Database Authentication and Configuration 141
- Configuration Wizard User Permissions 141
- Service Account Permissions 142
- Sites 143
- Stores 144

Storage Requirements 146
- Sample vDisk Storage Requirements 147

Provisioning Services Installation 148
- Provisioning Services Configuration 148
- Bootstrap Server Configuration 149
- Store Configuration 150

Provisioning Services Hosts in the Console 152
- Target Device Connections 152
- Server Properties 152
- Test Your Knowledge: Provisioning Services Installation and Configuration 152

Target Device Collection 154
- Target Device Template 154
- Target Device Properties 154
- Target Device Additions to the Database 155
- Auto-Add Wizard 155
- Test Your Knowledge: Target Devices 155

Active Directory Integration 157
- Domain Password Validation Process 158
- Automatic Password Renegotiation 158
- To Integrate Active Directory 159
- To Reset Computer Accounts for Target Devices 159
- Test Your Knowledge: Active Directory 160

Module 7: Managing vDisks and Target Devices 161

Overview 163

vDisk Image Modes 164
- Standard Image Mode 164
- Private Image Mode 165

Write Caches 166
- Write-Cache Types 166
- Cache on Server Disk 167
- Cache on Server Disk: Local Storage 168
- Cache on Server Disk: Shared Storage 168
- Cache on Target Device Hard Drive 169
- Cache on Target Device Hard Drive: Local Storage 169
- Cache on Target Device Hard Drive: Shared Storage 170
- Target Device-based RAM Cache 171
- Target Device RAM 171
- Cache Persistent on Server 172
- Test Your Knowledge: vDisk Image Modes and Write Cache 172

Microsoft Licensing for Provisioning Services 174

Master Target Device ... 175
Preparing a Master Target Device ... 175
Preparing the Hard Disk of the Master Target Device ... 175
Configuring the BIOS of a Master Target Device ... 176
To Configure the BIOS of a Master Target Device ... 176
Installing Master Target Device Software ... 176
Imaging a Windows Target Device with XenConvert ... 177
Test Your Knowledge: Master Target Device ... 178
vDisk File Creation ... 179
vDisk Lifecycle Operations ... 179
VHD Formats ... 180
Create vDisks Automatically Using Imaging Wizard ... 181
To Create a New vDisk Automatically Using the Imaging Wizard ... 181
Creating a vDisk Manually ... 182
vDisk Management ... 183
Assigning vDisks to Target Devices ... 183
To Assign a vDisk to a Single Target Device ... 184
vDisk Versions ... 184
vDisk Backup ... 184
vDisk Updates ... 185
Autoupdate Tool ... 186
vDisk Update Management Requirements ... 186
To Install Updates Automatically ... 187
Incremental Update Rollback ... 188
vDisk Replication ... 188
vDisk Inventory Service ... 188
High Availability Overview ... 189
Provisioning Services Failover ... 190
Configuring the Boot File for High Availability ... 191
Adding Provisioning Services Hosts to the Boot File ... 191
Enabling High Availability on vDisks ... 191
To Provide Provisioning Services Hosts with Access to Stores ... 192
Considerations for Offline Database Support ... 192
To Enable Offline Database Support ... 193
Stream Logging ... 194
To Enable Logging ... 194
Troubleshoot vDisk Images ... 195
Test Your Knowledge: Managing vDisks ... 197

Module 8: Implementing Resource Pools ... 199
Overview ... 201
Resource Pools ... 202
Pool Member Types ... 202
Resource Pool Communication ... 203
Homogeneous Pools ... 203
Heterogeneous Pools ... 204

Test Your Knowledge: Resource Pools 204
Shared Configurations 206
Adding a Host to a Resource Pool 206
To Add a Member to a Resource Pool 206
To Remove a Secondary Member from a Resource Pool 207
XenMotion 207
Migrating a Virtual Machine or Template 208
Test Your Knowledge: Pool Member Types 208
Pool Member Failure Overview 210
Failure Detection Process 210
To Promote a Secondary Member to a Master 210
Test Your Knowledge: Resource Pool Requirements 212

Module 9: Distributed Virtual Switching 213
Overview 215
Distributed Virtual Switching Features 216
Distributed Virtual Switching Components 217
Test Your Knowledge: Distributed Virtual Switch Features and Components 218
vSwitch Controller Virtual Appliance Deployment 219
Accessing the vSwitch Controller 219
Test Your Knowledge: vSwitch Controller Access 219
vSwitch Controller Configuration 220
High Availability 220
Test Your Knowledge: Distributed Virtual Switch Configuration 221
Monitoring Network Activity 222
Viewing Flow Statistics 223
Exporting NetFlow Statistics 223
Managing Address Groups and Virtual Machine Groups 224
Test Your Knowledge: Monitoring Networking Activity 224
Distributed Virtual Switching Policy Configuration Hierarchy 225
Access Control List Policies 225
Access Control Policies Hierarchy 226
Defining Access Control List Rules 226
Access Control List Rule Enforcement Order 227
Configuring Fail Modes 227
Test Your Knowledge: Access Control List Rule Enforcement 228
Setting Up Port Configuration Policies 229
Configuring QoS 229
Configuring RSPAN 230
To Configure RSPAN 230
Configuring MAC Address Spoof Checking 230
Test Your Knowledge: Port Configurations 231
Cross-Server Private Networks 232
Cross-Server Private Network Requirements 232
XenMotion with Cross-Server Private Networks 232
Recovering from a Failed vSwitch Controller 233

Module 10: Workload Balancing 235
Overview 237
Workload Balancing Overview 238
  Workload Balancing Key Concepts 238
  Workload Balancing Components 239
  To Download the Workload Balancing Virtual Appliance 239
Import Considerations for the Workload Balancing Virtual Appliance 240
Workload Balancing Configuration 241
  To Update Workload Balancing Credentials 241
  Test Your Knowledge: Importing Workload Balancing Considerations 241
Connecting to the Workload Balancing Virtual Appliance 243
  Workload Balancing Access Control Permissions 243
Workload Balancing Settings 245
  Critical Thresholds 245
  Critical Threshold Evaluation and Edit 245
  Metric Weighting 246
  Fixed and Scheduled Optimization Modes 246
  Optimization Recommendations 247
Workload Balancing Power Management 248
  Power Management Behavior 249
  Power Management and Virtual Machine Consolidation 250
  To Apply Optimization Recommendations Automatically 250
  Optimal Server Selection 250
  Test Your Knowledge: Power Management 251
  Host Exclusion from Recommendations 251

Module 11: Configuring High Availability 253
Overview 255
High Availability Requirements 256
  High Availability Considerations 256
Restart Priorities 258
  Server Failure Tolerance 258
  Overcommitting 259
  Host Fencing 259
  Test Your Knowledge: High Availability Requirements 259
Disaster Recovery 261
  Disaster Preparedness and Response 261
  Test Your Knowledge: Disaster Recovery 262
Backup and Recovery of XenServer Hosts and Virtual Machines 263
  Metadata 264
    Backing Up a Pool Installation 264
    To Restore a Pool Installation 265
    Backing Up Virtual Machines as Snapshots 265
      To Take a Snapshot of a Virtual Machine 266
    To Restore a Running XenServer Host 266
    Listing Snapshots Using the Command-line Interface 266

Deleting a Snapshot ... 267
Restoring a Virtual Machine from a Snapshot ... 267
Test Your Knowledge: Creating Snapshots ... 267
Test Your Knowledge: Backup and Restore ... 268

Module 12: Managing and Troubleshooting XenServer ... 269
Overview ... 271
Command-line Interface Overview ... 272
Command-line Interface Operations Targets ... 272
Command-line Interface Basics ... 272
Command-line Interface Shortcuts ... 273
XenServer Menu-Driven Text Console ... 274
Test Your Knowledge: XenServer Commands ... 274
Monitoring XenServer ... 276
Using the XenTop Utility to Monitor Host and Virtual Machine Performance ... 276
Logging to a Remote Syslog Server ... 278
To Log Using a Remote Syslog Server ... 278
Test Your Knowledge: Monitoring XenServer ... 278
Events and Alerts ... 279
Alert Customization ... 279
Persistent XenServer Performance Statistics ... 280
Performance Statistics in XenCenter ... 280
Test Your Knowledge: XenServer Alerts ... 281
XenServer Logs and Reports ... 282
XenCenter Event Logs ... 282
Server Status Reports ... 282
Xen-Bugtool Utility ... 282
XenServer Crashdump Overview ... 283
Generating the Crashdump File in the Command-Line Interface ... 283
Test Your Knowledge: Viewing and Generating Reports and Logs ... 284
Running a Network Trace and Triggering a Memory Dump ... 285
Test Your Knowledge: Troubleshooting XenServer ... 286
XenServer Updates and Hotfixes ... 287
Update or Hotfix Considerations ... 287
To Update XenServer Hosts Using XenCenter ... 288
To Update Hosts or a Pool Using the Command-line Interface ... 288
Rolling Pool Upgrade ... 290
Planning a Rolling Upgrade ... 290
Workload Balancing Reports Overview ... 291
Workload Balancing Reports ... 291

Notices

Citrix Systems, Inc. (Citrix) makes no representations or warranties with respect to the content or use of this publication. Citrix specifically disclaims any expressed or implied warranties, merchantability, or fitness for any particular purpose. Citrix reserves the right to make any changes in specifications and other information contained in this publication without prior notice and without obligation to notify any person or entity of such revisions or changes.

Citrix Systems, Inc.

851 West Cypress Creek Road

Fort Lauderdale, FL, 33309 USA

www.citrix.com

The following marks are service marks, trademarks or registered trademarks of their respective owners in the United States and other countries.

Mark	Owner
Active Directory®, Microsoft®,Microsoft Internet Explorer®, Windows®,	Microsoft Corporation
AMD®	Advanced Micro Devices, Inc.
Citrix®, Citrix Provisioning Services™, XenApp™, XenDesktop™, XenServer™, XenCenter™, XenMotion™,	Citrix Systems, Inc.
Dell™, Net EqualLogic™	Dell Inc.
Emulex®	Emulex Corporation
IBM®	International Business Machines Corporation
Intel®	Intel Corporation
Linux®	Linus Torvalds
NetApp®	Network Appliance, Inc.

Mark	Owner
PuTTY®	Simon Tatham, Open Source Certified
QLogic®	QLogic Corporation
Realtek™	Realtek Semiconductor Corporation
Red Hat®	Red Hat, Inc.
Sun™	Sun Microsystems, Inc.
Suse®	Novell, Inc.
Toolwire®	Toolwire
Unix®	The Open Group

Other product and company names mentioned herein might be the service marks, trademarks or registered trademarks of their respective owners in the United States and other countries.

Credits

Instructional Designers:	Rachel White, Orlando A. Martinez, Raymond Kung, Omid Mirshafiei
Product Specialist:	Andrew Garfield, George Komoto
Graphic Artists:	Nathan Jackson, Joshua Jack
Manager:	Gina Alesse
Editor:	Kathryn Morris
Translation Coordinator:	Yashica Burgess
CCI Stakeholder	Jeff Apsley
Subject Matter Experts:	Christopher Campbell, Joel Stocker, Mark Simmons, Shane Broomhall, Blaine Anaya, Nick Kieffer, Peter Svoboda, Patrick Carey, Elisabeth Teixeira

Module 1

Introduction to XenServer

Overview

Citrix XenServer is a complete server virtualization platform, optimized for both Windows and Linux virtual servers, with all the capabilities required to create and manage a virtual infrastructure.

XenServer is a hypervisor that runs on the physical or host server to provide a virtual computer environment. XenServer works by virtualizing the hardware. Hardware virtualization abstracts system components, such as hard drives, resources, and ports, and allocates them to the virtualized servers running on the system. These virtualized servers are known as virtual machines. They run operating systems and applications that are known as guest software.

Timings

- Module: 60 minutes
- Total time: 60 minutes

Objectives

After completing this module, you will be able to:

- Describe the XenServer virtualization platform.
- Describe the new features in XenServer 6.0.
- Identify XenServer architecture and key components.
- Describe resource pool communication within a XenServer environment.
- Describe the XenServer storage architecture.
- Describe the network architecture for XenServer.

XenServer Product Line

XenServer is available in four editions to meet the needs of any organization.

Free

XenServer Free is a downloadable virtualization platform that includes features like live migration, virtual machine disk snapshots, Active Directory integration, shared storage support, and centralized multiserver management, plus physical-to-virtual and virtual-to-virtual conversion tools.

Advanced

XenServer Advanced includes specific administration features within XenCenter, including advanced alerting and performance history, automated virtual machine protection and recovery, and dynamic memory control.

Enterprise

XenServer Enterprise adds Role-based Access Control (RBAC), integrated StorageLink configuration, live memory snapshots, automated Workload Balancing, and remote power management to the set of available features.

Platinum

XenServer Platinum adds features that are administered from outside of XenServer, such as Provisioning Services and Site Recovery with StorageLink.

For more information about the XenServer editions, see the *www.citrix.com* Web site.

Product Simplification

XenServer has been simplified for easier, faster setup:

- XenServer no longer requires Windows-based virtual machines for features such as StorageLink, Site Recovery, and Workload Balancing. In fact, for StorageLink and Site Recovery, no additional management infrastructure is required.
- Workload Balancing and its historical reporting features have been moved to a Linux-based virtual appliance for easy installation and management.
- The Linux Supplemental Pack has been removed, leaving only one base installation ISO. The Demo Linux virtual machine functionality has been moved to a virtual appliance format, so it

can be easily imported into a host or resource pool. You can download the Demo Linux virtual appliance from the *mycitrix.com* Web site.

Architectural Changes

XenServer 6.0 streamlines the product architecture and enhances performance:

- The XenServer 6.0 release is based on the Xen 4.1 hypervisor.
- The Open vSwitch is now the default network stack for the product. Improvements to Distributed Virtual Switching include a fail-safe option and various improvements based on customer feedback from XenServer 5.6 Feature Pack 1.
- General network performance has been improved, particularly aggregate host network throughput.
- Support for hardware-assisted Single Route I/O Virtualization network performance optimizations has been improved, particularly for use with the NetScaler VPX and SDX products.

Virtual Appliances and Broader Vendor Support

Within XenCenter, you can create multi-virtual-machine virtual appliances, with relationships between the virtual machines, such as start up sequence, for use with high availability and Site Recovery. Virtual appliances can be easily imported and exported using the Open Virtualization Format (OVF) standard.

Importing of the VMware Virtual Machine Disk (VMDK) and Microsoft Virtual Hard Disk (VHD) image formats is now integrated into XenCenter.

Microsoft System Center Integration

XenServer 6.0 supports System Center Virtual Machine Manager (SCVMM) 2012 for managing XenServer hosts and virtual machines. To enable these management capabilities, install a supplemental pack from Citrix.

For more information about SCVMM, visit: http://www.microsoft.com.

XenDesktop Integration

XenServer 6.0 is the first XenServer release to include High Definition User Experience (HDX) enhancements for an optimized end user experience with virtual desktops.

HDX technology is a set of capabilities that delivers a high definition desktop virtualization user experience to end users for any application, device, or network. HDX technology provides network and performance optimizations to deliver the best end user experience over any network, including low-bandwidth and high-latency WAN connections.

With version 6.0, a physical Graphics Processing Unit (GPU) can be assigned to a virtual machine, so the applications running in the guest operating system can use GPU instructions (GPU Pass-Through). This feature provides significant Total Cost of Ownership (TCO) benefits for the XenDesktop HDX 3D Pro technology used for the delivery of CAD and other graphical applications using virtual desktops. With GPU Pass-Through, either a single GPU card or a GPU on a multi GPU card can be assigned to a virtual machine.

Enhanced Operating System Support

XenServer 6.0 brings broader support for the following guest operating systems:

- Formal support for Ubuntu 10.04
- Updates for support of Red Hat Enterprise Linux (RHEL) 5.6, CentOS 5.6, and SUSE Linux Enterprise Server (SLES) 10 Service Pack 4
- Experimental virtual machine templates for Solaris and Ubuntu 10.10

Other Enhancements and Improvements

Enhancements have been made to reliability, capacity, and localization:

- A Rolling Pool Upgrade wizard is provided in XenCenter to enable more reliable upgrades from XenServer 5.6, 5.6 Feature Pack 1, and 5.6 Service Pack 2.
- High availability now supports Network File System (NFS) for storage of the heartbeat disk; the heartbeat disk provides a way to check for communication between hosts.
- Host RAM support has been increased to 1 TB.
- Virtual machines virtual CPU (vCPU) and virtual RAM (vRAM) support is increased, up to 16 vCPUs and 128 GB vRAM for Windows; increased Linux vCPU and vRAM support levels vary by Linux distribution.
- XenServer 6.0 improves Network Interface Card (NIC) bonding reliability and adds formal support for active/passive bonding.

Test Your Knowledge: New Features in XenServer 6.0

Match the following terms with the correct descriptions.

- Open vSwitch
- SCVMM
- HDX
- Physical GPU
- Workload Balancing (WLB)
- High Availability

- Storage Link, Site Recovery
- IntelliCache

Term	Description
Workload Balancing (WLB)	Is a Linux-based virtual appliance.
HDX	Provides an optimized end user experience with virtual desktops.
SCVMM	Manages XenServer hosts and virtual machines.
High Availability	Supports NFS for storage of the heartbeat disk.
IntelliCache	Requires a thin-provisioned, local storage repository.
Open vSwitch	Is the default network stack.
Physical GPU	Can be assigned to a virtual machine so apps running in the guest operating system can use GPU Pass-Thru.
Storage Link, Site Recovery	Require no additional management structure.

XenServer Architecture Overview

The XenServer virtualization platform uses two technologies:

Hardware-assisted virtualization

XenServer is designed to use hardware-assisted virtualization technologies delivered by both Intel and AMD. With hardware-assisted virtualization, the guest operating system on the virtual machine does not require modifications in order to have direct access to the server resources.

Paravirtualization

Paravirtualization is accomplished by allowing a guest operating system, such as Windows, to communicate with the hypervisor. This direct communication improves performance and is enabled on Windows virtual machine running on XenServer by installing XenServer Tools.

XenServer Architectural Components

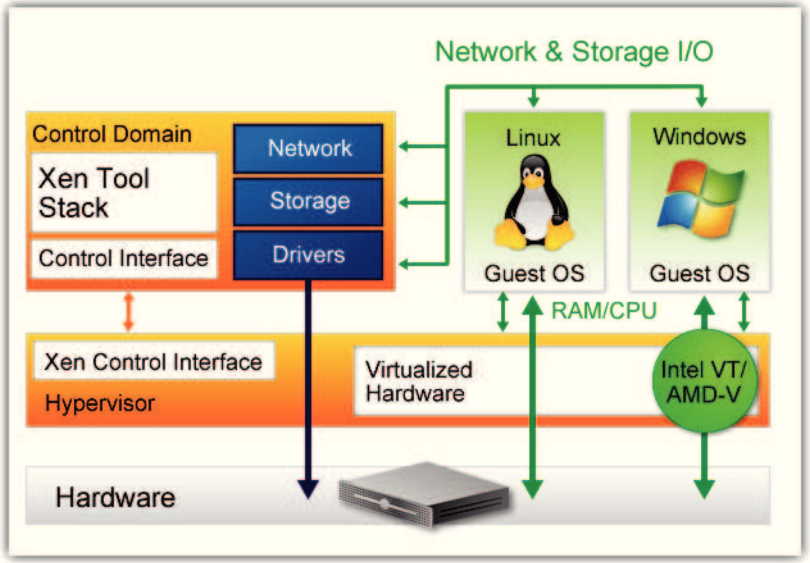

The following list provides descriptions of the XenServer architectural components:

Hardware Layer

The hardware layer contains the physical server components, including memory, CPU, and disk drives.

Xen Hypervisor

The Xen hypervisor is a thin layer of software that runs on top of the hardware. Xen provides an abstraction layer that allows each physical server to run one or more virtual machines, effectively decoupling the operating system and its applications from the underlying hardware.

Control Domain

The control domain is a Linux virtual machine with higher priority to the hardware than guest operating systems. The control domain manages the network and storage I/O of all virtual machines. Because the control domain uses Linux device drivers, a broad range of physical devices is supported.

Guest Operating System

The guest operating system is the operating system that is installed on the virtual machine.

Linux Virtual Machine

The Linux virtual machines include paravirtualized kernels and drivers. Storage and network resources are accessed through the control domain, while CPU and memory are accessed through Xen to the hardware.

Windows Virtual Machine

The Windows virtual machines use paravirtualized drivers to access storage and network resources through the control domain. XenServer is designed to use the virtualization of Intel VT- and AMD-V-enabled processors.

Test Your Knowledge: Architectural Components

Match the following terms with the correct descriptions.

- Control Domain
- Xen Hypervisor
- Hardware Layer
- Linux Virtual Machine
- Windows Virtual Machine

Description	Term
Uses hardware virtualization to enable the high-performance virtualization capabilities of the host OS kernel without using legacy emulation technology.	Windows Virtual Machine
Runs on top of the hardware as a thin abstraction layer of software, decoupling the OS and its applications from the underlying hardware.	Xen Hypervisor
Is a Linux virtual machine with higher priority to the hardware than the priorities of the guest operating systems to the hardware.	Control Domain
Contains paravirtualized kernels and drivers. Storage and network resources are accessed through the control domain, while CPU and memory are accessed through Xen to the hardware.	Linux Virtual Machine
Contains the physical server components, including memory, CPU, and disk drives.	Hardware Layer

XenCenter Overview

XenCenter is a graphical, Windows-based user interface. XenCenter allows you to manage XenServer hosts, resource pools and shared storage, and to deploy, manage and monitor virtual machines from your Windows desktop machine.

Multiple XenCenter consoles can be used to manage a single resource pool. The event mechanism keeps each client updated.

Resource Pools

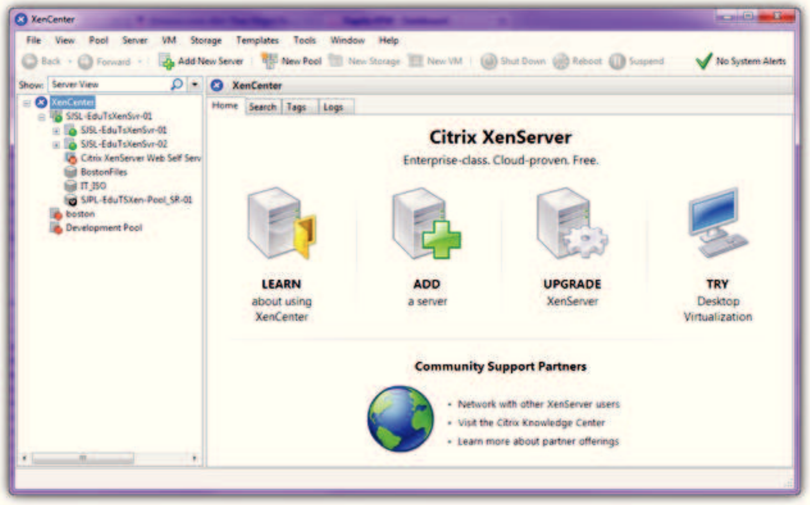
XenCenter
File View Pool Server VM Storage Templates Tools Window Help
Back Forward Add New Server New Pool New Storage New VM Shut Down Reboot Suspend
No System Alerts
Show: Server View
XenCenter
SJSL-EduTsXenSvr-01
SJSL-EduTsXenSvr-01
SJSL-EduTsXenSvr-02
Citrix XenServer Web Self Serv
BostonFiles
IT_ISO
SJPL-EduTSXen-Pool_SR-01
boston
Development Pool
Home Search Tags Logs
Citrix XenServer
Enterprise-class. Cloud-proven. Free.
LEARN about using XenCenter
ADD a server
UPGRADE XenServer
TRY Desktop Virtualization
Community Support Partners
Network with other XenServer users
Visit the Citrix Knowledge Center
Learn more about partner offerings

XenCenter requires a minimum of 1GB of RAM. 2 GB or more of RAM is recommended.

XenServer allows you to manage multiple XenServer hosts as a single entity through the use of resource pools. Resource pools provide you with the ability to move and run virtual machines on different XenServer hosts. This ability allows you to move virtual machines from one XenServer host to another:

- In the event of a host failure
- In preparation for upgrade a XenServer host
- To consolidate virtual machines to a select number of XenServer hosts in order to reduce power consumption

Topics to discuss:

- Virtual machine icons in XenCenter. Discuss the different states of a virtual machine and their corresponding icons displayed in XenCenter.
- The Logs tab. Demo the Logs tab and discuss the information displayed.

XenServer Storage Overview

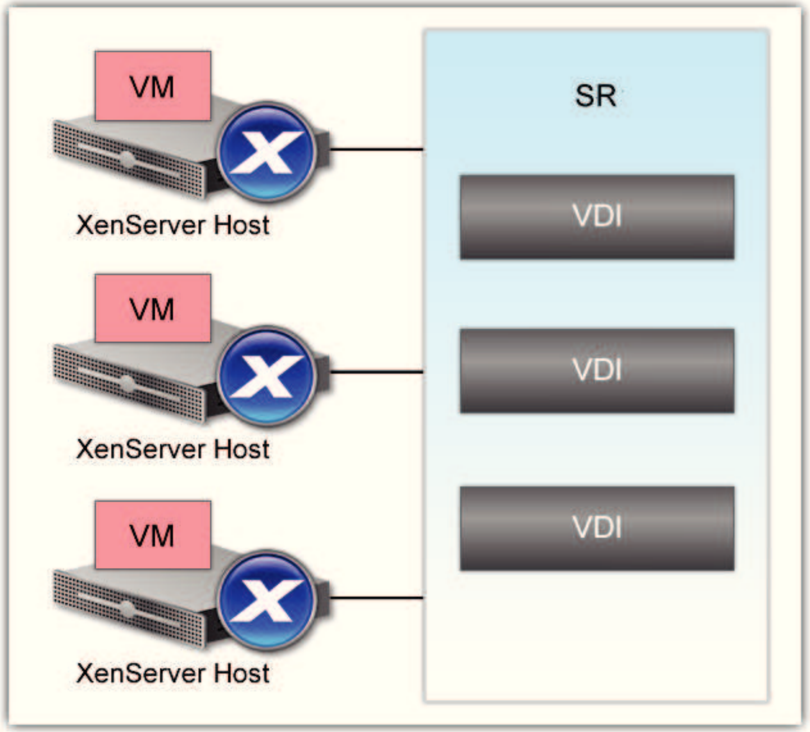

XenServer storage targets are called storage repositories. A storage repository stores virtual disk images, which contain the contents of a virtual disk. The virtual disk images are the fundamental unit of virtualized storage in XenServer.

XenServer hosts can have multiple storage repositories, and storage repositories can be shared between XenServer hosts.

XenServer Networking Overview

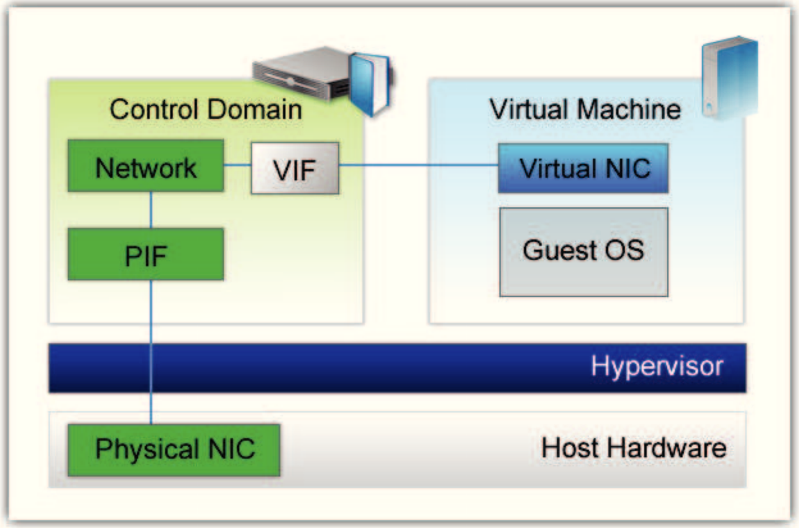

One network is created for each physical Network Interface Card (NIC) on the physical machine during the XenServer installation. These networks are used to provide communication between the physical network and virtual machines running on XenServer hosts.

There are three network objects in XenServer:

- Physical NIC, the physical network card on the XenServer host
- Virtual interface, the virtual network interface on a virtual machine running on the XenServer host
- Network, the virtual Ethernet switch used to route network traffic on a XenServer host

Provisioning Services Overview

The Provisioning Services infrastructure is based on software-streaming technology. This technology allows computers to be provisioned in real-time from a single shared-disk image. By streaming a single shared disk image (vDisk) rather than copying images to individual machines, Provisioning Services enables organizations to reduce the number of systems that they manage, even as the number of end users continues to grow. The administrators can completely eliminate the need to manage and update individual systems. Instead, all image management is done on the master image that is streamed.

Test Your Knowledge: XenServer Virtualization Platform

1. Which two technologies does XenServer use? (Choose two.)
   a. Emulation
   b. Full virtualization
   c. Paravirtualization
   d. Hardware-assisted virtualization

   Answer: c and d
2. Networking options can be configured with XenCenter but not the xe CLI.
   a. True
   b. False

   Answer: b
3. All three types of server-side networking entities have a globally unique UUID.
   a. True
   b. False

   Answer: a

Module 2

Installing and Configuring XenServer

Overview

The XenServer installer uses a graphical text interface and is designed to ask a minimum set of questions to configure the installation. You can install XenServer from the product CDs or from a network location.

XenCenter is a Windows application that provides a graphical management console for managing and deploying XenServer hosts and virtual machines. From XenCenter, you can assign shared storage and create resource pools.

Timings

- Module: 90 minutes
- Exercises: 20 minutes
- Total: 110 minutes

Objectives

After completing this module, you will be able to:

- Describe the XenServer installation process.
- Import the Citrix License Server Virtual Appliance.
- Integrate XenServer and Active Directory to support end user authentication.
- Secure a XenServer host by configuring Role-based Access Control.

XenServer Installation

The XenServer host consists of a Xen-enabled Linux operating system, a management agent, virtual machine templates, and a local storage repository reserved for virtual machines. The XenServer host must be installed on a dedicated 64-bit x86 server.

You can configure local disk storage repositories during the initial XenServer installation. When more than one physical disk is available, additional installation options are available. After installation, you can create additional local disk and remote storage repositories.

Do not install any other operating system in a dual-start configuration with the XenServer host; this configuration is not supported.

Virtual Machine Storage

The amount of storage required for virtual machines depends on the operating system and the options selected during installation of the virtual machine. During initial setup, only local and SAN storage options are available. Installing the XenServer host enables other virtual machine storage repository options.

With multiple physical disks, the following virtual machine storage repository options are available:

- The virtual machine storage repository on the same disk as the XenServer host, spanning multiple disks
- The virtual machine storage repository and the XenServer host on separate disks

Installing Single and Multiple Disk Systems

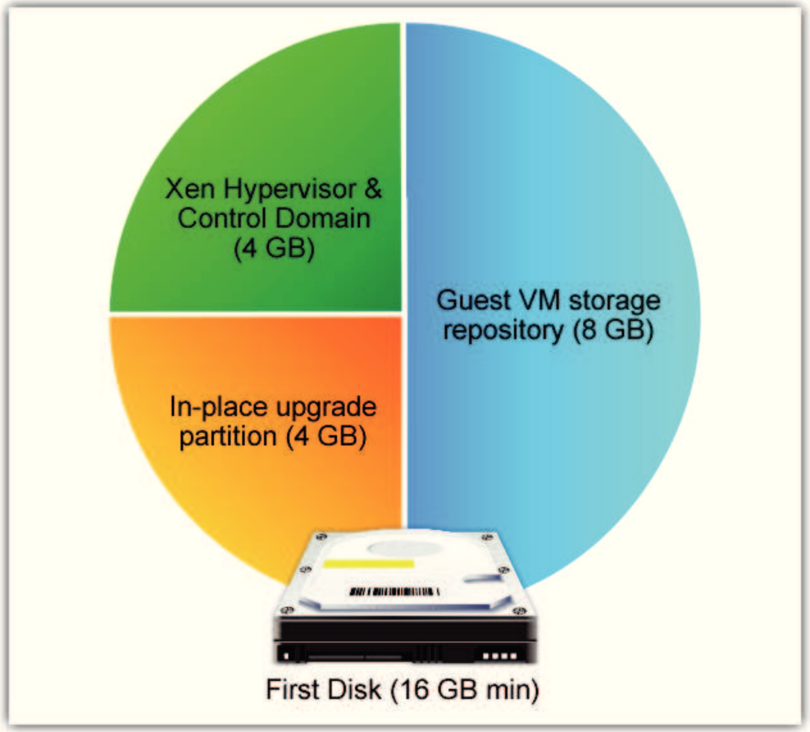
Xen Hypervisor & Control Domain (4 GB)
Guest VM storage repository (8 GB)
In-place upgrade partition (4 GB)
First Disk (16 GB min)

The XenServer host can be installed on either a single local or multiple local disks. In a single disk configuration, both the XenServer host and the local storage repository reside on the same disk.

When installed on multiple disks, the XenServer host is installed on one disk and the virtual machine storage repository is installed on a separate disk. The storage repository can be installed in full on one or multiple disks or on a SAN-attached LUN.

Two configurations are recommended for a multiple local disk environment: configuring the XenServer disk in a RAID 1 format and configuring the storage repository in a RAID 5 setup for performance and reliability.

Installing XenServer

There are three XenServer installation methods:

- Install from a CD.

  Download the installer (ISO file format) and burn it to a CD. To download the installer, visit *www.citrix.com/xenserver*.
- Set up a network-accessible Trivial File Transfer Protocol (TFTP) server to start up using a Pre-Boot Execution Environment (PXE).
- Install XenServer to a remote disk on a SAN to enable start up from SAN.

For more information about Installing XenServer, see Citrix article CTX130421 on *http://support.citrix.com.*

Pre-Installation Checklist

The XenServer host computer is dedicated entirely to the task of running virtual machines and is not used for other applications. The XenServer host runs an optimized and hardened Linux partition with a Xen-enabled kernel, which controls the interaction between the virtual devices, seen by virtual machines and the physical hardware.

Before installing XenServer:

- Verify that the server meets the CPU, memory, and networking requirements.
- Verify availability of a whole physical disk or Logical Unit Number (LUN) with at least 16 GB of space. The installer will list only those disks with enough space.
- If installing on a remote disk, test the connection before installing.

For more information about the Hardware Compatibility List (HCL), see *www.hcl.vmd.citrix.com*. You can also e-mail Citrix at *xenserver.hcl@citrix.com* with specific questions.

- The XenServer installation process destroys any existing operating system or data on the selected disk.
- The XenServer host gives an SSL error and fails to connect if there is a time difference between the XenCenter console and the XenServer host. Network Time Protocol (NTP) synchronizes the Linux system clock with an accurate time source. For information about how to set up an NTP server for a XenServer network, refer to Citrix article CTX116307 on *http://support.citrix.com.*

XenServer Installation Process

During the installation process, the installer:

1. Prepares the disks for startup data, the control domain, and virtual machine storage.
2. Installs an EXTLINUX boot loader.
3. Detects hardware.
4. Creates the control domain:
    a. Installs the XenServer host packages
    b. Configures the device drivers
    c. Configures the network/storage options
5. Initializes the virtual machine storage repository (optional).
6. Restarts the XenServer host.

If virtualization technologies are not enabled at the BIOS, you will receive an error message; however, the installation will proceed.

IntelliCache

Using XenServer with IntelliCache makes hosted XenDesktop deployments more cost effective by enabling the use of a combination of shared storage and local storage. IntelliCache is the intelligent management of write-cache.

IntelliCache works by caching data from a virtual machine's parent VDI in local storage on the virtual machine host. This local cache is then populated as data and read from the parent VDI. When many virtual machines share a common parent VDI, the data pulled into the cache from on virtual machine can be used by another virtual machine. Therefore, further access to the master image on shared storage is not required.

This feature is only supported when using XenServer with XenDesktop.

The requirements for using IntelliCache are:

- A thin-provisioned, local storage repository, which is configured during XenServer installation by selecting Optimized storage on XenDesktop for thin provisioning
- An NFS- or EXT-based shared storage to host the source virtual disk image

For more information about IntelliCache, see Citrix article CTX130421 on *support.citrix.com*.

Configuring NTP

You can use NTP to enable the server to determine local time.

Windows

If choosing NTP over manual entry to determine local time, either select **NTP is configured by my DHCP server** so that DHCP will set the time on the server, or at least enter one NTP server name or IP address. Click **OK**. Make sure port 123 using the UDP protocol is open on your firewall. For more information about configuring NTP, see Citrix article CTX130422 on *support.citrix.com*.

XenServer assumes that the time setting in the BIOS of the server is the current time in Universal Time Clock (UTC). XenServer compares the UTC time, so the XenServer hosts can each be in different time zones. To ensure that synchronization is correct, you can choose the same NTP servers for your XenServer resource pool and the Active Directory server.

Test Your Knowledge: XenServer Installation

Place the XenServer installation steps in the correct order.

1	Format the disks for startup data, control domain, and virtual machine storage
3	Detect hardware
5	Optionally initialize the virtual machine storage repository
2	Install a boot loader
6	Restart the XenServer host
4	Create the control domain

XenCenter Installation and Management Consoles

XenCenter is installed on a remote machine and connects to the XenServer host through the network to manage the XenServer environment. XenCenter is installed using the Base Installation ISO or by locating the latest XenCenter version from the *www.citrix.com* or *www.mycitrix.com* Web sites.

Uninstall any previous version before moving forward. If installing from a CD or DVD, open the client_install folder and run the `XenCenter.msi` file.

> XenCenter is backwards compatible.

XenCenter Consoles

XenServer provides XenCenter management consoles for Windows and Linux virtual machines.

Windows Console

XenCenter uses a graphical console to interact with Windows virtual machines. Virtual machines write their screens to a Cirrus VGA adapter, which is then converted by the control domain to a Virtual Network Computing (VNC) stream. This data is sent over the default SSL link on port 443 to XenCenter.

XenCenter includes a remote management feature that starts the built-in Windows Remote Desktop Protocol (RDP) client and sends it the virtual machine hostname if RDP and networking inside the virtual machine are enabled. This remote management method uses less bandwidth than the VNC console. The RDP data is carried over a separate network connection to the XenCenter client.

Linux Console

A text console is available with Linux virtual machines.

The graphical console requires a standard VNC server running within the virtual machine. XenCenter connects to the VNC server using the 5900 default port and the IP addresses returned from XenServer Tools. This connection is made outside of the SSL connection. Click Switch to X Console to access the graphical VNC console.

> If the Switch to X Console button is dimmed, then VNC is likely not configured on port 5900, a firewall is blocking traffic, or networking is not functional within the virtual machine.

XenServer Licensing Components

You must consider three components in order to deploy licensing correctly.

License Server The license server stores license files.

License File The license file keeps the license information for the product. It contains vital information such as the product edition and any applicable expiration dates.

License Administration Console The license administration console allows you to maintain the license server and license files over a Web-based interface.

Licensing Communication Overview

Citrix products depend on communication with the license server. You must perform the following tasks for a license server to accept connection and license requests:

- Add a license file to the license server
- Configure the product to use a specific license server

License File Management

Citrix requires each organization that uses Citrix products to purchase licenses for the product. The licenses allow for a connection to the product and enable the use of features in the product edition. License files store the company license information in a plain text format with authenticated content. Each license file can store information for one or more licenses.

During the licensing section, please refer to the Citrix article CTX128013 *XenDesktop Licensing FAQ*. *http://support.citrix.com*

Citrix License Server Virtual Appliance

XenServer 6.0 Advanced, Enterprise, and Platinum editions support a Citrix License Server Virtual Appliance (11.6.1 or higher) and a XenServer 6.0 license. A Citrix License Server Virtual Appliance is available for download from the Citrix XenServer 6.0 download page. To use an existing Citrix License Server Virtual Appliance, you must install a XenServer 6.0 license. For more information about downloading and installing the Citrix License Server Virtual Appliance, see Citrix article CTX124501 on *support.citrix.com*.

Obtaining License Files

The *www.mycitrix.com* Web site issues license files. You can allocate some or all of the licenses to one or more license servers. Therefore, you are not obligated to allocate all licenses simultaneously and can choose where to use the remainder at a later date.

This administration design allows companies to purchase licenses in bulk and distribute them as needed for various licenses servers, production farms, test farms, or other schema that fit the environment. For example, if you purchase a single 100-count license you could allocate it to several license files.

License Management Console

The license management console is a Web-based interface that allows you to maintain the license server and manage license files for that license server.

> You cannot install the license management console on a server other than the license server, but you can access it remotely through a Web browser.

The following list provides a brief description of the licensing features available using the license management console.

Tracking License Usage	Tracks concurrent license information.
Reporting	Creates reports based on current license usage.
Configuring Alerts	Creates and views alerts based on license usage and expiration dates.
Configuring Delegated Administrators	Assigns rights to administrators to limit capabilities and ensure proper license management.

Consider the following:

- Citrix recommends that you configure Secure Sockets Layer (SSL) and configure Secure HTTP(S) when accessing the license management console using a browser on a UNIX workstation, or in an unsecured environment.
- If the vendor daemon stops running, you can restart vendor daemon services in the license management console, which is less intrusive than restarting the server.
- Citrix recommends that you use a Virtual Private Network (VPN) when accessing the license management console from outside the network.

- Ask students if they know how to download a license file. If not demo this procedure.
- The license for XenServer Free is not downloaded from *www.mycitrix.com*. When you attempt to activate the license in XenCenter, you are taken to a Web form. The license for XenServer Free will be e-mailed to you when the form has been submitted.
- For more information about license files, see Citrix article CTX125301 on *http://support.citrix.com*

- The License Server is Web-based.

> Authentication is not required to view the Dashboard, but is required to administer a license. Installation creates a default "Admin" account and configures a password. If you forget the password, you must reinstall the license server.

Test Your Knowledge: Citrix License Server

1. Which three options does Citrix recommend for setting up secure access to a license server? (Choose three.)

   a. SSL

   b. HTTPS

   c. RDP

   d. VPN

   Answer: A, B, D

Managing XenServer Users

When you first install the XenServer host, a user account is added to the XenServer host automatically. This account is the local super user (LSU), or root, which is authenticated locally by the XenServer host computer. The LSU is used for system administration and has all rights and permissions.

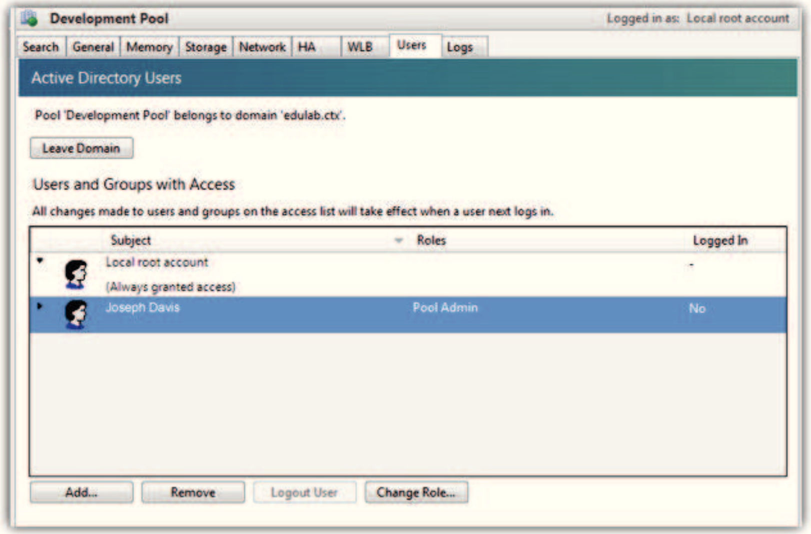

Pool 'Development Pool' belongs to domain 'edulab.ctx'.

Users and Groups with Access

All changes made to users and groups on the access list will take effect when a user next logs in.

Subject	Roles	Logged In
Local root account (Always granted access)		-
Joseph Davis	Pool Admin	No

All editions of XenServer can add user accounts from Active Directory. However, only XenServer Enterprise and Platinum Editions let you assign these Active Directory accounts different levels of permissions through the Role-based Access Control feature.

Key Benefits of Active Directory Integration

Key benefits of authenticating end users through Active Directory include:

- Easy access control to XenServer hosts
- Basic auditing control and enabled access revocation
- Access with the xe command-line interface using the appropriate -u and -pw arguments

Configuring Active Directory Integration

To grant an end user with access to the XenServer host, you must add a subject for that end user or group. To manage end user permissions in Active Directory, you can create a single group from which to add and remove end users. Alternatively, you can add and remove individual users from the XenServer host or a combination of users and groups as appropriate for authentication requirements. The subject list can be managed from XenCenter or through the command-line interface.

Credentials are first checked against the local root account when authenticating an end user, allowing you to recover a system in which an Active Directory server has failed. If the credentials do not match, then an authentication request is made to the Active Directory server. If the second request is successful, then the XenServer host retrieves the end user information and validates it against the local subject list. Validation against the subject list is successful if the end user or a group in the transitive group membership is in the subject list.

End-User Authentication Using Active Directory

If you are familiar with XenCenter, note that the XenServer host command-line interface uses slightly different terminology to refer to Active Directory and user account features.

XenCenter Term	XenCenter Command-line Interface Term
Users	Subjects
Add users	Add subjects

Active Directory authentication for a XenServer host requires that the same DNS servers are used for both the Active Directory server (configured to allow for interoperability) and the XenServer host. In some configurations, the Active Directory server might provide the DNS itself. This can be achieved either by using DHCP to provide the IP address and a list of DNS servers to the XenServer host, by setting values in the physical network interface (PIF) objects, or by using the installer if a manual static configuration is used.

Citrix recommends enabling DHCP to broadcast XenServer host names. In particular, you should not assign the reserved host names localhost or linux to XenServer hosts.

Active Directory Integration

Although the external authentication property is individual to each XenServer host, Citrix recommends enabling or disabling authentication by resource pool rather than by host. The XenServer host deals with failures that occur when enabling authentication on a particular host and performs any rollback changes that might be required. This ensures that a consistent configuration is used across a resource pool. The `xe host-param-list` command can be used to check host properties and to determine the external authentication status by checking related field values.

External Authentication Process

The external authentication process includes the following steps:

1. The XenServer host passes credentials to the remote authentication directory service for authentication.
2. The remote authentication service checks the credentials. If they are invalid, then the authentication immediately fails.
3. The external authentication directory service is queried to obtain the subject identifier associated with the credentials, if they are valid.
4. The authentication is successfully completed if the subject identifier matches the credentials stored in the XenServer host persistent metadata.

If the credentials of the end user are invalidated while the user is connected, the XenServer host invalidates the logon to XenCenter but does not invalidate any active SSH session.

Test Your Knowledge: Integrating XenServer and Active Directory

1. Which four steps must you take to configure varying levels of access for users? (Choose four.)
    a. Enable Active Directory
    b. Install XenCenter
    c. Create a subject entry for the person or group requiring access
    d. Add user accounts
    e. Assign roles to added user accounts

    Answer: A, C, D, E

Role-Based Access Control

The Role-based Access Control (RBAC) of the XenServer host allows you to assign users, roles, and permissions to control who has access to your XenServer host and which actions they can perform. The XenServer host RBAC system maps a user (or a group of users) to defined roles, which in turn have associated XenServer host permissions. RBAC depends on Active Directory for authentication services. Specifically, the XenServer host keeps a list of authorized users based on Active Directory user and group accounts.

Roles

XenServer features pre-established roles for role-based administration. Note that Pool Administrator, Local super user, and Read-only roles are only available in the Free and Advanced editions.

Pool Administrator

The resource pool administrator is equivalent to local root. The user can perform all operations.

Local super user (root)

The local super user will always have the Pool Admin role. The Pool Admin role has the same permissions as the local root.

Pool Operator

The resource operator can do everything except for adding/removing users and modifying their roles. This role is focused mainly on host and resource management.

Virtual Machine Power Administrator

The Virtual Machine Power Administrator creates and manages virtual machines. This role provides the ability to provision virtual machines for use by a virtual machine operator.

Virtual Machine Administrator

The Virtual Machine Administrator is similar to a Virtual Machine Power Administrator but cannot migrate virtual machines or perform snapshots.

Virtual Machine Operator

The Virtual Machine Operator is similar to Virtual Machine Administrator; however, while the Virtual Machine Operator can perform start/stop life-cycle operations, it cannot create or destroy virtual machines.

Read-only	The Read-only role can view resource pool and performance data.

Security Logs

The RBAC audit log will record any operation taken by a logged-in user.

- The log entry will explicitly record the Subject ID and user name associated with the session that invoked the operation.
- If an operation is invoked for which the subject does not have authorization, this will be logged.
- If the operation succeeds, it is recorded; if the operation fails, the error code is logged.

Test Your Knowledge: Role-Based Access Control

1. Which role is equivalent to the resource pool administrator with regard to permissions?
    a. Active Directory admin
    b. SR master
    c. (local) root
    d. an exempt user account

    Answer: c
2. The Read-only role can view performance data.
    a. True
    b. False

    Answer: a
3. Which three options does an audit log track during a session? (Choose three.)
    a. error codes
    b. Group ID of session that invoked the operation
    c. user name
    d. logon information

    Answer: a, b, c

Module 3

XenServer Networking

Overview

One of the goals of networking in XenServer is to make the physical network interface cards (NICs) in XenServer hosts available for networking in virtual machines. The relationship between physical NICs and virtual NICs adds another layer of complexity to networks. An initial network is setup when XenServer is installed. You can add additional networks after the installation. This module provides information on how to create and configure all networks in a XenServer environment.

Objectives

After completing this module, you will be able to:

- Describe the network components and architecture of XenServer.
- Connect virtual machines by creating a virtual network and assigning a virtual local area network (VLAN).
- Configure two physical NICs to function as one logical NIC by creating a NIC bond.
- Configure a XenServer host to use a specific network by configuring a management interface.

Timings:

Module: 90 minutes

Exercises: 10 minutes

Total Time: 100 minutes

For more information about XenServer networking, view the following Knowledge Center articles:

- CTX128502
- CTX130924

For students how are familiar with a previous version of XenServer, be sure to inform the student that this module covers the Open vSwitch, and not the Linux networking stack.

XenServer Networking Overview

XenServer hosts provide the network communication between virtual machines running on one or more XenServer hosts, as well as communication with the physical network.

You can configure four different virtual networks in XenServer:

- Single-Server Private networks, which are a type of internal network
- External networks
- Bonded networks, which are a type of external network
- Cross-Server Private networks, which are a type of internal network

Cross-Server private networks are accomplished by having a "switching host" establish GRE tunnels (in a star topology) to each of the other hosts (which have an active virtual machine running on the private network) in the pool.

Cross-Server private networks will be discussed in the Distributed Virtual Switching module.

Network Stacks Supported by XenServer

There are two networking stacks in XenServer 6.0:

- The Open vSwitch, the default networking stack
- The Linux bridge

The Open vSwitch is a software switch running on XenServer. The Open vSwitch supports Open Flow, a network protocol used to manage and direct traffic among routers and switches.

The Open vSwitch can be used alone, or it can be used in conjunction with a separately installed Distributed vSwitch Controller. When the Open vSwitch and Distributed Virtual Switch are used together you gain additional functionality and features, such as Remote Switched Port Analyzer (RSPAN) and quality of service (QoS). The table provides a list of supported features for the Open vSwitch network when used alone and the Open vSwitch and the Distributed vSwitch Controller used together.

	Open vSwitch Only	Open vSwitch and Distributed vSwitch Controller
Active/Active	X	X
Active/Passive	X	X
Private Network	X	X
QoS		X
Jumbo frames	X	X
NetFlow		X

The change from the Linux stack to the Open vSwitch is seamless to the administrator.

Previous versions of XenServer used standard Linux bridging code for building virtual switches. If required, you can revert back to the Linux stack after XenServer installation by running the following command :

```
xe-switch-network-backend bridge
```

You must restart your server after running this command. The Linux network stack is not open flow enabled, does not support Cross Server Private Networks, and cannot be managed by the XenServer vSwitch Controller.

	Open vSwitch Only	Open vSwitch and Distributed vSwitch Controller
Cross-Server Private Network		X
Access Control Lists		X

The default Open vSwitch is the networking stack that will be discussed throughout the rest of this module.

Network Components and Architecture

One network is created for each physical network interface card during XenServer installation. The XenServer host performs all the required configurations of the physical network interface cards. When you add a XenServer host to a resource pool, the default networks are merged so that all physical NICs within the same device name are attached to the same network.

The Linux bridge and the Open vSwitch have the same network architecture in relation to XenServer. The change from the Linux bridge to the Open vSwitch is transparent to the administrator.

Network Example

All hosts in a pool with an eth0 NIC will have a corresponding physical interface plugged into the pool-wide Network 0 network. The hosts with eth1 network interface cards will have a corresponding physical interface plugged into the pool-wide Network 1 network.

Network Architecture Diagram

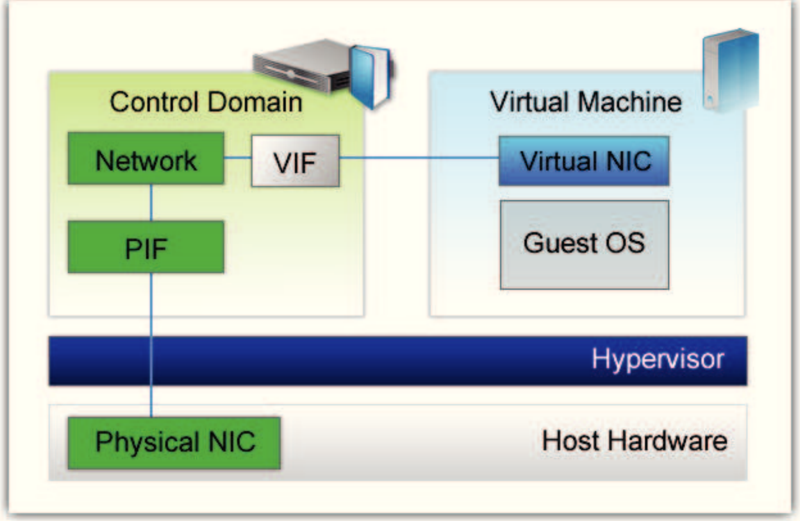

The components that are related to networking on XenServer are:

Physical interface (PIF) A PIF represents a physical network interface card for each XenServer host. The XenServer supports up to 16 physical network interfaces (or up to 8 bonded network interfaces) per XenServer host.

Virtual interface (VIF)

A VIF is a server-side software object that is a virtual representation of a computer network interface. A virtual machine connects to a virtual interface to provide network connectivity to other virtual machines and the physical network.

Virtual NIC

A virtual NIC is the virtual representation of a NIC on a virtual machine. The virtual NIC uses paravirtualized drivers to connect to the VIFs in the control domain.

Network

The control domain contains one or more virtual switches. A virtual switch is a software switch able to bridge multiple virtual network interfaces to a physical interface.

The control domain uses standard Linux device drivers to connect to the physical NICs in the host, which allows XenServer to support a broad range of physical devices.

Network Adapter Drivers

In a Linux guest operating system, virtual NICs are always displayed as standard Linux network devices and use the high speed Xen paravirtualized (PV) network driver.

In a Windows guest operating system, the initial Windows installation has an emulated network device that uses a built-in driver. Windows sees the device as a RealTek Fast Ethernet NIC. After XenServer Tools--which includes PV network drivers--is installed, Windows sees a XenServer PV Ethernet Adapter. When the high-speed drivers are installed, any network settings set during or after the Windows installation for the RealTek adapter are copied over to the XenServer PV Ethernet Adapter.

The XenServer PV Ethernet Adapter reports a speed of 2 GB per second in Windows virtual machines. This speed is a hardcoded value and is not relevant in a virtual environment because the virtual NIC will perform at the same rate as the physical NIC. The most important factor in the network speed of a virtual machine is the speed of the physical NIC on the host.

Private-Server and External Networks

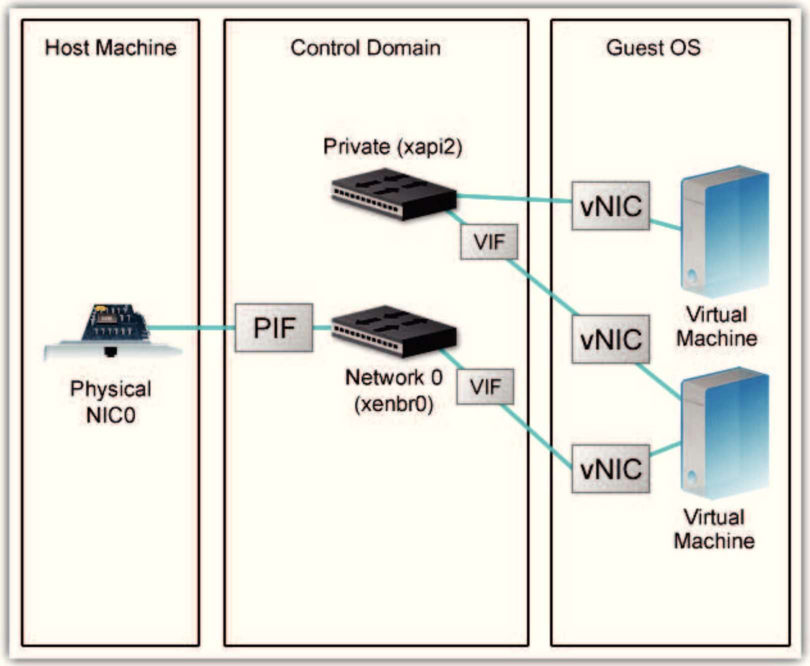

Networks without an association to a PIF are considered single-server private networks and can only be used to provide connectivity between virtual machines on a given XenServer host. With a single-server private network, no connection is made to a physical network interface card.

Networks with a PIF association are considered external and provide connectivity between virtual machines and the physical network.

The following list provides an overview of the elements of single-server private networks and external networks:

Virtual interface (VIF) The VIF transfers data between the virtual machine and the network.

Internal switch External networks form a bridge with the physical network and support VLAN trunking if multiple VLANs are associated with a single network interface card.

Physical interface (PIF) The PIF connects the physical network to Network0.

Physical NIC 0 (external networks)	The NIC on the physical XenServer host.

Test Your Knowledge: Network Components

Match the following terms with the correct descriptions in the courseware.

- PIF
- Virtual NIC
- VIF
- Network

Description	Term
The __________ bridges multiple virtual interfaces to a physical interface.	network
The __________ connects the physical network to the internal network.	PIF
A __________ uses paravirtualized drivers to connect to the virtual interfaces in the control domain.	virtual NIC
The __________ transfers data between the virtual machine and the network.	VIF

VLAN Support and Components

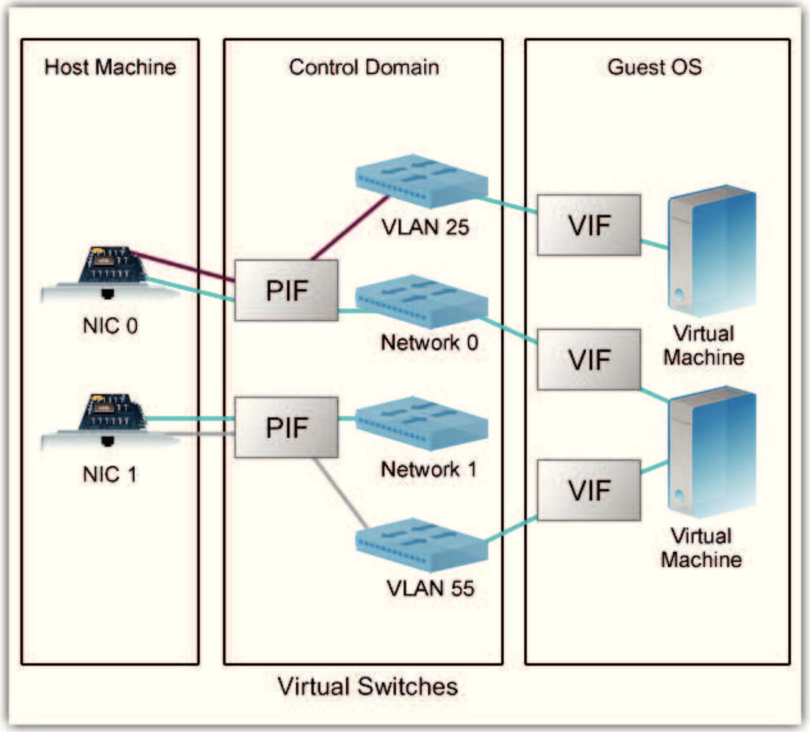

XenServer supports the use of multiple VLANs to mapped physical network interfaces on the host server. When creating a network, you must properly configure the VLAN tag for the virtual NIC in XenServer to correspond with the VLANs on the virtual switches.

VLANs allow a single physical network to support multiple logical networks. XenServer supports the use of multiple VLANs to mapped physical network interfaces on the host server.

Network

A new network is configured for each VLAN, and VLAN tags are added to packets and stripped off at the network. It is not necessary to configure a virtual machine for the VLAN. The virtual machine needs only to be connected to the switch for the VLAN.

VLAN tagging

A VLAN ID is a tag added on every packet. Incoming VLAN traffic tags are stripped off at network and added on outgoing packets.

VLAN Example

You have configured four different VLANs labeled VLAN 100, 200, 300, and 400 that are available through a trunked connection. You must verify that the naming of the VLAN tags is consistent. To ensure proper communication configure all VLANs at the resource pool level.

You can configure up to 4092 VLANs. Demo VLANs for the class by adding a network with a VLAN.

Initial Network Setup

You can designate one physical NIC and IP address for the management NIC during the installation of XenServer. Management NICs are used for management traffic between the XenServer host and XenCenter, as well as networks for storage and other distributed functions such as XenMotion. During installation:

- PIFs are created for each physical NIC on the host.
- The PIF of the physical NIC selected for use as the management interface is configured with the IP addressing options specified during installation.
- A network is created for each PIF (network 0, network 1).
- The IP addressing options of all other PIFs are left unconfigured.

After installation, you can configure non-management physical NICs, which are recommended for virtual machine network connections.

Test Your Knowledge: VLANs

1. You have three VLANs labeled 25, 26, and 27. To ensure that there is proper communication on your physical network, you must:
    a. Configure VLAN 25 on the physical switch as VLAN tag 25.
    b. Configure all VLANs on the physical switch with the proper VLAN tags.
    c. Configure VLAN 25 on the virtual switch as VLAN tag 25.
    d. Configure all VLANs on the virtual switch with the proper VLAN tags.

   Answer: b
2. You are installing XenServer on a physical server that has two NICs. Which two options will you need to configure when completing the XenServer installation? (Choose two.)
    a. A management interface
    b. One virtual switch
    c. Two PIFs
    d. The IP addresses for both NICs

   Answer: a, c

NIC Bond

A NIC bond can improve the XenServer host resiliency by using two physical NICs as if they were one physical NIC. If one physical NIC within the bond fails, the network traffic on the XenServer host will automatically be routed over the second physical NIC. NIC bonds can work in:

- Active/Active mode, with traffic balanced between the bonded NICs
- Active/Passive mode, in which traffic is only passed over one of the active NICs

NIC Bonding Architecture

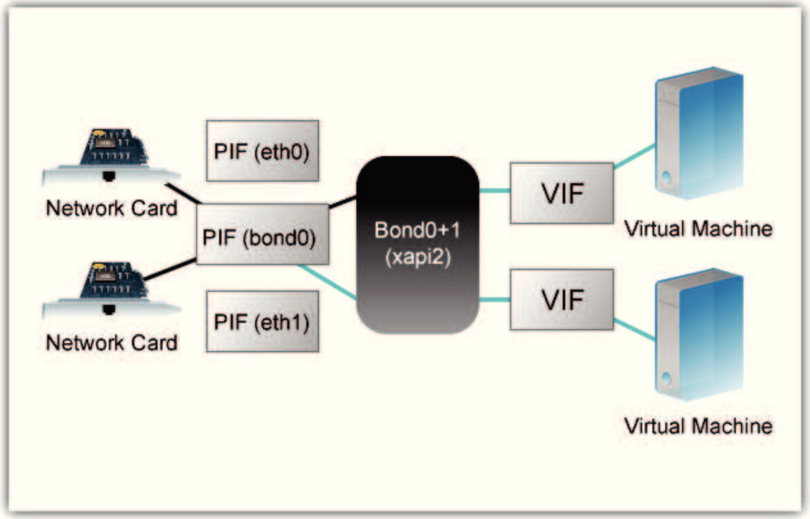

Load balancing is at source MAC granularity.

NIC bonds are represented by additional physical interface objects, such as Bond0. The bonded physical interface is connected to vSwitch.

When the bond is used for non-guest traffic, one IP configuration is required for each bond.

Load Balancing

XenServer supports Source Level Balancing (SLB) and Active-Passive NIC bonding.

Test Your Knowledge: NIC Bond Modes

1. A NIC bond configured in active-active mode:
   a. Supports network traffic over only one NIC at a time

b. Requires switch support for EtherChannel

c. Sends traffic based on the source MAC address

d. Supports network traffic over both NICs

Answer: C, D

NIC Bond Configuration for Resource Pools

Citrix recommends creating NIC bonds as part of the initial resource pool creation prior to joining additional hosts to the pool. Doing so allows the bond configuration to be automatically replicated to hosts as they are joined to the pool and reduces the number of steps required. Adding an NIC bond to an existing pool requires creating the bond configuration manually on the master and each of the members of the pool. Creating an NIC bond using physical NICs that are in use is a disruptive operation.

Do not attempt to create NIC bonds that are currently in use for high availability.

Module 4

XenServer Storage Repositories

Overview

Virtual machines frequently require large amounts of storage, and they typically have to share that storage with other virtual machines. Within XenServer, the virtual machines behave like physical machines with locally attached disks. In reality, the XenServer host has allocated a section of physical disk space and has made this space available as a disk resource to the virtual machine.

Objectives

After completing this module, you will be able to:

- Determine the features of different XenServer storage technologies.
- Describe the storage options for a XenServer storage repository.
- Configure and manage a local storage location for the storage of virtual disk images.
- Configure and manage a shared storage location for the storage of virtual disk images.
- Create a new storage repository using Advanced StorageLink technology.

Timings:

Module: 105 minutes

Exercises: 25 minutes

Total Time: 130 minutes

View the following Citrix TV videos on XenServer storage:

- http://www.citrix.com/tv/#videos/3673
- http://www.citrix.com/tv/#videos/106

Storage Technologies

In a XenServer environment, the virtual machines behave like physical machines with locally attached disks. XenServer supports several physical storage types:

Local physical disk
Logical volume management (LVM) on a physical drive directly attached to a machine

NFS
Virtual hard drive on the Network File System (NFS) that allows access to the file system over the network

iSCSI SAN
LVM over Internet Small Computer System Interface (iSCSI) SAN

Fibre Channel SAN
LVM over Fibre Channel using SAN

XenServer hosts support Fibre Channel SANs and uses the Emulex or QLogic host bus adapter (HBA).

Virtual Disk Image Formats

There are three methods for mapping physical storage to a virtual disk image:

File-based VHD on a file system
Virtual machine images are stored as thin-provisioned VHD format files on either a local non-shared File system (EXT) or a shared NFS target.

Logical Volume-based VHD on a Logical Unit Number (LUN)
The default XenServer block device-based storage inserts a Logical Volume Manager (LVM) on a disk: either a locally attached device or a SAN attached LUN over either Fibre Channel, iSCSI, or SAS.

LUN per Virtual Disk Image
LUNs are directly mapped to virtual machines as virtual disk images by storage repository types that provide an array-specific plugin (NetApp, EqualLogic, or StorageLink). The array storage abstraction matches the virtual disk image storage abstractions for environments that manage storage provisioning at an array level.

Storage Technology Comparison

The following tables compare the features of different storage types that are supported by XenServer.

	Storage Repository Type	Shared	Thin Provisioning
Local Physical Disk	LVM		
Local File System	EXT		X
HW iSCSI	LVMoHBA	X	
LVM over FC LUN	LVMoFC	X	
Software iSCSI	LVMoISCSI	X	
NFS Based	NFS	X	X
StorageLink (NetApp, Dell EqualLogic)	NETAPP,EQUAL	X	X

	Disk Resize	NIC Bonding	Multipathing	Fast Clone
Local Physical Disk	X			
Local File System				X
HW iSCSI	X		X	
LVM over FC LUN	X	X	X	
Software iSCSI	X	X	X	
NFS Based		X		X
StorageLink (NetApp, Dell EqualLogic)	X	X		

Define thin provisioning during this topic. Mention that NFS path is case-sensitive.

Test Your Knowledge: Storage Technologies

Choose the correct description for each of the storage technologies. A storage technology can match more than one description, and a description can have more than one storage technology.

1. Local physical disk
    a. Fast Clone
    b. Disk Resize
    c. NIC bonding
    d. Multipathing
    e. Thin Provisioning

   Answer: b
2. Local file system (Choose two.)
    a. Fast Clone
    b. Disk Resize
    c. NIC bonding
    d. Multipathing
    e. Thin Provisioning

   Answer: a and e
3. NFS-based (Choose three.)
    a. Fast Clone
    b. Disk Resize
    c. NIC bonding
    d. Multipathing
    e. Thin Provisioning

   Answer: a, c, and e
4. Software iSCSI (Choose two.)
    a. Fast Clone
    b. Disk Resize
    c. NIC bonding
    d. Multipathing
    e. Thin Provisioning

   Answer: b, c and d

XenServer Storage Architecture

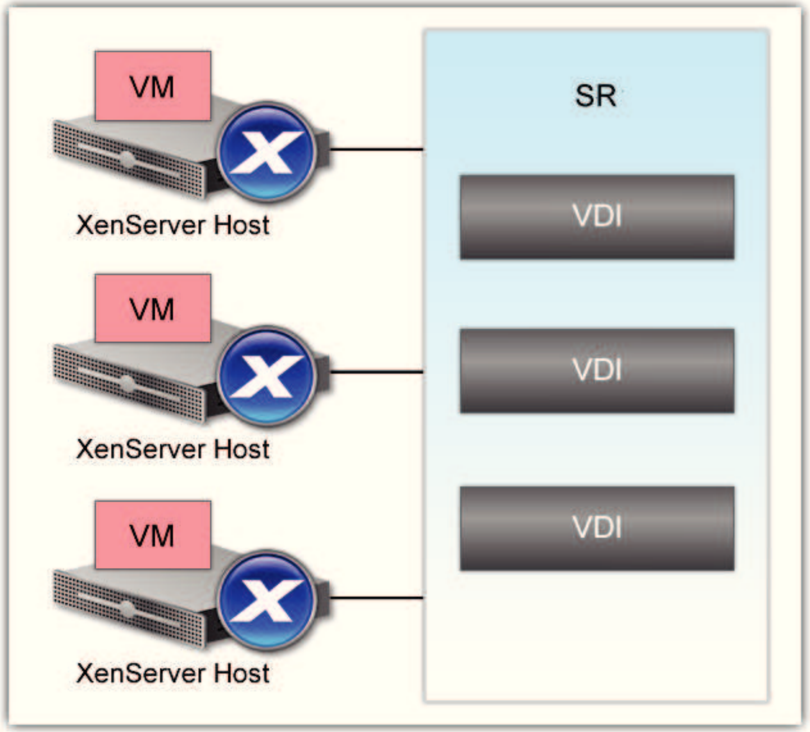

XenServer defines a container called a storage repository—a persistent, on-disk data structure—to describe a particular storage target in which virtual disk images are stored. A virtual disk image is a disk abstraction that contains the contents of a virtual disk.

The interface to storage hardware allows virtual disk images to be supported on a large number of storage repository types. The XenServer storage repository is very flexible, with built-in support for IDE, SATA, and locally connected SCSI and SAS drives, and iSCSI, NFS, and remotely connected SAS and Fibre Channel. The storage repository and virtual disk image abstractions allow advanced storage feature—such as sparse provisioning, virtual disk image snapshots, and fast cloning—to be exposed on storage targets that support them. For storage subsystems that do not inherently support advanced operations, a software stack is provided based on the Microsoft VHD specification, which implements these features.

Each XenServer host can use multiple and different storage repositories simultaneously, which can be local or remote. If the storage repositories are remote, they can be shared between hosts or dedicated to particular hosts. Shared storage is pooled between multiple hosts within a defined resource pool. Each host must be able to access a shared storage repository on the network. All hosts in a single resource pool must have at least one shared storage repository in common.

Storage Repository Architecture

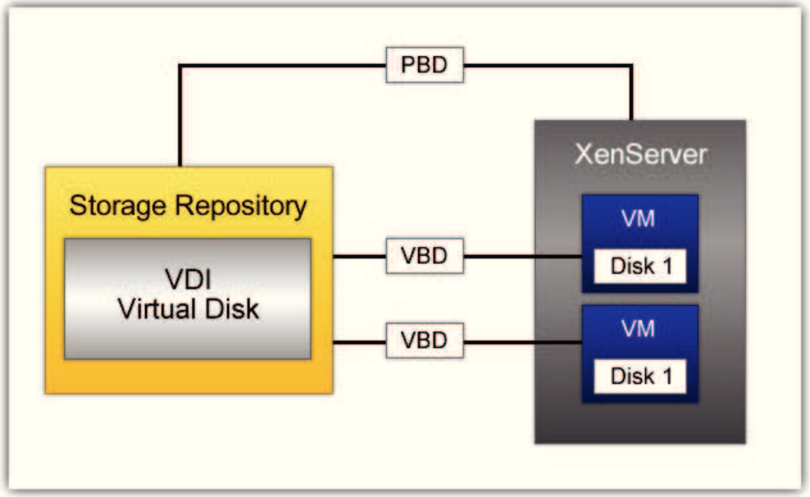

XenServer storage repository architecture consists of the following components:

- Physical block device
- Virtual disk image
- Virtual block device

Physical Block Device

A physical block device (PBD) represents the interface between a physical host and an attached storage repository. PBDs store the device configuration fields that are used to connect to and interact with a given storage target. In the case of NFS, for instance, this device configuration includes the IP address of the NFS server and the associated mount path. PBD objects manage the run-time attachment of a given storage repository to a given host.

Virtual Disk Image

A virtual disk image is an on-disk representation of a virtual disk provided to a virtual machine. The virtual disk image is the fundamental unit of virtualized storage in XenServer. The virtual disk image is a virtual disk drive. The format of the virtual disk image depends on the type of storage repository in which it is contained.

Virtual Block Device

A virtual block device (VBD) is a connector object that is similar to a PBD that allows mappings between virtual disk images and virtual machines. In addition to providing a mechanism to attach a virtual disk image to a virtual machine, VBDs allow the fine-tuning of parameters regarding QoS, statistics, and the ability of a given virtual disk image to start.

Multiple Storage Repositories

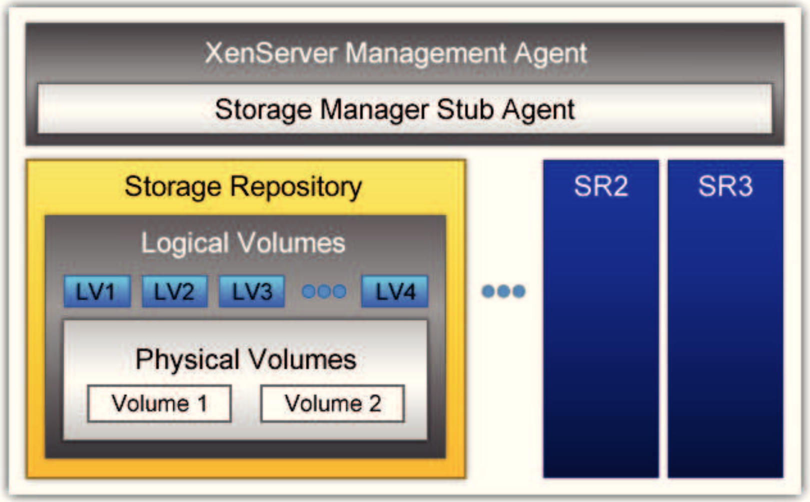

The storage manager subsystem on the server deals with all storage repository management. You can attach and use multiple storage repositories on a XenServer.

Virtual disk images can be stored in different storage repositories, and a single virtual machine can have virtual disk images in different storage repositories. Moving virtual disk images from a local storage to remote storage enables the use of XenMotion. Virtual machines should be moved from local to remote storage repositories by copying or moving the virtual machine.

Local Storage Repositories

XenServer supports:

- LVM, which represents the disks within a locally attached Volume Group.
- EXT3, which represents disks as VHD files stored on a local path.

By default, XenServer uses the local disk on the physical host on which it is installed. The Linux LVM is used to manage virtual machine storage. A virtual disk image is stored in VHD format in an LVM-managed logical volume.

A local disk EXT storage repository must be configured using the XenServer command-line interface.

> Storage repositories created on the local storage are not available to share with other XenServer hosts.

Virtual Disk Size

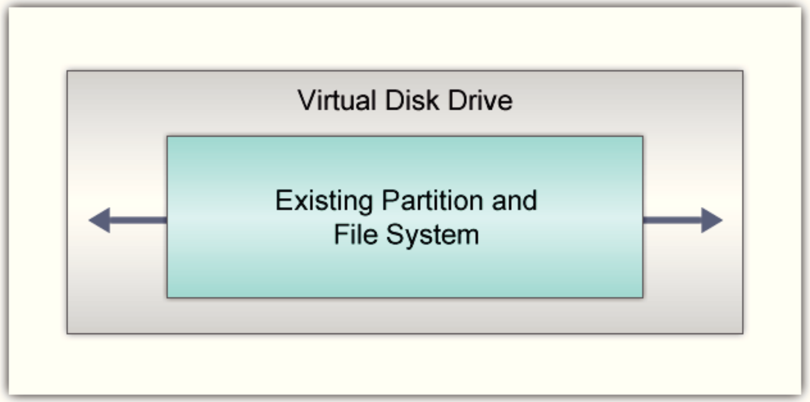

Storage changes can be made after the XenServer installation. When the virtual disk size is increased in XenCenter, the guest operating system views the file system and partitioning as they were originally created and views the additional space added to the virtual disk as free space, if specific tools in the guest operating system can be used to expand the file system to use the free space.

A Windows virtual disk can be changed using the `diskpart` command-line interface tool or the Windows Disk Manager utility. Linux uses the `lvextend` command-line interface tool to change virtual disk size.

The virtual machine to which a virtual disk is assigned must be shut down before modifying the virtual disk size.

Test Your Knowledge: Storage Repositories

Match the following methods with their correct descriptions. A method can match more than one description.

- Virtual block device
- Virtual disk image
- Storage repository
- Physical block device

Method	Description
Is a collection of volumes and disk devices.	Storage repository
Is a connector object that allows mapping between virtual disk images and virtual machines.	Virtual block device
Represents the interface between a physical host and an attached storage repository.	Physical block device
Is an on-disk representation of a virtual disk provided to a virtual machine.	Virtual disk image

Test Your Knowledge: Local Storage

Indicate whether each statement is true or false.

Statement	True or False
NFS-based storage allows for thin provisioning.	True
Both LVM and EXT3 are supported local storage types.	True
Virtual disk drive size can be decreased.	False

Statement	True or False
Additional local storage can be configured after installation by using the xe command-line interface.	True

Storage Capability Comparison

All storage types can store virtual machine images; however, each storage type has different capabilities. Shared remote storage offers more features than local storage options.

	Local Disk	Fibre Channel	iSCSI Hardware	iSCSI Software	NFS-based
Store Virtual Machine	x	x	x	x	x
Automatic Virtual Machine Placement		x	x	x	x
XenMotion Virtual Machines		x	x	x	x
Resize Disks	x	x	x	x	
Fast Clone	x				x
Thin Provision					x

Citrix XenServer natively supports thin provisioning, dynamic virtual disk image resizing, native snapshot functionality and fast cloning for NetApp and Dell EqualLogic storage solutions. For more information about configuring XenServer for use with NetApp and EqualLogic storage, see Citrix articles CTX122737 and CTX118841 on *support.citrix.com*.

As shown in the table in the courseware, each storage type has different capabilities.

All storage types can store virtual machine images.

Shared remote storage offers more features than local storage options.

NFS Storage Overview

NFS is a common storage infrastructure that is available in many deployments. XenServer allows existing NFS servers that support NFS V3 over TCP/IP to be used as a storage repository for virtual disks. Virtual disks are stored in the Microsoft VHD format.

Any XenServer can access NFS storage because XenServer includes a tuned NFS client.

Consider the following points when configuring NFS-based remote storage:

- A high performance hardware NFS solution is recommended for large deployments.
- A software-based NFS server can be used for smaller deployments.
- A NFS storage can be set up after the product installation in XenCenter or in the command-line interface.

Any mapping function to an NFS share is case-sensitive.

NFS Architecture

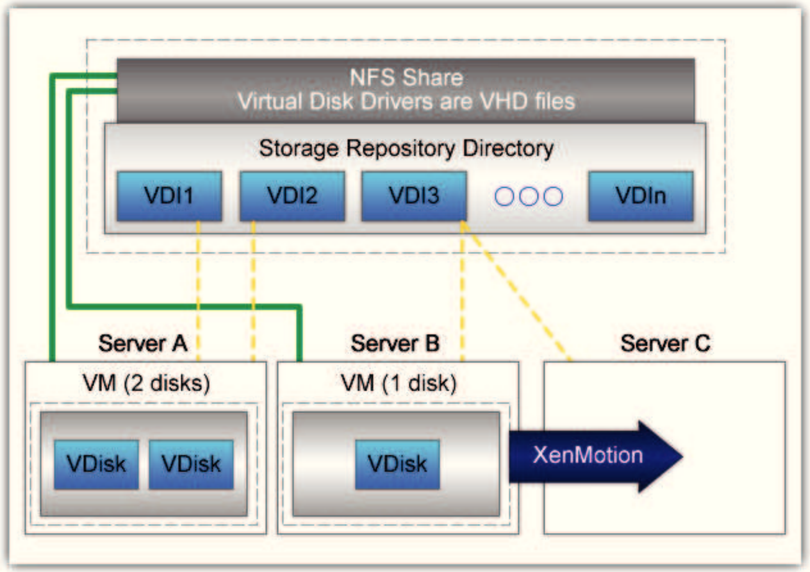

Each virtual disk of each virtual machine is represented by a file in the NFS repository directory. All servers connect to the same NFS Share; virtual disk drives are virtual hard drive files on the NFS share. Virtual hard drive files for each virtual machine are stored using a unique ID. Only the server running a virtual machine connects to the individual virtual disk for that virtual machine.

NFS Virtual Disks

Two advantages to using the Microsoft VHD format on an NFS storage include:

- Dynamic growth of the virtual disks can dynamically grow
- Fast cloning, the default behavior in XenCenter and which sets:
  - The user templates to read-only base images for clones
  - The virtual disk to track changes from the base image

NFS Remote Storage Configuration and Recommendations

Citrix recommends that you:

- Use NFS, the preferred storage for XenServer deployments.
- Set up a dedicated network for NFS traffic using dedicated NICs on servers and dedicated switches and cabling.
- Use dedicated NAS-based NFS servers, which are better optimized with high-speed cache and faster than most Linux-based solutions.

Test Your Knowledge: NFS Storage

1. Which two statements about NFS storage are correct? (Choose two.)
    a. XenServer supports NFS V3.
    b. NFS storage is recommended for all XenServer deployments.
    c. NFS storage is similar to FC/iSCSI storage architecture.
    d. NFS storage is configured during the product installation.

    Answer: A, B

iSCSI Storage Overview

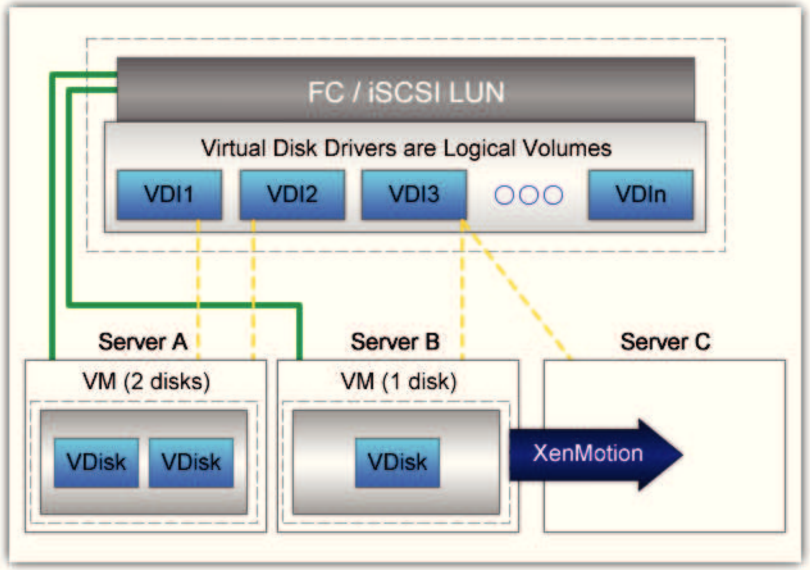

XenServer provides support for shared storage repositories on iSCSI LUNs. iSCSI is supported using the open-iSCSI software iSCSI initiator or using a supported iSCSI HBA.

iSCSI-based storage repositories enable virtual machine agility using XenMotion. Virtual machines can be started on any XenServer host in a resource pool and migrated between hosts.

Both Fibre Channel and iSCSI storage share the following features:

- The servers connect to a LUN.
- The virtual disk drives are mapped to logical volumes on the LUN.
- Only the server running the virtual machine can connect to the virtual disk for that virtual machine.

iSCSI Architecture

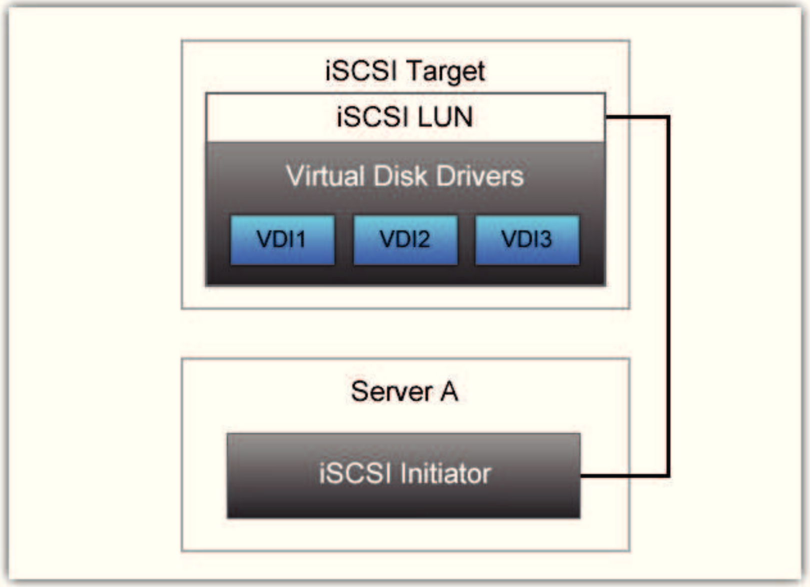

The iSCSI Storage architecture uses several unique components.

iSCSI Initiator A system that connects to iSCSI-based storage - Each system gets a unique initiator name automatically and the unique initiator name can be changed after installation. Example: Iqn.2007-07.example.com:97fdaca0

iSCSI Target An IP address on an iSCSI-based storage array - A unique name is provided during array installation and setup. Example: iqn.myiscsitarget.iscsi0

iSCSI LUN A logical disk drive on an iSCSI array - A unique name is provided during LUN creation. Example: 0

An IP address or an iSCSI Qualified Name (IQN) are needed to connect to an iSCSI storage repository. With the IP, you can perform a discovery to find the name.

iSCSI Adapters

iSCSI HBA offloads iSCSI processing from the server, which allows for higher performance and lower burden on server CPU resources.

Shared iSCSI support is implemented based on the LVM and provides the same performance benefits provided by LVM in the local disk case.

iSCSI Setup Configuration and Recommendations

An iSCSI target:

- Should have a dedicated network for iSCSI traffic.
- Uses dedicated NICs on servers.
- Uses dedicated switches and cabling.

iSCSI arrays:

- Allow all initiator names to connect during setup.
- Ensure that the target allows multiple servers (initiators) to access the same LUN (often an option when setting up a LUN).

Test Your Knowledge: iSCSI Storage

1. Which two options do you need in order to configure software iSCSI storage in a XenServer environment? (Choose two.)
    a. IP address
    b. IQN
    c. Share name
    d. Number of LUNs

    Answers: A, B

Resizing a Storage Repository

If more storage is needed for virtual disks, the LUN on which an iSCSI or HBA storage repository is based can be resized with the array tools of the storage repository to increase the amount of storage.

In previous versions of XenServer, explicit commands were required to resize the physical volume group of iSCSI and HBA storage repositories. These commands are no longer required.

iSCSI and HBA storage types have a different process for resizing a storage repository.

Resizing a Storage Repository - iSCSI

1. Shut down all virtual machines on the SR.
2. Use the `xe sr-list` command to find and record the Universally Unique Identifier (UUID) of the storage repository.
3. Use the `# xe sr-param-list uuid=<SR UUID>|grep PBD` command to find and record the UUID for the PBD. Replace <SR UUID> the storage repository UUID that was recorded in the previous step.
4. Use the `# xe pbd-unplug uuid=<PBD UUID>` command to unplug the PBD. The storage size is increased.
5. Use the `# xe pbd-plug uuid=<PBD UUID>` command to plug the PBD.
6. Turn on the virtual machines.

Resizing a Storage Repository - HBA

1. Turn off the virtual machines. You can also use XenMotion to move the virtual machines to another host in the pool.
2. Restart the XenServer host.
3. Turn on the virtual machines.
4. Repeat this procedure for each host in the resource pool.

Fibre Channel SANs Overview

XenServer supports Fibre Channel SANs using the HBA in which LUNs are mapped to XenServer disk devices.

A startup from SAN:

- Works with HBAs that support startup from LUN.
- Requires configuration of the HBA firmware to start from LUN setup before launching the XenServer installation.

XenServer support for Fibre Channel has the following restrictions:

- Dual Multipath (DM) is not currently supported.
- Direct mapping of a LUN to a virtual machine is not supported. A LUN must be added to an LVM volume group before one or more volumes can be assigned to a virtual machine.

Make sure your hardware is listed in the hardware compatibility list (HCL). For more information about the HCL, visit the *hcl.vmd.citrix.com* Web site.

Fibre Channel Architecture

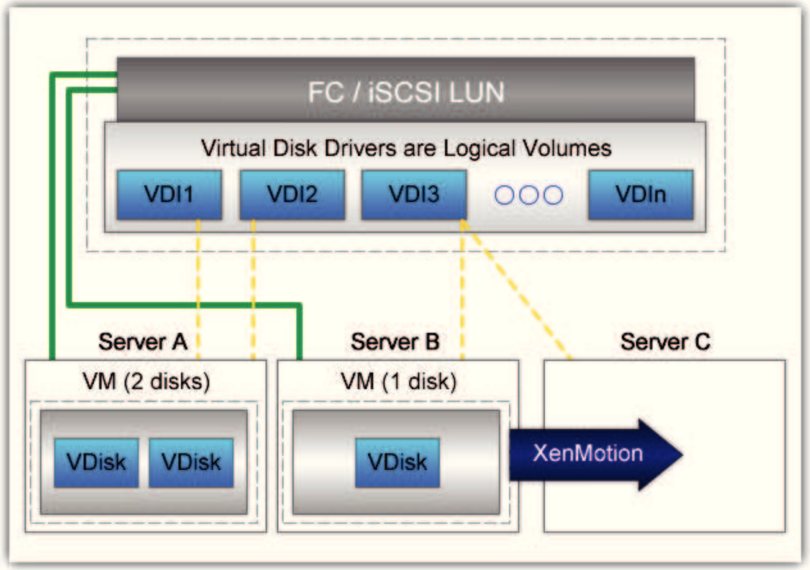

With Fibre Channel storage:

- All servers connect to the same LUN.

© Copyright 2012 Citrix Systems, Inc.

- Virtual disk drives are individual, logical volumes on the LUN similar to local disk setup.
- Only the server running a virtual machine connects to the individual virtual disk for that virtual machine.

LUN Device Path

You can use one of several methods to identify the global device path to the new LUN:

- Use the storage repository driver probe feature to return a list of valid, unused by XenServer, and attached SCSI devices.
- Identify the LUN to use based on one or all of the following values:
  - Vendor
  - LUN size
  - LUN serial number
  - SCSI bus LUN ID
- Use the globally unique device path that is returned using the storage repository probe corresponding to the specific LUN entry.

Triggering the HBA to scan for new LUNs might require specific vendor operations; consult the relevant documentation for the HBA.

Test Your Knowledge: Fibre Channel HBA Management

Indicate whether each statement is true or false.

Statement	True or False
Each server connects to a different LUN, which is identified by scanning for the LUN.	False
LUNs appear as SCSI devices.	True
Direct mapping of a LUN to a virtual machine is supported by XenServer.	False
A LUN must be added to a LVM volume group before it can be made available to a virtual machine.	True

Dedicated NIC Bonds for Remote Storage

A NIC can be configured and dedicated to specific functions by assigning a specific IP address to the bond. After a management interface has an IP address, it can be dedicated to remote storage and will not be available for any other purpose. For example, to dedicate a NIC to storage traffic, the NIC, storage target, and switch must be configured so that the target is only accessible over the assigned NIC. This configuration allows for the use of standard IP routing to control how traffic is routed between multiple NICs within a XenServer.

The bonded interfaces can act in various ways including: hot standby or link integrity monitoring. As soon as the bond is assigned an IP address, it will be available immediately upon starting; it will not have to be manually plugged in.

Dedicated Remote Storage Interface Configuration

Before dedicating a network interface as a storage interface for use with iSCSI or NFS storage repositories, you must ensure that the dedicated interface uses a separate IP subnet that cannot be routed from the main management interface. If this is not enforced, then storage traffic might be directed using the main management interface after a host restart, due to the order in which network interfaces are initialized.

To ensure that there is a non-routable subnet:

1. Configure the interface (or NIC bond) to management mode.
2. Assign an IP address to the interface.
3. Dedicate the IP address to remote storage traffic.

Storage Multipathing

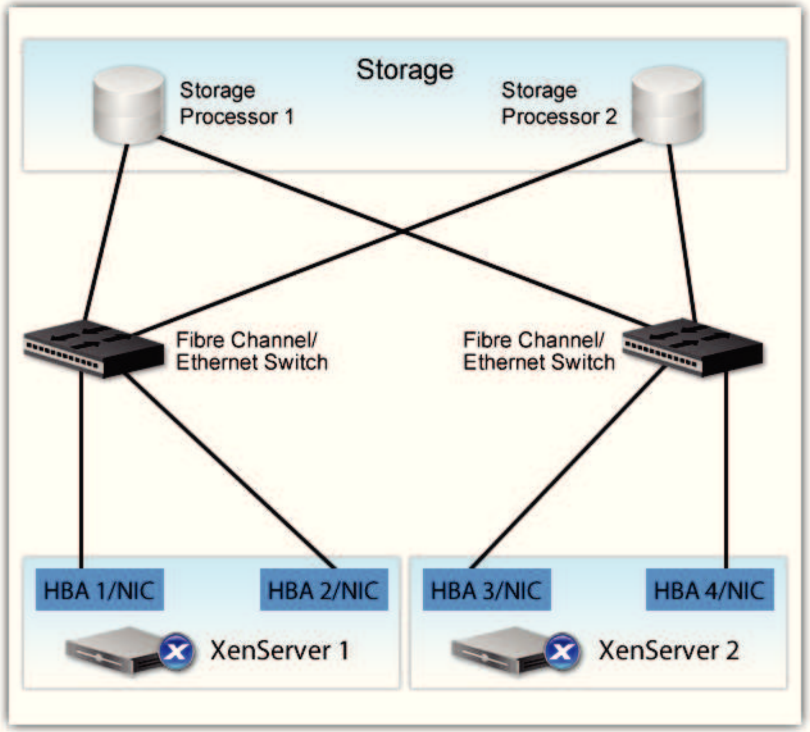

Storage multipathing for XenServer is a networking feature which:

- Allows the use of redundant paths to a storage device.
- Increases availability in case of hardware failure on one of the paths.
- Is controlled on each host, allowing varying degrees of support in a resource pool.
- Is available for all storage repository types that use raw LUNs.

Dynamic multipathing support is available for Fibre Channel and iSCSI on both hardware- and software-based storage back-ends. By default, multipathing uses the round robin load balancing method. Both routes will have active traffic during normal operation.

Multipathing does not benefit from being stacked with NIC bonding. Make sure each iSCSI channel is on a unique subnet.

Citrix StorageLink Overview

StorageLink technology lets your virtual server infrastructure take full advantage of all the resources and functionality of existing storage systems. StorageLink supports all third-party storage architectures and delivers deep integration with leading storage platforms, allowing you to switch seamlessly between Citrix XenServer and Microsoft Hyper-V platforms. StorageLink provides the following advantages:

- Simplifies virtual machine creation by automating storage resource management
- Enables hardware-assisted virtual machine cloning, as well as rapid cloning and provisioning of virtual machines
- Improves storage utilization through de-duplication and thin provisioning
- Lowers operating cost and reduces error-prone manual tasks
- Improves virtual machine storage visibility and control with reduced total cost of ownership (TCO)

Citrix StorageLink Storage Repository

The Citrix StorageLink storage repository provides direct access to native array APIs to offload intensive tasks such as LUN provisioning, snapshot, and cloning data. The StorageLink provides a number of supported adapter types to communicate with array management APIs.

By using StorageLink, all provisioning and mapping of storage to XenServer hosts is handled on demand. Data path support for LUNs includes both Fibre Channel and iSCSI, if these are supported by the hardware.

StorageLink storage repositories can co-exist with other storage repository types on the same storage array hardware, and multiple StorageLink storage repositories can be defined in the same resource pool.

The exact features available for a given StorageLink storage repository depends on the capabilities of the array. All StorageLink storage repositories use the LUN-per-VDI model in which a new LUN is provisioned for each virtual disk.

StorageLink Supported Array Types

Citrix StorageLink in XenServer 6.0 supports the following array types:

- NetApp/IBM N Series
- Dell EqualLogic PS Series

XenServer Advanced edition or above is required to use the integration with the NetApp or Dell EqualLogic storage repository types. These storage repositories can be used as standard iSCSI, Fibre Channel, or NFS storage without the benefit of direct control of the hardware features with the free XenServer.

Upgrading to StorageLink with XenServer 6.0

With the release of XenServer 6.0, a new StorageLink storage repository is easily created using the New Storage wizard in XenCenter. Upgrading an existing StorageLink repository from a previous version of XenServer to XenServer 6.0 is not completed using the wizard. The upgrade process from a previous version of XenServer depends on the storage adapter. There are three standard scenarios that can occur when upgrading:

- Upgrading to XenServer 6.0 with an adapter that is supported
- Upgrading to XenServer 6.0 with an adapter that is no longer supported
- Upgrading to XenServer 6.0 with a legacy adapter

For more information about upgrading StorageLink, see Citrix article CTX130421 on *support.citrix.com*.

Test Your Knowledge: Storage

1. You need to create a XenServer resource pool. All hosts have access to two Fibre Channel networks. Shared storage needs to be configured and protected against storage area network failures. What should you do to use the shared storage?
    a. Enable high availability
    b. Enable multipathing
    c. Bond the network connections
    d. Create a Fibre Channel network

    Answer: B

2. Which two adapters are not supported Advanced Storage Link technology adapters? (Choose two.)
    a. EMC CLARiiON (SMI-S)
    b. Dell EqualLogic
    c. NetApp/IBM N Series
    d. QLogic

    Answer: A, D

Module 5

Creating and Managing Virtual Machines

Overview

Virtual machines are created from templates. A template is an image that contains all the configuration settings necessary to create an instance of a specific virtual machine. XenServer is equipped with a base set of templates, which range from generic, raw virtual machines that can turn on an operating system vendor installation CD or run an installation from a network repository to complete, pre-configured operating system instances.

Different operating systems require slightly different settings in order to run at their best. XenServer templates are tuned to maximize operating system performance.

XenServer 6.0 introduces the concept of virtual appliances. A virtual appliance is logical group of one or more related virtual machines.

Timings:

- Module: 180 minutes
- Exercises: 75
- Total: 255 minutes

Objectives

After completing this module, you will be able to:

- Create an ISO library to use in the creation of a Windows virtual machine.
- Use the Linux Demo template to create a Linux virtual machine in XenCenter.
- Install XenServer Tools on a virtual machine to enhance performance through the paravirtualized drivers.
- Use XenCenter to perform lifecycle operations of a virtual machine.
- Determine which virtual machine template to use for a given scenario.
- Create a new virtual machine by using a template, cloning an existing virtual machine, and importing a virtual machine.
- Summarize the process that occurs when XenConvert is used for a physical-to-virtual conversion.
- Outline the steps required to increase host server utilization.

Windows Virtual Machine Architecture

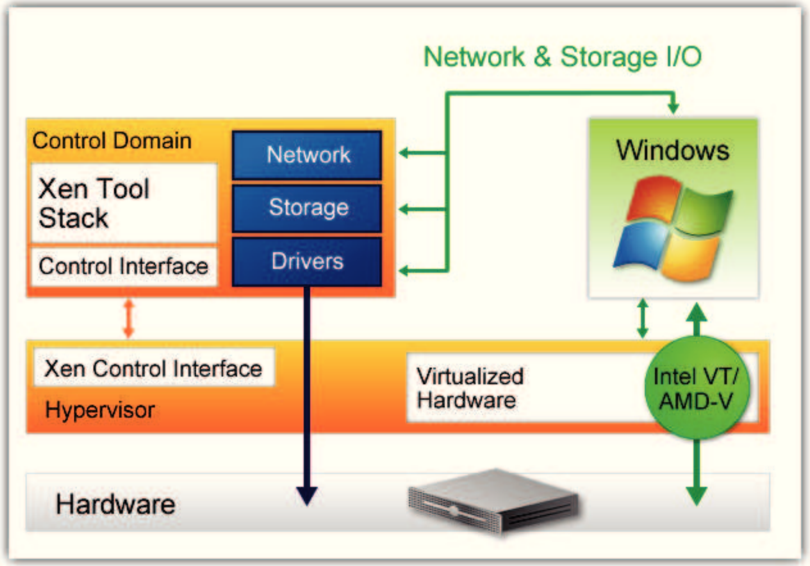

The Windows virtual machines use paravirtualized drivers to access storage and network resources through the control domain. XenServer is designed to use the virtualization capabilities of Intel VT and AMD-V processors.

Virtual Memory and Disk Size Limits for Windows Virtual Machines

When installing virtual machines, be sure to follow the memory and disk space guidelines of the operating system and any relevant applications that you want to run when allocating resources such as memory and disk space. Note that individual versions of the operating systems might also impose their own maximum limits on the amount of memory supported; for example, the amount of memory supported might be limited for licensing reasons.

Setting a memory maximum that is greater than the operating system supported limit might lead to stability problems within your virtual machine.

For more information about supported virtual memory and disk size limits for each operating system, see Citrix article CTX130420 on *support.citrix.com*.

Virtual Device Support for Windows Virtual Machines

The current version of the XenServer product family has the following general limitations on virtual devices for virtual machines. Note that specific guest operating systems might have lower limits for certain features. These limitations are noted in the individual guest installation section.

Virtual devices	Windows virtual machines
Number of virtual CPUs	8
Number of virtual disks	7 (including virtual CD-ROM)
Number of virtual CD-ROM drives	1
Number of virtual NICs	7

Installation from an ISO

An ISO image acts as a virtual CD. ISO images can be mapped to a virtual machine and used for product installations. You can use a CD-burning program to create an ISO image of the installation media. You can then copy it to an ISO storage repository that is stored on a Windows CIFS or NFS share.

Citrix recommends that you create an ISO image library on a network share that several XenServer hosts can access.

ISO Libraries

An ISO library stores ISOs in a single network share so that they can be made available to multiple XenServer hosts. All ISOs should be in the root of the share because XenServer does not look in subdirectories. XenServer includes NFS (Linux) and CIFS (Windows) clients, so you can store the ISO library on either type of file share.

Once you transfer ISOs to the share, you can use either the wizard in XenCenter or the command-line interface to map a server to the ISO library.

Create an ISO library to store virtual CDs and to create your virtual machines from those ISOs. This can save an administrator the work of from copying data to new servers in order to create virtual machines or to install applications.

To Create a Windows Virtual Machine

1. Determine the installation source:
   - A physical CD
   - An ISO file
   - Network installation (the CD/DVD image on the network server)
2. Select an appropriate template.
3. Install the Windows operating system.
4. Install XenServer Tools.

Test Your Knowledge: Creating an ISO Library for a Windows Virtual Machine

1. An ISO library is useful for __________. (Choose three.)
   a. Storing images on a network and making them accessible to multiple XenServer hosts

b. Installing without needing the physical CD or DVD
c. Backing up images on a network.
d. Uninstalling XenServer hosts

Answer: a, b, c

Linux Distributions

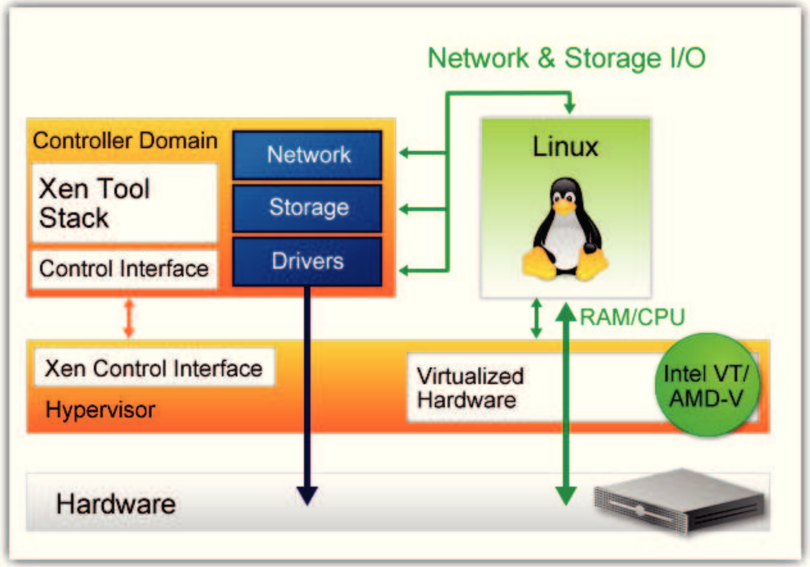

Most modern Linux distributions support Xen paravirtualization directly, but have different installation mechanisms and some kernel limitations.

Xen-enabled Distributions

- Red Hat Enterprise Linux (RHEL) 4 & 5
- SUSE Linux Enterprise Server (SLES) 10 SP1, SP2 32-bit/64-bit
- CentOS 5
- Oracle Enterprise Linux 5

Use vendor-supplied kernels and drivers when creating new virtual machines. XenServer creates a paravirtualized virtual machine from the vendor installation files.

Non-Xen enabled Distributions

- RHEL 3
- SLES 9 SP1 and SP2

Creating 32-bit Linux virtual machines on a host with more than 128 GB of memory is not supported.

Vendor-supplied kernels are only changed where the kernel touches the hardware. The changes affect less than one percent of the kernel. XenServer uses the latest distribution version to ensure up-to-date security fixes and hardware drivers. For example, a RHEL 4.4 virtual machine uses the RHEL 4.6 distribution. RHEL 4.4 and lower are non-Xen-enabled.

For more information about Linux distributions, see Citrix article CTX130422 on *support.citrix.com*.

Creating Linux Virtual Machines

The XenServer guest installer allows you to install an operating system from a network-accessible copy of vendor media onto a virtual machine. To prepare for installing from vendor media, make an expanded network repository of your vendor media, not ISO images, and export it over NFS, HTTP, or FTP so that it is accessible from the XenServer host administration interface.

Verify that you can access the network repository from the control domain of the XenServer host. The URL must be directed to the base of the CD/DVD image on the network server and be of one of the following forms:

HTTP	`http://<server>/<path>`
FTP	`ftp://<server>/<path>`
NFS	`nfs://<server>/<path>`

See the vendor installation instructions for information about how to prepare for a network-based installation such as where to unpack the ISO.

For a SUSE-based distribution, use `nfs://<server>/<path>`. For an RHEL distribution, use `nfs:<server>:/<path>`.

When creating virtual machines from templates, the XenCenter New VM wizard prompts you for the repository URL. When using the command-line interface, you can install the template as normal using `vm-install` and then point to the URL with the `other-config:installrepository` parameter. When the virtual machine starts, it initiates the network installation process.

When you install a new Linux-based virtual machine, finish the installation and restart the virtual machine before performing any other operations. When using the NFS installation method from XenCenter, use `nfs://`. XenCenter modifies this path into the correct form when passing it to the server.

Test Your Knowledge: Using a Linux Demo Template

1. Which two parameters should be used to identify the URL of the network repository containing the vendor media? (Choose two.)

   a. locate-vendormedia

   b. vm-install

   c. vm-install-repo

   d. other-config:installrepository

   e. vm-add

   Answer: b, d

XenServer Tools

The XenServer Tools installer runs in the virtual machine. You must be logged on to the virtual machine before installing XenServer Tools.

Without XenServer Tools, you cannot:

- Restart a virtual machine cleanly.
- Suspend a virtual machine.
- Migrate a running virtual machine (XenMotion).
- Use the checkpoint and roll back feature.
- Adjust dynamically the number of virtual CPUs assigned to a running Linux virtual machine; Windows virtual machines require a restart for this to take effect.

XenServer Tools for Windows Virtual Machines

New Windows virtual machines initially use emulated storage and network drivers. The emulated drivers create overhead, which dramatically impacts performance. Running a virtual machine without installing XenServer Tools is not supported.

XenServer Tools:

- Provides custom Windows drivers and a guest agent for XenServer virtual machines.
- Replaces emulated storage and network drivers with high-speed paravirtualized versions.

Updating XenServer Tools for Windows Virtual Machines

The XenServer Tools are available in XenCenter on the built-in `xs-tools.iso`.

1. On the virtual machine menu, select Install XenServer Tools; this attaches the CD image containing the XenServer Tools to the virtual machine.
2. If Autoplay is not enabled, double-click the CD drive and select xensetup.exe to begin the XenServer Tools installation. If Autoplay is enabled for the virtual machine disk, installation will be started automatically.
3. Follow the on-screen prompts to install the new drivers.

XenServer Tools for Linux Virtual Machines

Although all supported Linux distributions are natively paravirtualized and, therefore, do not need special drivers for full performance, XenServer includes a guest agent that provides additional information about the virtual machine to the host.

This additional information includes:

- Linux distribution name and version (major, minor revision)
- Kernel version (`uname`)
- IP address of each Ethernet interface
- Total and free memory within the virtual machine

Install this agent and keep it updated as you upgrade your XenServer host.

> For more information about specific Linux distributions, see Citrix article CTX130422 on *support.citrix.com.*

To Install XenServer Tools

1. Initiate the installation in XenCenter.
    1. Right-click the virtual machine and select **Install XenServer Tools.**
    2. Click **OK** on the message dialog box to go to the VM console.
2. Click the virtual machine, click the **Console** tab, and click in the Console.
3. Log on with root level access to the Linux instance in the virtual machine.
4. Mount the image onto the virtual machine:

    ```
    mount /dev/xvdd /mnt
    ```

5. Execute the installation script as the root user:

    ```
    /mnt/Linux/install.sh
    ```

6. Start the virtual machine.

Test Your Knowledge: XenServer Tools

1. Which three options are enabled by XenServer Tools when working with a virtual machine? (Choose three.)
    a. Copy
    b. Suspend
    c. Migrate
    d. Use the checkpoint and rollback feature

    Answer: b, c, d

Life Cycle Operations

You can use XenCenter to control the life cycle of any virtual machine. All virtual machines support the following operations:

Start	Turns on the virtual machine
Shutdown	Initiates an orderly shutdown of the virtual machine
Force Shutdown	Imitates pulling the power plug on a physical machine
Reboot	Restarts the running virtual machine
Force Reboot	Follows the force shutdown operation with a reboot operation

Suspend and Resume

The suspend operation freezes the virtual machine to disk. Suspending a virtual machine:

- Stores the memory image outside the virtual disk of the virtual machine.
- Frees up host memory and CPU resources.
- Changes the machine state in XenCenter.
- Reduces the virtual machine startup time.

If the virtual machine is located on shared storage, the virtual machine can be suspended on one server and restarted on another. If the virtual machine is located on non-shared storage, it can only be resumed on the same machine.

- Suspend and resume operations of Windows virtual machines are not supported until the paravirtualized drivers are installed on the virtual machine.
- See Workload Balancing for more information about suspend operations.

Deleting a Virtual Machine

Deleting a virtual machine removes the virtual machine, its configuration, and its file system from the server.

Deleting a virtual machine cannot be undone.

When you delete a virtual machine, you can choose to delete or preserve:

- Attached virtual disks
- Virtual machine snapshots

Citrix recommends that when you delete a virtual machine, you should also delete any virtual disks attached to it, because deleting the virtual machine creates orphan virtual disk drives that consume space.

Preserved snapshots can still be exported, deleted, or used to create new virtual machines and templates.

Test Your Knowledge: Life Cycle Operations

Match the following terms with the correct descriptions in the courseware.

- Shutdown
- Force Shutdown
- Shared Storage
- Uninstall

Term	Description
Shutdown	Requires paravirtualized drivers for Windows virtual machines.
Force shutdown	Is similar to unplugging a physical machine.
Shared storage	Allows you to resume a virtual machine on a different server.
Uninstall	Erases all configurations and virtual disks.

Virtual Machine Templates

A virtual machine template can contain installation metadata, including:

- The operating system information.
- The settings for optimum storage, CPU, and memory.
- The network configuration information.

XenServer includes templates that can help you to quickly deploy new virtual machines. Custom templates can be created from virtual machines created on XenServer. The three types of templates available in XenCenter are basic, full, and custom.

Basic templates

Basic templates are skeleton templates that include recommended settings for new virtual machines. These templates do not include an operating system. Use this template in conjunction with vendor-installed media to create new Windows and Linux virtual machines.

Full templates

Full templates are full copies of the operating system with files and recommended settings. Additional packages can be added to these templates using the apt-get utility after installing the operating system.

Custom templates

Custom templates can be created from virtual machines created on XenServer. Any virtual machine can be turned into a custom template. You can duplicate custom templates and make basic changes.

Importing and exporting templates allows for movement of configurations between XenServers located in different resource pools.

For Windows virtual machines, the Microsoft sysprep utility should be used before creating the template.

> Converting a virtual machine to a template is a one-way operation. Templates become read-only and can no longer be started as normal virtual machines.

Virtual Machine Snapshots

XenServer 6.0 introduces Live Memory Snapshots, which can capture the memory state of a running virtual machine along with the disk. When reverting to the snapshot, the virtual machine will go back to the exact point in time when the snapshot was taken. Without a Live Memory Snapshot, the restoration will be from a shut down state.

Live memory snapshots:

- Create a record of a virtual machine at a point in time.
- Do not require the virtual machine to be shut down.
- Capture all the storage and configuration information for the original virtual machine, including networking information.

> Snapshots are not to be confused with templates. Templates are the most flexible copy method. There are templates suitable for backing up, exporting to another resource pool, or creating additional virtual machines.

To Create a Template from a Windows Virtual Machine

1. Create, install, and configure the Windows virtual machine as desired.
2. Apply all relevant Service Packs and updates.
3. Install the XenServer Tools.
4. Install any applications and perform any other configuration.
5. Run sysprep. This will shut down the virtual machine when it completes.
6. Convert the virtual machine into a template using XenCenter.
7. Clone the newly created template into new virtual machines as required.
8. When the cloned virtual machine starts, it will get a new SID and name, run a setup to prompt for configuration values as necessary, and restart.

> The original virtual machine that sysprep was run on should not be restarted again after the sysprep stage; it should be converted to a template immediately after sysprep is completed. If the source virtual machine is restarted, sysprep must be run on it again before the virtual machine can be used to make additional clones. For more information about sysprep, visit the Microsoft support page on the *support.microsoft.com* Web site.

Creating a Template from a Linux Virtual Machine

1. Create the virtual machine for your target operating system using XenCenter or the command-line interface.
2. Install the operating system using vendor installation media.
3. Install the XenServer Tools.
4. Configure the correct time and time zone on the virtual machine and VNC as you would in a normal, non-virtualized environment.

Test Your Knowledge: Determining Which Template to Use

Match the following terms with the correct descriptions.

- Basic templates
- Full templates
- Custom templates

Term	Description
Basic templates	Include recommended settings for new virtual machines, but do not include an operating system.
Custom templates	Can be created from any virtual machine template.
Full templates	Include copies of the entire operating system with settings.

Exporting a Virtual Machine

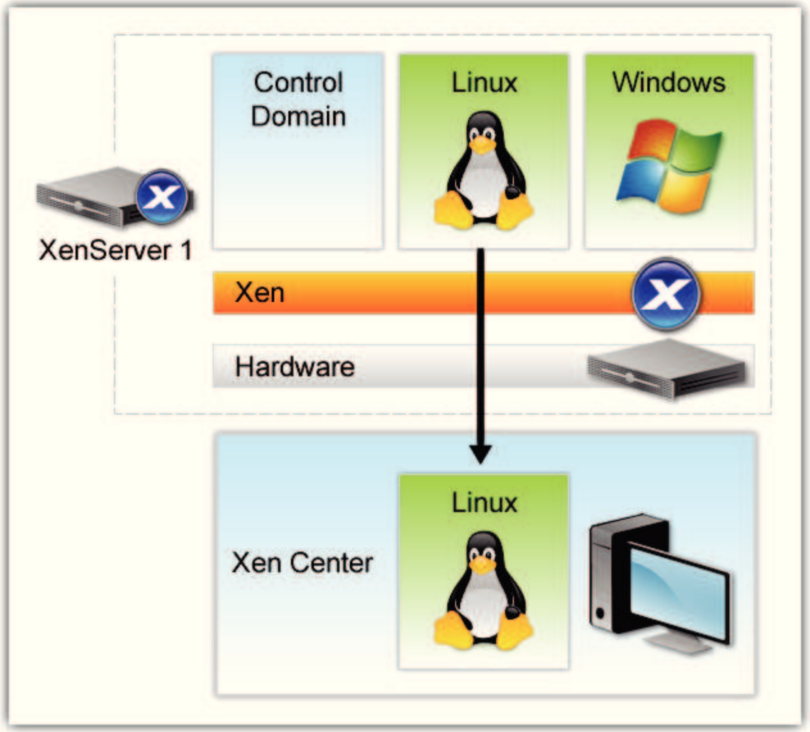

You can import and export virtual machines as an appliance package, which is a collection of one or more virtual machines, as well as disk images, using the Import and Export wizards in XenCenter or the XenServer command-line interface.

When you export virtual machines as an appliance, they are exported as configuration data along with the virtual hard disks of each virtual machine.

A virtual machine can be exported using the `xe vm-export` command.

Exporting a virtual machine allows for:

- The movement of virtual machines.
- The creation of a library of virtual machine images.
- A full backup of virtual machines.
- The archiving of old virtual machines that are being taken out of service.

When exporting a virtual machine:

- The virtual machine must be shut down.
- The virtual machine is exported by XenServer as a single, uncompressed binary file.

- The virtual machine data is sent over the SSL link to the system running XenCenter or the command-line interface.
- Bandwidth between the administration station and the server must be adequate.

Demo how to copy a virtual machine.

Copying a Virtual Machine

There are several ways to copy a virtual machine. Existing environment and virtual machine availability are key factors in determining the copy method.

Copy (Cloning)

The copy function is best for creating a duplicate virtual machine within the same resource pool.

- The copy function is limited to the same resource pool.
- The copy types include fast clone and full copy.
- The virtual machine needs to be shut down.

When copying a virtual machine, you have the option to create a full copy or a fast clone. The fast clone mode only writes modified blocks to disk and is only supported for file-backed virtual machines. Fast clone is designed to save disk space and allow fast clones, but will slightly slow down normal disk performance. A template can be fast-cloned multiple times without performance degradation.

Convert to Template

The convert to template function is best used for the initial creation of base virtual machines.

- The convert to template function is used to create a base image of a virtual machine.
- The virtual machine needs to be shut down.
- The source virtual machine is no longer available.

If you create a virtual machine using Quick Create and convert the virtual machine back into a template, disk performance can decrease depending on the number of times the conversion has been performed. In this event, you can use the `vm-copy` command to perform a full copy of the disks and restore disk performance.

Export as Backup

The export as backup function creates an XVA file that you can use to back up, export, or create additional virtual machines.

- The export as backup function creates a single file with all the virtual machine settings.
- The virtual machine needs to be shut down.
- The source virtual machine is still available after export. Exporting does not remove the virtual machine from the storage repository.

Snapshot

The snapshot function is the most flexible copy method, providing a template suitable for backing up, exporting to another resource pool, or creating additional virtual machines. Unlike the export as backup function, the snapshot function does not require you to shut down the source virtual machine.

- The snapshot function creates a record of a virtual machine at a point in time.
- The virtual machine does not need to be shut down.
- A template created by the snapshot function is similar to a normal virtual machine template but contains all the storage and configuration information for the original virtual machine, including networking information.

Assigning Resources to a Virtual Machine

Virtual machines can be assigned physical resources of the XenServer host.

GPU

XenServer 6.0 allows you to assign a physical graphics processing unit (GPU) in a XenServer host machine to a Windows virtual machine running on the same host. This GPU Pass-Through feature is intended for graphics power end users, such as CAD designers, who require high performance graphics capabilities. It is supported only for use with XenDesktop.

GPU Pass-Through is supported for specific machines and GPUs. In all cases, the IOMMU chipset feature, known as VT-d for Intel models, must be available and enabled on the XenServer host. Before enabling GPU Pass-Through, visit *www.citrix.com/ready/hcl* to check the hardware compatibility list. For any further questions, e-mail *xenserver.hcl@citrix.com*.

> GPU Pass-Through is only available to Windows virtual machines on Citrix XenServer Enterprise Edition and higher. It can be enabled using XenCenter or the xe command-line interface.

Network Resources

For Windows virtual machines, the XenServer host communicates with the virtual machine through the vSwitch Controller, which is initially assigned an IP by DHCP. However, Citrix recommends assigning the vSwitch Controller a static IP.

For more information about Networking, see Citrix article CTX130420 on *support.citrix.com*.

> If you want to revert to the default Linux network stack, run the `xe-switch-network-backend bridge` command and restart your XenServer host.

Storage

You can create one or more storage repositories, each containing multiple virtual disk images.

For more information about Storage Configuration and Managing Storage Repositories, see Citrix articles CTX130420 and CTX130422 on *support.citrix.com*.

Importing a Virtual Machine

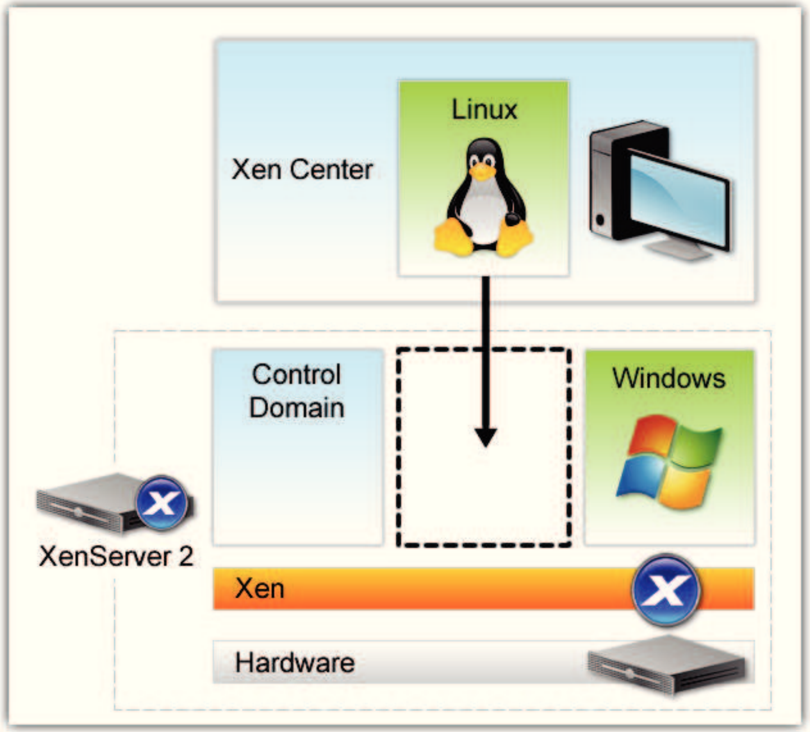

You can create a virtual machine by importing an exported virtual machine. Like cloning, exporting and importing a virtual machine allows you to create additional virtual machines of a certain configuration so that you can increase the speed of your deployment. Importing is more efficient than creating a virtual machine from scratch.

When a virtual machine is imported:

- Data is sent over SSL.
- The virtual machine is automatically assigned a new UUID.
- A new MAC address is assigned.
- XenServer re-attaches existing virtual network interfaces to any network with the same name as the network on the server from which the virtual machine was exported. If no matching network can be found, the virtual machine network interface, also known as VIF, is added to a new private network.

- Citrix recommends that you verify the virtual NIC settings after it has been imported.
- A virtual machine cannot be exported on an Intel-based server then imported on an AMD-based server. A virtual machine also cannot be exported on an AMD-based server then imported on an Intel-based server.
- To maintain the original MAC address, import using the command line and add the `preserve=true` argument.

A virtual machine can be imported by using the `xe vm-import` command.

XenConvert V2V is now done through importing. Importing an exported virtual machine might take some time, depending on the size of the virtual machine and the speed and bandwidth of the network connection between the XenServer host and XenCenter.

Modifying Virtual Machine or Template Resources

You can adjust resource settings after creating a virtual machine. XenCenter allows you to adjust most resource settings, like disk capacity and network performance, regardless of whether the virtual machine is powered on or shut down.

Adjustable virtual machine resource settings include:

- Virtual disks and disk sizes
- Virtual disk resizing
- Virtual NICs
- Virtual CPUs
- Memory

Virtual Disks

Virtual disks can be added and removed in Linux and Windows virtual machines because virtual disks look like raw storage to the virtual machine. You can add or remove virtual disks from a running virtual machine without restarting, which allows you to make adjustments to storage resources while a workload is running.

Virtual Disk Sizes

You can increase the size of virtual disk files on every type of storage except NFS. NFS-based virtual disks use sparse allocation, so you can allocate extra disk space during virtual disk creation without allocating too much space. You can add additional disks on an NFS-based virtual disk if needed. You cannot reduce virtual disk sizes.

Increasing the size of a virtual disk in XenCenter increases the size of the disk seen by the virtual machine. You should expand the file system and partitions within the virtual machine to cover the added space.

You must use utilities such as Disk Manager in Windows and fdisk in Linux to partition and format new virtual disk drives. Diskpart is used to expand an existing NTFS partition using the command-line interface.

Virtual NICs

Virtual NICs can be added, removed, and edited in the Network tab of a virtual machine in XenCenter.

Virtual CPUs

Increasing the number of virtual CPUs does not require any virtual machine kernel changes if it is performed on a Linux or Windows 2003/XP virtual machine. Windows systems still need to be restarted to change to the multiprocessor Hardware Abstraction Layer.

Windows 2000 virtual machines configured with a single vCPU need to be manually reconfigured.

Virtual CPUs

Increasing the number of virtual CPUs does not require any virtual machine kernel changes if it is performed on a Linux or Windows 2003/XP virtual machine. Windows systems still need to be restarted to change to the multiprocessor Hardware Abstraction Layer.

Windows 2000 virtual machines configured with a single vCPU need to be manually reconfigured.

Memory

When a virtual machine is first created, it is allocated a fixed amount of memory. XenServer Dynamic Memory Control works by automatically adjusting the memory of running virtual machines, keeping the amount of memory allocated to each virtual machine between specified minimum and maximum memory values.

Dynamic Memory Control provides the following benefits:

- Memory can be added or removed without restarting the virtual machine, providing a seamless experience to the end user.
- When servers are full, dynamic memory control allows you to start more virtual machines on these servers, reducing the amount of memory allocated to the running virtual machines proportionally.

Virtual Appliance Packages

Both Open Virtualization Format (OVF) and Open Virtualization Appliance (OVA) are appliance package standards defined by the Distributed Management Task Force (DMTF). This section provides an overview of the OVF and OVA formats and describes the scenarios most appropriate for each type.

Open Virtualization Format

OVF is an open standard for packaging and distributing software to be run in virtual machines. This standard describes the metadata of one or more virtual machines. OVF packages also supports other non-metadata related capabilities, such as encryption, compression, archiving, EULA attachment, and annotations, among other capabilities. An OVF package consists of a descriptor file (*.ovf) and any other files representing the following attributes of the package:

Signature	Digital signature used by a public key certificate in the X.509 format to authenticate the producer of the package
Manifest	An SHA-1 digest of every file in the package to verify its contents by detecting any corruption
Virtual disks	Files comprising virtual disks in the format defined by the virtualization product that exported the virtual disks - VMware products export a virtual disk in the Stream-Optimized VMDK format for an OVF package. XenServer products export a virtual disk in the Dynamic VHD format for an OVF package.

Open Virtualization Appliance

An OVA package is a single archive file in the Tape Archive (.tar) format, containing the files that comprise an OVF package.

Selecting the Appropriate Package

OVF packages contain a series of uncompressed files, which makes it easier for those who want to access individual disk images in the file. OVA packages are one archive file of an OVF package. An OVA file can be compressed, but it does not offer the flexibility of a series of files, unlike OVF. OVA is useful for specific applications for which it is beneficial to have just one file, such as creating packages for Web downloads. Consider using OVA only as an option to make the package easier to handle, because using this option lengthens both the export and import processes.

To Create Virtual Appliances

1. Select **Pool** > **Virtual Appliances** .
   This displays the Manage Virtual Appliances window.
2. Click **New Appliance**.
3. Type a name for the Virtual Appliance, and optionally a description, and then click **Next.**
4. Select the virtual machines that you want to group together into a Virtual Appliance, and then click **Next.**
5. Specify the start order and the delay between each virtual machine in the appliance and click **Next**.
   The details of the appliance are displayed.
6. Click **Finish** to create the appliance.

A virtual appliance can span multiple servers in a single resource pool, but cannot span several resource pools.

Exporting an Appliance Package

The Export and Import wizards in XenCenter have been enhanced to allow appliance packages in OVF/OVA format to be imported and exported.

Disk images in VHD and VMDK formats can also be imported using this wizard.

When you export virtual machines in an appliance, they are exported as the configuration data along with the virtual hard disks of each virtual machine. The process creates an OVF or OVA file that describes the way these virtual hard disks join together to form a virtual machine. The file also describes the resource settings (CPU, memory, and disk space) associated with that virtual machine.

To Export an Appliance Package

1. Click the VM menu and click **Export**.
2. Type a name for the export file and specify the folder in which you want it to be saved and click **Next**.
3. Select the virtual machines you want to export and click **Next**.
4. Accept the license agreement and click **Next**.
5. Configure the Advanced Options and click **Next**.
6. Configure the appropriate settings and click **Next**.
7. Click **Finish**.

Importing an Appliance Package

When you import the appliance package, the wizard reads the OVF, OVA, or disk image to determine the resource requirements for the virtual machines in the appliance—and which virtual disks are associated with each virtual machine—and then reconstitutes them.

To Import an Appliance Package

1. Right-click the resource pool or XenServer host in which you want to place one or more new virtual machines and select **Import**.
2. Click **Browse** and locate the file you want to import. Click **Next**.
3. Select the appropriate host and click **Next**.
4. Select the appropriate storage option and click **Import**.
5. Select the appropriate network and click **Next**.
6. Click **Finish**.

XenConvert Overview

As a physical-to-virtual conversion tool, Citrix XenConvert can convert a server or desktop workload from an online physical machine running Windows to a XenServer virtual machine or Provisioning Services vDisk.

For more information about host machine system and virtual machine/virtual disk requirements, see Citrix article CTX130945 on *support.citrix.com*.

> Virtual-to-virtual conversions have been removed in order to align with new appliance import features in XenServer 6.0.

XVA is the default format of a Xen Virtual Appliance. There are two versions of XVA.

Version 1	XVA V1 includes: • The ova.xml file, which is a meta-data file that defines the properties of a Xen virtual machine • The dha folder, which is a folder containing one or more compressed chunks of virtual hard disk
Version 2	The XVA V2 format is a single file archive of the files that comprises a Xen Virtual Appliance.

Physical Machine Conversion

You can convert a physical machine, including up to four volumes, to one of the following destinations:

- Physical machine to Provisioning Services vDisk
- Physical machine to VHD
- Physical machine to XenServer
- Physical machine to OVF Package

All conversions from a physical machine include preparing the source machine and choosing volumes.

Source Machine Preparation

Before converting from a physical machine:

- Enable Windows Automount on Windows Server operating systems.

- Disable Windows Autoplay.
- Remove any virtualization software.
- Verify that adequate free space exists at the destination drive.

Volume Selection

When converting from a physical machine to a virtual machine or virtual disk, XenConvert provides options to:

- Select up to four volumes to include.
- Resize the volumes on the virtual disk by changing the amount of free space.
- Resize the entire virtual disk by changing the allocated space (for destinations other than Provisioning Services).

These options become available after the Welcome page appears in the wizard.

For more information about physical machine conversions, see Citrix article CTX130945 on *support.citrix.com*.

Converting from Physical to Virtual Machine

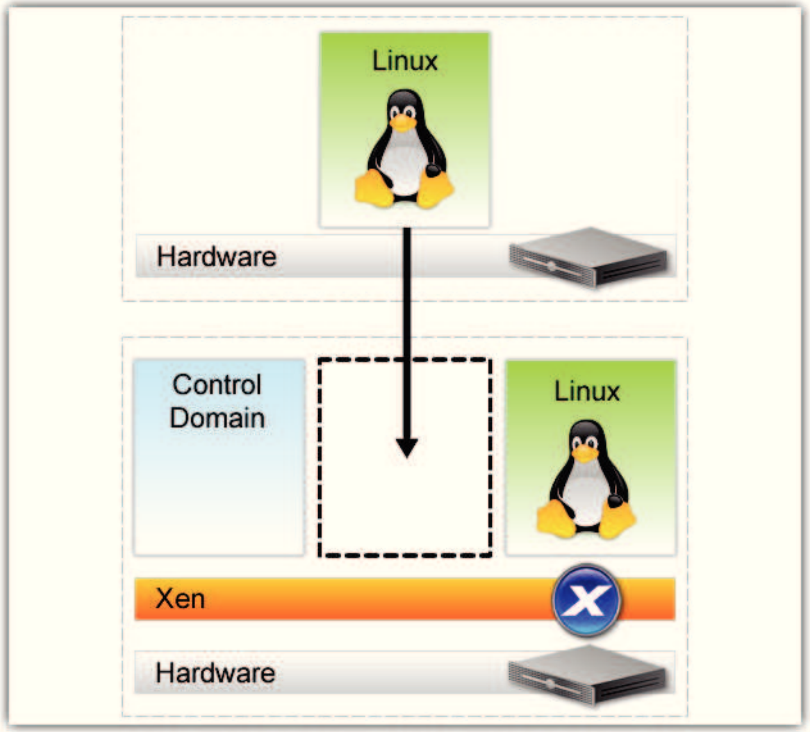

XenConvert is a physical-to-virtual (P2V) conversion tool that converts a workload from a server or desktop machine to a virtual machine, virtual appliance or virtual disk.

XenConvert produces:

- A virtual appliance in OVF that XenCenter can import into XenServer.
- A virtual disk in the Dynamic Virtual Hard Disk format that is compatible with XenServer and Provisioning Services.

Performing a P2V Conversion

1. Start the physical server.
2. Start XenConvert.
3. Choose the source and destination for the conversion.
4. Convert the physical server to a virtual server.
5. Review the results of the conversion.

Dynamic Memory Control

Dynamic Memory Control (DMC) works by automatically adjusting the memory of running virtual machines. This keeps the amount of memory allocated to each virtual machine between specified minimum and maximum values, guaranteeing performance and permitting a greater density of virtual machines for each XenServer host. Without DMC, when a server is full, starting further virtual machines will fail with "out of memory" errors. With DMC enabled, even when the XenServer host is full, XenServer will attempt to reclaim memory by automatically reducing the current memory allocation of running virtual machines within their defined memory ranges.

You can use DMC to operate a virtual machine in Target mode or Dynamic Range mode.

- If you set the dynamic min/max to the same value, XenServer will ensure that this exact amount of memory is allocated to the virtual machine.
- DMC is only available for XenServer Advanced or higher editions. For more information about XenServer Advanced or higher editions and to find out how to upgrade, see Citrix article CTX130421 on *support.citrix.com* and *XenServer features by edition* on the *www.citrix.com* Web site.

Target Mode

You can specify a memory target for the guest operating system. XenServer adjusts the guest operating system's memory allocation to meet the target. Specifying a target is particularly useful in virtual server environments and when you know exactly how much memory you want a guest operating system to use. XenServer will adjust the virtual machine's memory allocation to meet the target you specify.

Dynamic Range Mode

You can specify a dynamic memory range for the guest operating system. XenServer chooses a target from within the range and adjusts the guest operating system's memory allocation to meet the target. It is useful to specify a dynamic range in virtual desktop environments. XenServer chooses a target from within the range and adjusts the memory allocation to meet the target.

Memory Constraints

You can use any memory control operations with any guest operating system. XenServer allows you to change virtual machine memory properties to any values subject to validation checks. Citrix supports only certain guest operating system memory configurations for each supported operating system. For more information about configuring virtual machine memory for select operating systems, see Citrix article CTX130420 on *support.citrix.com*.

Dynamic Memory Control Optimization

Optimizing DMC requires configuring the following values:

Dynamic Memory Range The range within which memory can be added or removed from the virtual machine while it is running, without requiring a restart. XenServer always guarantees that the amount of memory allocated to the virtual machine within the dynamic range will be available.

Dynamic Minimum Memory A lower memory limit that you assign to the virtual machine - Allocating only a small amount of memory to a virtual machine can negatively impact it, leading to long startup time and compromised performance.

Dynamic Maximum Memory An upper memory limit that you assign to the virtual machine.

DMC Example

If the Dynamic Minimum Memory was set at 512 MB, and the Dynamic Maximum Memory was set at 1024 MB, the virtual machine would have a Dynamic Memory Range (DMR) of 512 through 1024 MB, within which it would operate. XenServer guarantees that each virtual machine will be assigned memory within its specified memory range at all times.

Static Memory Range

Static memory is the maximum amount of memory that a virtual machine will ever be asked to consume, declared at the time that the virtual machine starts. This allows the virtual machine to size its page tables and other memory management structures accordingly. The static memory range cannot be adjusted while the virtual machine is running. The dynamic range is constrained so that it will always be contained within this static range.

The static minimum (the lower bound of the static range) protects the administrator and is set to the lowest amount of memory with which an operating system can run on XenServer.

When configuring the memory for the virtual machine, be careful not to exceed the maximum amount of physical memory addressable by your operating system. Setting a memory maximum that is greater than the operating system supported limit might lead to stability problems.

DMC Behavior when Launching New Virtual Machines

If a new virtual machine is required to start on a full server, the required extra memory is obtained by reducing the memory of running virtual machines proportionally within their pre-defined dynamic ranges.

Automatic virtual machine memory adjustment scenarios include:

DMC disabled, host full	New virtual machines fail to start with an "out of memory" error.
DMC enabled, host full	XenServer attempts to reclaim sufficient memory. If enough memory can be reclaimed, the virtual machine starts. All existing virtual machines receive at least their Dynamic Minimum Memory level. If insufficient memory was recovered at the minimum memory level, the new virtual machine fails to start.
DMC enabled, plenty of memory	All running virtual machines receive their Dynamic Maximum Memory level.

Test Your Knowledge: Dynamic Memory Control

1. Which three DMC values must you change to optimize host server utilization? (Choose three.)
    a. Dynamic Terminal Memory
    b. Dynamic Memory Range
    c. Static Memory Range
    d. Dynamic Minimum Memory

    Answer: b, c, d
2. Which three resources can you modify to increase the XenServer host utilization?(Choose three.)
    a. Memory
    b. Number of virtual disks
    c. Virtual NICs
    d. IP addresses

    Answer: a,b,c

Module 6

Installing and Configuring Provisioning Services

Overview

XenServer is the simplest and most effective way to virtualize and provision servers. With the addition of Provisioning Services, on-demand provisioning of both physical and virtual servers is enabled, providing the ability to:

- Stream a virtual server workload image to a virtual machine.
- Create a common workload image that can provision both physical and virtual servers.

Timings:

Module: 60 minutes

Exercises: 90 minutes

Total Time: 150 minutes

Objectives

After completing this module, you will be able to:

- Configure Provisioning Services.
- Determine which options to select during a Provisioning Services installation.
- Identify the key components and services of Provisioning Services architecture.
- Integrate Provisioning Services with Active Directory to enable Active Directory password management.
- Prepare a target device for vDisk assignment by adding the device to the Provisioning Services database.

Provisioning Services Technology

For more information about Provisioning Services 6.0, read the Citrix blog: *http://blogs.citrix.com/2011/09/28/provisioning-services-6-0-from-single-image-delivery-to-single-image-management/*

Provisioning Services provides the ability to provision the operating system of a computer and re-provision it in real-time from a single shared-disk image. In doing so, you can completely eliminate the need to manage and update individual systems.

Provisioning Services is based on software-streaming technology. After installing and configuring Provisioning Services components, a vDisk is created from the hard drive of a device by taking a snapshot of the operating system and application image, and then storing that image as a vDisk file on the network. A device that is used during this process is referred to as a Master target device. The devices that use those vDisks are called target devices.

vDisks can exist on a Provisioning Services host, file share, or in larger deployments, on a storage system with which the Provisioning Services host can communicate, such as iSCSI, SAN, NAS, and CIFS. vDisks can be assigned to a single target device in private image mode, or to multiple target devices in standard image mode.

When a target device is turned on, it is set to start from the network and to communicate with a Provisioning Services host. The target device downloads the startup file from a TFTP server, and then the target device starts up. Based on the device start up configuration settings, the appropriate vDisk is located, then mounted by a Provisioning Services host. The software on that vDisk is streamed to the target device, as needed.

Instead of immediately pulling all the vDisk content down to the target device, the data is brought across the network in real-time, as needed. The Provisioning Services host provides files from the vDisk as they are requested by the operating system, in the same way that the operating system would normally request them from its hard drive. This approach allows a target device to load a completely new operating system, and software from the vDisk in the time it takes to restart. This approach dramatically decreases the amount of network bandwidth required by traditional disk imaging tools; making it possible to support a larger number of target devices on your network without impacting overall network performance.

Provisioning Services for XenServer

Each licensed XenServer host can provision:

- Unlimited virtual machines on the XenServer Advanced, Enterprise, or Platinum host.
- An additional three physical servers from a XenServer Platinum host.

Citrix Provisioning Services Components

During a Provisioning Services installation, you install the components and services necessary to provide software streaming. It is essential to know the functions of each component and how each service facilitates communication between these components. The Citrix Provisioning Services components include:

Provisioning Services Host A Provisioning Services host streams a vDisk to a target device. The Provisioning Services host acts as a proxy between the target device and the vDisk store by using the Stream Service to stream content from the vDisk to the target device.

Provisioning Services Database A Provisioning Services database stores all Provisioning Services hosts, vDisk, target devices, and system configuration settings that exist within a farm. Only one database can exist within a farm and all Provisioning Services hosts in that farm must be able to communicate with the database. Provisioning Services supports Microsoft SQL 2005, 2008, and 2008 R2.

Stores A store is the logical name given to a physical or virtual storage location for vDisks. A store can be placed on a local drive on a Provisioning Services host, a SAN, a CIFS share, a NAS, or a UNC path. In this way, a store can be used for an entire farm or for a particular site or server.

vDisk A vDisk is a file that contains an image of the hard drive of a device, including operating system and any installed applications. A Provisioning Services host streams the image to target devices. vDisks are housed in a store, which can be located either locally on a Provisioning Services host or on a shared storage device.

vDisk Pools A vDisk pool is a collection of all vDisks available to a site. A site can contain only one vDisk pool.

Target Device A target device is any desktop or server system that receives a streamed operating system and applications from a vDisk. Each target device continues to have its own identity on the network and within the existing network directory services.

Master Target Device A master target device is used to create and test a golden image and represents the pristine state of a system, including all operating system and application updates and configurations. The master target device is used to create a vDisk that will be shared by multiple end users.

Citrix License Server

The license server stores Provisioning Services licenses. You download licenses from the Citrix.com web site to the license server, which then checks the license out to target devices as requested. For more information about Citrix licensing, see Citrix eDocs at *edocs.citrix.com*.

Citrix Provisioning Services Infrastructure

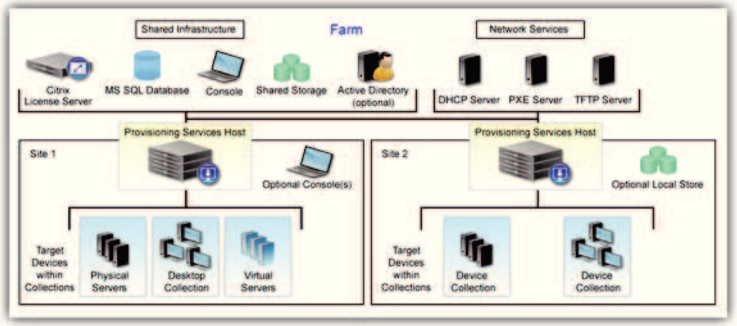

The Citrix Provisioning Services infrastructure is organized into the following hierarchy:

Farm

A farm represents the top level of a Provisioning Services infrastructure and provides you with a method of defining and managing logical groups of Provisioning Services components into sites. A farm contains a Citrix License Server and Microsoft SQL database.

Site

A site contains one or more Provisioning Services hosts, device collections, views, vDisk pools, and can contain local shared storage. A site is an administrative unit that can correspond with a physical location, such as a branch office or floor of a building, an IP range, or other logical grouping. Provisioning Services hosts within a site communicate with the farm components to obtain information necessary to start target devices and stream vDisks. If vDisks are located on shared storage at the farm level, Provisioning Services hosts within the site must have access to the store.

Device Collection

A device collection is a logical grouping of target devices. A device collection could represent a physical location, a subnet range, or a logical grouping of target devices. Organizing target devices into collections simplifies management because tasks can be performed at the collection level, rather than on a device-by-device basis. A target device can only be a member of one device collection.

View

Views allow you to quickly manage a group of target devices. Views are typically created according to business needs, such as a physical location, or user type. Unlike device collections, a target device can be a member of multiple views.

vDisk Pool

A vDisk pool is a collection of all vDisks available to a site. There is one vDisk pool for each site.

Provisioning Services Administration

Provisioning Services can be administered from the Provisioning Services Console or from the Management command-line interface.

Provisioning Services Console

The Provisioning Services Console is a utility that is used to manage Provisioning Services and to create and configure vDisks and target devices. The Console is an MMC-based administration console that includes the following functionalities:

- Integrated Windows authentication using local groups or Active Directory groups
- Remote administration of servers and devices in a farm You can also take advantage of the MMC taskpad view to configure lists of common administration tasks.

Example: Taskpad View

You might add the Active Directory Users and Computers snap-in to the Provisioning Services Console. You could then use the taskpad view to create a list of tasks in Provisioning Services and Active Directory that must be completed when a vDisk and target device are provisioned for a new employee.

Management Command-line Interface

The Management command-line interface allows you to manage Provisioning Services through a command-line interface and to automate common tasks, such as adding or deleting a vDisk with scripts. For more information about Management command-line interface, see Citrix eDocs at *edocs.citrix.com*.

During installation, the vbox loads the management API. You can deselect this option to remove the management command-line interface.

Administrative Roles

The ability to view and manage objects within a Provisioning Services implementation is determined by the administrative role assigned to any object in Active Directory. Provisioning Services makes use of Windows and Active Directory groups that already exist within the network. The following administrative roles can be assigned to a group:

Farm Administrator

A farm administrator can view and manage all objects within a farm. Farm administrators can also create new sites and manage role memberships throughout the entire farm. When a farm is first configured, the administrator that creates the farm is automatically assigned the farm administrator role.

Site Administrator

A site administrator has full management access to all objects within a site. For example, a site administrator can manage Provisioning Services, site properties, target devices, device collections, vDisk assignments, and vDisk pools. If a farm administrator assigns a site as the owner of a particular store, the site administrator can also manage the store. The site administrator can also manage device administrator and device operator memberships.

Device Administrator

A device administrator manages device collections. Device collections consist of a logical grouping of target devices.

Device Operator

A device operator has administrator privileges to perform the following tasks within a device collection:

- Start and restart a target device.
- Shut down a target device.
- View target device properties.
- View vDisk properties for those vDisks assigned to a target device.

Key Services

Provisioning Services uses the following services.

Explain the PXE and DHCP option requirements.

For more information about load balancing TFTP servers, see Citrix article CTX116337 on *http://support.citrix.com*.

Citrix License Service

The Citrix license service retrieves the product license from the license server.

Stream Service

The Stream Service makes vDisk streaming possible by providing a vDisk and its contents to target devices. The Stream Service streams the contents of that vDisk on demand, eliminating the need to stream the entire contents of a vDisk during startup. The Stream Service also transfers data from a target device to a vDisk when the target device uses the vDisk in private image mode and from the device to a write cache when the vDisk is in standard image mode.

SOAP Service

The SOAP service provides a framework to enable external or existing solutions to interface with and manage Provisioning Services. The Provisioning Services Console and the Management CLI both use this service.

The Stream Service and SOAP Service are fundamental Provisioning Services technologies and required installation components. These services cannot run as standalone services.

In addition, Provisioning Services also uses boot services during the startup process to communicate and exchange information between the Provisioning Services components. The boot services include PXE, BOOTP, DHCP, and TFTP.

PXE

The Pre-Execution Environment (PXE) protocol is a BIOS extension that enables target devices to start up from a network interface card (NIC), regardless of the availability of local data storage devices or operating systems. The PXE protocol is made up of a generic component common to all devices and a vendor-specific component. PXE combines either BOOTP or DHCP and TFTP to locate IP address of the target device, the location of the Provisioning Services host, and download the bootstrap file. Target devices must support PXE in order to start up from the network.

When a target device is turned on, it sends a DHCP broadcast that identifies the target device as PXE compatible.

PXE receives data on UDP port 67 and sends data to UDP port 68.

> Provisioning Services supports PXE .99j or later. Running multiple PXE servers in the same environment might result in performance problems.

BOOTP

The Bootstrap Protocol (BOOTP) is a precursor to DHCP, and like DHCP, it is a UDP protocol that target devices use to request and obtain IP addresses from a BOOTP server. BOOTP can also deliver the bootstrap file location and file name to a target device. The BOOTP server receives requests on UDP port 67 and sends data to UDP port 68 on a target device. While BOOTP is no longer common, it can be used when DHCP does not meet the requirements of an environment.

DHCP

The Dynamic Host Configuration Protocol (DHCP) is used by the target device to request and obtain an IP address from the DHCP service. DHCP uses Options 66 and 67 to specify the bootstrap file location and file name to a target device. The DHCP service receives requests on UDP port 67 and sends data to UDP port 68 on a target device.

DHCP reservations can be used when one or more target devices must reliably receive the same IP address from the DHCP service.

TFTP

Target devices use the Trivial File Transfer Protocol (TFTP), which is a simple file transfer protocol, to request and receive a bootstrap file from the TFTP service. The TFTP service receives requests on UDP port 69 and sends data to UDP port 69 on a target device.

DHCP Deployment Options

Provisioning Services can be deployed with one of three DHCP configurations:

- Standalone DHCP
- Co-hosted DHCP and Proxy DHCP
- Separated DHCP and Proxy DHCP

Standalone DHCP is by far the most common configuration.

You can have only one TFTP per subnet.

Standalone DHCP

In a standalone DHCP configuration, the DHCP service performs all non-TFTP communications between the Provisioning Services host and target devices.

You must configure DHCP options 66 and 67 when a standalone DHCP configuration is used in order to communicate the IP address of the TFTP server and the bootstrap file name to a target device.

Standalone DHCP configurations are most commonly used in production environments and require assistance from DHCP or network administrators.

Co-hosted DHCP and Proxy DHCP

A proxy DHCP service is a PXE service running either on a server other than the server that hosts the DHCP service, or on the same server using port 4011 instead of port 67. Whereas PXE scope options might not be enabled on typical DHCP services, the appropriate DHCP options are enabled on the proxy DHCP service that is installed as part of the Provisioning Services host PXE service. This allows the proxy DHCP to respond appropriately to target device requests.

If you do not have access to the DHCP configuration, you can choose to set up a proxy DHCP server that hosts the PXE service. The following table lists the information the target device receives from each service in an environment in which the DHCP and proxy DHCP services are hosted on the same server.

Service	Port	Information Sent to Target Device
DHCP	67	IP address of the target device

Service	Port	Information Sent to Target Device
Proxy DHCP	4011	• IP address of TFTP service • Name of bootstrap file

Separated DHCP and Proxy DHCP

If you do not have access to the DHCP configuration, you can choose to set up a proxy DHCP server that hosts the startup services. The following table lists the information that the target device receives from each service in an environment in which the DHCP and proxy DHCP services are located on separate servers.

Service	Port	Information Sent to Target Device
DHCP	67	IP address of the target device
Proxy DHCP	67	• IP address of TFTP service • Name of bootstrap file

DHCP Configuration

You must manually configure DHCP PXE scope options to provide the bootstrap file information to a target device at the same time the DHCP service sends the target device IP address. If DHCP PXE options are not configured, the DHCP service provides only the IP address to the target device.

The following information describes the DHCP PXE options that are necessary in a Provisioning Services environment.

60 Identifies the target device as a PXE client. The default setting is PXEClient.

66 Identifies the FQDN or IP address of the TFTP service. When the FQDN of the TFTP service is used in DHCP option 66, DNS resolves the request and returns the IP address of the TFTP service to the target device.

67	Identifies the name of the bootstrap file. The default setting is `ardbp32.bin`.

Depending on the design of the farm and sites, you might want to set DHCP options at the global or scope level. Global options apply to all DHCP responses made by the server. Scope options apply only to those IP addresses that are part of the scope.

Spanning Tree Protocol is a link management protocol that prevents network loops in a bridged LAN and provides path redundancy. Spanning tree can cause PXE requests to time out. You can prevent PXE request timeouts using one of the following methods:

- Disabling spanning tree on Provisioning Services host switch ports
- Enabling spanning tree portfast mode on all Provisioning Services host switch ports

Test Your Knowledge: Key Services

Match the following terms with the correct descriptions.

- Citrix License Server
- Stream Service
- SOAP Service
- DHCP
- PXE
- TFTP

Term	Description
Stream Service	Provides a vDisk and its contents to target devices.
DHCP	Assigns an IP address to target devices.
Citrix License Server	Retrieves the product license.
PXE	Enables target devices to start up from a network interface.
TFTP	Delivers start up information to target devices.
SOAP Service	Provides a framework to enable external or existing solutions to interface with and manage Provisioning Services.

Installation Planning

A number of decisions must be made before installing and configuring Provisioning Services, including:

- Is proxy DHCP required?
- Which SQL server edition to use for the Provisioning Services database?
- Which administrative permissions are necessary?
- What is the optimal farm layout of databases, sites, and high availability?
- Which server will be used as the Citrix License Server?
- Does the environment meet the hardware and software requirements for a Provisioning Services installation?

Provisioning Services Farm Design

You should plan a Provisioning Services farm design prior to installing Provisioning Services to make decisions during the installation and configuration process. Consider the following items when planning a farm design:

- Farm name
- SQL database server
- SQL database authentication
- Sites
- Stores
- Device collections
- Role-based administration
- High availability

> Farm name, site, store, and device collection objects are used for administrative purposes only and do not correlate to farm structures used by XenDesktop, XenApp, XenServer, or any other product.

Farms

When planning the number of farms required to support a given environment, you should consider the communication that takes place between Provisioning Services hosts and the database. Provisioning Services hosts communicate with the database constantly to access information that is used to stream vDisks to target devices. The database should be located in physical proximity to the Provisioning Services hosts in the farm to minimize latency and ensure optimal target device performance.

For example, a large enterprise with offices around the world should consider creating a farm for each major business region to avoid latency. If target devices are located in North America, Europe, and South America, consider creating three farms.

SQL Database Server

Only one Provisioning Services database is associated with a farm.

The Provisioning Services database can be created on an existing SQL database server, provided the server can communicate with all Provisioning Services hosts within the farm.

> In some production environments, your database administrator might prefer to create a Provisioning Services database for you. In this case, provide the Microsoft SQL database administrator with the file created using the `DBScript.exe` utility, which is installed with the Provisioning Services software.

The database does not grow significantly as more objects are added to it. The database typically grows by 10 MB with the addition of 10,000 target devices. For example, the database in a Provisioning Services farm with 50,000 target devices would occupy approximately 50 MB of disk space. For more information about scalability statistics, see Citrix eDocs at *edocs.citrix.com*.

SQL Database Authentication and Configuration

Provisioning Services uses Windows authentication. All Provisioning Services components, including the Configuration Wizard and services that access the database, must run in the context of the logged-on user.

Services, such as the Stream Service and SOAP service, require minimal privileges in the end-user configuration.

> Provisioning Services supports Windows authentication, as recommended by Microsoft.
>
> Microsoft SQL Server authentication is not supported, except when running the Configuration Wizard.

Configuration Wizard User Permissions

The following Microsoft SQL permissions are required for the user that is running the Configuration Wizard:

- dbcreator, which is required for creating the database
- securityadmin, which is required for creating the SQL logons for the stream and SOAP services

If the end user does not have sufficient SQL permissions, a dialog box prompts for a SQL Server end user who has the appropriate permissions (dbcreator and securityadmin).

If using SQL Express in a test environment, you can choose to provide the end user who is running the Configuration Wizard sysadmin permissions (the highest database privilege level).

Alternatively, if the database administrator has provided an empty database, the end user running the Configuration Wizard must be the owner of the database and have the "View any definition" permission. These settings are set by the database administrator when the empty database is created.

Service Account Permissions

The user context for the Stream and SOAP services requires the following database permissions:

- db_datareader
- db_datawriter
- execute permissions on stored procedures

The Configuration Wizard assigns these permissions, provided the user has securityadmin permissions.

In addition, the service user must have the following system privileges:

- Run as service
- Registry read access
- Program Files\Citrix\Provisioning Services
- Read/write access to any vDisk location

The Stream and SOAP services can run under one of the following supported user accounts:

- Network service account, which is a minimum privilege local account that authenticates on the network as computers domain machine account
- Specified user account (required when using a Windows Share), which is a workgroup or domain user account
- Local system account (for use with SAN)

Because authentication is not common in workgroup environments, minimal privilege user accounts must be created on each server, and each instance must have identical credentials.

Installing SQL Server and Provisioning Services on the same server can cause poor distribution during load balancing.

The security option you select for a Provisioning Services farm impacts Role-Based Administration and user groups. You can choose only one of the following options:

- Use Active Directory groups for security (default), which is selected if a Windows Domain running Active Directory - this option enables you to take advantage of Active Directory for Provisioning Services' roles
- Use Windows groups for security, which is selected if a single server or in a Workgroup - this option enables you to use the Local User/Groups on that particular server for Provisioning Services' roles

Sites

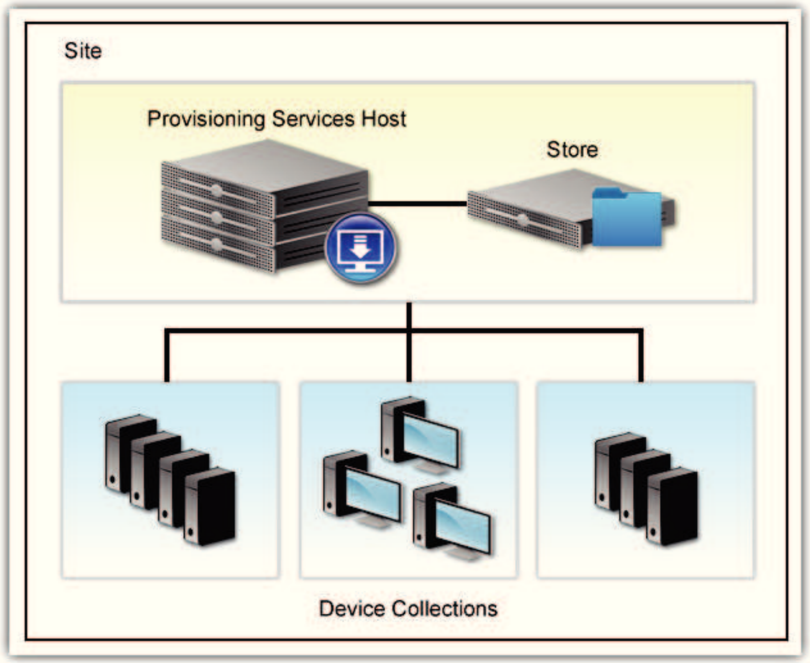

A Provisioning Services site provides both a site administrator and farm administrator with a method of representing and managing logical groupings of Provisioning Services hosts, device collections, and local shared storage.

You can create additional sites at any time, but it is helpful to plan how sites will be used to logically group Provisioning Services components and determine appropriate naming conventions before initial configuration.

Stores

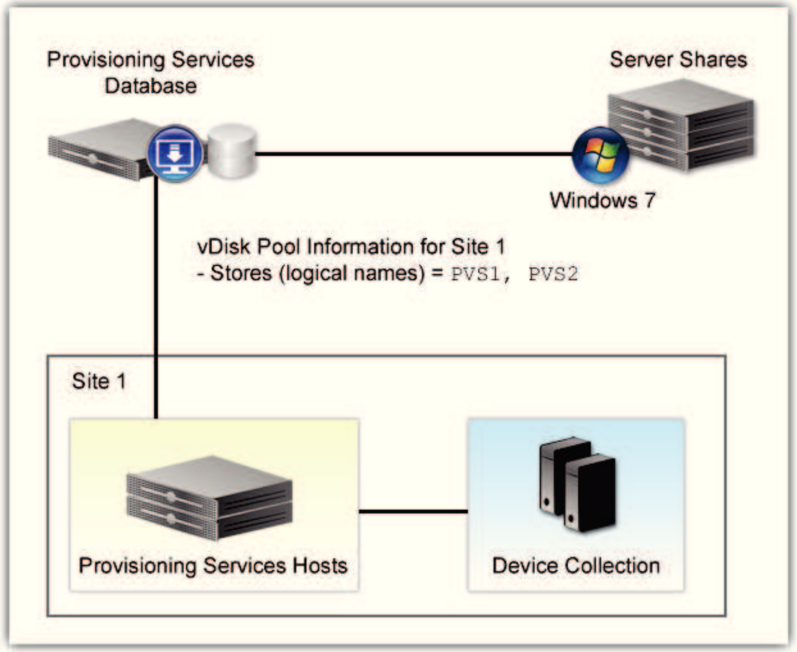
Provisioning Services Database
Server Shares
Windows 7
vDisk Pool Information for Site 1
- Stores (logical names) = PVS1, PVS2
Site 1
Provisioning Services Hosts
Device Collection

When vDisk files are created in the Console, they are assigned to a store. One or more Provisioning Services hosts within a site are given permission to access a store in order to serve vDisks to target devices.

Several different types of stores can be created based upon how configurations are made, but typical store configurations include:

Farm Store A farm store is available to all Provisioning Services hosts within a farm.

Site Store A site store is restricted to Provisioning Services hosts within a specified site.

Distributed Server Store

A distributed server store is comprised of vDisk storage locations hosted on several Provisioning Services hosts. Any vDisks that will have failover protection in this store must be manually copied to each Provisioning Services host. In this configuration, each Provisioning Services host is configured with a path to the storage location that overrides the path configured in the store. In this way, failover protection is shared among several Provisioning Services hosts.

A storage cluster file system is required when using SAN shared among multiple Provisioning Services hosts. Additionally, all servers need simultaneous read/write access to the SAN storage and database.

Single Server Store

A single server store can only be accessed by a single Provisioning Services host. Single server stores can be local drives on the Provisioning Services host or a SAN that has been configured as a local drive.

Storage Requirements

Hard disk size and free space are crucial to Provisioning Services performance if the Provisioning Services host is going to store write-cache files or vDisks, which can be very large in size. A RAID array, SAN, or NAS might improve streaming performance.

The hard disk space requirement varies depending upon the following options.

Static or Dynamic vDisks

Static vDisks require the allocation of a specific amount of hard disk space when the vDisk is created. The space allocated to a static vDisk cannot be changed once the vDisk has been created. A dynamic vDisk does not need allocated space because it expands as additional data is added to the vDisk file.

Standard or Private vDisks

Storage requirements will vary significantly depending upon the need for shared or private vDisks. An environment that requires a large number of private image vDisks will require significantly more storage space than an environment that can run target devices on relatively few standard image mode vDisks.

Write-Cache Storage Location

If the write-cache for a target device is stored on a Provisioning Services host, sufficient space must be available. Write-cache sizes vary depending upon how long the cache file has been left open between starts or restarts, so you should examine the needs of the environment when planning space for write-caches. Additionally, you should consider the number of target devices that will run simultaneously when configuring the write-cache.

vDisk Storage Location

The size of a vDisk varies depending upon the operating system and application stack in use, but they are typically large files that take up several GB at a minimum. You should ensure that the amount of hard drive space available on the desired storage location is sufficient to support existing vDisks and any planned future growth.

vDisk Backup Copies

You should maintain a backup copy of all production vDisks. Therefore, you should plan for enough free hard drive space to accommodate two copies of a vDisk.

Future Growth When planning space requirements for a static vDisk, you should plan to allow space for future growth, which could include additional applications and updates.

Sample vDisk Storage Requirements

The following table lists typical hard disk storage sizes for static vDisks running various operating systems and applications.

- All hard disk sizes listed are estimates.
- Additional space is required for vDisks based on the application stack.

Operating System	Hard Disk Size
Windows Server 2003	10-15 GB, based on system functionality
Windows Server 2008 R2	16-20 GB
Windows XP	2 GB for operating system
Windows Vista	15 GB for operating system
Windows 7	15 GB for operating system
Linux	5 GB for operating system
XenApp 5 (Windows Server 2008)	16-20 GB, based on system functionality
XenApp 6 (Windows Server 2008 R2)	16-20 GB, based on system functionality

Provisioning Services supports the use of several storage technologies.

For more information about the benefits and configurations of different storage technologies, see Citrix article CTX125126 on *support.citrix.com*.

Provisioning Services Installation

If you have time, show the students this video about Installing Citrix Provisioning Services, see Citrix TV at *www.citrix.com/tv/#videos/1605*.

Prior to beginning the installation process for Provisioning Services, it is important that you first install any Windows service packs, drivers, and updates.

A basic Provisioning Services implementation includes a single Provisioning Services host on which all server components are installed. This installation also includes various product utilities, such as the Configuration Wizard, the following components include and the Provisioning Services Console and the Stream Service.

However, complex environments require you to plan out the farm configuration and Provisioning Services installation carefully.

The Provisioning Services installation package includes the following components:

- Provisioning Services Console
- Stream Service
- Network boot services
- Product documentation
- Management Application Programming Interface (API)

> Installation of network services is optional. These services include DHCP, BOOTP, PXE, and TFTP.

For more information about installing Provisioning Services, see Citrix eDocs at *edocs.citrix.com*.

Provisioning Services Configuration

If you have time, show the students this video about Configuring Citrix Provisioning Services, see Citrix TV at *http://www.citrix.com/tv/#videos/1604*.

Provisioning Services can be configured using the Provisioning Services Configuration Wizard. The Configuration Wizard specifies settings such as the location of the DHCP and license server. The Configuration Wizard starts automatically after the Provisioning Services installation is completed and is available at any time from the Provisioning Services menu from the Start button.

> Running the Configuration Wizard restarts all services for Provisioning Services, which can be helpful when troubleshooting.

Farm Configuration

The Configuration Wizard allows you to configure the Provisioning Services farm by selecting to:

- Create a new farm or join an existing farm.
- Create a new database or use an existing database.
- Create a new site or join an existing site.

- Create a new device collection or use an existing device collection.

Bootstrap Server Configuration

When configuring the bootstrap server, select the appropriate options to enable for the Provisioning Services host:

Verbose Mode

Select the Verbose Mode option if you want to monitor the startup process on the target device (optional) or view system messages.

Interrupt Safe Mode

Select Interrupt Safe Mode if you are having trouble with your target device failing early in the startup process. This enables debugging of target device drivers that exhibit timing or start up behavior problems.

Advanced Memory Support

This setting enables the bootstrap to work with newer Windows operating system versions and enabled by default. Only disable this setting on older XP or Windows Server OS 32-bit versions that do not support PAE, or if your target device stops responding or behaves erratically in the early boot phase.

Network Recovery Method

There are two options for this method:

- Restore Network Connections: Selecting this option results in the target device attempting indefinitely to restore it is connection to the Provisioning Services host.
- Reboot to Hard Drive: Selecting this option instructs the target device to perform a hardware reset to force a restart after failing to re-establish communications for a defined number of seconds. The user determines the number of seconds to wait before restarting. Assuming the network connection cannot be established, PXE will fail and the system will restart to the local hard drive. The default number of seconds is 50, to be compatible with high availability configurations.

A hard drive must exist on the target device

Logon Polling Timeout

Enter the time, in milliseconds, between retries when polling for Provisioning Services hosts. Each Provisioning Services host is sent a log on request packet in sequence. The first Provisioning Services host that responds is used. In non-high availability configurations, this time-out simply defines how often to retry the single available Provisioning Services host with the initial log on request. This time-out defines how quickly the round-robin routine will switch from one Provisioning Services host to the next in trying to find an active Provisioning Services host. The valid range is from 1,000 to 60,000 milliseconds.

Logon General Timeout

Enter the time-out, in milliseconds, for all logon associated packets, except the initial logon polling timeout. This time-out is generally longer than the polling time-out, because the Provisioning Services host needs time to contact all associated servers, some of which might be down and will require retries and time-outs from the Provisioning Services host to the other Provisioning Services hosts to determine if they are indeed online or not. The valid range is from 1,000 to 60,000 milliseconds.

Store Configuration

Store Accessibility

If Server A has been configured to access Store A and resides in Site 1, all the vDisks that are located in Store A will appear in the vDisk pool for Site 1.

Store accessibility refers to the servers within a farm which are granted access to a store. Only specified servers are able to access vDisks within the store, regardless of whether a UNC path has been specified for the store. A store can be made accessible to multiple sites.

The type of store in use is selected largely on how the path to the vDisk storage location is specified. For example, farm or site stores must use a path that all servers can use to access the storage location, while private or distributed stores can use a path to a local hard drive. If vDisks are stored in a Windows file share, you should point the store to the UNC path of the share in order for several Provisioning Services hosts to have access to the store.

vDisks are either created or added to a store after the store has been configured. When a server has been given access to a store, the vDisks in that store appear in the vDisk pool for the site in which the server resides. vDisks that appear in the vDisk pool are available for assignment to any target devices in the site but can only be streamed by servers that have been given access in the Store properties.

Store Path

The file path provided for the storage location determines which Provisioning Services hosts within the farm are able to reach the store. For example, a local file path could direct a Provisioning Services host to look for vDisks within a folder stored locally or could indicate a SAN mapped to a local drive. UNC paths can be used by all Provisioning Services hosts within a farm provided that they can recognize the shared store, and the storage location has been configured to allow sharing.

Store Administration

Farm administrators can delegate store configuration to site administrators by configuring an optional site owner for the store. A store that does not have a specified site owner can be configured only by farm administrators.

Stores with specified site owners can be configured by both farm administrators as well as the site administrators within the designated site. Site owners are used primarily with site-based stores to allow delegated administration to site administrators.

Provisioning Services Hosts in the Console

In addition to streaming vDisks to target devices, Provisioning Services hosts also retrieve and provide configuration information to and from the Provisioning Services database. Provisioning Services host configuration options are available to ensure high availability and load-balancing of target device connections.

If time permits, show the video, *How to: Update a Pooled Desktop Group vDisk without Downtime*. It demonstrates how to update a Provisioning Services vDisk associated with a pooled desktop group without requiring end users to log off or prevent end users from accessing their desktop. *http://www.citrix.com/tv/#videos/1951*

If time permits, share the blog, *Provisioning Services or Machine Creation Services...Big Picture Matters*. This discusses using Machine Creation Services or Provisioning Services for single image desktop management. *http://virtualfeller.com/2011/02/15/provisioning-services-or-machine-creation-services%E2%80%A6-big-picture-matters/*

Target Device Connections

Target device connections to the Provisioning Services host are viewed and managed in the Console. The following tasks can be performed on one or more target devices.

- Shut down target devices that are highlighted in the dialog box. When selecting Shutdown or Reboot, a dialog box opens providing the option to type a message that displays on the affected devices.
- Restart target devices that are highlighted in the dialog box.
- Open the Edit Message dialog box to type and send a message to target device or devices highlighted in the dialog box.
- View a list of target devices that are currently connected to the host.

The Shutdown or Reboot options can be delayed by entering a delay time setting. If a message appears confirming that the target device was successfully turned off or restarted, but the icon in the Console window does not change accordingly, select the Refresh button.

Server Properties

You can modify Provisioning Services host configuration settings in the Provisioning Server Properties dialog box. To view the existing properties of a Provisioning Services host, choose one of the following methods:

- Highlight a Provisioning Services host, then select Properties from the Action menu.
- Right-click a Provisioning Services host, then select Properties.
- If the details pane is open, highlight a Provisioning Services host, then select the Properties menu item from the list of actions.

Provisioning Services displays a message if a change made on a Provisioning Server Properties dialog box requires that the server be restarted.

Test Your Knowledge: Provisioning Services Installation and Configuration

Indicate whether each statement is true or false.

Statement	True or False
You should install the Microsoft SQL database and Provisioning Services on separate servers.	True
As a general rule, you can plan to create a static vDisk that is 10% larger than the requirements for the operating system and application stack to allow for future growth.	False
Only one Provisioning Services database is associated with a farm.	True

Target Device Collection

A device collection can be used to simplify management by performing actions on collections, rather than on individual target devices.

A target device becomes a member of a device collection when it is added to the farm. A target device can only be a member in one device collection. However, a target device can exist in any number of views. If a target device is removed from the device collection, it is automatically removed from any associated views.

When a target device is added to a device collection, the device properties are stored in the Provisioning Services database. Target device properties include information such as the device name and description, start method, and vDisk assignments. You can use the Console to create a new device collection or to move a target device from one collection to another. The Console supports drag-and-drop functionality for this action.

In the Console, actions can be performed on:

- An individual target device
- All target devices within a collection
- All target devices within a view

Target Device Template

A target device can be set as the template for new target devices that are added to a device collection. This allows you to quickly add new devices to the device collection by using the template to imprint properties on the new device.

To set a target device as the template device for a collection, right-click the target device and select Set device as template.

The application of template properties is a one-time action. A new target device will not inherit changes made to the template target device after application.

Target Device Properties

Target device properties affect the performance of a Provisioning Services host environment. Target device settings can be updated in the Console by right-clicking a target device and selecting Properties.

Target device properties can be copied to one or more target devices in the Console by copying and pasting the properties to the appropriate target devices.

For more information about target devices, see Citrix eDocs at *edocs.citrix.com*.

Target Device Additions to the Database

You can create new target device entries in the Provisioning Services database by using one of the following methods:

- Using the Console to manually create target device entries
- Using the Auto-Add Wizard to create target device entries
- Importing target device entries

After the target device exists in the database, you can assign a vDisk to the target device.

Auto-Add Wizard

The Auto-Add Wizard automates the configuration of rules for automatically adding new target devices to the Provisioning Services database using the Auto-Add feature.

The Auto-Add Wizard can be started at the farm, site, collection, or device level. When started at a level lower than farm, the wizard uses that choice as the default choice. For example, if it is started on a particular target device, it will:

- Select the site for that device as the **Default Site** choice in the combo box.
- Select the collection for that device as the **Default Collection** choice in the combo box.
- Select that device as the **Template Device** choice in the combo box.

Each page is displayed with choices pre-selected based on the location that the Auto-Add Wizard was started from.

A farm administrator has the ability to turn Auto-Add on or off and to select the default Site.

A site administrator can only select the default site if appropriate permissions have been assigned. If the site administrator is not the administrator of the currently selected default site, then that administrator can only configure the sites they have access to.

Test Your Knowledge: Target Devices

1. What is the maximum number of target device collections that a target device can belong to?
   a. 1
   b. 2
   c. 3
   d. 4

   Answer: A
2. You need to add several target devices to a collection, but all target devices will have the same properties. How can you simplify the target device creation process?
   a. Designate a set up target device as the template for the collection.

b. Set up a script to automatically create target devices.

c. Import target devices from other target device collections.

d. Use the Auto-Add Wizard to add new target devices to a specific collection.

Answer: A

Active Directory Integration

Each target device that logs on to a domain requires a computer account in Active Directory. Target devices that access a vDisk in private image mode do not require any additional configurations to enable Domain Password management. However, Domain Password management must be configured on standard image mode vDisks to join target devices to a domain. Configuring Domain Password management ensures that target devices sharing the same vDisk image have unique domain accounts.

Demonstrate how to allow Provisioning Services to manage computer account password in Active Directory.

Active Directory Integration Prerequisites

The following prerequisites must be met prior to integration with Active Directory:

- Provisioning Services must be installed, configured, and running.
- The master target device that will be used to build the shared vDisk image for domain targets must be added to the Provisioning Services database.

Benefits of Active Directory Integration

Integrating Provisioning Services hosts and Active Directory allows you to:

- Select the Active Directory organizational unit in which the Provisioning Services host should create a target device computer account.
- Take advantage of Active Directory management features, such as delegation of control and group policies.
- Configure the Provisioning Services host to automatically manage the computer account passwords of target devices.

Domain Password Validation Process

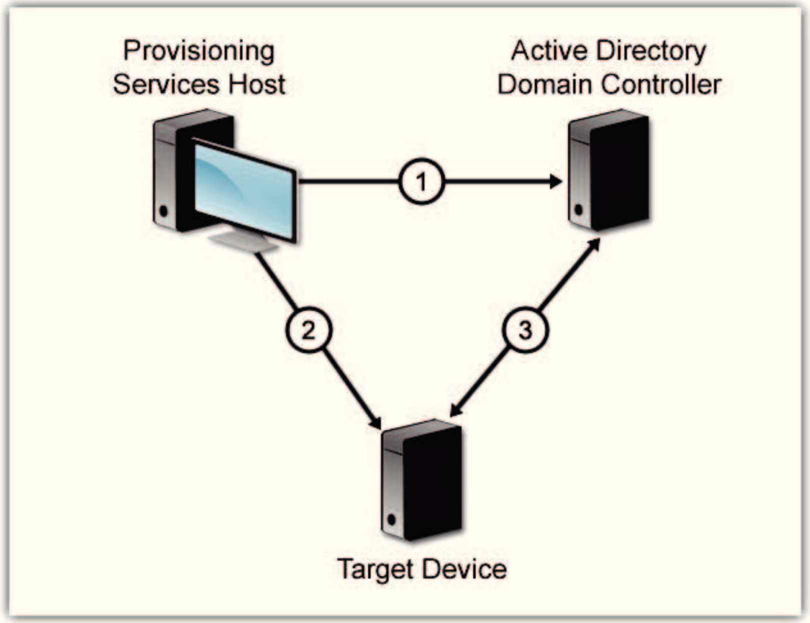
Provisioning Services Host
Active Directory Domain Controller
1
2
3
Target Device

1. An Active Directory account for a target device is created in the database.
2. The Stream service provides the account name to the target device.
3. The domain controller validates the password provided by the target device with the Active Directory password.

Automatic Password Renegotiation

Automatic password renegotiation can be configured at the domain level using the **Domain member: Disable machine account password changes** policy.

While target devices starting from vDisks no longer require Active Directory password renegotiation, configuring a policy to disable password changes at the domain level applies to any domain members starting from local hard drives. This might not be desirable. A better option is to disable machine account password changes at the local level. This can be done by selecting the **Optimize** option when building a vDisk image. The setting will then be applied to any target devices that start from the shared vDisk image.

To Integrate Active Directory

1. Verify that the vDisk file is in private image mode and assign the vDisk to the target device.
2. Set the target device to start from the vDisk.
3. Enable **Active Directory machine account password management** in vDisk properties.
   a. Open the Provisioning Services Console.
   b. Right-click a vDisk and select **File Properties**.
   c. Click the **Options** tab.
   d. Select **Active Directory machine account password management**.
   e. Click **OK** to close the vDisk file properties.
   f. Restart the Streaming Service.

Perform a demo showing students how to integrate Active Directory.

> This step has to be completed for each new vDisk that will stream to domain members. This option is disabled by default.

4. Enable **Automatic password support** in server properties.
5. Create the clean, golden image on the master target device hard drive.
6. Update the golden image with any necessary updates or drivers and prepare the system using sysprep.
7. Add the master target device to the domain.
8. Install the target device software.
9. Restart the target device and configure the BIOS to PXE boot.
10. Run the Image Optimization Wizard and verify that **Disable Machine Account Password Changes** is selected.
11. Build the image to the vDisk file, then shut down the target device.
12. Change the vDisk file to standard image mode.
13. Configure a new target device and assign the vDisk.
14. Create a machine account for the new target device using the Console or the `Run AddDeviceToDomain` command.
15. Turn on the new target device from the vDisk and log on to the domain.

To Reset Computer Accounts for Target Devices

1. Right-click one or more target devices in the Console window, then select **Active Directory Management**.
2. Select **Reset machine account**.
   The Active Directory Management dialog box appears.

Use this procedure if you did not disable password management.

3. In the target device table, highlight those target devices that should be reset, then click **Reset devices**.
4. Click **Close** to exit the dialog box.
5. Disable Windows Active Directory automatic password re-negotiation. To do this, on your domain controller, enable the following group policy: `Domain Member: Disable machine account password changes`.

> To make this security policy change, you must be logged on with sufficient permissions to add and change computer accounts in Active Directory. You have the option of disabling machine account password changes at the domain level or local level. If you disable machine account password changes at the domain level, the change applies to all members of the domain. If you change it at the local level (by changing the local security policy on a target device connected to the vDisk in Private image mode), the change applies only to the target devices using that vDisk.

6. Start each target device.

Test Your Knowledge: Active Directory

1. Which two prerequisites must you meet prior to integrating Provisioning Services with Active Directory? (Choose two.)
   a. Provisioning Services must be configured.
   b. The master target device must be added to the Provisioning Services database.
   c. Active Directory machine account password management must be enabled.
   d. The domain password must be set on the domain controller.

   Answers: A and B

2. Why must a Provisioning Services host manage the domain passwords for target devices that share a vDisk?
   a. To ensure that machine account password changes are disabled.
   b. In order to renegotiate the Active Directory password when the target device logs on.
   c. In order to be able to reset the target device password in case a target device is unable to log on.
   d. To ensure that the name and password assigned to the target device matches the computer account within the domain.

   Answer: D

Module 7

Managing vDisks and Target Devices

Overview

Proper vDisk setup is essential when a single vDisk is shared across multiple servers in order to avoid performance issues. It is important to know not only how to configure a vDisk, but also how different configuration settings affect the use and performance of a vDisk.

Timings:

Module: 180 minutes

Exercises: 95 minutes

Total Time: 275 minutes

Objectives

After completing this module, you will be able to:

- Determine the best vDisk image mode for a given scenario.
- Determine the best write cache location for a given scenario.
- Prepare for the creation of vDisk images by configuring a master target device.
- Deploy a vDisk image by assigning the vDisk image to a target device.
- Update and create a new version of a vDisk by using the Auto-update feature.
- Configure and test high availability for Provisioning Services to ensure server availability.
- Troubleshoot the availability of the vDisks by identifying common issues that can occur with the server-side streaming service and build process.
- Troubleshoot the availability of a Provisioning Services target device by identifying common pre-logon and logon issues.

vDisk Image Modes

Prior to class, if you are unfamiliar with RAM and storage requirements for Provisions Services hosts and target devices, please review the white paper, *Advanced Memory and Storage Considerations for Provisioning Services. http://support.citrix.com/article/CTX125126*

A vDisk is a file that contains a snapshot of the hard drive of a device, including the operating system.

A Provisioning Services vDisk can be configured as one of the following modes:

- Standard image mode
- Private image mode

Standard Image Mode

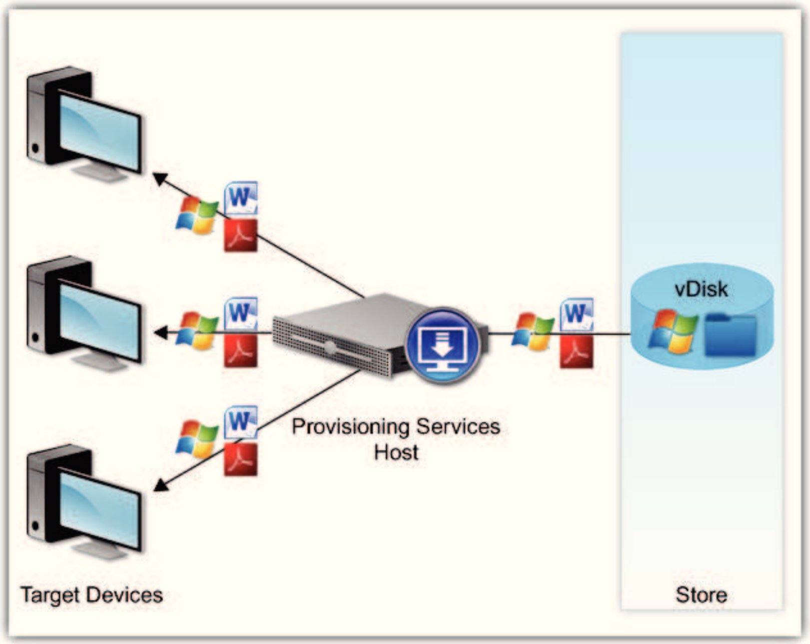

Standard image mode vDisks are read only, which allows multiple target devices to use a single vDisk at the same time. Any changes made by the target device are stored in a write-cache file for the duration of the session. Standard image mode is the most cost-effective mode, and it uses the least amount of disk space.

Private Image Mode

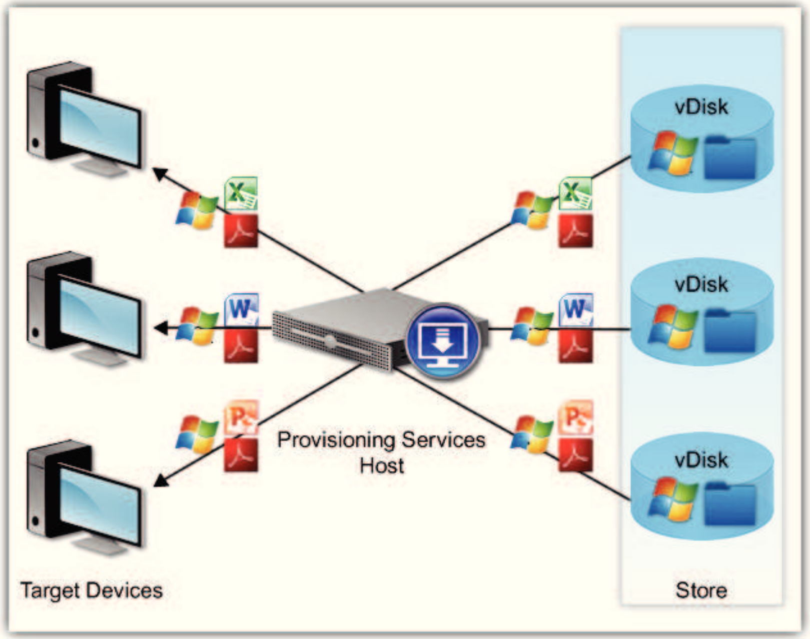

Private image mode closely models how a computer uses a regular hard drive by allowing only one target device to access a private image vDisk at a time. Provisioning Services performs read or write requests directly to the vDisk.

Write Caches

A write-cache file stores any writes that an operating system makes while a target device streams a vDisk. The size of a write cache varies depending upon the type of tasks and operations that are performed. For example, end users who perform repetitive tasks might only require a small cache size, while knowledge workers might require much larger cache sizes. You should consider both the size requirement of the write cache and the location where it will reside when designing your Provisioning Services implementation.

A write cache can be placed on shared storage and configured to use a UNC path by choosing the server disk option.

For more information about write-cache types, read the Citrix blog: *http://blogs.citrix.com/2011/10/06/pvs-write-cache-sizing-considerations/*

Write-Cache Types

Provisioning Services supports the following write-cache types:

- Cache on a Provisioning Services host
- Cache persistent on Provisioning Services host
- Cache on target device hard drive
- Cache in target device RAM

Write Cache Benefits and Considerations

The benefits and considerations for each cache location are listed in the following table.

Write Cache Location	Benefits	Considerations
Server disk	• Cache on Provisioning Services host disk allows for diskless target devices. • Cache size can be large.	Lowest performance: network utilization is high due to the amount of data requests to the Provisioning Services host.
Device hard drive	Good performance: network utilization is reduced the longer the target device is running because more data is stored in the cache and fewer requests are sent to the Provisioning Services host.	Cache size limit: if the cache exceeds the limit, the device will fail.

Write Cache Location	Benefits	Considerations
Device RAM	Best performance: accessing data from the target device RAM is faster than accessing data from the target device disk.	Small cache size: RAM cannot hold significant amounts of data. If cache exceeds the limit, errors can occur.

Cache on Server Disk

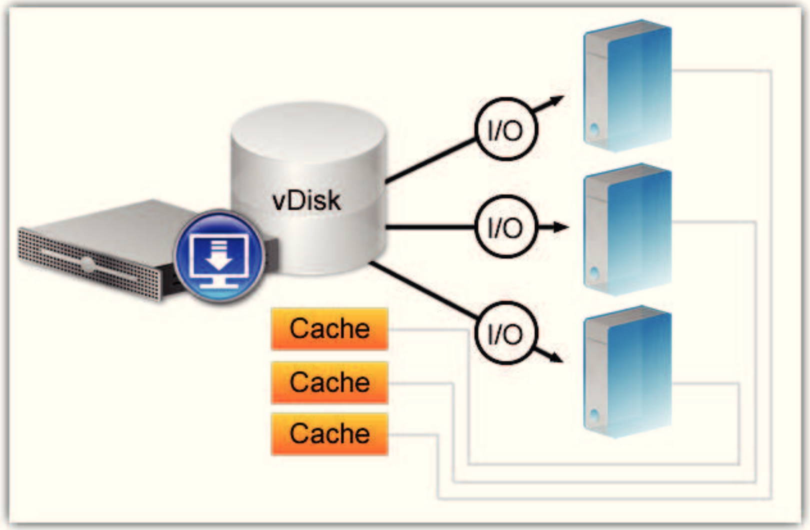

If the write cache is located on the server, all changes made to the vDisk image during a session are stored as a temporary file on the Provisioning Services host. The Provisioning Services host handles all writes in this configuration, which can increase disk I/O and network traffic.

The Provisioning Services host can be configured to encrypt write-cache files for additional security. The data will be encrypted in the event a hard drive is stolen because of the presence of the write-cache file on the hard drive.

When the cache on server disk option is selected, you can choose to store the write cache in one of the following locations:

- Local storage on the Provisioning Services host
- Shared storage attached to the Provisioning Services host

Cache on Server Disk: Local Storage

To configure cache on server disk for local storage, you should place the write cache on the physical disks of the Provisioning Services host.

Benefits	Considerations
• Simplest option to set up • No additional resources or configuration within the environment required • Inexpensive disk space	• Performance could be impacted due to requests to/from the write cache traversing the network between the target device and Provisioning Services host. • Provisioning Services host scalability is reduced because the Stream Service must also service the write-cache requests. • Provisioning Services high availability is unavailable because the write-cache storage is not accessible by other Provisioning Services hosts. This can be mitigated by implementing a third-party cluster file system. • Provisioning Services host will fail if the local storage space is exceeded.

Cache on Server Disk: Shared Storage

To configure the cache on server disk for shared storage, you should place the write cache on shared storage that is connected to the Provisioning Services host.

Benefits	Considerations
• Provisioning Services high availability is possible because all Provisioning Services hosts attached to shared storage can access the write cache. • Shared storage devices typically hold a large amount of data, which mitigates storage size concerns.	• Network congestion could impact performance because requests traverse the network twice. • Provisioning Services host scalability is reduced because the Stream Service must also service the write-cache requests. • Setup and configuration of a robust shared storage solution is required, if one is not already in place.

Cache on Target Device Hard Drive

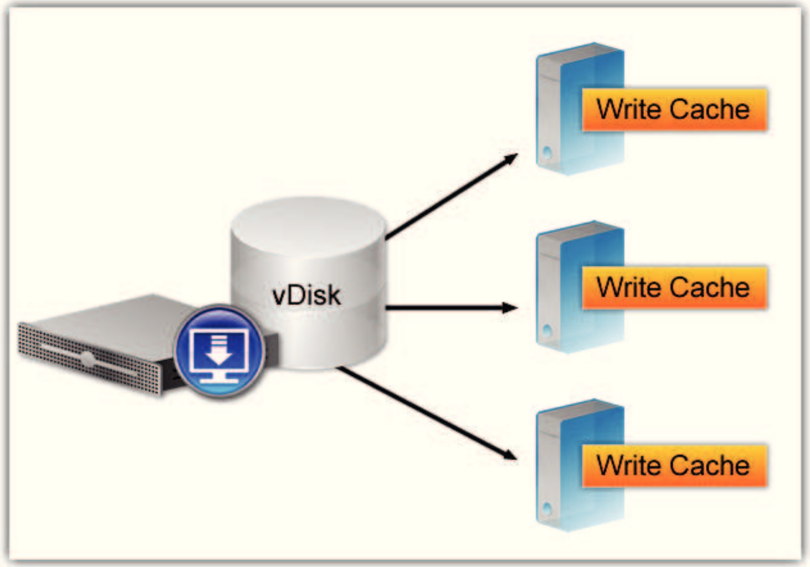

If the write cache is located on the hard disk of a target device, all changes made to the vDisk image during a session are stored as a temporary file on the hard drive of the target device. The target device hard disk does not require additional software to support this write-cache configuration. Storing the disk cache on the target device allows Provisioning Services hosts to use processing resources for other critical tasks.

When the cache on device hard drive option is selected, you can choose to store the write cache in one of the following locations:

- Local storage on the target device
- Shared storage attached to the target device

Cache on Target Device Hard Drive: Local Storage

To configure a cache on device hard drive for local storage, place a write cache on the physical disks of the target device. The local storage can be either a physical or virtual disk drive. This type of write cache is used mainly with physical target devices.

Benefits	Considerations
• Additional resources are not required if local disks are installed and unused on physical target devices. • Response times are fast because the read/write to and from the write cache is performed locally. • Local storage typically provides more than enough space for the write cache, minimizing risk of underestimating disk requirements. • Network I/O is reduced, which increases scalability on the Provisioning Services host.	• Live migration is not possible if the write cache is stored on a virtual infrastructure server local hard drive on virtual target devices. In this configuration the storage is not shared among virtual infrastructure servers. • Local storage configuration is slower than target device RAM cache. • Determining the size of the write cache is critical to prevent server failure.

Cache on Target Device Hard Drive: Shared Storage

To configure cache on device hard drive for shared storage, place the write cache on shared storage that is connected to the target device. This type of write cache is usually only valid in environments that use virtual target devices, such as those with Citrix XenServer. The storage is assigned to each virtual machine from a shared storage repository.

Benefits	Considerations
• Response times are faster. • Storage costs are significantly cheaper than purchasing RAM. • Live migration is possible because the target device cache storage is accessible from multiple virtual machines.	• This method is slower than target device RAM or local disk cache. • Setup and configuration of a shared storage solution is required if one is not already in place.

Target Device-based RAM Cache

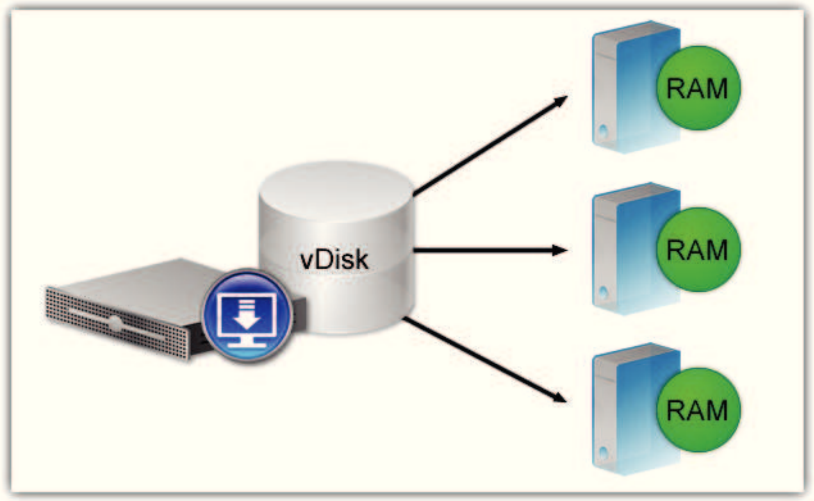

If the write cache is located in the target device RAM, all changes made to the vDisk image during a session are stored as a temporary file in the target device RAM. A portion of RAM is reserved for Provisioning Services use and any remaining RAM is available for the operating system.

Target Device RAM

No further configuration is necessary once the cache on device RAM write-cache option is selected.

Benefits	Considerations
This type of write cache is fastest.	• RAM is diverted from workload use. • The cost is greater than the cost of using storage. • Determining the amount of RAM required for the write cache is difficult yet critical to the stability of the environment. • Target device fails when the allocated write-cache space reaches capacity.

Explain that this was formerly known as Differencing Disk.

Cache Persistent on Server

The cache persistent on server option provides the ability to save changes between restarts. After restarting, a target device is able to retrieve changes made from previous sessions that differ from the read-only vDisk image. If a vDisk is set to cache persistent on server, each target device that accesses the vDisk automatically has a device-specific, writable disk file created. Any changes made to the vDisk image are written to that file, which is not automatically deleted upon shutdown.

The file name uniquely identifies the target device by including the MAC address and disk identifier of the target device. A target device can be assigned to multiple vDisks and therefore multiple cache files can be associated to it.

To restore a vDisk that uses cache persistent on server, be sure to back up all vDisk files and associated end user cache files prior to making any vDisk modifications.

Benefits	Considerations
• Saves target device-specific changes that are made to the vDisk image • Provides the same benefits as standard image mode	The cache file is saved provided that the file remains valid. Any changes made to the vDisk force the cache file to be automatically deleted. For example, if the vDisk is set to private image mode, all associated cache files are deleted. Invalidating changes include: • vDisk is placed in maintenance. • vDisk mode is changed to private image mode. • The drive is mapped from the Console. • The location of the write-cache file is changed. • Automatic updates are used.

Test Your Knowledge: vDisk Image Modes and Write Cache

Match the following terms with the correct descriptions.

- Private image mode
- Standard image mode
- Cache on server disk
- Cache on device hard drive
- Target device-based RAM cache

Description	Term
Allows for multiple target devices to use a single vDisk.	Standard image mode
The Provisioning Services host handles all writes in this configuration.	Cache on server disk
End users are able to personalize their desktops and all applications.	Private image mode
This configuration provides the fastest access to the write cache.	Target device-based RAM cache
All changes made to a vDisk image during a session are stored as a temporary file on the target device hard drive.	Cache on device hard drive

Microsoft Licensing for Provisioning Services

Provisioning Services supports Microsoft Key Management Service (KMS) or Multiple Activation Key (MAK) volume licensing.

Configuring Microsoft KMS Volume Licensing

Microsoft provides two mechanisms for administering volume licenses.

KMS volume licensing uses a centralized activation server that runs in the datacenter and serves as a local activation point. The tasks involved in configuring a vDisk image to use KMS volume licensing and managing that vDisk in a Provisioning Services farm include:

- Enabling KMS licensing on the vDisk being created. This is done by selecting the KMS menu option on the Microsoft Volume Licensing tab when running the Imaging Wizard.
- Preparing the new base vDisk image for KMS volume licensing. This is done by using the rearm command to reset the vDisk to a non-activated state. This operation must be performed on a vDisk in private image mode.
- Maintaining or upgrading a vDisk image that uses KMS volume licensing. This should be done from the Master Target Device and the original Provisioning Services host.

For more information about configuring Microsoft KMS volume licensing, see Citrix eDocs at *edocs.citrix.com*.

It might take a few minutes for the KMS licensing to be activated. Until then, the end user might see an error that the license is not authentic.

Microsoft MAK Volume Licensing Support

Another mechanism for administering Microsoft volume licenses is called Multiple Activation Keys (MAKs). A MAK corresponds to a certain number of purchased operating system (OS) licenses. The MAK is entered during the installation of the OS on each system, which activates the OS and decrements the count of purchased licenses centrally with Microsoft. Alternatively, a process of 'proxy activation' is done using the Volume Activation Management Toolkit (VAMT). This allows activation of systems that do not have network access to the internet. Provisioning Services uses this proxy activation mechanism for standard image mode vDisks that have MAK licensing mode selected when the vDisk is created.

In order for MAK licensing to work, the VAMT must be installed on all Provisioning Services hosts within a farm.

For more information about configuring Microsoft MAK volume licensing, see Citrix eDocs at *edocs.citrix.com*.

Master Target Device

A target device with an operating system which is imaged to create a vDisk is called a master target device. As a benefit of using Provisioning Services, you can manage a single vDisk rather than an individual workstation; however, for this reason, the initial image must be prepared properly. A pristine vDisk image created from a master target device can also be called a golden image.

A master target device can be either a virtual or physical machine.

Preparing a Master Target Device

Provisioning Services streams the contents of a vDisk created from the master target device to other target devices.

Four steps are needed to prepare a master target device:

- Preparing the hard disk of the master target device
- Configuring the BIOS of the master target device
- Configuring the Network Adapter BIOS
- Installing the master target device software

Preparing the Hard Disk of the Master Target Device

The master target device is typically different from subsequent target devices because it initially contains a hard disk. This is the hard disk that will be imaged to the vDisk. If necessary, after imaging, the hard disk can be removed from the master target device.

To support a single vDisk that is shared by multiple target devices, those devices must have certain similarities to ensure that the operating system has all required drivers. The three key components that must be consistent include the:

- Motherboard
- Network card, which must support PXE
- Video card

However, the Provisioning Services Common Image Utility allows a single vDisk to simultaneously support different motherboards, network cards, video cards, and other hardware devices.

If target devices will be sharing a vDisk, the master target device serves as a template for all subsequent diskless target devices as they are added to the network. It is crucial that the hard disk of the master target device be prepared properly and all software is installed on it in the following order:

1. Windows operating system
2. Device drivers
3. Service pack updates

4. Target device software

If target devices will be members of a domain, and will share a vDisk, additional configuration steps must be completed.

For more information about managing domain accounts, see Citrix eDocs at *edocs.citrix.com*.

Configuring the BIOS of a Master Target Device

The following steps describe how to configure the BIOS and BIOS extension provided by the network adapter of a target devices system to start from the network. Different systems have different BIOS setup interfaces. If necessary, consult the documentation that came with your system for further information on configuring these options.

To Configure the BIOS of a Master Target Device

1. If the target device BIOS has not yet been configured, restart the target device and enter the system's BIOS setup.
2. Set the network adapter to **On with PXE**.

> Depending on the system vendor, this setting might appear differently.

3. Configure the target device to start from **LAN** or **Network first**. Optionally, select the Universal Network Driver Interface; select **UNDI first** if using a NIC with Managed Boot Agent (MBA) support.
4. Save changes, then exit the BIOS setup program.
5. Start the target device from its hard drive over the network to attach the vDisk.

Installing Master Target Device Software

Provisioning Services target device software must be installed on a master target device prior to building a vDisk image using the installation wizard.

Provisioning Services target device software components include:

Provisioning Services Virtual Disk The Provisioning Services Virtual Disk is the virtual media used to the disk components of the operating system and applications.

Provisioning Services Network Stack
The Provisioning Services Network Stack is a proprietary filter driver that is loaded over the NIC driver, allowing communications between the target devices and the Provisioning Services host.

Provisioning Services SCSI Miniport Virtual Adapter
The Provisioning Services SCSI Miniport Virtual Adapter allows the vDisk to be mounted to the operating system on the target device.

Provisioning Services Imaging Wizard
The Provisioning Services Imaging Wizard is used to create the vDisk file and image the master target device.

Virtual Disk Status Tray Utility
The Virtual Disk Status Tray Utility provides general vDisk status and statistical information. This utility includes a help system.

Target Device Optimizer Utility
The Target Device Optimizer Utility is used to change target device settings to improve performance.

> Before installing the product software on a master target device, turn off any BIOS-based virus protection features. To include antivirus software on the vDisk image, be sure to turn the antivirus software back on prior to running the Imaging Wizard.

Imaging a Windows Target Device with XenConvert

As a physical-to-virtual conversion tool, XenConvert can convert a server or desktop workload from an online physical machine running Windows to a XenServer virtual machine or Provisioning Services vDisk.

1. Start XenConvert.
2. Select the location that will be used to create the vDisk image.
3. Select the destination where the vDisk image will be stored.
4. Select the volumes that should be included in the vDisk image.
5. Configure the amount of free space to include on the vDisk image.
6. Configure a log of the files copied during the image build.
7. Optimize the vDisk for Provisioning Services if the vDisk will be used in standard image mode.
8. Build the vDisk image.

Test Your Knowledge: Master Target Device

1. Which four steps must you complete to create a master target device? (Choose four.)
   a. Configuring the Network Adapter BIOS.
   b. Install the master target device software.
   c. Configure the BIOS of the master target device.
   d. Prepare the hard disk of the master target device.
   e. Attach the golden image to the master target device.
   f. Install all applications before installing the master target device software.

   Answer: A, B, C, D

2. In which order must you install software on the master target device hard disk to ensure it functions correctly?
   a. Windows operating system, service pack updates, device drivers, and target device software.
   b. Windows operating system, target device software, service pack updates, and device drivers.
   c. Windows operating system, service pack updates, target device software, and device drivers.
   d. Windows operating system, device drivers, service pack updates, and target device software.

   Answer: D

3. Which Provisioning Services target device software component provides general vDisk status and statistical information?
   a. Provisioning Services Virtual Disk
   b. Provisioning Services Wizard
   c. Virtual Disk Status Tray Utility
   d. Target Device Optimizer Utility

   Answer: C

vDisk File Creation

A vDisk file is the file in which an operating system image taken from a master target device is stored. All vDisk files created with Provisioning Services use Microsoft's Virtual Hard Disk (VHD) format. When a vDisk is created, two files are created within the designated storage location: a `.vhd` file that contains the vDisk image and a `.pvp` file that contains property configurations for the vDisk. These files are automatically named with the same file name as the vDisk.

Properties File

The `.pvp` and `.vhd` files must always be stored in the same directory. If a vDisk is moved to a different location, the corresponding `.pvp` file must also be moved to the same folder as that of the vDisk. If you want to duplicate an existing vDisk, you must also create a copy of the existing `.pvp` file.

If the `.pvp` file is deleted, missing, or becomes corrupt, Provisioning Services will automatically generate a new file. However, the new file will contain default vDisk settings and all previously configured settings will be lost. The `.pvp` file should be backed up as part of the normal Provisioning Services backup process.

vDisk Lifecycle Operations

vDisks are managed throughout the vDisk lifecycle. Provisioning Services provides support for a full image lifecycle that takes a vDisk from initial creation, through deployment and subsequent updates, and finally to retirement. The lifecycle of a vDisk consists of four stages:

1. Creating
2. Deploying
3. Updating
4. Retiring

Creating a vDisk

Creation of a vDisk requires preparing the master target device for imaging, creating, and configuring a vDisk file where the vDisk will reside, and then imaging the master target device to that file, resulting in a new base vDisk image. This process can be performed automatically, using the Imaging Wizard, or manually. Provisioning Services also provides the option to create a common image for use with a single target platform or for use with multiple target platforms.

Deploying a vDisk

After a vDisk base image is created, it is deployed by assigning it to one or more devices. A device can have multiple vDisk assignments. When the device starts, it starts from an assigned vDisk.

There are two mode options; Private image mode (single device access, read/write) and Standard image mode (multiple device access, write cache options).

Updating a vDisk

It is often necessary to update an existing vDisk so that the image contains the most current software and updates. Updates can be made manually, or the update process can be automated using vDisk Update Management features. Each time a vDisk is updated a new version is created. Different devices can access different versions based on the type of target device and version classification.

Retiring a vDisk

Retiring a vDisk is the same as deleting. The entire VHD chain including differencing and base image files, properties files, and lock files are deleted.

VHD Formats

When creating a vDisk you must select a VHD format—fixed or dynamic.

Fixed

Using a fixed VHD format allocates a specified amount of space to a vDisk file that cannot be changed once it has been configured. This allotted space must be large enough to hold the operating system, any required applications, and any applications that might be installed at a later time. If you format the vDisk to use NTFS, the size limit is approximately 2 TB. The limit is 4095 MB if you format the vDisk to use a FAT32 file system.

> While you cannot change the size of a fixed vDisk file using the Console, several third-party tools are available for expanding fixed VHD vDisks. If a third-party tool is not used, a vDisk can be expanded using the reverse imaging process. For more information about expanding fixed VHD vDisks, see Citrix article CTX124792 on *support.citrix.com*.

Dynamic	Using a dynamic VHD format allows a vDisk file to expand as changes to the vDisk are made. The file size associated with a dynamic vDisk is the maximum size that the file will be allowed to reach. Dynamic vDisks make planning for vDisk sizes easier, particularly for private image mode vDisks, which can grow at their own pace as end users install applications and add data. The dynamic VHD format adds overhead to the disk-write process. As a result, the write process takes longer.

Create vDisks Automatically Using Imaging Wizard

When using the Imaging Wizard to automatically create the base vDisk image from a master target device, you must:

- Enable Windows Automount on Windows Server operating systems.
- Disable Windows Autoplay.
- Verify that adequate free space exists in the vDisk store, which is approximately 101% of used space on the source volumes.
- Make note of which NIC the master target device was bound to when the Provisioning Services software was installed on the target device. This information is necessary during the imaging process.

To Create a New vDisk Automatically Using the Imaging Wizard

1. Select **Citrix** > **Provisioning Services** > **Imaging Wizard** from the Windows Start menu of the master target device.
   The wizard's Welcome page appears.
2. Click **Next.**
   The Connect to Farm page appears.
3. Type the name or IP address of a Provisioning Services host within the farm to connect to and the port to use to make that connection.
4. Use the Windows credentials, or enter different credentials, then click **Next**. If using Active Directory, enter the appropriate password information.
5. Select the volume license option to use for target devices or select **None** if volume licensing is not being used.
6. Select to create a new vDisk, or use an existing vDisk by entering that vDisks name, then click **Next.**
   The Add Target Device page appears.

7. Select the target device name, the collection to add this device to, and the MAC address associated with one of the NICs that was selected when the target device software was installed on the master target device. Click **Next**.
   If the target device is already a member of the farm, the Existing Target Devices page appears.
8. Click **Next**.
   The Summary of Farm Changes appears.
9. Verify all changes, then click **Next.**
   A confirmation message displays.
10. Click **Yes** on the confirmation message to start the imaging process.

Creating a vDisk Manually

1. Create a vDisk file and provide the following information:
   a. Site that will contain the vDisk
   b. Store where the vDisk will reside
   c. File name
   d. Description
   e. Size
   f. VHD format
2. Mount the vDisk (if formatting from the Provisioning Services host).
3. Assign the vDisk to the master target device.
4. Format the vDisk and provide the following information:
   a. File system type
   b. Volume label
   c. Format options
5. Build the vDisk image.
6. Unmount the vDisk (if formatting from the Provisioning Services host).

vDisk Management

A vDisk acts as a hard disk for a target device. You should consider the following information when creating a vDisk image file:

- For large implementations with many target devices, spreading the I/O across multiple disks can increase efficiency.
- The number of vDisk image files that can be created is unlimited. The only constraint is the space available on the Provisioning Services host, or on the storage device containing the vDisk image files.
- vDisk files use FAT or NTFS file systems. EXT2 and EXT3 can be used for Linux.
- Depending upon the file system used to store the vDisk, the maximum size of a vDisk is 2 TB (NTFS) or 4096 MB (FAT).
- A vDisk can be shared (Standard Image) by one or more target devices, or it can exist for only one target device to access (Private Image).
- vDisks can be started directly from a Windows Virtual Server or Hyper-V without needing to stream to a target device.
- The vDisk image is created using the Imaging Wizard Utility, and the vDisk file is created and configured using the Console.

vDisks in the Console

In the Console, a new vDisk can be created by right-clicking the vDisk Pool or the Store and then selecting the Create new vDisk menu option. vDisks are displayed in the details pane when a site vDisk pool is selected or when a store in the farm is selected.

The administrator role determines which displays and which tasks you can perform in the Console. For example, you can view and manage vDisks in sites in which you are a site administrator. However, unless the farm administrator sets a site as the owner of a store, the site administrator cannot perform store management tasks.

Assigning vDisks to Target Devices

A vDisk can be assigned to a single target device or to all devices within a target device collection. If more than one vDisk is assigned to a target device, a list of vDisks displays when the target device starts, allowing the end user to select the appropriate vDisk to start.

> If one or more versions exist for a vDisk, the version target devices use in production is either the highest numbered production version or an override version.

vDisks can be assigned to a single target device using:

- Drag-and-drop

- Target Device Properties dialog box

To Assign a vDisk to a Single Target Device

1. Expand the Device Collections folder in the Console tree, then click the collection folder where the target device is a member.
   The target device displays in the details pane.
2. Right-click the target device, then select **Properties**.
   The Target Device Properties dialog box appears.
3. Select the startup method that this target device should use from the **Boot from** drop-down menu options on the General tab.
4. Select the **Add** button within the **vDisk for this Device** section of the vDisks tab.
   The Assign vDisks dialog box appears.
5. Select a specific store or server under the **Filter** options to locate vDisks to assign to the target device, or accept the default settings, which include **All Stores** and **All Servers**.
6. Highlight the vDisk to assign in the **Select the desired vDisks** list, then click **OK**, then **OK** again to close the Target Device Properties dialog box.

vDisk Versions

Versioning simplifies vDisk update and management tasks, providing a more flexible and robust approach to managing vDisks.

A vDisk consists of a Virtual Hard Disk (VHD) base image file, any associated side-car files, and if applicable, a chain of referenced VHD differencing disks. Differencing disks are created to capture the changes made to the base disk image, leaving the original base disk unchanged.

> vDisk Versioning
>
> A new version of a vDisk is created each time a vDisk is placed in Maintenance and changes are made to the base disk. The base disk is represented by version 0. Each subsequent disk will have an incrementing version number. For example, a base disk might be named XYZ.VHD. A subsequent update would lead to a new version of the vDisk named XYZ.1.AVHD.

vDisk Backup

The Provisioning Services host treats a vDisk image file like a regular file, but the target device treats it as a hard drive. The procedure for backing up a vDisk image file is the same as backing up any other file on your server. If a vDisk image file becomes corrupt, restoring it requires simply replacing the corrupted file with a previous, functional version.

Do not back up a vDisk while it is in use or while it is locked. Integrate the backup of vDisks into your normal Provisioning Services host backup routine.

vDisk Updates

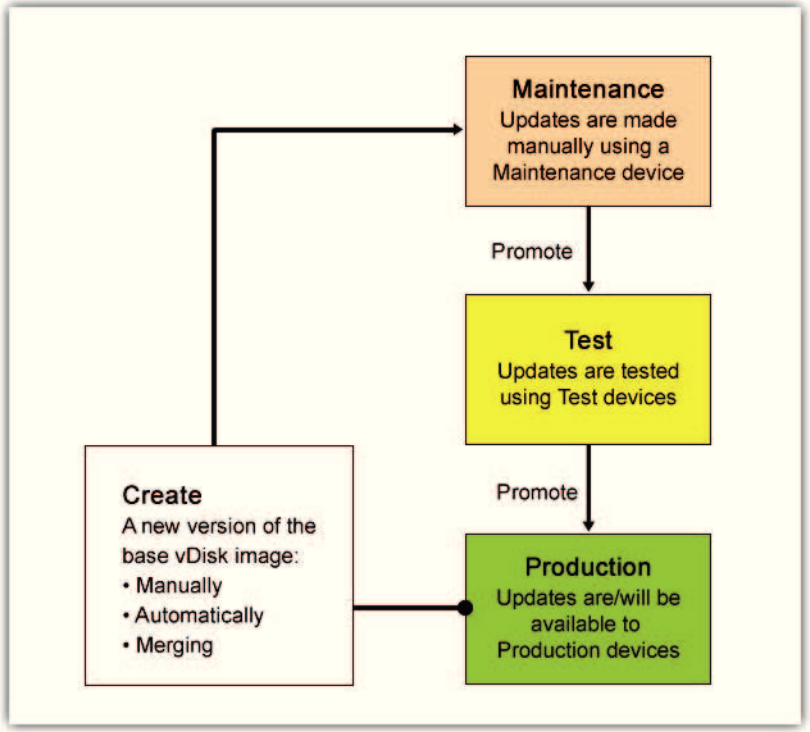
Maintenance
Updates are made manually using a Maintenance device
Promote
Test
Updates are tested using Test devices
Promote
Production
Updates are/will be available to Production devices
Create
A new version of the base vDisk image:
• Manually
• Automatically
• Merging

It is often necessary to update an existing vDisk so that the image contains the most current software and updates. Each time the vDisk is to be updated, a new version of that vDisk is created (VHD file) to capture the changes without changing the base vDisk image.

Updating a vDisk involves:

- Creating a new version of the vDisk, manually or automatically.
- Starting the newly created version from a device, make and save changes to the vDisk, then shut down the device.
- Promoting the new version to Production.

Below are the vDisk update scenarios that are supported.

Manual Update

You can choose to update a vDisk manually by creating a new version of that vDisk and then using a Maintenance device to capture updates to that version. Manual updates are initiated by selecting the New button. The Access column on the vDisk Versioning dialog box displays that the newly created version is currently under maintenance. While under maintenance, this version can only be accessed and updated by a single Maintenance device. Multiple Maintenance devices can be assigned to a vDisk. However, only one device can start and access that version of the vDisk at any given time. During that time that Maintenance device will have exclusive read/write access.

Automated Update

Creating automated updates saves administration time and physical resources. Updates are initiated on demand or from a schedule and are configured using vDisk Update Management. If updating automatically, the Access column on the vDisk Versioning dialog box displays that the newly created version is currently under maintenance. While under maintenance, this version can only be accessed and updated by the one Update device to which it is assigned (only one Update Device exists for each vDisk).

> vDisk Update Management is intended for use with standard image mode vDisks only. Private image mode vDisks can be updated using normal software distribution tool procedures. Attempting to register a private image mode vDisk for vDisk update management, or switching a vDisk that is already registered, will cause errors.

Merge

Merging VHD differencing disk files can save disk space and increase performance, depending on the merge option selected. A merge update is initiated manually by selecting the Merge button, or automatically when the maximum vDisk versions count is reached.

Autoupdate Tool

In the Console, the vDisk Update Management feature is used to configure the automation of vDisk updates using virtual machines. Automated vDisk updates can occur on a scheduled basis, or at any time that you invoke the update directly from the Console. This feature supports updates detected and delivered from Windows Server Update Services (WSUS) and System Center Configuration Manager (SCCM) Electronic Software Delivery (ESD) servers.

When the Site node is expanded in the Console tree, the vDisk Update Management feature appears. When expanded, the vDisk Update Management feature includes the following managed components:

- Hosts
- vDisks
- Tasks

> For more information about enabling automatic vDisk updates, configuring virtual host connections for automated vDisk updates, creating and configuring ESD update VMs, and configuring managed vDisk for automated updates, see Citrix eDocs at *http://edocs.citrix.com*.

vDisk Update Management Requirements

vDisk Update Management requires completing the following high-level tasks:

1. Designate a Provisioning Services host within the site to process updates.
2. Configure a Virtual Host Pool for Automated vDisk updates.
3. Create and configure ESD virtual machine that will be used to update the vDisk.
4. Configure the vDisks for automated updates.
5. Create and manage update tasks.
6. Run the update task by right-clicking on the task object in the Console, and then select the **Run update now** menu option.

After vDisk Update Management is configured, managed vDisks can be updated using the following methods:

- Scheduled: the Image Update Service automatically updates a vDisk, on a scheduled basis as defined in the Update Task.
- User-Invoked: you can select a managed vDisk to be updated from the Console.

The Update virtual machine will start, install updates, and restart as necessary. After the update task successfully completes, the virtual machine is automatically shut down. The update status can be checked from the Console tree under **vDisk Update Management** > **vDisks** > **vDisks** > **vDisk name** > **Completed Update Status**. The status can also be checked using the event viewer or in WSUS.

To Install Updates Automatically

1. Under the vDisk Update Management node in the Console tree, right-click vDisks, then select the **Add vDisks** option.
   The Managed vDisk Setup Wizard Welcome page appears.
2. Click **Next** to begin.
   The vDisk page appears.
3. Select the default search options or use the filtering options to display the vDisks to be managed. vDisks that are not already managed will display in the vDisk selection box.
4. Select one or more vDisks to be managed, then click Next.
5. Select the type of connection to use when hosting the virtual machine.
6. Select the virtual machine device to use to process the vDisk update from the drop-down list.
7. Click **Next**.
   The Active Directory page appears.
8. If using Active Directory, enter a Domain and Organizational Unit to create an Active Directory machine account that will be used by the Update Device that is created exclusively for updating the vDisk, then click **Next**.
   The Confirmation page appears.
9. Review all settings, then click **Finish.**

Incremental Update Rollback

Provisioning Services automatically creates a special rollback file when an incremental vDisk update occurs. The rollback file is a delta file - a file that contains the changes from one point to the next - that is used to reverse the update process and revert the new vDisk image to the original image. The vDisk update process creates a subfolder in the vDisks folder, named Rollback, where the original vDisk resides. The rollback file is given the same name specified for the delta file with an `.rbk` extension appended to the filename.

For more information about rolling back a vDisk update, see Citrix article CTX124791 on *support.citrix.com*.

> Each incremental update must be rolled back individually if several incremental updates have been applied to a vDisk. For example, if you apply three incremental updates to a vDisk and you want the vDisk state to return to the original, each of the three updates must be rolled back beginning with the most recent and working backward sequentially.

vDisk Replication

Provisioning Services allows you to safely use replication solutions like Microsoft Distributed File System (DFS) Replication to distribute vDisks across multiple servers or geographic locations. Provisioning Services hosts will maintain an inventory of available vDisks and versions and adjust load balancing as needed to ensure sessions are only assigned to servers that have access to the required vDisk version.

For more information about using Microsoft DFS Replication with Provisioning Services, read the Citrix blog:*http://blogs.citrix.com/2010/06/25/using-microsofts-dfs-replication-with-provisioning-services-ha/*

vDisk Inventory Service

The vDisk inventory service keeps track of every vDisk version that is found in the file system. In addition, the service allows you to choose between four different replication methods:

- Microsoft DFS
- PeerSync
- Robocopy
- Scripts

High Availability Overview

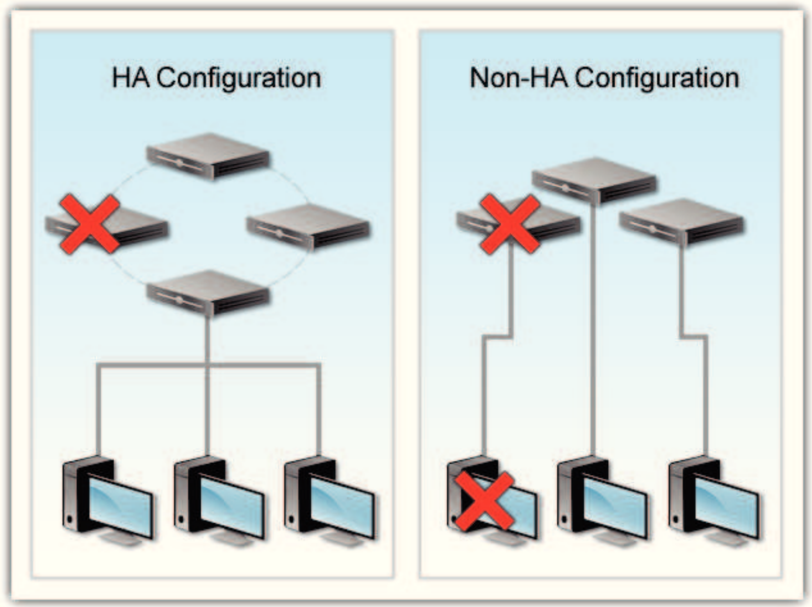

High availability refers to an implementation in which at least two Provisioning Services hosts are configured to provide a vDisk to one or more target devices. Should the primary Provisioning Services host fail for any reason, and high availability is enabled, the connection will fail over to the secondary Provisioning Services host.

In order to provide maximized operational time, high-availability-enabled implementations use a shared storage architecture. Multiple Provisioning Services hosts access the same physical files located on shared storage, which allows a target device to establish a connection on an alternate Provisioning Services host if the connection to the active Provisioning Services host is interrupted for any reason. A target device does not experience any disruption in service or loss of data when failover occurs.

When failover occurs, a target device attempts to connect to the next available Provisioning Services host. If unable to make a connection, the target device continues to try different Provisioning Services hosts until it can successfully connect.

The Provisioning Services host to which a target device accesses for logon does not necessarily become the Provisioning Services host that accesses the vDisk on behalf of the target device. In addition, once connected, if one or more Provisioning Services hosts can access the vDisk for this target device, the server that is least busy is selected.

To purposely force all target devices to connect to a different Provisioning Services host in a high-availability configuration, while preventing targets from timing out and attempting to reconnect to the current Provisioning Services host, stop the Stream Service on that Provisioning Services host.

For more information about implementing Provisioning Services high availability, see Citrix article CTX121090 on *http://support.citrix.com*.

Upon shutdown, the Stream Service will notify each target device to log on again to another Provisioning Services host.

Provisioning Services Failover

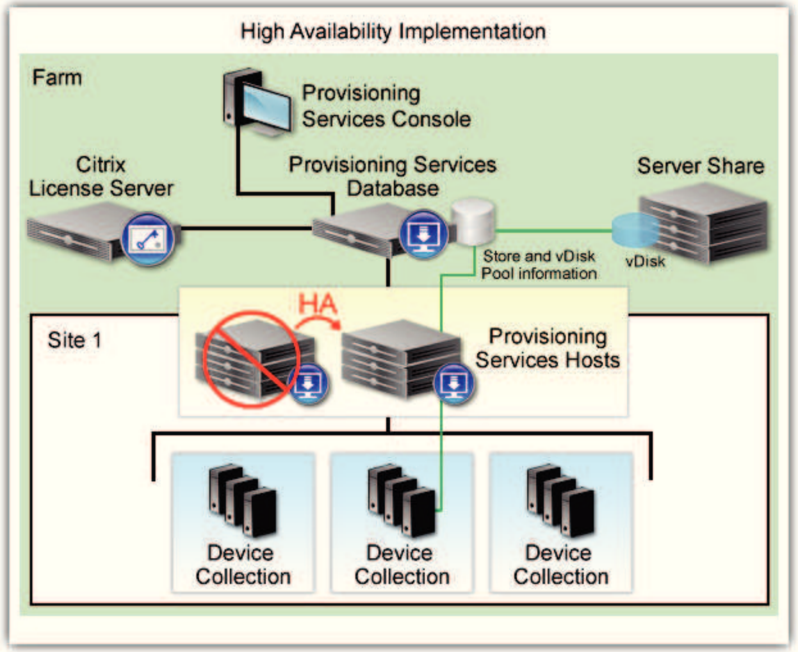
High Availability Implementation
Farm
Provisioning Services Console
Citrix License Server
Provisioning Services Database
Server Share
Store and vDisk Pool information
vDisk
Site 1
HA
Provisioning Services Hosts
Device Collection
Device Collection
Device Collection

By default, all Provisioning Services hosts within a site that can access a vDisk can provide that vDisk to target devices. Multiple Provisioning Services hosts can access the same physical files located on shared storage, which allows a target device to establish a connection on an alternate Provisioning Services host if the connection to the active Provisioning Services host is interrupted for any reason. A target device does not experience any disruption in service or loss of data when failover occurs.

For more information about planning and implementing Provisioning Services High Availability, see Citrix article CTX121090 on *http://support.citrix.com*.

For information on configuring Provisioning Services to automatically balance the target device load between servers, refer to Balancing the Target Device Load on Provisioning Servers in the Provisioning Services Administrator's Guide.

For implementations that use vDisk replication, if a server failover occurs, only those Provisioning Services hosts with access to an identical replicated vDisk can provide that vDisk to target devices. For example, if a vDisk is replicated across three Provisioning Services hosts' hard drives and then one of the vDisks is updated, that vDisk is no longer identical and will not be considered if a server failover occurs. Even if the same exact update is made to two of the vDisks, the timestamps on each will differ, therefore the vDisks are no longer identical.

If load balancing is enabled for the vDisk and a Provisioning Services host providing that vDisk should fail, Provisioning Services automatically balances the target device load between the remaining Provisioning Services hosts. If the load balancing option is not enabled, a single

Provisioning Services host is assigned to provide the vDisk to target devices; therefore failover will not occur.

Configuring the Boot File for High Availability

The boot file of a target device contains the IP addresses of up to four logon Provisioning Services hosts, as well as other configuration information. The boot file lists the Provisioning Services hosts that a target device can contact to get access to the Provisioning Services farm. The server that is contacted can hand the target device off to a different Provisioning Service host that is able to provide the target device with its vDisk.

A target device initiates the boot process by first loading a bootstrap program. A bootstrap program is a small program that runs before the operating system is loaded. Provisioning Services uses a special bootstrap program which initializes the streaming session between the target device and the Provisioning Services host. After this session starts, the operating system is streamed and loaded from the vDisk that was initiated.

> A shared storage system ensures the availability of the Provisioning Server vDisks. Depending on the type of shared storage, the vDisks use either the Universal Naming Convention (UNC) or the usual DOS naming convention.

Adding Provisioning Services Hosts to the Boot File

You must add Provisioning Services hosts to the boot file to provide a target device with the information necessary to make contact with the Stream Service.

During configuration, you can configure a Provisioning Services host to provide TFTP services. If all target devices are on one network segment, there will typically be one TFTP server for each farm. If target devices are on multiple network segments, and each segment is configured as an independent site, then one TFTP server for each site (network segment) can be used.

Provisioning Services hosts can also be configured as logon servers in the Console using the Configure Bootstrap dialog box.

For more information about adding Provisioning Services hosts to a boot file, see Citrix eDocs at *edocs.citrix.com*.

Enabling High Availability on vDisks

After the bootstrap file has been configured, the high availability feature must be enabled on the vDisk.

To enable high availability on vDisks:

1. Right-click the vDisk and select the **File Properties** menu option.
2. Select the **Options** tab.

3. Select the **High availability (HA)** check box.
4. Click **OK** to save this vDisk property change and continue.
5. Configure load balancing in the properties of the vDisk.

To Provide Provisioning Services Hosts with Access to Stores

For each store, select the Provisioning Services hosts that can access the store:

1. Right-click the Store, then select the **Properties** menu option.
   The Store Properties dialog box appears.
2. Select the location of Provisioning Services hosts that should be able to access this store.
3. Enable the checkbox next to each Provisioning Services host that can provide vDisks in this store, then click **OK.**

Considerations for Offline Database Support

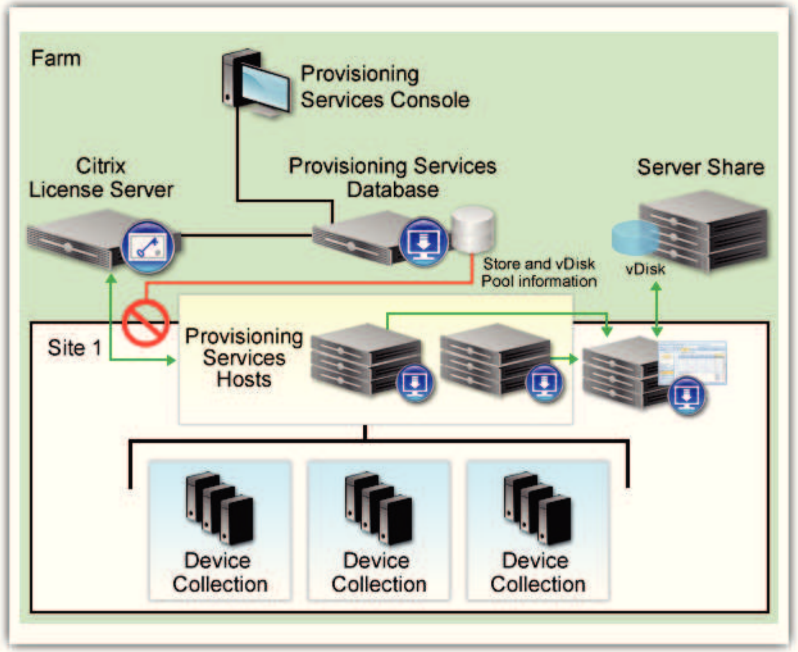
Farm
Provisioning Services Console
Citrix License Server
Provisioning Services Database
Server Share
Store and vDisk Pool information
vDisk
Site 1
Provisioning Services Hosts
Device Collection
Device Collection
Device Collection

The Offline Database Support option allows Provisioning Services hosts to use a snapshot of the Provisioning Services database in the event that the connection to the database is lost.

> This option is disabled by default and is only recommended for use with a stable farm running in production. Only a farm administrator can set this option.

When offline database support is enabled on the farm, a snapshot of the database is created and initialized when the Provisioning Services host is started. The Provisioning Services host is then continually updated by the Stream Service. If the database becomes unavailable, the Stream Service uses the snapshot to get information about the Provisioning Services host and the target devices available to the Provisioning Services host; this functionality allows Provisioning Services hosts and target devices to remain operational. However, when the database is offline, Provisioning Services management functions and the Console become unavailable.

When the database connection becomes available, the Stream Service synchronizes any Provisioning Services hosts or target device status changes made to the snapshot back to the database.

It is important to note that the following features, options, and processes remain unavailable when the database connection is lost, even if the Offline Database Support option is enabled:

- AutoAdd target devices
- vDisk updates
- vDisk creation
- Active Directory password changes
- Stream Process startup
- Image Update service
- Management functions such as PowerShell, MCLI, SoapServer, and the Console

To Enable Offline Database Support

1. Right-click the Farm, then select **Properties**.
   The Farm Properties dialog box appears.
2. Check the checkbox next to **Offline Database Support** on the Options tab.
3. Restart the Stream services.

Stream Logging

The files used in Provisioning Services to manage logging, as well as any generated log files, are located in: `%APPDATA%\Citrix\Provisioning Services\logs`.

For example: `C:\Documents and Settings\All Users\Application Data\Citrix\Provisioning Services\logs`

Stream log files include:

- `Stream_log.config`
- `Stream.log`

The StreamProcess.exe, Manager.dll, and Streamdb.dll all write to the `Stream.log file`.

> The `Stream_log.config` file should not be edited manually. Logging levels should be set through the Console. Any edits made to this file manually are lost when the Provisioning Server restarts, or when logging levels are changed using the Console.

The content of a log file includes:

- Timestamp
- Logging Level
- Component and method used to perform logging
- Provisioning Services host and target device identity (name, IP, or MAC)
- Logging message with supporting data of Windows error codes, when appropriate

To Enable Logging

1. Right-click the Provisioning Services host, then select the **Properties** menu option.
2. Select one of the following logging levels on the Logging tab:
    - OFF
    - FATAL
    - ERROR
    - WARN
    - INFO
    - DEBUG

> The logging levels are listed from the minimum level (OFF) to the maximum level (DEBUG) of logging information that you can collect. Logging levels are inclusive of previous levels. For example, if you select INFO, log information will include WARN, ERROR, and FATAL.
>
> For more information about Provisioning Services log properties, see Citrix eDocs at *edocs.citrix.com*.

3. In the **Max File Size** text box, scroll to select the maximum size that a log file can reach.
   When the max file size is reached, the file is closed and an index number is appended to the file name, then a new file is created.
4. Scroll to select the maximum number of backup files to retain in the **Max Backup Files** text box, then click **OK**.
   The oldest log file is automatically deleted when the maximum number of backup files is reached.
5. Enable Log events to the Windows Event Log of the Provisioning Services host that is communicating with the target device. This log includes errors that might occur after the early start phase as well as any critical error reporting. Click **OK.**

Troubleshoot vDisk Images

When troubleshooting vDisk images, there are several steps you can take to resolve the problem. Below are a few common troubleshooting issues.

Troubleshooting vDisk High Availability Issues

If you have multiple Provisioning Services hosts in your farm yet you encounter clients that stop responding when one of the Stream services shut down, you might have a configuration issue within the Provisioning Services Console. After verifying that all vDisks are available and accessible from all Provisioning Services hosts, verify these additional configuration options:

- Verify that the Store is set to service all Provisioning Services hosts.
- Verify that the actual vDisk is set for high availability.
- Verify that the vDisk is set to use a load balancing algorithm and not assigned to a particular Provisioning Services host.
- Check the Bootstrap configuration for all Provisioning Services hosts to make sure they are listed correctly. Incorrect settings in the Bootstrap configuration will affect failover behavior.

Troubleshooting and Viewing Replication Status for a Particular vDisk

Provisioning Services allows users to view the availability of replicated vDisks to Provisioning Services hosts within a farm.

1. Right-click a vDisk in the Console, then select the **Versions** menu option.
2. Highlight a version in the dialog box, then click the Replication button. The vDisk Version Replication Status dialog box displays showing the replication status availability for each server that can provide this version of the vDisk.
    - If a version is in **Maintenance** (hammer icon), **Test** (magnifying glass), or **Pending** (hour glass) states, that state displays in the first row.
    - A green checkmark indicates that the server has access to the highlighted version.
    - A yellow warning indicates that a server currently does not have access to one or more versions of the highlighted vDisk. The version that is missing, or has an issue, has a yellow warning under the version column.

Releasing a vDisk Lock

Multiple target devices and Provisioning Services hosts can gain access to a single vDisk image file. Therefore, it is necessary to control access to prevent corruption of the image. If multiple target devices are configured to start from a private image, a corrupt image would result. Therefore, the image becomes locked appropriately for a given configuration. A small lock appears over the vDisk icon to indicate that the vDisk is locked.

> Ensure that the vDisk is not in use before removing a lock. A vDisk image can become corrupted if the lock is released while a target device is still connected.

Using the Status Tray on a Target Device

The purpose of the Virtual Disk Status Tray tool is to aid in the management and troubleshooting of vDisks. The Virtual Disk Status Tray provides key troubleshooting information such as:

- Status, which indicates whether the target device is accessing the vDisk (active) or not (inactive).
- Server, which indicates the IP address and port of the Provisioning Services host providing access to the vDisk.
- Boot from, which indicates if the vDisk is set to start from a local hard drive or from a vDisk.
- Virtual Disk, which provides the name of the vDisk that is being accessed by the target device.
- Mode, which indicates the current access mode for the vDisk.
- Version, which indicates the edition and provides version and server-pack information.
- Boot Statistics, which provides information on the start time, retries, bytes read, bytes written, and throughput.
- Session Statistics, which provides information on uptime, retries, bytes read, bytes written.

Test Your Knowledge: Managing vDisks

Indicate whether each statement is true or false.

Statement	True or False
Microsoft KMS volume licensing uses a centralized activation server that runs in the datacenter, and serves as a local activation point for Provisioning Services.	True
If you want to duplicate an existing vDisk, you must simply create a copy of the existing .VHD file.	False
Once you have created a fixed .VHD file, you can never expand its size.	False
When creating a vDisk manually, the last step in the process is to unmount the vDisk.	True

Module 8

Implementing Resource Pools

Overview

A resource pool consists of one or more XenServer hosts, managed as a single entity. When combined with shared storage and a shared network, a resource pool enables virtual machines to be started on any XenServer host. In addition, virtual machines that are running within the resource pools can also be dynamically moved between hosts without perceived downtime.

Timings:

Module: 60 minutes

Exercises: 20 minutes

Total Time: 80 minutes

Objectives

After completing this module, you will be able to:

- Explain the purpose and functionality of resource pools.
- Manage several XenServer hosts as a single entity through the creation and configuration of a resource pool.
- Replace an existing pool master by promoting a secondary host in the resource pool.
- Migrate a virtual machine between hosts in a resource pool using XenMotion.

If you have time, show the students a Citrix TV video about Creating a Resource Pool *http://www.citrix.com/tv/#videos/107*.

Resource Pools

XenServer uses resource pools to provide enhanced administration, flexibility, and reliability. If an individual XenServer host suffers a hardware failure, then you can restart the failed virtual machines on another XenServer host in the same resource pool. If high availability is enabled on the resource pool, virtual machines will be started automatically started on another XenServer host.

A resource pool always has at least one host, known as the master. The pool master forwards commands executed in the command-line interface to individual pool members as necessary.

Resource pools:

- Allow several XenServer hosts to be treated as a single entity from a management perspective, creating consolidated management of server resources.
- Allow authentication to be performed in one place, eliminating the need to log on to each XenServer host individually.
- Allow all servers to share a common framework for network and storage, which facilitates features such as automatic placement of virtual machines and XenMotion.
- Use a master/secondary server model for management.

Pool Member Types

There are two types of resource pool members: master and secondary members. While a resource pool can only have one master, it can have up to 15 secondary members.

A master pool member:

- Allows for a single control point for all servers in the pool.
- Maintains all configuration data for the pool.
- Handles distributed locking for shared storage.
- Configures secondary servers in the pool.

Secondary pool members:

- Are controlled through the master.
- Keep a backup of all configuration data on the master.
- Can be promoted, in the event that the master becomes unavailable.

Resource Pool Communication

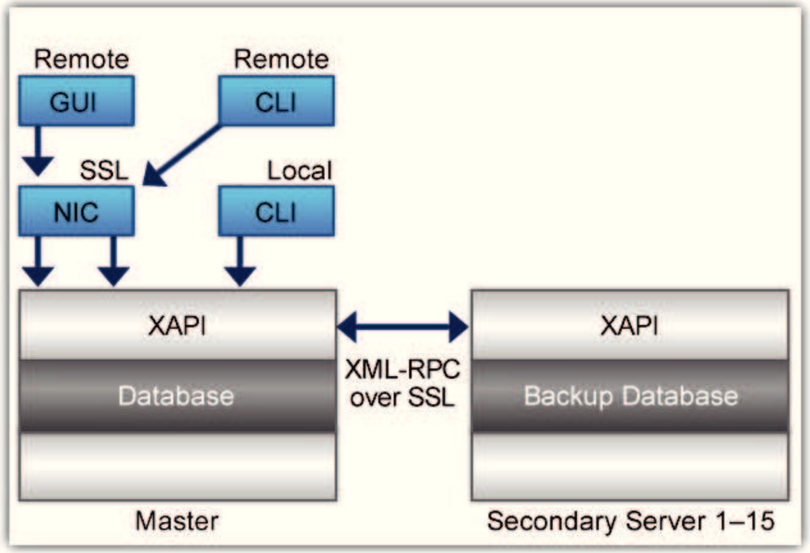

When communicating with a resource pool:

- The graphical user interface and remote command-line interface always connect to the master.
- The local command-line interface runs on any system in the resource pool. However, unless the UUID of the local server is included, the commands are passed to the master.

Homogeneous Pools

Homogeneous pools consist of one or more XenServer hosts and are defined as homogeneous if:

- the CPUs on the XenServer host joining the resource pool are the same (in terms of vendor, model, and features) as the CPUs on hosts already in the pool.
- the XenServer host joining the resource pool is running the same version of XenServer software, at the same upgrade level, as XenServer hosts already in the pool

Pool Member Requirements

In a homogeneous pool, members are required to have similar hardware configurations. The following XenServer host resources must be consistent throughout the resource pool:

- CPU
- Networking settings
- XenServer Software

Explain the importance of the XAPI service and how to restart the service.

CPU Requirements

All members of a resource pool require the same CPU vendor, model, and features, though the speed of the CPU can vary. For example, AMD-V and Intel VT CPUs cannot be mixed in the same resource pool.

Network Settings

Network settings should be similar among all pool members, including:

- NICs across all members must connect to the same networks.
- NICs must be in the same binding order on each system.
- NICs should be the same speed but can be from different vendors.

Virtual machines do not need the same number of NICs.

Heterogeneous Pools

Most system vendors add and discontinue CPU offerings within the life cycle of a server model, which makes it difficult to purchase servers with identical CPUs. However, with XenServer 6.0, the heterogeneous pools feature allows you to expand your existing XenServer environment with different hardware configurations. Heterogeneous resource pools use the technology in recent Intel (FlexMigration) and AMD (Extended Migration) CPUs that allow CPU masking or leveling. This feature enables you to configure a CPU so that it appears to provide a different make, model, or functionality than it actually does. With CPU masking, you are able to create resource pools consisting of XenServer hosts with different CPUs, while still supporting live migrations.

Using XenServer to mask the CPU features of a new server so that it will match the features of the existing XenServer hosts in a resource pool requires that:

- The CPUs of the XenServer host joining the resource pool must be of the same vendor (for example, AMD or Intel) as the CPUs on XenServer hosts already in the resource pool, though the specific type (for example, family, model and stepping numbers) need not be.
- The CPUs of the XenServer host joining the resource pool must support either Intel FlexMigration or AMD Enhanced Migration.
- The features of the older CPUs must be a sub-set of the features of the CPUs of the XenServer host joining the resource pool.
- The XenServer host joining the resource pool is running the same version of XenServer software, with the same hotfixes installed, as XenServer hosts already in the resource pool.
- XenServer Advanced Edition or higher is licensed.

Test Your Knowledge: Resource Pools

1. What is the purpose of using a resource pool in XenServer?

a. Resource pools allow for automatic failover protection in case a XenServer host within the pool fails.

b. Using a resource pool allows you to enable workload balancing within a resource pool, ensuring optimal virtual machine placement.

c. Resource pools eliminate the need to update and maintain multiple XenServer hosts, as only the master server needs to be managed within the pool.

d. Using a resource pool provides enhanced administration, flexibility, and reliability by allowing several XenServer hosts to be treated as a single entity from a management perspective.

Answer: D

2. At a minimum, which XenServer edition must you use for each XenServer host in a heterogeneous resource pool?

   a. Free

   b. Advanced

   c. Enterprise

   d. Platinum

   Answer: B

3. The master pool member is responsible for configuring secondary servers in the pool.

   a. True

   b. False

   Answer: A

Shared Configurations

Once a group of XenServer hosts is configured as a resource pool, any new shared configuration settings are reflected immediately across the resource pool. New resource pool members automatically inherit the shared configuration settings.

Shared configuration data includes:

- Virtual machine metadata
- Storage settings
- Networking settings

However, it is important to note that while the structural details of NICs, VLANs, and bonded interfaces are all inherited, policy information is not. The policy information that must be reconfigured includes:

- The IP addresses of management NICs, which are preserved from the original configuration.
- The location of the primary management interface, which remains the same as the original configuration. For example, if the primary management interfaces of other pool hosts are on a bonded interface, then the joining host must be explicitly migrated to the bond once it has joined.
- Re-assigning dedicated storage NICs to the joining host from XenCenter or the CLI and re-plugging the physical block devices (PBD), which act as PIFs to the physical storage, to route the traffic accordingly. This is because IP addresses are not assigned as part of the pool join operation, and the storage NIC is not useful without this configured correctly.

Adding a Host to a Resource Pool

To join a resource pool, a XenServer host must:

- Have a static management IP address.
- Synchronize its clock to the same time source as the pool master (for example, through NTP).
- Not be a member of an existing resource pool.
- Not bond its management interface.
- Not be running or suspending its virtual machines.
- Not place its virtual machine operations in progress.

To Add a Member to a Resource Pool

Because the secondary member inherits much of the configuration from the master, you need to follow this process to prepare and add the secondary member to a resource pool:

1. Disconnect any shared storage.
2. Halt all running virtual machines.

3. Ensure that physical NICs are connected to the same networks, in the same resource pool, in the same order.
4. Ensure that the secondary member has access to the shared storage in the existing resource pool.

> The resource pool join operation will fail if the conditions in steps 3 and 4 are not satisfied.

5. Click **Add member to the existing pool** from the Pool menu.

To Remove a Secondary Member from a Resource Pool

> Removing a member from a resource pool destroys all the virtual machines on the local storage of the member being removed.

1. Shut down any virtual machines on the member or use XenMotion to move the virtual machines to another pool member.
2. Export virtual machines that are on the local storage or copy them to shared storage.

> Copy any other data stored on the secondary member to a shared storage repository in the resource pool.

3. Remove the server from the pool.

XenMotion

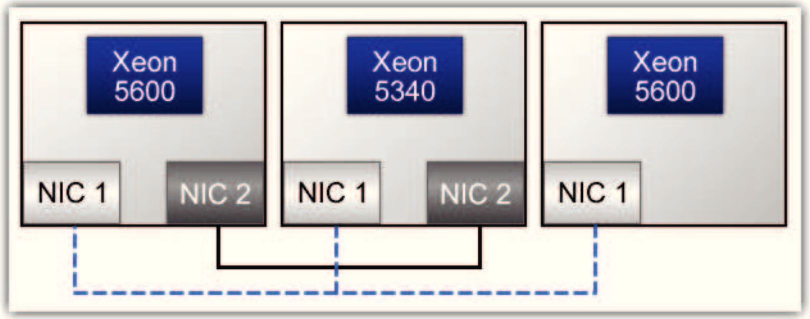

Resource pools enable the movement of virtual machines to distribute the load across XenServer hosts in a resource pool. To move virtual machines from one XenServer host to another, the source and the target servers:

- Must have access to the same shared storage.
- Must be connected to the same networks.

A virtual machine cannot be moved if the virtual machine:

- Is located on local or remote non-shared storage.
- Needs to be connected to a single-server private network.
- Does not have XenServer Tools installed.

Migrating a Virtual Machine or Template

Using XenMotion, you can migrate a running virtual machine from one XenServer host to another in the same resource pool.

If you have time, show this Citrix TV video about Migrating a Virtual Machine *http://www.citrix.com/tv/#videos/109*. In addition, demo XenMotion for the class as this is usually a unique feature that students want to see.

To migrate specified virtual machines to specified XenServer hosts, use the command:

```
xe vm-migrate
```

By using the vm-migrate command, you will have full control over the distribution of migrated virtual machines to other XenServer hosts in the resource pool. To automatically live-migrate all virtual machines to other XenServer hosts in the resource pool, use the command:

```
xe host-evacuate
```

There are several reasons for migrating virtual machine, including:

- Gaining better performance on another XenServer host.
- Keeping the virtual machines running when a XenServer host needs to be shut down for maintenance.
- Keeping the virtual machines running if a XenServer host needs to be removed or ejected from a resource pool.

By using the host-evacuate command, you leave the distribution of migrated virtual machines to XenServer.

Test Your Knowledge: Pool Member Types

Match the following terms with their correct descriptions. A term can match more than one description.

- Master
- Secondary

Description	Category
Allows for a single control point.	Master
Can be promoted.	Secondary

Description	Category
Configures other pool members.	Master
Handles distributed locking for shared storage.	Master
Keeps back up of all configuration data.	Secondary
Keeps track of all configuration data for the pool virtual machines.	Master

Pool Member Failure Overview

If the resource pool's master fails, master re-election will only take place if high availability is enabled. If high availability is not enabled, you cannot perform any virtual machine lifecycle operations.

XenServer uses synchronized resource pool configuration data to minimize the effect of a member failure. In a resource pool environment, the master XenServer host provides an authoritative database that is synchronously mirrored to all the member hosts in the resource pool. This provides a degree of built-in redundancy to the resource pool. The synchronized data enables every member of the resource pool to:

- Continue operating in the event of a master server failure.
- Take over the role of the master server if necessary.

Due to the synchronized resource pool configuration, every member of a resource pool contains the information necessary to take over the role of master if required. When a master node fails, one of the following events occur:

- If high availability is enabled, another master is elected automatically.
- If high availability is not enabled, each member will wait for the master to return or until you manually select a new pool master.

If the master server can be reliably recovered, secondary members will reconnect and synchronize configuration data. The secondary members will also exit emergency mode and return to normal operations.

Failure Detection Process

Pool configuration data is synchronized to all secondary systems every minute. When the secondary members do not receive this update from the master, the secondary members:

1. Retry the connection to the master for 1 minute.
2. Go into emergency mode if the master remains unreachable.
3. Retry the connection to the master every 3-5 minutes afterward.

To Promote a Secondary Member to a Master

If the master server cannot be recovered, you can promote a secondary member to the new master and instruct the new master to update the other members in the resource pool.

1. Connect to the local command-line interface of a secondary member.
2. Type the following command and press **Enter**.

   ```
   xe pool-emergency-transition-to-master
   ```

   The secondary member becomes the new master.
3. Type the following command and press **Enter**.

```
xe pool-recover-slaves
```

The secondary members now recognize the new master.

4. Verify that the default pool storage repository is set to an appropriate value by typing the following command and press **Enter**.

```
xe pool-param-list
```

Test Your Knowledge: Resource Pool Requirements

1. Which three tasks should you perform before adding a secondary server to a resource pool? (Choose three.)
    a. Disconnect local disk.
    b. Disconnect any shared storage.
    c. Shut down all running virtual machines.
    d. Ensure that NICs are connected to the same networks in the same pool in the same order.

    Answer: B, C, D
2. Which two statements accurately describe resource pools in a XenServer environment? (Choose two.)
    a. Provide a process for disaster recovery
    b. Use a master/secondary server model for management.
    c. Use heartbeats on the network and a storage device to determine the state of the servers.
    d. Allow multiple virtualization servers to be treated as a single entity from a management perspective.

    Answer: B, D
3. Which two statements regarding resource pool architecture are true? (Choose two.)
    a. The local command-line interface runs on the master.
    b. The GUI and remote command-line interface always connect to the master.
    c. The master and secondary servers communicate using the same XML-RPC protocol.
    d. The GUI and command-line interface communicate using XML-RPC protocol over HTTP to call the XAPI.

    Answer: B, C
4. Which three requirements are necessary to ensure that a virtual machine can be moved from one XenServer host to another? (Choose three.)
    a. Same processor class.
    b. XenServer Tools is installed.
    c. Connected to the same network.
    d. Access to the host's local storage.
    e. Access to the same shared storage.

    Answer: A, C, and E

Module 9

Distributed Virtual Switching

Overview

Distributed Virtual Switching creates a multitenant, highly secure, and extremely flexible network fabric that is the basis for building any public or private cloud. The vSwitch allows for centralized management of multiple virtual switches by decoupling the control and data planes of the virtual switch. The control plane manages control traffic such as routing protocols, while the data plane handles the forwarding of traffic. This configuration enables a distributed virtual switch to access the control plane of individual network switches—separately from the data plane—in order to act as a singular point of configuration for all virtual switch modules within a resource pool. In this manner, a distributed virtual switch creates an abstraction layer consisting of a single network switch spanning multiple hypervisors.

Objectives

After completing this module, you will be able to:

- Describe the key features of Distributed Virtual Switching.
- Explain the relationship between the vSwitch and the Distributed Virtual Switch Controller.
- Install and configure distributed virtual switching.
- Monitor resource status and network behavior through the distributed virtual switch interface.
- Configure a Cross-Server Private network.
- Apply access control lists and port configuration policies to customize the hierarchical policy model within the Distributed Virtual Switch.

Timings:

Module: 110 minutes

Exercises: 35 minutes

Total time: 145 minutes

If you have time, show the following Citrix TV videos

- http://www.citrix.com/tv/#videos/4219
- http://www.citrix.com/tv/#videos/4020
- http://www.citrix.com/tv/#videos/3705a

Ensure that when discussing the vSwitch, you are describing the network described in earlier modules.

Distributed Virtual Switching Features

A Distributed Virtual Switching solution provides the following features:

True network portability

Distributed Virtual Switching provides portable access control lists (ACLs) that associate with a virtual machine and move with that virtual machine. Distributed Virtual Switching bridges subnets to provide stateful migration of virtual machines between networks and / or between on-premise and cloud networks without manual intervention.

Network Fault Tolerance

Multiple vSwitches can be tied to each virtual machine to ensure network availability through redundant routes to each application or virtual machine.

Improved Network Security

Distributed Virtual Switching provides detailed traffic isolation. This feature enables customers to manage traffic, support multi-tenancy, and ensure packet isolation.

Transparently Support Network Compliance

Distributed Virtual Switching supports Remote Switched Port Analyzer (RSPAN) to transparently mirror all network traffic to a virtual machine. RSPAN is configured directly on the vSwitch ensuring separation of network compliance from server and virtual machine administration.

Industry Standard Network Monitoring

Distributed Virtual Switching integrates seamlessly with all network monitoring solutions supporting the NetFlow standard.

Distributed Virtual Switching Components

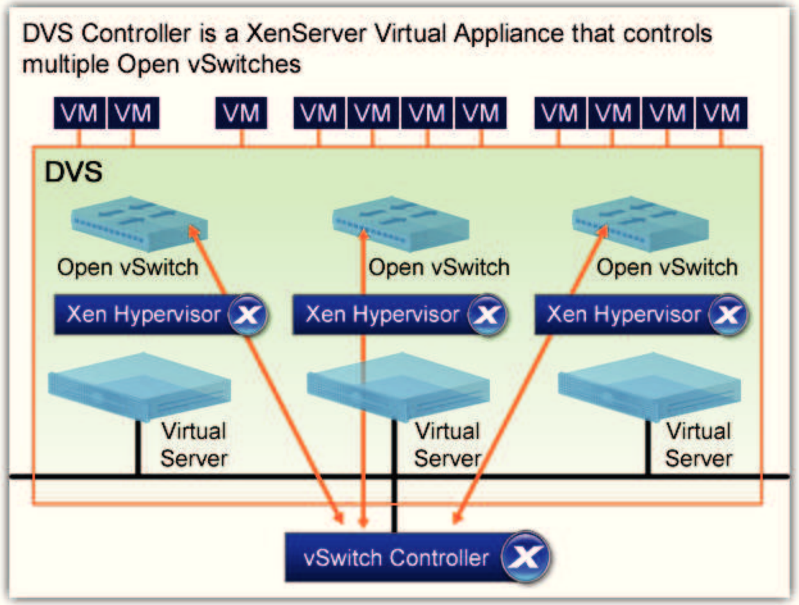
DVS Controller is a XenServer Virtual Appliance that controls multiple Open vSwitches

Distributed Virtual Switching consists of a virtualization-aware switch (the Open vSwitch) running on each XenServer host and the vSwitch Controller, a centralized server that manages and coordinates the behavior of each vSwitch to provide the appearance of a single distributed virtual switch.

Open vSwitch

Open vSwitch is an open source project aimed at developing a production-quality switch for virtual machine environments that supports distribution across multiple physical servers. The Open vSwitch enables detailed management over traffic flows with per-flow admission control, forwarding rule control, and isolation between tenants or applications. It enables you to dynamically reconfigure the network state for each virtual machine as it is deployed or migrated. As a virtual machine moves about the physical infrastructure, all of the policies associated with the virtual interface move with it.

The Open vSwitch includes the following features:

- Detailed access control lists and quality-of-service policies
- Visibility into the flow of traffic through NetFlow
- Traffic mirroring
- Port bonding
- Individual virtual machine traffic policing

vSwitch Controller

The vSwitch Controller provides a virtualized central server capable of applying network configuration policies on different scales ranging from global to specific virtual interfaces. The controller also enables a deep level of visibility within virtual machine traffic through an embedded NetFlow visualizer. A single vSwitch controller can manage up to 64 hosts.

Test Your Knowledge: Distributed Virtual Switch Features and Components

1. Which two choices describe key features offered by Distributed Virtual Switching? (Choose two.)
   a. Accesses real-time network traffic statistics.
   b. Decentralizes management of multiple virtual switches.
   c. Enhances local security by tunneling IP traffic between separate private networks.
   d. Enables virtual machines connected to a private network to live on multiple hosts within the same pool.

   Answers: A and D

2. Distributed Virtual Switching is comprised of a(n) ________ and a(n) __________.
   a. vSwitch, NetFlow analyzer
   b. vSwitch controller, Open vSwitch
   c. sFlow, vSwitch controller
   d. vSwitch controller, NetFlow analyzer

   Answer: B

vSwitch Controller Virtual Appliance Deployment

You can import the vSwitch Controller using the virtual appliance that is shipped with XenServer 6.0. During the import, the single VIF of the imported virtual appliance is attached to a network through which the vSwitch can communicate with the XenServer host or XenServer pool.

The vSwitch Controller can be imported into an existing XenServer pool, or as a stand-alone virtual appliance. The performance of the configurations is similar, with the exception of a Controller migration or restart. If you are running the vSwitch Controller within a resource pool, it might take slightly longer to connect to all vSwitches running on the virtual machines based on the differences in how the individual vSwitches route control connections.

The vSwitch Controller must be configured after it is imported. You can configure the vSwitch Controller by using:

- XenCenter
- A Web browser
- An SSH client

Accessing the vSwitch Controller

When the vSwitch Controller virtual appliance starts, the text console within XenCenter will display a message indicating the IP address that you can use to access the user interface remotely. If the virtual machine did not receive an IP address, the user interface cannot be used locally or remotely until one is assigned. The text console will provide instructions on setting the IP address locally in the command-line interface.

> The default password for the user interface is `dvscadmin`.

The vSwitch Controller can be managed using:

- The vSwitch Controller Graphical User Interface (GUI) locally from within XenCenter or remotely using a Web browser with the default password `dvscadmin`.
- The vSwitch Controller command-line interface locally from within XenCenter or remotely using an SSH client with the default password `admin`.

Test Your Knowledge: vSwitch Controller Access

Indicate whether each statement is true or false.

Statement	True or False
The remote command-line interface connects to the vSwitch controller using a self-signed certificate.	False
The local command-line interface connects to the vSwitch controller through the text console of the controller appliance virtual machine.	True
The remote graphical user interface connects to the vSwitch controller using the SSH client.	False
The local graphical user interface connects to the vSwitch controller through XenCenter.	True

vSwitch Controller Configuration

When the vSwitch Controller is started for the first time, it will attempt to obtain an IP address using DHCP; however, you can assign a static IP address through the Settings in the GUI. If DHCP is configured, resource pools cannot be set to Fail-Safe mode.

After the IP address has been configured, you can specify which resource pools the vSwitch Controller manages. When a resource pool is added, the vSwitch Controller automatically begins managing all XenServer hosts in that pool. You can add a resource pool through the Visibility & Control screen in the GUI.

A username and password is specified when adding a resource pool. The vSwitch Controller uses this account to communicate with the pool master server using the XAPI protocol. When communication is established, the new resource pool is added to the resource tree, along with the associated resources.

- Typically the user account, `root`, is specified when adding resource pools; however, this account could be different if Role-based Access Control is implemented in the XenServer environment.
- The user must have full management capabilities in the resource pool. The vSwitch Controller will not be able to properly manage the pool if the account has restricted capabilities.

High Availability

To ensure that the vSwitch Controller is always available, Citrix recommends enabling high availability for the vSwitch Controller virtual appliance.

Because Access Control List (ACL) rule enforcement depends on continuous operation of the vSwitch Controller, the virtual machine `restart-priority` should be set to 1 and `ha-always-run` should be set to true. These settings ensure that XenServer hosts can always reach an active vSwitch Controller. For more information about enabling high availability, see Citrix CTX130423 article on *support.citrix.com.*

Test Your Knowledge: Distributed Virtual Switch Configuration

1. Which action would cause an error based on the configuration of Distributed Virtual Switching?
   a. Setting the vSwitch Controller virtual machine restart-priority to 1
   b. Overriding the existing vSwitch Controller configuration while adding a new resource pool
   c. Adding a resource pool with user-level logon credentials
   d. Enabling ha-always-run while setting the vSwitch Controller virtual machine restart-priority to 1

   Answer: C

Monitoring Network Activity

The vSwitch Controller interface presents summary statistics and information about events within the virtual network environment. The information is automatically updated every few seconds. You can access the following network data by clicking Dashboard at the top of the vSwitch Controller interface.

Server Statistics

The Server Statistics section provides general information on the up time and CPU load of the vSwitch Controller.

Network Statistics

The Network Statistics section provides an inventory of the network elements for resource pools, XenServer hosts, networks, and virtual machines.

Recent Network Events

The Recent Network Events section lists the most recent events that have occurred within the managed virtual networks since the vSwitch Controller was last restarted. Over time, older events are automatically deleted from the list.

Recent Administrative Events

The Recent Administrative Events section lists events that have occurred within the vSwitch Controller itself, often as a result of changes made to the configuration within the graphical user interface. Over time, older events are automatically deleted from the list.

Throughput, Flows, and Bit Rate Graphs

The Throughput, Flows, and Bit Rate graphs section displays information about the behavior of the most active virtual machines and protocols.

Viewing Flow Statistics

By default, the vSwitch on each managed XenServer host sends NetFlow data to the vSwitch Controller, which uses this data to generate Flow Statistics tables and charts. NetFlow records are generated for all IPv4 flows after five seconds of inactivity or 60 seconds of total activity.

The data rate of a flow is represented as the total traffic of the flow averaged across the duration of the flow.

Because NetFlow uses UDP datagrams to transport NetFlow records between the vSwitch and the vSwitch Controller, the collector is not able to determine if or why a NetFlow record is not received. Dropped records can result in inaccurate data with Flow Statistics tables or charts.

Citrix recommends to disabling flow visibility in deployments of more than 100 virtual machines to avoid overloading the vSwitch Controller and the network used to send NetFlow records.

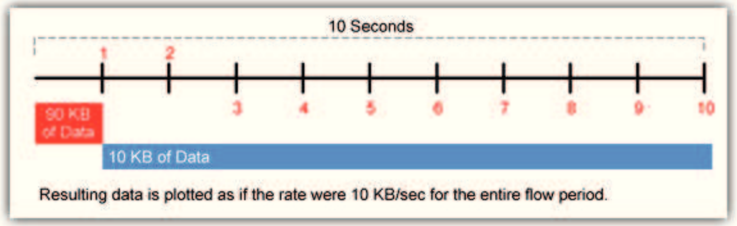

Resulting data is plotted as if the rate were 10 KB/sec for the entire flow period.

A flow lasts 10 seconds with 90 KB sent in the first second and 10 KB sent in each of the nine remaining seconds. The resulting data is plotted as if the rate were 10 KB/second for the entire flow period.

Exporting NetFlow Statistics

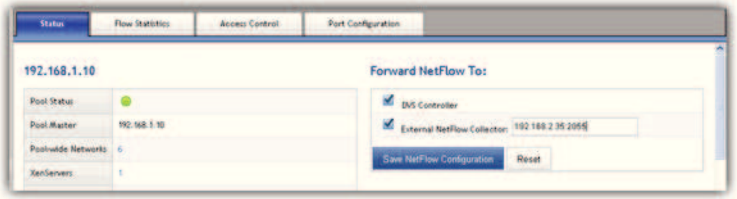

192.168.1.10

Pool Status	●
Pool Master	192.168.1.10
Poolwide Networks	6
XenServers	1

Forward NetFlow To:

☑ DVS Controller

☑ External NetFlow Collector: 192.168.2.35:2055

NetFlow statistics can be forwarded to an external NetFlow collector. To export the data, you must specify the IP address and port number for the NetFlow Controller.

If NetFlow is not being forwarded to the vSwitch Controller, a warning blue status text displays under the Flow Statistics tab. Click the blue text to display the available resource pool and reconfigure forwarding.

Managing Address Groups and Virtual Machine Groups

In a XenServer environment, flow statistics can be managed by Address Groups or Virtual Machine Groups.

Address Group

Address groups represent a collection of IPv4 addresses, specified as a single IP address, or they can use a network prefix notation, such as 10.0.0.0/8, which represents the network 10.0.0.0 with a netmask of 255.0.0.0. Address groups can be used to limit the scope of an ACL rule or to use the Flow Statistics settings to view only traffic coming to or from a set of IP addresses.

Virtual Machine Group

A Virtual Machine Group is a collection of virtual machines from any pool managed by the vSwitch Controller. Virtual Machine Groups provide a simple way to limit the data displayed by the Status or Flow Statistics tabs to an arbitrary subset of virtual machines.

You can create Address Groups and Virtual Machine Groups under the Visibility & Control tab in the vSwitch Controller interface.

Flow Statistics can be managed by:

- Address groups, which represent a collection of IPv4 addresses.
- Virtual Machine Groups, which are a collection of virtual machines from any pool managed by the vSwitch Controller.

Test Your Knowledge: Monitoring Networking Activity

1. You are one of two XenServer administrators of a enterprise environment. You were out of the office yesterday. When you came in this morning, you were informed that VLAN traffic is no longer being monitored. In which section can you find information that will help you determine a potential cause of this issue?
    a. Recent Administrative Events
    b. Server Statistics
    c. Recent Network Events
    d. Network Statistics

    Answer: A

Distributed Virtual Switching Policy Configuration Hierarchy

Although all policies are applied at the virtual interface level, the vSwitch Controller exposes a hierarchical policy model that supports default policy declarations across a collection of VIFs, such as a resource pool. You can override this default policy by creating detailed exceptions when needed, such as exempting a particular virtual machine from the default resource pool policy.

The policy hierarchy contains the following levels:

Global (most general level)	Includes all VIFs in all resource pools
Resource pools	All VIFs in a particular resource pool
Networks	All VIFs attached to a particular network
Virtual Machines	All VIFs attached to a particular virtual machine
Virtual Interfaces (most specific level)	A single VIF

Access Control List Policies

ACL list policies allow or deny virtual machine traffic based on traffic attributes in the Access Control tab. All ACLs are enforced as sets of rules on VIFs on the vSwitch.

Each ACL policy consists of the following set of rules:

Action	Indicates whether traffic matching the rule should be permitted or dropped.
Protocol	Indicates the network protocol to which the rule applies. You can apply the rule to all protocols, choose from an existing protocol list, or specify a new protocol.

Direction Indicates the direction of traffic to which the rule applies. The text of the rules is meant to be read from left to right. **To** means traffic outbound from the virtual machine, while **From** means traffic inbound to the virtual machine.

Remote Addresses Indicates whether the rule is limited to traffic to/from a particular set of remote IP addresses.

Access Control Policies Hierarchy

You can specify policies at any supported level of the access control policies hierarchy. At each level, rules are organized as follows:

Mandatory rules These rules are evaluated before any child policy rules. The only rules that take precedence over them are mandatory rules of parent policies. Mandatory rules are used to specify rules that cannot be overridden by child policies.

Child rules The child policy placeholder indicates the location in the rule order at which rules in child policies will be evaluated. It divides the mandatory rules from the default rules.

Default rules These are evaluated last, after all mandatory rules and all child policy default rules. They only take precedence over default rules of parent policies. They are used to specify behavior that should only be applied if a more specific child policy does not specify conflicting behavior.

Defining Access Control List Rules

A new ACL is defined using the resource tree in the vSwitch Controller GUI by selecting the node at the appropriate level in the policy configuration hierarchy. At each level, you can add rules for that level and higher levels. New rules can be added by using:

- The gear icon in the header bar for the level
- The gear icon for the entry

The new rule is added with the following default settings:

- Action - Allow

- Protocol - Any
- Directions - To/From
- Remote address - Any
- Description - None

Access Control List Rule Enforcement Order

While ACLs can be defined at different levels of the policy configuration hierarchy, ACLs are enforced on an individual virtual interface basis. For actual enforcement, the hierarchy is applied in the following order and applied to each virtual interface:

1. Mandatory rules at the global level
2. Mandatory rules for the resource pool containing the virtual interface
3. Mandatory rules for the network containing the virtual interface
4. Mandatory rules for the virtual machine containing the virtual interface
5. Rules for the virtual interface containing the virtual interface
6. Default rules for the virtual machine containing the virtual interface
7. Default rules for the network containing the virtual interface
8. Default rules for the resource pool containing the virtual interface
9. Default rules for the global level containing the virtual interface

The first rule that matches is executed and no further rules are evaluated.

To see the currently applied rules on a virtual interface along with the associated statistics, select the virtual interface in the resource tree and view the ACL in the Status tab.

Configuring Fail Modes

Under normal operation, the vSwitch maintains connection to its configured vSwitch Controller to exchange network management and status information. In the vSwitch Controller GUI you can configure how a vSwitch in the resource pool enforces ACL rules when it is unable to connect with the configured vSwitch Controller. The Fail Mode section allows you to configure the following fail modes:

Fail-open	All traffic is allowed, previously defined ACLs are not applied until the vSwitch is able to reconnect with the vSwitch Controller.
Fail-safe	Traffic is routed based on existing ACLs, all ACLs are applied.

If the vSwitch Controller becomes unavailable due to network disruption or controller restart, the vSwitch enters a period of inactivity during which network traffic is dropped. When vSwitch controller is not detected, the vSwitch enters the configured fail mode.

In fail-safe mode, existing ACLs are applied after the vSwitch loses connectivity to its configured vSwitch Controller. Traffic that does not match existing ACLs is denied. Traffic is denied in fail-safe mode when:

- A new VIF is plugged in.
- A virtual machine is migrated using XenMotion or Workload Balancing.
- Virtual machines on a host are added to a pool.

> If the vSwitch is restarted in fail-safe mode and the controller is unavailable after the vSwitch has started, all ACLs are lost, which means all traffic is denied. The vSwitch stays in fail-safe mode until connectivity with the vSwitch Controller is re-established and ACLs are pushed down to the vSwitch by the vSwitch Controller.

Test Your Knowledge: Access Control List Rule Enforcement

Put the following steps in order to illustrate the process of access control list enforcement.

Step	Description
4	Mandatory rules for the virtual machine containing the virtual interface
1	Mandatory rules at the global level
3	Mandatory rules for the network containing the virtual interface
2	Mandatory rules for the resource pool containing the virtual interface

Setting Up Port Configuration Policies

Port configuration policies are applied to VIF ports to support the following policy types:

QoS Quality of service (QoS) policies control the maximum transmit rate for a virtual machine connected to a distributed virtual switch port.

Traffic Mirroring RSPAN policies support mirroring traffic sent or received on a virtual interface to a VLAN in order to support traffic monitoring applications.

MAC Spoof Policy MAC address spoof check policies control whether MAC address enforcement is performed on outbound traffic from a VIF.

Enabling RSPAN without correct configuration of your physical and virtual network can cause a network outage.

Configuring QoS

QoS can be configured on the vSwitch Controller by specifying a QoS limit and a burst size. The QoS policy for a VIF sets a hard limit on the rate at which traffic can be sent from the VIF. The rate limit includes traffic to other virtual machines and traffic exiting the XenServer host. At a minimum, the burst size must be larger than the Maximum Transmission Unit (MTU) of the local network.

Setting the burst rate too small relative to the rate limit can prevent a VIF from being able to send enough traffic to reach the rate limit, especially with protocols that perform congestion control such as TCP.

Setting QoS to an inappropriately low burst size, such as 1KB, on any interface on which the vSwitch Controller sits might result in losing all communication with the vSwitch Controller and forcing an emergency reset situation.

QoS port policies are configured at the global, resource pool, network, virtual machine, and VIF levels by using the Port Configuration tab in the user interface. When you select a node in the resource tree, the configured value for each parent level in the hierarchy is shown; however, you can only change the configuration for the selected policy level.

To prevent any virtual machine from inheriting QoS policy configurations, disable the QoS policy at the virtual machine level.

Configuring QoS

You select a virtual machine in the resource tree and click the Port Configuration tab. The global, resource pool, network and virtual machine configured values are displayed; however, you are only able to change the value at the virtual machine level.

Configuring RSPAN

When RSPAN is enabled on a VIF, each packet sent to and from the VIF is copied and tagged with a VLAN value called the target VLAN. A host performance monitoring device is connected on the switch port that is configured to use the target VLAN. If the monitoring host interface uses promiscuous mode, it can see all traffic sent to and from the VIFs configured to use RSPAN.

To configure RSPAN policies on the vSwitch Controller, RSPAN must first be configured correctly on the physical switch. It is critical to correctly configure the physical network to be aware of the RSPAN traffic to avoid network outages. RSPAN should only be enabled if the physical switching infrastructure connecting all RSPAN-enabled VIFs can be configured to disable learning on the target VLAN.

To Configure RSPAN

To configure RSPAN you must:

1. Specify the available target VLAN IDs on the Status tab. You can specify available target VLAN IDs at the resource pool, network, or server level.

> When target VLANs are added at a level of the hierarchy, the VLANs are available at that level and lower levels of the hierarchy when configuring RSPAN. The correct level at which to specify a target VLAN depends on how widely you have configured your physical infrastructure to be aware of that target VLAN.

2. Select the appropriate VLAN ID on the Port Configuration tab to create the policy.

Configuring MAC Address Spoof Checking

MAC address enforcement can only be configured on a VIF basis and does not inherit or override parent configurations. When MAC address enforcement is configured, the vSwitch Controller detects and drops packets with an unknown MAC address from a VIF. All subsequent traffic from the VIF is also dropped.

MAC address spoof check policies are on by default and should be disabled on VIFs running software such as Network Load Balancing on Microsoft Windows servers. You can disable MAC address enforcement by selecting the MAC address spoof checkbox. The policy takes effect immediately after the configuration changes are saved.

Test Your Knowledge: Port Configurations

Match the following terms with the correct descriptions.

- QoS
- Traffic Mirroring
- MAC Spoof Policy

Term	Description
MAC Spoof Policy	Controls whether MAC address enforcement is performed on traffic outbound from a VIF
Traffic Mirroring	Controls the maximum transmit rate for a virtual machine connected to a distributed virtual switch port
QoS	Supports mirroring traffic sent or received on a virtual interface to a VLAN in order to support traffic monitoring applications

Cross-Server Private Networks

Cross-server private networks provide the following benefits without requiring a physical switch:

- The isolation properties of single-server private networks
- The ability to span a resource pool, enabling virtual machines connected to a private network to live on multiple hosts within the same pool
- Compatibility with features such as XenMotion and Workload Balancing

Cross-Server Private Networks must be created on a management interface, as they require an IP-enabled PIF. Any IP-enabled PIF can be used as the underlying network transport. If you choose to put cross-server-private network traffic on a second management interface, then this second management interface must be on a separate subnet.

Cross-Server Private Network Requirements

To create a cross-server private network, the following conditions must be met:

- Each of the servers in the pool must be using XenServer 6.0 or greater.
- Each of the servers in the pool must be using the vSwitch for networking.
- The pool must have a vSwitch Controller configured that handles the initialization and configuration tasks required for the vSwitch connection.
- Each server must be have an IP-enabled PIF.

XenMotion with Cross-Server Private Networks

Cross-server private networks are able to use XenMotion without requiring a physical switch. When XenMotion is implemented the following events occur:

1. The policies of the virtual machine being migrated are copied to the target host.
2. The virtual machine is moved to the new host.
3. The rules are applied to the virtual machine on the new host.
4. The policies are removed from the original host.

Recovering from a Failed vSwitch Controller

In the event that a vSwitch Controller failure occurs, the vSwitch continues to pass traffics; however, it will not enforce ACLs if it is in the default mode of fail-open.

If the vSwitch is restarted and the vSwitch Controller is not available, ACLs are not enforced because the vSwitch is not able to obtain the ACLs from the vSwitch.

To recover from a failed vSwitch Controller, you must import a new Distributed Virtual Switch Controller and restore the configuration from a backup. By default, a backup is created every 12 hours.

For more information about recovering the vSwitch Controller from a backup, see Citrix article CTX130423 on *support.citrix.com*.

The vSwitch Controller is not required to use the vSwitch network. The vSwitch Controller is used only to apply ACLS, QoS, RSPAN, MAC spoofing and collect metrics.

Module 10

Workload Balancing

Overview

Workload Balancing works with XenServer to provide optimal placement for a virtual machine in a XenServer resource pool. When Workload Balancing locates a virtual machine, it determines the best host on which to start a virtual machine or rebalances the workload across hosts in a pool. When Workload Balancing and high availability are working at the same time, high availability takes precedence.

Timings:

Module: 30 minutes

Exercises: 70 minutes

Total Time: 100 minutes

Objectives

After completing this module, you will be able to:

- Describe the key concepts and Workload Balancing components.
- Configure Workload Balancing by connecting a resource pool to the Workload Balancing virtual appliance.
- Ensure the balance of virtual machines in a resource pool by configuring Workload Balancing to optimize workloads.
- Deploy power management tools to optimize resource pool workload by understanding power management behavior.

Workload Balancing Overview

Workload Balancing provides recommendations for optimal virtual machine placement and virtual movement as well as provides a series of features that work together to reduce power consumption during off-peak work hours. These features automatically adjust the Workload Balancing placement strategy during off-peak periods, rebalances virtual machines to consolidate workloads onto hosts as densely as possible, and turns off lightly loaded hosts after their virtual machines have been relocated. The Workload Balancing virtual appliance is a single, pre-installed virtual machine designed to run on a XenServer host.

Use Workload Balancing to:

- Select the optimal host for each virtual machine when placing a host out of service for maintenance.
- Provide recommendations to help restart virtual machines on the optimal host, when host machines are taken offline.
- Help determine when to turn off hosts at certain times of the day, if necessary.

Workload Balancing Key Concepts

Workload Balancing for XenServer 6.0 is a new Linux-based version that is packaged as a virtual appliance that no longer needs to be installed as a separate application; you only need to import the virtual appliance and configure it. Also, you do not need to install an SQL Server database for Workload Balancing. A PostgreSQL database is now included in the virtual appliance.

Workload Balancing configuration preferences include settings for performance or density placement, virtual CPUs, and performance thresholds.

Workload Balancing makes recommendations to rebalance the virtual machine workload in your environment based on a strategy for placement you select known as the optimization mode.

Workload Balancing has two optimization modes:

Maximize Performance (Default)

Workload Balancing attempts to spread workload evenly across all physical hosts in a resource pool. The goal is to minimize CPU, memory, and network pressure for all hosts. When Maximize Performance is your placement strategy, Workload Balancing recommends optimization when a virtual machine reaches the High threshold.

Maximize Density

Workload Balancing attempts to fit as many virtual machines as possible onto a physical host. The goal is to minimize the number of physical hosts that must be online. When you select Maximize Density as your placement strategy, you can specify rules similar to the ones in Maximize Performance. However, Workload Balancing uses these rules to determine how it can pack virtual machines onto a host. When Maximize Density is your placement strategy, Workload Balancing recommends optimization when a virtual machine reaches the Critical threshold.

Workload Balancing Components

Workload Balancing consists of the following components.

Workload Balancing Server

The Workload Balancing server collects data from the virtual machines and their hosts and writes the data to the data store. This service is also called the data collector.

Data Store

The data store is a PostgreSQL database that is included in the virtual appliance.

To Download the Workload Balancing Virtual Appliance

1. Browse to `www.mycitrix.com` and click **Downloads**.
2. Log on to **MyCitrix**.
3. Click **Search Downloads by Product** > **XenServer** .
4. Click **XenServer Advanced Management Services**.
5. Click **XenServer 6.0 Dynamic Workload Balancing Virtual Appliance**.
   The Citrix Download Manager dialog box will appear
6. Click **Download Now**.
   The Download Manager will appear.
7. Select the download location and click **Save**.
   The Citrix Download Manager will display the download progress.

Import Considerations for the Workload Balancing Virtual Appliance

The Workload Balancing virtual appliance is designed to run on XenServer 5.6 Feature Pack 1 and later. It is capable of monitoring resource pools running XenServer 5.5 hosts and higher. Citrix recommends using XenCenter 6.0 to import the virtual appliance. The Workload Balancing virtual appliance follows the same import procedures as all virtual appliances. The Workload Balancing virtual appliance requires a minimum of 1 GB of RAM and 4 GB of disk space to run. Consider the following information before importing:

Communications Port

Before launching the Workload Balancing Configuration wizard, determine the port over which the Workload Balancing virtual appliance will communicate for configuration. The default port is 8012.

> Do not set to port 443. The Workload Balancing virtual appliance cannot accept connections over port 443 (the standard SSL/HTTPS port) because HTTPS is bound to the port making the pool.

Account for Workload Balancing

The Workload Balancing Configuration wizard requires a user name and password for the Workload Balancing account and the database account. These accounts will be created during configuration.

Resource Pools Monitoring

The Workload Balancing virtual appliance can monitor a resource pool while residing in another resource pool. The Workload Balancing virtual appliance requires that the time matches on the host running the Workload Balancing virtual appliance with the resource pool master in the resource pool that is being monitored.

> There is no way to manually change the time of the Workload Balancing virtual appliance. You can assign the XenServer hosting the virtual appliance and the resource pool master it is monitoring to the same Network Time Protocol (NTP) server.

XenServer and Workload Balancing Communication over HTTPS

During Workload Balancing configuration, Workload Balancing automatically creates a self-signed certificate. You can change this certificate to one from a certificate authority to configure XenServer to verify the certificate.

Workload Balancing Configuration

The Workload Balancing virtual appliance must be configured and enabled on each resource pool to be monitored before Workload Balancing can gather data for a resource pool.

You can configure the Workload Balancing virtual appliance using the configuration wizard in XenCenter found in the Console tab.

After initial configuration, evaluate the performance thresholds. Workload Balancing must be set to the correct thresholds for the environment or its recommendations might not be appropriate.

To Update Workload Balancing Credentials

Run the WLBConfig service command in the Workload Balancing appliance to change the Workload Balancing user name and password or the PostgreSQL password. After you execute this command, the Workload Balancing services are restarted.

1. Change the directory by running: `cd /opt/citrix/wlb`
2. Run `mono WlbConfig.exe`
   The screen displays a series of questions guiding you through changing your Workload Balancing user name and password and the PostgreSQL password. Follow the prompts on the screen to change these items.
3. Run `service postgresql-9.0 restart`, and `service workloadbalancing restart`. After modifying either the Workload Balancing user name, password, or the PostgreSQL password, you must restart both the PostgreSQL and Workload Balancing services.

Test Your Knowledge: Importing Workload Balancing Considerations

Indicate whether each statement is true or false.

Statement	True or False
A user name and password for both the Workload Balancing account and the database account will be created during configuration.	True
The port over which the Workload Balancing virtual appliance will communicate for configuration is port 443.	False

Statement	True or False
There is no way to manually change the time of the Workload Balancing virtual appliance to match the host and the resource pool.	True
You only need to configure Workload Balancing once for the first resource pool and Workload Balancing will enable every additional resource pool to be monitored automatically.	False

Connecting to the Workload Balancing Virtual Appliance

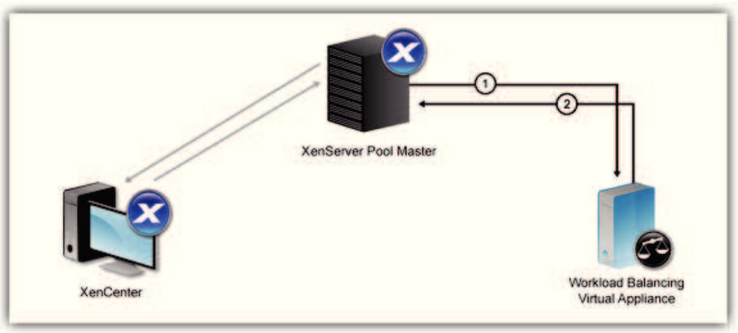

You can connect the Workload Balancing virtual appliance to the resource pool to be monitored using the WLB tab at the pool node in XenCenter.

To connect the Workload Balancing virtual appliance, you will need the:

- Host name (or IP address) and port of the Workload Balancing virtual appliance.
- Credentials for the resource pool (the resource pool master) you want Workload Balancing to monitor.
- Credentials for the account you created on the Workload Balancing virtual appliance during Workload Balancing configuration. This is often known as the Workload Balancing user account. XenServer uses this account to communicate with Workload Balancing.

If you want to specify the FQDN in Workload Balancing virtual appliance when connecting to the Workload Balancing server, you must first manually add its host name to your DNS.

When you first connect to the Workload Balancing virtual appliance, it uses the default thresholds and settings for balancing workloads. Automatic features, such as Automated Optimization Mode, Power Management, and Automation are disabled by default.

Workload Balancing Access Control Permissions

When RBAC is implemented in your environment, all user roles can access the WLB tab. However, not all roles can perform all operations. The following table lists the minimum roles that are required to use Workload Balancing features.

Task	Minimum Required Role
Configure, Initialize, Enable Disable WLB	Resource Pool Operator

Task	Minimum Required Role
Apply WLB Optimization Recommendations	Resource Pool Operator
Modify WLB Report Subscriptions	Resource Pool Operator
Accept WLB Placement Recommendations	VM Power Admin
Generate WLB Reports, including the Pool Audit Trail report	Read Only
Display WLB Configuration	Read Only

If a user tries to use Workload Balancing and does not have sufficient permission, a role elevation dialog box appears.

Workload Balancing Settings

After you connect the resource pool to the Workload Balancing virtual appliance, you can configure the following settings:

- Placement strategy
- Automatic optimizations and power management
- Critical thresholds and metric weightings
- Host exclusions

Workload Balancing settings apply collectively to all virtual machines and hosts in the resource pool. However, you must configure individual settings for each resource pool in your environment.

Critical Thresholds

Workload Balancing evaluates CPU, Memory, Network Read, Network Write, Disk Read, and Disk Write utilization for physical hosts in a resource pool. Workload Balancing recommendations are triggered when the High threshold in Maximum Performance mode or Low and Critical thresholds for Maximum Density mode are violated. After you specify a new Critical threshold for a resource, Workload Balancing resets the other thresholds of the resource relative to the new Critical threshold.

You can adjust Workload Balancing settings to create recommendations by changing critical thresholds, metric weighting factors, and resource settings.

When evaluating utilization, Workload Balancing compares its daily average to four thresholds: low, medium, high, and critical. After you specify or accept the default critical threshold, Workload Balancing sets the other thresholds relative to the critical threshold on a resource pool.

> The critical threshold is the only threshold that you can change through XenServer.

For more information about critical thresholds, see Citrix article CTX130420 on *support.citrix.com*.

Critical Threshold Evaluation and Edit

You can edit the critical threshold and modify the importance or weight that Workload Balancing gives to a resource in a resource pool. Citrix recommends to initially use the defaults in the Workload Balancing Configuration wizard. However, specific hardware might align better with different network and disk thresholds.

After Workload Balancing is enabled for a period of time, Citrix recommends evaluating your critical thresholds and determining if you need to edit them. For example, consider if you are:

- Getting optimization recommendations when they are not yet required. If this is the case, try raising the thresholds until Workload Balancing begins providing suitable optimization recommendations.
- Not getting recommendations when you think your network has insufficient bandwidth. If this is the case, try lowering critical thresholds until Workload Balancing begins providing optimization recommendations.

You can use the Pool Health report or the Pool Health History report to evaluate the effectiveness of your critical thresholds.

Demo maximum performance and density bullet points for students.

Metric Weighting

Metric weights are used to indicate the priority for individual resources in Workload Balancing recommendation calculations. Moving the slider towards Less Important indicates that a small amount of available resource for this pool is acceptable. That is, if you set memory as a less important factor in placement recommendation, Workload Balancing might still recommend placing virtual machines on a server with high memory utilization. The effect of the weighting varies according to the selected placement strategy. For example, if you select:

- Maximum performance with network writes as less important, and if the network writes on that server exceed the critical threshold, then Workload Balancing still makes a recommendation to place a virtual machine workload on a server with the goal of ensuring performance for the other resources.
- Maximum density as the placement recommendation and the network writes as less important, then Workload Balancing still recommends placing workloads on that host if the network writes exceed the critical threshold. However, the workloads are placed in the densest possible way.

By default, all metric weighting is set to More Important.

You can edit metric weighting factors by using the WLB tab of the XenCenter console by adjusting the slider next to the appropriate resource.

Fixed and Scheduled Optimization Modes

Workload Balancing also lets you specify when the optimization modes will be applied.

Fixed Optimization

Fixed optimization modes set Workload Balancing to one of the following specific optimization behaviors at all times:

- Try to create the best performance
- Try to create the highest density

Scheduled Optimization

Scheduled optimization modes allow different optimization modes to be applied depending on the time of day. Workload Balancing automatically applies the optimization mode at the beginning of the specified time period. Scheduled optimization can be configured for Everyday, Weekdays, Weekends, individual days, and by the hour.

Optimization Recommendations

The optimization recommendations display the following information:

- Name of the virtual machine that Workload Balancing recommends relocating
- Host it currently resides on
- Host that Workload Balancing recommends as the new location for a virtual machine
- Reason Workload Balancing recommends moving the virtual machine

After Apply Recommendations is selected, XenServer relocates all virtual machines listed as recommended for optimization and automatically displays the Logs tab. View this tab to monitor the progress of the virtual machine migration.

> Workload Balancing applies all configured recommendations during the optimization. Single recommendations cannot be applied.

Workload Balancing Power Management

Power management refers to the ability to turn on or off the power for physical hosts in a resource pool based on the total workload of the resource pool.

Configuring power management on a host requires that:

- The host server hardware can be turned on or off remotely.
- The Host Power On feature is configured for the host.
- The optimization mode of the resource pool is set to Maximum Density or Maximum Performance, either as a Fixed mode or a Scheduled mode.
- Workload Balancing is configured to apply Optimization recommendations automatically.
- Workload Balancing is configured to apply Power Management recommendations automatically.
- The host has been explicitly selected as a host to be managed under Power Management.

In Maximum Density mode, if Workload Balancing detects that there is insufficient host capacity in the resource pool to turn off virtual machines, it will recommend leaving the virtual machines on until the workload of the resource pool decreases enough to turn them off. When you configure Workload Balancing to turn off extra servers, it applies these recommendations automatically and behaves in the same way.

When a host is set to participate in Power Management, Workload Balancing makes turn on and turn off recommendations as needed. If you configure Workload Balancing so that it can automatically turn on the host when you are running in Maximum Performance mode, Workload Balancing will turn on one or more hosts if the utilization of a resource in the hosts exceeds the High threshold. However, when running in Maximum Performance mode, Workload Balancing never turns off hosts that it has turned on.

Power management is enabled at the resource pool level; however, you can specify the individual hosts from the resource pool for which you want to enable Power Management.

Power Management Behavior

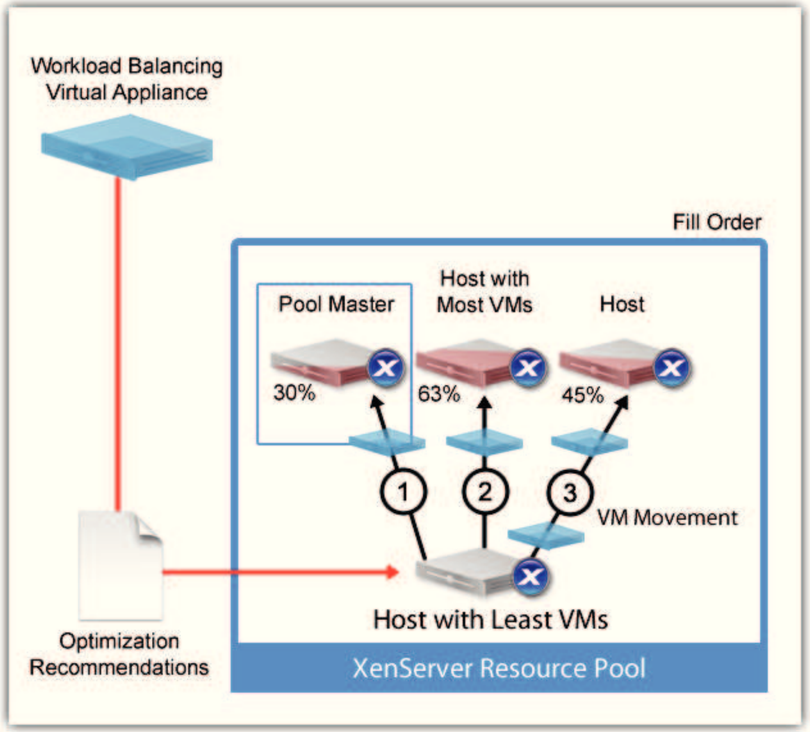

Before Workload Balancing turns servers on or off, it selects the hosts to which the virtual machines will be transferred based on which hosts have the most virtual machines running. It does so in the following order:

1. Filling the resource pool master because it is the host that cannot be turned off.
2. Filling the host with the most virtual machines.
3. Filling subsequent hosts according to which hosts have the most virtual machines running.

If Workload Balancing detects a performance issue while the resource pool is in Maximum Density mode, it attempts to address the issue by recommending migrating workloads among the hosts that are on. If Workload Balancing cannot resolve the issue using this method, it attempts to turn on a host. Workload Balancing determines which hosts to turn on by applying the same criteria it would if the optimization mode was set to Maximum Performance.

When Workload Balancing is running in Maximum Performance mode, Workload Balancing recommends turning on hosts until the resource utilization on all hosts in the resource pool falls below the High threshold.

Workload Balancing will turn on hosts automatically, or recommend to turn on hosts, if Workload Balancing determines that increasing capacity would benefit the overall performance of the resource pool while one or more virtual machines are migrating.

Workload Balancing never recommends turning on a host unless it was turned off by Workload Balancing.

Power Management and Virtual Machine Consolidation

When planning a XenServer implementation, consider your workload design with automatic virtual machine consolidation and power management. For example, you might want to place different types of workloads in separate resource pools or exclude a host from Workload Balancing.

If you have an environment with distinct types of workloads—for example, end user applications instead of domain controllers or types of applications that perform better with certain types of hardware—consider if you need to locate the virtual machines hosting these workloads in different resource pools.

Because power management and virtual machine consolidation are managed at the resource pool level, you should design resource pools so they contain workloads that you want consolidated at the same rate.

To Apply Optimization Recommendations Automatically

1. Select the **Automatically apply Optimization recommendations** check box under the WLB tab.
2. Specify the number of minutes Workload Balancing waits before applying the recommendation.
3. Select the lowest level of optimization to apply.
4. Select the hosts that Workload Balancing will control in the Power Management section.

Optimal Server Selection

When Workload Balancing is enabled and an offline virtual machine is restarted, XenCenter provides recommendations to help determine the optimal physical host in the resource pool on which to start the virtual machine. Workload Balancing makes these placement recommendations by using performance metrics that it previously gathered for that virtual machine and the physical hosts in the resource pool. Likewise, when Workload Balancing is enabled and migration of a virtual machine to another host occurs, XenCenter recommends which virtual machines to move to the host. This Workload Balancing enhancement is also available for the Initial (Start On) Placement and Resume features.

Workload Balancing functions more effectively and makes better, less frequent optimization recommendations if you start virtual machines on the servers it recommends.

When using these features with Workload Balancing enabled, host recommendations appear as star ratings beside the name of the physical host. Five empty stars indicates the least optimal server. An X appears beside host names that are not available.

Starting and Resuming a Virtual Machine

You can start or resume a virtual machine on an optimal server by selecting the virtual machine and specifying the Optimal Server under Resources > VM in the XenCenter console.

Test Your Knowledge: Power Management

Put the following steps in the order that Workload Balancing uses to select the hosts to which the virtual machines will be transferred before turning servers on or off:

Step	Description
2	Filling the host with the most virtual machines.
3	Filling subsequent hosts according to which hosts have the most virtual machines running.
1	Filling the resource pool master because it is the host that cannot be powered off.

Host Exclusion from Recommendations

Workload Balancing can be configured to exclude specific hosts from Workload Balancing optimization and placement recommendations, including Start On placement recommendations.

Situations in which you might exclude hosts from recommendations include when:

- Maximum Density mode and consolidating and shutting down hosts is desired, but specific hosts should be excluded from this behavior.
- Two virtual machine workloads always need to run on the same host.
- The workloads, such as SQL Server, should not be moved.
- Maintenance is required on a host, and the host should remain on the network.
- The performance of the workload is so critical that the cost of dedicated hardware is irrelevant.
- Specific hosts are running high-priority workloads, and prioritizing using the high availability feature is not desired.

- The hardware in the host is not optimal for the other workloads in the resource pool.

Excluded hosts remain excluded even when the optimization mode changes. To prevent Workload Balancing from shutting off a host automatically, consider disabling Power Management for that host instead.

You can exclude hosts from Workload Balancing optimization by Selecting Exclude Host in the WLB tab properties.

Module 11

Configuring High Availability

Overview

When high availability is enabled, XenServer continually monitors the health of the XenServer hosts in a resource pool. High availability automatically moves protected virtual machines to a healthy host if the current resource host fails. Additionally, if the XenServer host that fails is the pool master, high availability selects another host to take over the master role automatically, so that you can continue to manage the XenServer resource pool. If working concurrently, high availability will take precedence over Workload Balancing.

Timings:

Module: 120 minutes

Exercises: 35 minutes

Total Time: 155 minutes

Overview

After completing this module, you will be able to:

- Identify the requirements for high availability.
- Identify the steps needed to prepare for disaster recovery.
- Use a snapshot to create and restore a virtual machine.
- Run the XenServer commands to back up and restore a virtual machine.

High Availability Requirements

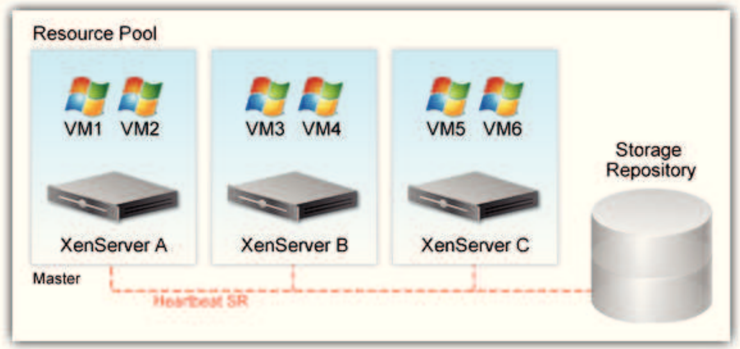

Requirements for high availability are:

- Shared storage, including at least one iSCSI, NFS, or Fibre Channel LUN of a recommended size of 356 MB or greater for the heartbeat storage repository.
- A XenServer resource pool.
- XenServer Advanced Edition or higher on all hosts.
- Static IP addresses for all hosts.

For maximum reliability, Citrix strongly recommends that you use a separate NFS or iSCSI storage array as your high availability heartbeat disk and not for any other purpose.

For a virtual machine to be protected by high availability, it must:

- Have its virtual disks on shared storage.
- Not have a connection to a local DVD drive configured.
- Have its virtual network interfaces on pool-wide networks.

A minimum of three hosts is recommended for high availability to work. For more information about high availability, see Citrix article CTX129721 on *support.citrix.com*.

High Availability Considerations

Explain to students why high availability won't work with two virtual machines. See Citrix article CTX129721.

The high availability mechanism creates two volumes on the heartbeat storage repository:

- 4 MB heartbeat volume is used for heartbeats.
- 256 MB metadata volume stores resource pool master metadata to be used in the case of master failover.

When high availability is enabled, some operations that would compromise the plan for restarting virtual machines, such as removing a server from a pool, are disabled. To perform these operations, high availability must be temporarily disabled.

Citrix recommends using the bonded management interface on the servers in the resource pool if high availability is enabled and using multipathed storage for the storage repository heartbeat.

Restart Priorities

Virtual machines can be assigned a restart priority and a flag to indicate whether they will be protected by high availability. When high availability is enabled, every effort is made to keep protected virtual machines live. If a restart priority is specified, any protected virtual machine that is halted will be started automatically. If a server fails then the running virtual machine will be started on another server. The restart priorities determine the order in which XenServer attempts to start virtual machines when a failure occurs.

Restart priorities can be set at:

- 0 - the first group of virtual machines to attempt to restart
- 1 - the second group of virtual machines to attempt to restart
- 2 -the third group of virtual machine to attempt to restart
- 3 - the fourth group of virtual machines to attempt to restart
- best-effort - the last group of virtual machines to attempt to restart

High availability Always Run flag can be set as:

- True, which includes the virtual machines in the restart plan.
- False, which does not include virtual machines in the restart plan.

Discuss scenarios for how to tag.

Server Failure Tolerance

XenServer dynamically maintains a failover plan that details what to do if any of the hosts in a pool fail at any given time. You can define the number of host failures that XenServer will tolerate as part of high availability configuration.

In a given configuration in which a number of host failures are greater than zero:

- Virtual machines that have restart priorities 0, 1, 2, or 3 are guaranteed to be restarted given the stated number of host failures.
- Virtual machines with a best-effort priority setting are not part of the failover plan and are not guaranteed to be kept running, since capacity is not reserved for them.

If the resource pool experiences host failures and enters a state in which the number of tolerable failures drops to zero:

- The protected virtual machines will no longer be guaranteed to be restarted.
- A system alert will be generated.
- If an additional failure occur in this state, all virtual machines that have a restart priority set will behave according to the best-effort priority behavior.
- If a protected virtual machine cannot be restarted at the time of a host failure, further attempts to start this virtual machine will be made as the state of the resource pool changes.

No running virtual machine will ever be stopped or migrated in order to free resources for a virtual machine with `always-run=true` to be restarted.

Overcommitting

A resource pool is overcommitted if the virtual machines that were running cannot be restarted in the resource pool, such as when there is not enough memory available. For example, if a resource pool consists of 16 XenServer hosts, and the tolerated failures is set to 3, the resource pool calculates a failover plan that enables any 3 XenServer hosts to fail and then restart virtual machines on other hosts. If a plan cannot be found, then the resource pool is considered to be overcommitted.

Explain that dynamic plan is updated when information is received from the heartbeat.

If a protected virtual machine cannot be restarted because the resource pool is overcommitted, then further attempts to start the virtual machine are made as extra capacity becomes available in a pool.

If you attempt to start or resume a virtual machine and that action causes the resource pool to be overcommitted, a warning alert appears. The alert explains that the pool will not be able to tolerate any further server failures and that the high availability configuration will no longer be able to restart protected virtual machines. You are then allowed to cancel the operation or you can proceed by first turning off high availability.

Host Fencing

Host fencing occurs when the XenServer host acts to ensure that the virtual machines are not running on two hosts simultaneously when a host failure occurs. host failure examples include:

- A loss of network connectivity.
- A problem with the control stack.

When a fence action is taken, the host immediately is restarted, causing all virtual machines running on it to stop. The other hosts detect that the virtual machines are no longer running, and the virtual machines are restarted according to the assigned priorities. The fenced host enters a restart sequence, and when it has restarted, it attempts to rejoin the resource pool.

Test Your Knowledge: High Availability Requirements

Indicate whether each statement is true or false.

Statement	True or False
A high availability heartbeat disk should be shared between two storage arrays.	False

Statement	True or False
High availability can be used in all XenServer editions.	False
Virtual machines protected by high availability must have their virtual network interfaces on pool-wide networks.	True
Virtual machines protected by high availability must have a connection to a local DVD drive.	False

Disaster Recovery

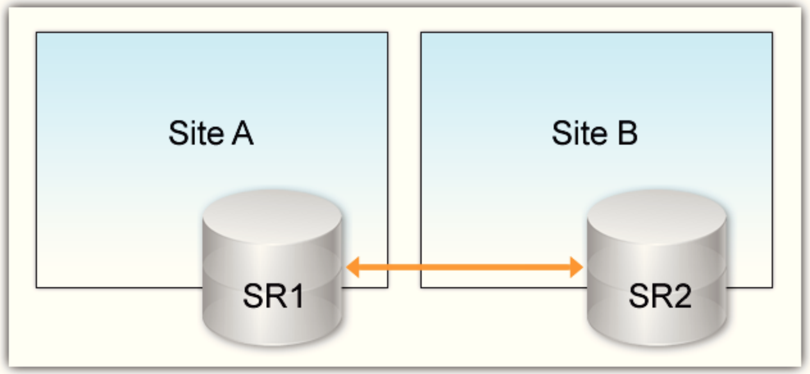

With disaster recovery enabled, XenServer automatically stores all metadata for virtual machines in a resource pool on one or more dedicated storage repositories.

In the event of a disaster, the Disaster Recovery wizard in XenCenter is used to access this storage and import chosen virtual machines into a recovery pool. When the virtual machines are running in the recovery pool, the recovery pool information is also saved to allow any changes to virtual machine settings to be replicated back to the primary pool should the primary pool be recovered.

If the Disaster Recovery wizard finds information for the same virtual machine present in two or more places--for example, storage from the primary site, storage from the disaster recovery site, and in the pool that the data will be imported into--then the wizard will ensure that only the most recent information for the virtual machine is used.

For more information about disaster recovery, see Citrix article CTX130420 on *support.citrix.com*.

Disaster Preparedness and Response

You can take measures to prepare your environment before a disaster and steps to recover from the disaster after it has happened. The following is an example of how to prepare for an Active/Passive disaster recovery.

Prepare Your Environment

- Note how your virtual machines and virtual appliances are mapped to your storage repositories and the storage repositories to your LUNs. Ensure that the name_label and name_description are named appropriately so you can recognize the storage repository later.
- Arrange replication of the LUNs using the native vendor tools of your storage systems.
- Enable resource pool metadata replication to one or more storage repositories on these LUNs.

Demonstrate how to obtain the name_label lists.

Discuss with the students that this is an example of how to prepare for an active/passive disaster recovery.

Disaster Response

- Break any existing storage mirrors so that the recovery site has read and write access to the shared storage.
- Ensure that the LUNs you designate to recover virtual machine data are not attached to any other pool, or corruption might occur.
- On the recovery site, re-synchronize any storage mirrors and shut down the virtual machines or virtual appliances that you want to move back to the primary site.
- On the primary site, select virtual machines or virtual appliances to fail back the primary site.

Test Your Knowledge: Disaster Recovery

1. Which step is not necessary before or after disaster recovery?
    a. Arranging replications of the LUNs
    b. Breaking existing storage mirrors
    c. Noting how virtual machines are mapping to the storage repositories
    d. Disabling pool metadata replication to the storage repositories on LUNs

    Answer: D

Backup and Recovery of XenServer Hosts and Virtual Machines

You can choose to back up and restore the metadata of the resource pool or virtual machine, the virtual machine itself, or the XenServer host.

Backup

XenServer hosts use a database on each host to store metadata about virtual machines and associated resources such as storage and networking. When combined with storage repositories, this database forms the complete view of all virtual machines available across the pool.

To back up virtual machine metadata only, run the command:

```
xe vm-export vm=<vm_uuid> file-name=<backup> metadata=true
```

To back up a virtual machine, choose a remote host with enough disk space and use the command:

```
xe vm-export vm=<vm_uuid> file-name=<backup>
```

To back up a XenServer host, run the following command:

```
xe host-backup file-name=<filename> -h <hostname> -u root -
pw <password>
```

Recovery

To restore virtual machine metadata, run the following command:

```
xe pool-restore-database file-name=/var/backup/pool-database-*
```

To restore the XenServer host, run the following command:

```
xe host-restore file-name=<filename> -h <hostname> -u root -
pw <password>
```

After recovery:

- Re-synchronize any storage mirrors.
- On the recovery site, shut down cleanly the virtual machines or virtual appliances that you want to move back to the primary site.
- On the primary site, follow the same procedure as for the failover above to fail back selected virtual machines or virtual appliances to the primary.

For more information about back and recovery, see Citrix article CTX130420 on *support.citrix.com*.

Metadata

Illustrate altering the command-line in a script environment. Script samples are available in the SDK.

XenServer uses a database on each host to store metadata about virtual machines and associated resources, such as storage and networking. When combined with storage repositories, this database forms the complete view of all virtual machines available across the resource pool. It is important to understand how to back up this database in order to recover from physical hardware failure and other disaster scenarios.

The metadata back up/restore feature is supported by two commands:

xe -backup-metadata This command provides an interface to create the backup virtual disk images (with the `-c` flag) and also regularly performs the metadata backup and examines its contents.

xe -restore-metadata This command is used to probe for a backup virtual disk image on a newly attached storage repository and also selectively re-imports virtual machine metadata to recreate the associations between virtual machines and their disks.

Go through the different flags with above commands.

For more information about metadata, see Citrix article CTX130420 on *support.citrix.com*.

Backing Up a Pool Installation

In the unfortunate event that your entire resource pool fails, you will need to recreate the pool database from scratch.

Use the following command to backup the pool database:

```
xe pool-dump-database file-name=<backup>
```

To verify that there is enough free space on the disk to perform the backup you can use the following command:

```
xe pool-restore-database file-name=<backup>
          dry-run=true
```

To back up metadata only, run the following command: `xe vm-export vm=<vm_uuid> filename=<backup> metadata=true`.

For more information about backing up the resource pool database, see Citrix article CTX130420 on *support.citrix.com*.

To Restore a Pool Installation

1. Run the `xe pool-restore-database` on the host designated to be the new master.
2. Run the `xe host-forget` on the new master to remove the old member machines.
3. Run the `xe pool-join` on the member hosts to connect them to the new resource pool.

Backing Up Virtual Machines as Snapshots

There are three options to perform a backup of a virtual machine; these backup options are also called snapshots:

- Regular snapshots
- Quiesced snapshots
- Snapshots with memory

Regular snapshots are fail-consistent and can be performed on all virtual machine types, including Linux virtual machines.

Quiesced snapshots take advantage of the Windows Volume Shadow Copy Service (VSS) to generate application-consistent, point-in-time snapshots. Quiesced snapshots are therefore safer to restore but can have a greater performance impact on a system while they are being taken. They might also fail under load, so more than one attempt to take the snapshot might be required.

Snapshots with memory can be useful if you are upgrading or updating software, or want to test a new application but want the option to get back to the current, pre-change state (RAM) of the virtual machine. Reverting to a snapshot with memory does not require restarting the virtual machine.

Before taking a snapshot, see the section called Preparing to Clone a Windows VM and Preparing to Clone a Linux VM in the *XenServer 6.0 Virtual Machine Installation Guide*. See Citrix article CTX130422 on *support.citrix.com*.

Explain to students that the following can also be done in XenCenter using the GUI. Demo backups from a batch file.

To Take a Snapshot of a Virtual Machine

1. Verify that the virtual machine is running or suspended so that the memory status can be captured.
2. Use the following command to take a regular snapshot:

```
xe vm-snapshot vm=<vm uuid> new-name-label=<vm_snapshot_name>
```

3. Use the following command to capture the RAM state (snapshot with memory).

```
xe vm-checkpoint vm=<vm uuid> new-name-
label=<name of the checkpoint>
```

When XenServer has completed creating the snapshot with memory, its `UUID` will be displayed. For example:

Snapshot with Memory

```
xe vm-checkpoint vm=2d1d9a08-e479-2f0a-69e7-24a0e062dd35
new-name-label=example_checkpoint_1
b3c0f369-59a1-dd16-ecd4-a1211df29886
```

To Restore a Running XenServer Host

1. Use the following command to restore the compressed image to the hard disk of the XenServer host on which the command is run (not the host on which `<filename>` resides).

```
xe host-restorefile-name=<filename> -h <hostname> -u root -pw
                <password>
```

2. Restart the XenServer host using the XenServer installation CD and select the **Restore from backup** option.
3. Use the following command to restore the virtual machine metadata:

```
xe pool-restore-databasefile-name=/var/backup/pool-database
```

Restoring from a backup as described here does not destroy the backup partition.

Listing Snapshots Using the Command-line Interface

You can use the following command to view a list of virtual machine snapshots:

```
xe snapshot-list
```

If you want to view the snapshot for a particular virtual machine, you must use the UUID of the virtual machine. Use the following command to view the UUID of the virtual machine:

```
xe vm-list
```

Deleting a Snapshot

Each snapshot is assigned a UUID when it is created. You must use the UUID of the snapshot to delete it. You can use the `xe snapshot-list` to find the UUID of the snapshot.

You can use the following command to delete the snapshot after you have identified the UUID:

```
xe snapshot-uninstall snapshot-uuid=<snapshot_uuid>
```

You can remove the metadata of a checkpoint or snapshot by using the following command:

```
xe snapshot-destroy snapshot-uuid=<snapshot_uuid>
```

Restoring a Virtual Machine from a Snapshot

Use the following command to list the virtual machines.

```
xe vm-list
```

Note the UUID of the snapshot or checkpoint to which you want to revert.

Use the following command to revert to a specific snapshot. Use the UUID listed with the previous command to determine the snapshot for the reversion.

```
xe snapshot-revert  snapshot-uuid=<snapshot_uuid>
```

After reverting to a snapshot, the virtual machine will be suspended.

Test Your Knowledge: Creating Snapshots

1. Which three types of snapshots are available in XenServer 6.0? (Choose three.)
    a. Acquiesced
    b. Full
    c. Regular
    d. Quiesced
    e. Partial
    f. Snapshot with memory

    Answers: C, D, F

Test Your Knowledge: Backup and Restore

1. Which two commands support the metadata backup/restore feature in the control domain? (Choose two.)
    a. xe-backup-vmdata
    b. xe-backup-metadata
    c. xe-restore-vmdata
    d. xe-restore-metadata

    Answers: B, D

Module 12

Managing and Troubleshooting XenServer

Overview

XenServer includes a number of tools for monitoring hosts, managing XenServer, performing updates, and troubleshooting.

Objectives

After completing this module, you will be able to:

- Monitor the health of a XenServer host by viewing reports and logs.
- Troubleshoot the health of a XenServer host by running a network trace and memory dump.
- Locate the XenServer log files that are requested when calling Technical Support.
- Use the XenCenter and the command-line interface to access tools to troubleshoot Role-based Access Control.
- Describe the process for upgrading XenServer hosts.

Timings:

Module: 60 minutes

Exercises: 85 minutes

Total Time: 145 minutes

Command-line Interface Overview

The XenServer command-line interpreter:

- Is automatically installed on XenServer hosts.
- Is included with XenCenter.
- Is available as a standalone remote command-line interface for Linux.
- May be installed as a Windows application.
- Uses SSL to secure remote communications.

Command-line Interface Operations Targets

The target for command-line interface operations depends on the existing XenServer environment. When the XenServer host is in a standalone server environment, command-line interface commands run directly against the server.

When the XenServer host is in a resource pool environment with other hosts, the command-line interface:

- Runs against the pool master by default.
- Runs on any member of the pool by specifying the host UUID.

Command-line Interface Basics

XenServer command-line interface commands use the following syntax:

```
xe command-name argument=value
```

Descriptions of the syntax components are as follows:

xe Specifies the XenServer binary

Command-name Specifies the XenServer command and is comprised of the object name and specific action associated with the object.

Tell students to go to the command-line appendix of the Administrator's Guide. Explain to the students that they should clear the history or restrict access to keep others from accessing previously used CLI commands.

Argument=value

Specifies the values used for the command. Combinations of argument=value can be placed in any order after the command. A command will often assume default values when optional arguments are not included in the command. For values containing spaces, type:

```
argument="value with spaces"
```

Each command contains its own set of arguments that are of the form argument=value. Some commands have required arguments, and most have a set of optional arguments. Typically a command will assume default values for some of the optional arguments when invoked without them.

If the `xe` command is executed remotely, additional connection and authentication arguments are used. These arguments also take the form `argument=argument_value`.

For more information about the command-line interface, see Citrix article CTX130420 on *support.citrix.com*

Command-line Interface Shortcuts

The XenServer command-line interface provides the command recall and command completion shortcuts.

Command Recall

The command recall shortcut provides you with access to recently entered commands in the command-line interface using the up arrow key.

Command Completion

The command completion shortcut provides you with the ability to fill in the rest of a command using the tab key.

Command Completion

If you type `xe vm-l` and then press the **TAB** key, the rest of the command will be displayed when it is unambiguous. If more than one command begins with `vm-l`, pressing **TAB** a second time will list the possibilities. This is particularly useful when specifying object UUIDs in commands.

XenServer Menu-Driven Text Console

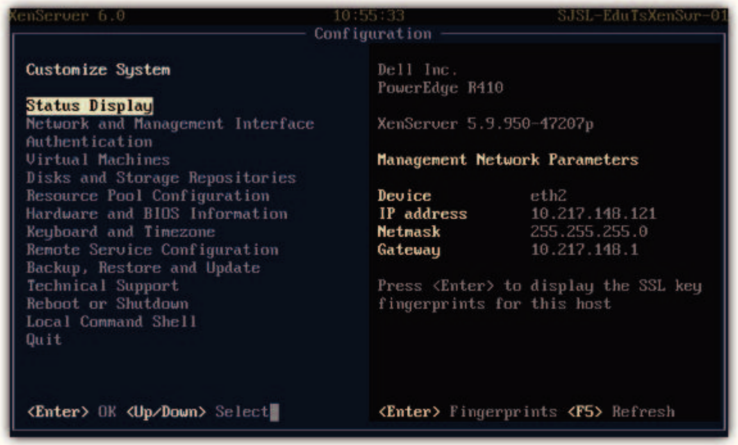

```
XenServer 6.0                    10:55:33                     SJSL-EduTsXenSvr-01
                              Configuration
 Customize System                         Dell Inc.
                                          PowerEdge R410
 Status Display
 Network and Management Interface         XenServer 5.9.950-47207p
 Authentication
 Virtual Machines                         Management Network Parameters
 Disks and Storage Repositories
 Resource Pool Configuration              Device          eth2
 Hardware and BIOS Information            IP address      10.217.148.121
 Keyboard and Timezone                    Netmask         255.255.255.0
 Remote Service Configuration             Gateway         10.217.148.1
 Backup, Restore and Update
 Technical Support                        Press <Enter> to display the SSL key
 Reboot or Shutdown                       fingerprints for this host
 Local Command Shell
 Quit

 <Enter> OK <Up/Down> Select              <Enter> Fingerprints <F5> Refresh
```

The menu-driven text console enables access to common XenServer functions without entering command-line interface commands or installing XenCenter. The updated text console of the XenServer host can be used when there is limited or no network connectivity. The menu-driven text console is automatically enabled on the main screen of the physical XenServer host computer. To use the updated console using XenCenter or an SSH connection, type `xsconsole` at the command-line interface prompt.

Test Your Knowledge: XenServer Commands

Match the following commands with their correct descriptions. A command can match more than one description.

- xe vm-param-set uuid=uuid param-name=param-
- xe host-shutdown host=host
- xe sr-param-set uuid=uuid parameter=value
- xe network-param-set uuid=uuid param=value
- xe host-disable host=host

Description	Command
Sets storage repository parameters	xe sr-param-set uuid=uuid parameter=value
Sets network parameters	xe network-param-set uuid=uuid param=value
Sets the properties of a virtual machine	xe vm-param-set uuid=uuid param-name=param-

Description	Command
Shuts down a server	xe host-shutdown host=host and xe host-disable host=host

Define what a tag is. Go to XenCenter and show the students how to add a tag and how to use tags to search.

Monitoring XenServer

XenServer and XenCenter provide access to alerts that are generated when noteworthy events occur. XenCenter provides various mechanisms of grouping and maintaining metadata about managed virtual machines, hosts, and storage repositories.

> Full monitoring and alerting functionality is only available with XenServer Advanced edition and above.

Alerts generated from XenServer:

- Are visible from the XenCenter GUI.
- Can be automatically e-mailed to the resource pool administrator.
- Can support custom fields and tags.
- Can support custom searches.

Additionally, the XenTop utility can be used to monitor a XenServer system in real-time, while historical logs can be tracked using XenServer to write logs to a remote syslog server.

Using the XenTop Utility to Monitor Host and Virtual Machine Performance

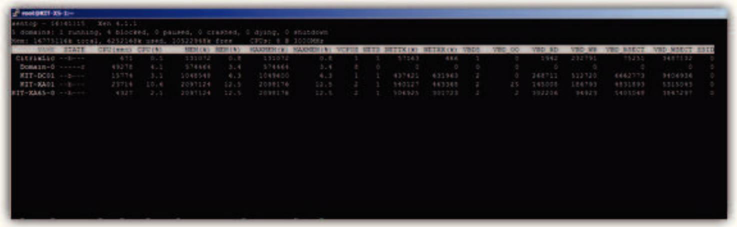

```
xentop - 16:41:15   Xen 4.1.1
5 domains: 1 running, 4 blocked, 0 paused, 0 crashed, 0 dying, 0 shutdown
Mem: 16775116k total, 6252148k used, 10522968k free    CPUs: 8 @ 3000MHz
      NAME  STATE   CPU(sec) CPU(%)     MEM(k) MEM(%)  MAXMEM(k) MAXMEM(%) VCPUS NETS NETTX(k) NETRX(k) VBDS   VBD_OO   VBD_RD   VBD_WR  VBD_RSECT  VBD_WSECT SSID
 CitrixLic --b---        471    0.1     131072    0.8     131072       0.8     1    1    57143      466    1        0     1942   232791      75251    3487132    0
  Domain-0 -----r      49278    4.1     574664    3.4     574664       3.4     8    0        0        0    0        0        0        0          0          0    0
  KIT-DC01 --b---      15776    3.1    1048548    6.3    1049600       6.3     1    1   437421   631963    2        0   268711   512720    6662773    9404936    0
  KIT-XA01 --b---      23714   10.6    2097124   12.5    2098176      12.5     2    1   540127   463385    2       25   145008   186793    4831893    5315043    0
KIT-XA65-0 --b---       4327    2.1    2097124   12.5    2098176      12.5     2    1   506935   301723    2        2   302236    96925    5405048    3867297    0
```

Define what a domain is according XenTop.

XenTop can be used to view the daemon and the resources being used on a domain if, for example, you have high spiking use of kernel or exhausted memory for the control domain. Demonstrate XenTop in progress.

Go to the command-line interface and demo XenTop. Define what is shown on the output.

The XenTop utility displays real-time information about a XenServer system and running domains in a level of detail not available from XenCenter. It uses a partially graphical interface to display all details in a friendly format. The XenTop utility is included in all versions of XenServer and can be accessed from the command-line interface using the `xentop` command.

> The `xentop` command does not need the xe.exe interpreter; therefore the command is not preceded by `xe`.

Column Name	Column Description
CPU(sec)	Domain CPU usage in seconds

Column Name	Column Description
CPU(%)	CPU percentage statistic
VCPUS	Number of virtual CPUs
NETS	Number of virtual networks
MEM	Current memory
MAXMEM(k)	Maximum domain memory statistic in KB
MAXMEM(%)	Memory percentage statistic, ratio of current domain memory to total node memory
NETTX	Number of total network tx bytes statistic/1024
NETRX	Number of total network rx bytes statistic/1024
VBDS	Prints number of virtual block devices
VBD OO	Prints number of total VBD OO requests.
VBD_RD	Number of read requests
VBD_WR	Number of write requests

Possible Host and Virtual Machine States	Description
d	domain is dying
s	domain shutting down
b	blocked domain
c	domain failed
p	domain paused
r	domain is actively running on one of the CPU

Logging to a Remote Syslog Server

XenCenter can be used to gather XenServer host information as logs. You can configure a XenServer host to log to a remote server. The syslog application must be running on the remote server to receive the logs and aggregate them correctly.

To Log Using a Remote Syslog Server

1. Set the syslog-destination parameter to the hostname or IP address of the remote server where you want the logs to be written by using the following command:

   ```
   xe host-param-
   set uuid=<xenserver_host_uuid>
   logging:syslog_destination=<hostname>
   ```

   Replace <xenserver_host_uuid> with the UUID for the host and <hostname> with the name of the syslog server.
2. Log to a remote server by using the following command:

   ```
   xe host-syslog-reconfigure uuid=<xenserver_host_uuid>
   ```

   Replace <xenserver_host_uuid> with XenServer host UUID.

Test Your Knowledge: Monitoring XenServer

Match the following methods with their correct descriptions. A tool can match more than one description.

- XenTop Utility
- XenServer Alerts
- Point to a Syslog Server

Description	Category
Generates Noteworthy Events	XenServer Alerts
Monitors Real-time Performance	XenTop Utility
Gathers Historical Logs	Point to a Syslog Server

Events and Alerts

XenServer supports configurable alerts for defined events, such as:

- The resource pool is overcommitted.
- A tolerable host failure occurrs, such as a multipathing failure or a failed NIC bond.
- The XenCenter is out of date or a new XenCenter version is available.

You can set alerts on hosts and virtual machines by accessing the **General tab** > **Properties**.

Alert Customization

Triggers for alerts are checked using perfmon. Perfmon requests and reads updates of performance variables from XenServer and runs once every five minutes. These default variables can be changed in `/etc/sysconfig/perfmon` but cannot be set lower than the minimum interval of five minutes.

Explain that alert customization has to be in XML format.

Default performance variables:

- Are averaged over one minute.
- Are separated into groups:
  - For the host itself.
  - For each virtual machine running on the host.

For more information see *The XenServer 6.0 Administrator's Guide.*

Persistent XenServer Performance Statistics

XenServer records statistics about the performance of various aspects of a XenServer installation. Metrics are stored as persistent data for long-term access and for analysis of historical trends. The amount of data recorded stays the same over time; however data is averaged during the initial five minutes, averaged again during the next hour, and averaged in decreasing time intervals thereafter. Metric data is stored on each XenServer host and can be requested by other XenServer hosts in the resource pool. Metrics are self-maintaining.

This is done because some metrics lose value over time.

Performance Statistics in XenCenter

The Performance tab in XenCenter:

- Is available for managed servers and virtual machines.
- Displays graphs showing:
  - CPU usage
  - Memory usage
  - Network usage
  - Disk I/O for virtual machines only

Go to XenCenter and show the students the Performance tab; demo how to change views in XenCenter.

Performance Graphs in XenCenter

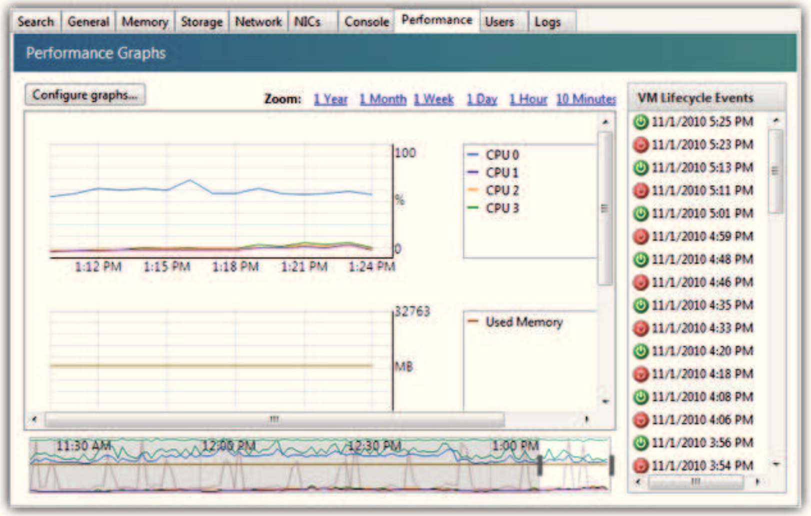

11/1/2010 5:25 PM
11/1/2010 5:23 PM
11/1/2010 5:13 PM
11/1/2010 5:11 PM
11/1/2010 5:01 PM
11/1/2010 4:59 PM
11/1/2010 4:48 PM
11/1/2010 4:46 PM
11/1/2010 4:35 PM
11/1/2010 4:33 PM
11/1/2010 4:20 PM
11/1/2010 4:18 PM
11/1/2010 4:08 PM
11/1/2010 4:06 PM
11/1/2010 3:56 PM
11/1/2010 3:54 PM

Performance graphs are easy to use and allow for dynamic review of performance metrics. You can add and alter items in the performance graphs by clicking Configure Graphs on the Performance tab in XenCenter.

Test Your Knowledge: XenServer Alerts

Indicate whether each statement is true or false.

Statement	True or False
XenServer supports configurable alerts for defined events.	True
Alerts can only be set at the host level and not the virtual machine level.	False
Triggers for alerts are checked using perfmon.	True
Perfmon is set at a default minimum interval of 1 minute.	False

XenServer Logs and Reports

XenCenter provides access to various logs and reports, which can be helpful with troubleshooting. The Logs tab within XenCenter provides detailed information about the host, virtual machine, and other resources. A consolidated log file for all events is available through XenCenter or the command-line interface. A XenServer Status Report groups related logs together.

XenCenter Event Logs

XenCenter creates two log files to track client-side events:

- `XenCenter.log` records every action, information, alert, and error event.
- `XenCenter Audit Trail.log` records the operations-related events from XenCenter.

Both log files are located in the following profile folder: `%appdata%\citrix\xencenter\logs\XenCenter.log`.

You can access the log files from XenCenter by clicking View Application Log Files from the Help menu and double-clicking either the `XenCenter.log` or `XenCenter Audit Trail.log` file to view the contents.

Server Status Reports

The Server Status Report wizard within XenCenter provides a convenient way to collect and package a comprehensive snapshot of a specific XenServer installation for troubleshooting. Options for including or excluding a range of different configuration files and log files for selected servers are available.

The Server Status Report wizard produces a single zip file that can be stored or e-mailed. The size of the report varies, depending on the files selected.

For more information about common log messages used in XenServer reports, see Citrix article CTX116371 on *support.citrix.com*.

Xen-Bugtool Utility

The XenServer host has several command-line interface commands to collate the output of logs and various system information using the xen-bugtool utility.

Use the xe command: `host-bugreport-upload` to collect the appropriate log files and system information and upload them to the Citrix Support ftp site.

For more information about the xen-bugtool utility, see *The XenServer 6.0 Administrator's Guide.*

XenServer Crashdump Overview

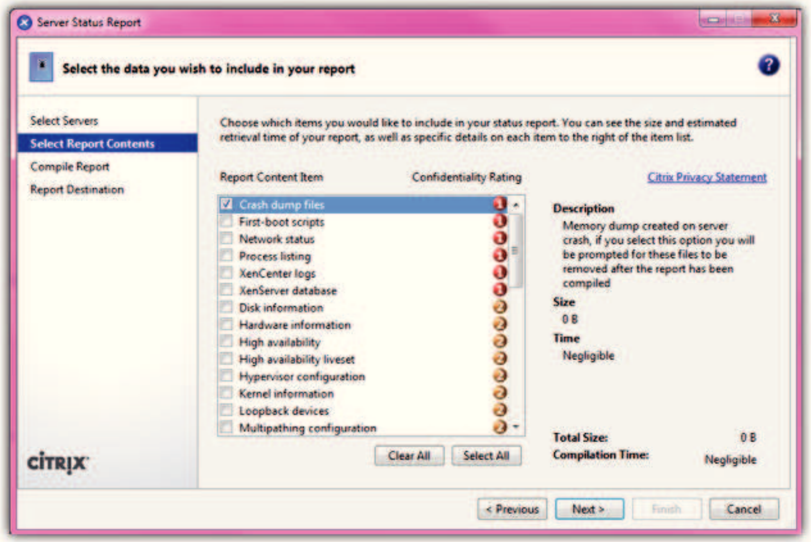

Server Status Report

Select the data you wish to include in your report

- Select Servers
- **Select Report Contents**
- Compile Report
- Report Destination

Choose which items you would like to include in your status report. You can see the size and estimated retrieval time of your report, as well as specific details on each item to the right of the item list.

Citrix Privacy Statement

Report Content Item	Confidentiality Rating
☑ Crash dump files	
□ First-boot scripts	
□ Network status	
□ Process listing	
□ XenCenter logs	
□ XenServer database	
□ Disk information	
□ Hardware information	
□ High availability	
□ High availability liveset	
□ Hypervisor configuration	
□ Kernel information	
□ Loopback devices	
□ Multipathing configuration	

Description
Memory dump created on server crash, if you select this option you will be prompted for these files to be removed after the report has been compiled

Size
0 B

Time
Negligible

Total Size: 0 B
Compilation Time: Negligible

The XenServer crashdump file can be used to troubleshoot XenServer issues. The crashdump file is available through XenCenter and the command-line interface. After the crashdump is obtained through XenCenter, it is up to you to send the file to Citrix Support. Using command-line interface `host-crashdump-upload` to obtain the crashdump file will automatically send the file to the Citrix Support FTP site.

Generating the Crashdump File in the Command-Line Interface

You can use the following command to view the crashdump file for all hosts using the command-line interface:

```
xe host-crashdump-list
```

To obtain a dmesg file from a specific host, add the host selector value at the end of the command.

For example:

```
xe host-crashdump-upload UUID=<host_UUID>
```

Replace host UUID with the UUID of the XenServer host system. For more information about available host selector values or selecting an alternate upload location, see *The XenServer 6.0 Administrator's Guide.*

Test Your Knowledge: Viewing and Generating Reports and Logs

1. What is the main difference between using a server status report and using XenCenter for event logs?
   a. The XenCenter event logs present current status, while server status reports present past data.
   b. The server status report gathers XenServer host information while XenCenter event logs are records of client-side operations and errors.
   c. The server status reports are for assessing current status and event logs are for past events.
   d. The XenCenter event logs gather XenServer host information, while the server status report keep track of client-side operations and errors.

   Answer: B

Running a Network Trace and Triggering a Memory Dump

Citrix Technical Support might request a network trace or memory dump when assisting you with a XenServer issue.

Running a Network Trace from a XenServer Host

For troubleshooting purposes, Citrix Technical Support might ask you to capture a network trace from the XenServer PIF, virtual bridge, the virtual machine, or the VIF.

The `tcpdump` command is used to capture the network traces.

For more information about Network Tracing, see Citrix article CTX120869 on *support.citrix.com.*

Triggering a Memory Dump from a Virtual Machine Running on XenServer

To find the root cause of an issue, such as frozen or unresponsive Windows virtual machines running on XenServer, an analysis of a full memory dump from the system might be necessary.

The preferred method by Microsoft for generating a Windows memory dump is to hold down the right CTRL key and pressing the Scroll Lock key twice.

For more information about Triggering a Memory Dump, see Citrix article CTX123177 on *support.citrix.com.*

Test Your Knowledge: Troubleshooting XenServer

1. Which tool is best for locating the root cause of an unresponsive Windows virtual machine?
    a. Network Trace
    b. Memory Dump
    c. bugtool
    d. perfmon

    Answer: B

XenServer Updates and Hotfixes

Between releases of XenServer, Citrix occasionally releases updates and hotfixes.

Updates	Contain accumulated bug fixes and, occasionally, small feature improvements for XenServer.
Hotfixes	Fix one or more specific issues with XenServer.

Updates and hotfixes can often be applied by:

- Migrating virtual machines away from each host as the hotfix or update is applied
- Apply the hotfix or update
- Updating one host at a time

Migrations of virtual machines away from each host must be performed manually if you are using the xe command-line interface to apply a hotfix or update. Otherwise, using XenCenter to perform an update migrates virtual machines automatically as the update is applied.

Update or Hotfix Considerations

It is important to be aware of the following before applying an update or hotfix:

- Follow the Release Notes that come with each update file. Each update file has unique instructions for installation, particularly with regard to preparatory and post-update operations.
- It is best to update all hosts in a resource pool within a short period. Running a resource pool that includes updated and non-updated hosts for general operation is not supported.
- Citrix recommends that you restart any hosts that you plan to update to ensure that the hosts are healthy and configurations are correct. If there are any pre-existing configuration issues, any updates will fail. XenCenter will restart each host automatically before applying the update file. If you are using the xe command-line interface, then you will have to restart hosts manually.
- Citrix also strongly recommends that you perform a backup of the state of the resource pool, virtual machines metadata, and the virtual disk images in the storage repositories or hosts that you want to update.
- Log on with a user account that has full access permissions.
- Empty the CD/DVD drives of any virtual machines you plan to suspend.
- Disable high availability, if applicable.

To Update XenServer Hosts Using XenCenter

1. Download the update file with the `.xsupdate` file extension to a known location on the computer running XenCenter.
2. Shut down or suspend any virtual machines on the hosts that you want to update.
3. Select **Install New Update** in the Tools menu.
   The Install Update wizard opens.
4. Read the information in Before You Start and click **Next**.
5. Select **Add**, browse to, and select the update file.
6. Click **Open**.
7. Click **Next** to continue.
8. Select the hosts or the pool to update and click **Next**.
9. Follow any recommendations to resolve any update prechecks that have failed.
   Click **Resolve All** for XenCenter to automatically resolve all failed prechecks.
10. Click **Next** to continue.
11. Choose between automatic or manual update mode.
    If you choose automatic, XenCenter will perform any required post-update actions that might be required. If you choose manual, you will need to perform the actions manually. The required post-update actions are listed in a text box. If you want to save the listed actions to a text file for reference, click **Save to File**.
12. Select **Install update** to proceed with the installation.
    The Install wizard shows the progress of the update, printing the major operations that XenCenter performs while updating each host.
13. Click **Finish** to close the Install Update wizard.
14. Perform any post-update actions if you chose manual update mode.

To Update Hosts or a Pool Using the Command-line Interface

1. Download the update file with the `.xsupdate` file extension to a known location on the computer running the xe command-line interface.
2. Shut down or suspend any virtual machines on the host that you want to update by using the `vm-shutdown` or `vm-suspend` commands.
3. Use the following command to upload the update file to the host or the pool which is to be updated:

```
xe -s <server> -u <username> -pw <password> patch-upload file-
name=<filename>
```

XenServer assigns a UUID to the update file, which this command prints. Note the update file UUID.

4. Perform any preparatory alerts that XenServer provides and proceeds when there are no more alerts.
5. Use the `vm-migrate` or the `host-evacuate` commands to migrate virtual machine to other XenServer hosts.
6. Use the following command to update the host or the pool:
   - To update hosts:

     ```
     xe patch-apply host-uuid=<UUID_of_host> uuid=<UUID_of_file>
     ```

   - To update the pool:

     ```
     xe patch-pool-apply uuid=<UUID_of_file>
     ```

7. Verify that the update has been successfully applied by using the `patch-list` command. If the update has been successful, the host field contains the host UUID.
8. Perform any post-update operations as necessary.

The update is applied to the master and then replicated down to all hosts in the pool.

Rolling Pool Upgrade

XenServer allows you to upgrade a pool of XenServer hosts to the next major version while keeping virtual machines on that pool running, thus avoiding service downtime. This is achieved by upgrading on a host-by-host basis, with only one XenServer host offline at a time.

It is not possible to migrate virtual machines located on a XenServer host with a newer XenServer version to one running an older version.

When updating a pool of XenServer hosts in XenCenter, the Rolling Pool Upgrade wizard performs the update path and virtual machine migration automatically. If you need to control the update path and virtual machine migration manually, you can update each host individually.

Start by updating the pool master.

When performing a rolling upgrade, Citrix strongly advises not to perform other virtual machine actions until the upgrade has been completed and recommends that you back up the state of your existing pool using the `pool-dumpdatabase` command-line interface command. This allows you to revert a partially complete rolling upgrade back to its original state without losing any virtual machine data.

Planning a Rolling Upgrade

You should plan your upgrade path carefully. Citrix strongly advises against running a mixed-mode pool, because the pool will be operating in a degraded state during the upgrade. While running in a degraded state, all virtual machines will continue to function as normal, but control operations other than migration might be unavailable. Operations such as `vm-copy`, `vm-start`, and `vm-export` are unavailable. In particular, storage-related operations such as adding, removing, or resizing virtual disks, are unsafe in this mode.

Performing storage-related operations, such as adding, removing, or resizing virtual disks in mixed-mode, might lead to data corruption or loss.

Workload Balancing Reports Overview

Workload Balancing provides reporting on three object types: physical host, resource pools, and virtual machines. Workload Balancing provides two types of reports:

- Historical reports that display information by date
- Roll-up style reports

Workload Balancing also provides some reports for auditing purposes, so you can determine, for example, the number of times a virtual machine moved.

Workload Balancing provides the following reports:

- Performance Optimizations Performance History Report
- Pool Audit Log History
- Pool Health Report
- Pool Health History Report
- Pool Optimization History
- Pool Optimization Performance History Report
- Virtual Machine Motion History Report
- Virtual Machine Performance History Report

Workload Balancing Reports

Utilization Analysis

This report can determine how much of a resource a specific department within your organization has used. Specifically, the information about all the virtual machines in your pool, including their availability and resource utilization. It can help you demonstrate Service Level Agreements compliance and availability, or implement a simple chargeback solution and facilitate billing.

Host Health History

This report displays the performance of resources (CPU, memory, network reads, and network writes) on specific hosts in relation to threshold values.

Pool Optimization Performance History

This report displays optimization events against that pool's average resource usage. Specifically, it displays resource usage for CPU, memory, network reads, and network writes. This report displays average resource usage for the day but does not display peak utilization. This report can help you determine if Workload Balancing is working successfully in your environment. Use this report to see what led up to optimization events or to see how a resource pool is performing if Workload Balancing is not making optimization recommendations.

Pool Audit Trail

This report displays the contents of the XenServer Audit Log, a feature designed to log attempts to perform unauthorized actions and select authorized actions, including import/export, host, and pool backups, and guest and host console access. By default, the Audit Log is always enabled.

Pool Health

This report displays the percentage of time a resource pool and its hosts spent in four different threshold ranges: Critical, High, Medium, and Low. Use this report to evaluate the effectiveness of your performance thresholds.

Pool Health History

This report provides a line graph of resource utilization on all physical hosts in a pool over time. It lets you see the trend of resource utilization. You can evaluate the effectiveness of your performance thresholds by monitoring trends of the data points in this report. Although similar to the Pool Health report, the Pool Health History report displays the average utilization for a resource on a specific date rather than the amount of overall time the resource spent in a threshold.

Pool Optimization History

This report provides chronological visibility into Workload Balancing optimization activity.

Virtual Machine Motion History

This line graph displays the number of times virtual machines moved on a resource pool over a period of time. It indicates if a move resulted from an optimization recommendation and to which host the virtual machine moved. This report also indicates the reason for the optimization. You can use this report to audit the number of moves on a pool.

Virtual Machine Performance History

This report displays performance data for each virtual machine on a specific host for a time period you specify. The performance data is based on the amount of virtual resources allocated for the virtual machine. This report displays data for CPU Usage, Free Memory, Network Reads/Writes, and Disk Reads/Writes. For more information about Workload Balancing Reports, see the *Citrix XenServer 6.0 Workload Balancing Administration Guide* on *http://support.citrix.com*.

851 West Cypress Creek Road Fort Lauderdale, FL 33309 USA (954) 267 3000 www.citrix.com
Rheinweg 9 8200 Schaffhausen Switzerland +41 (0) 52 63577 00 www.citrix.com