

lenovo[®]

Implementing Lenovo Client Virtualization with Citrix XenDesktop

**Introduces Lenovo x86 servers and
Citrix XenDesktop offerings**

**Reviews design, planning, and
deployment considerations**

**Provides step-by-step configuration
guidance**

**Describes VMware vSphere and
Microsoft Hyper-V implementation
scenarios**

Ilya Krutov

David Blair

Reza Fanaei Aghdam

Gica Livada





Implementing Lenovo Client Virtualization with Citrix XenDesktop

May 2015

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Last update on May 2015

This edition applies to System x, Flex System, and RackSwitch offerings that were announced prior to 4Q/2014, and Citrix XenDesktop version 7.5.

© Copyright Lenovo 2015. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Part 1. Introduction to Virtual Desktop Infrastructure	1
Chapter 1. Lenovo Client Virtualization overview	1
1.1 Virtual desktop infrastructure overview	2
1.2 Lenovo Client Virtualization	2
1.3 Citrix XenDesktop	5
Chapter 2. Components of the virtual desktop infrastructure	9
2.1 Planning for Lenovo VDI components	10
2.2 VDI servers	11
2.2.1 Flex System Enterprise Chassis	11
2.2.2 Flex System compute nodes	12
2.2.3 System x3550 M4	16
2.2.4 System x3650 M4	17
2.3 Networking components	18
2.3.1 Flex System networking I/O modules	18
2.3.2 RackSwitch offerings	22
2.3.3 Virtual Fabric adapters	27
2.4 Storage components	29
2.4.1 Fibre Channel connectivity	29
2.4.2 Converged fabrics	31
2.4.3 Solid-state drives in the VDI solution	32
2.4.4 RAID considerations	33
2.4.5 Flex System Storage Expansion Node	33
2.4.6 IBM Storwize V7000	34
2.4.7 IBM Storwize V3700	36
2.5 Management components	39
2.5.1 Integrated Management Module II	39
2.5.2 Chassis Management Module	39
2.5.3 Flex System Manager	40
2.5.4 Introduction to Upward Integration	41
Part 2. VDI design considerations	45
Chapter 3. VMware vSphere design considerations	47
3.1 ESXi and vSphere features	48
3.1.1 ESXi hypervisor	48
3.1.2 VMware vCenter Server	49
3.1.3 vMotion	49
3.1.4 Distributed Resource Scheduler	51
3.1.5 High Availability	52
3.1.6 vSphere licensing considerations	53
3.1.7 System x integration with VMware	54

3.2 Networking considerations	54
3.2.1 Virtual switch	54
3.2.2 Ports and port groups	55
3.2.3 Uplink ports	55
3.3 Storage considerations	56
3.3.1 Local or shared storage	56
3.3.2 Tiered storage	56
3.3.3 Redundancy	57
Chapter 4. Microsoft Hyper-V design considerations	59
4.1 Hyper-V virtualization and management features	60
4.1.1 Hyper-V overview	60
4.1.2 Hyper-V management	61
4.1.3 Lenovo management devices and tools for Hyper-V	61
4.1.4 VM and Storage Migration	62
4.1.5 VM placement	62
4.1.6 High Availability with Hyper-V clusters	63
4.2 Networking considerations	63
4.2.1 Hyper-V Network Virtualization	64
4.3 Storage Considerations	64
4.3.1 Local or shared storage	64
4.3.2 Tiered storage	64
4.3.3 Usage of flash based storage	65
4.3.4 Redundancy and load balancing	65
Chapter 5. Citrix XenDesktop design considerations	67
5.1 Citrix XenDesktop components	68
5.2 Desktop and application delivery	70
5.3 Citrix XenDesktop provisioning	71
5.3.1 Provisioning Services solution	72
5.3.2 Machine Creation Services	74
5.3.3 Personal vDisk	75
5.3.4 Image assignment models	76
5.4 Storage configuration	76
5.5 Network configuration	79
5.6 Operational model and sizing guidelines	80
5.6.1 VDI server configuration	80
5.6.2 Shared storage	84
Part 3. VDI deployment and management	89
Chapter 6. Citrix XenDesktop lab environment	91
6.1 Lab environment	92
6.2 Use case for the lab environment	93
6.3 Component model	95
6.4 Operational model	95
6.5 Logical design	97
6.5.1 Ethernet segment	97
6.5.2 Storage disk and host mapping	99
Chapter 7. Deploying Flex System	101
7.1 Initial configuration of the Chassis Management Module	102
7.1.1 Connecting to the Chassis Management Module	102

7.1.2	Using the initial setup wizard.	104
7.1.3	Configuring IP addresses for the chassis components.	112
7.2	Firmware updates and basic hardware configuration	113
7.3	Configuring Active Directory Integration for CMM.	121
7.4	Configuring the EN4093 10Gb Ethernet Switch	122
7.5	Enabling UFP on the x240 compute node	130
7.6	Configuring iSCSI on the x240 compute node	131
7.7	V7000 configuration	132
7.7.1	V7000 initial configuration.	132
7.7.2	V7000 Storage Node setup wizard	134
7.7.3	Configuring storage volumes	140
7.7.4	Configuring hosts	144
Chapter 8.	Deploying Citrix XenDesktop	147
8.1	Configuring utility services	148
8.2	Provisioning VMs for Citrix XenDesktop components.	149
8.2.1	Installing the Citrix License Server	149
8.2.2	Configuring the licenses	151
8.3	Installing Citrix XenDesktop Controller	153
8.3.1	Installing the XenDesktop Controller.	154
8.3.2	Advanced settings.	161
8.4	Installing Citrix XenApp	163
8.5	Installing Citrix StoreFront.	167
8.5.1	Configuring the StoreFront	169
8.6	Installing Citrix Provisioning Services	171
8.6.1	Installing the Citrix Provisioning Console	177
Chapter 9.	Operating Citrix XenDesktop.	179
9.1	Introduction	180
9.2	Configuring the gold image	180
9.2.1	Preparing the gold image for streaming services	180
9.2.2	Preparing the gold image for persistent desktops.	194
9.3	Configuring desktop distribution	198
9.3.1	Configuring streamed desktops	198
9.3.2	Configuring streaming desktops with personal VDisk	208
9.3.3	Configuring persistent desktops	216
9.3.4	Assigning a catalog to a group	222
9.4	Roaming profiles and folder redirection	225
9.4.1	Configuring the roaming profile.	226
9.4.2	Configuring folder redirection	236
9.4.3	Configuring the Citrix Receiver	239
9.4.4	Group Policy Object link	240
9.4.5	Configuring application distribution.	241
Chapter 10.	Managing System x and Flex System hardware in a VDI environment	243
10.1	Managing a vSphere environment with UIM	244
10.1.1	Enabling UIMs for a newly added ESXi host.	244
10.1.2	Collecting system inventory with UIM	245
10.1.3	Monitoring hardware status.	250
10.1.4	Using PFA alert to move VMs to another ESXi host.	251
10.1.5	Rolling firmware upgrades	255
10.1.6	Changing IMM and UEFI configuration.	257
10.2	Managing a Windows Server environment with UIM.	261
10.2.1	Enabling Hardware Monitoring on the Flex System	262

10.2.2 Deploying System Center agents for hardware monitoring.	268
10.2.3 Monitoring hardware status in SCOM.	273
10.2.4 Lenovo Hardware Performance and Resource Optimization Pack.	276
10.2.5 Rolling firmware upgrades by using UIM for System Center VMM.	277
Abbreviations and acronyms	283
Related publications	285
Lenovo Press publications	285
Online resources	285

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www.lenovo.com/legal/copytrade.html>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

BladeCenter®	RackSwitch™	UpdateXpress System Packs™
Flex System™	Lenovo(logo)®	VMready®
Lenovo®	ServerProven®	vNIC™
Omni Ports™	System x®	

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

The Lenovo® Client Virtualization offers robust, cost-effective, and manageable virtual desktop solutions for a wide range of clients, user types, and industry segments. These solutions help to increase business flexibility and staff productivity, reduce IT complexity, and simplify security and compliance. Based on a reference architecture approach, this infrastructure supports various hardware, software, and hypervisor platforms.

The Lenovo Client Virtualization solution with Citrix XenDesktop that is running on System x® rack and blade servers offers tailored solutions for every business, from the affordable all-in-one Citrix VDI-in-a-Box for simple IT organizations to the enterprise-wide Citrix XenDesktop. XenDesktop is a comprehensive desktop virtualization solution with multiple delivery models that is optimized for flexibility and cost-efficiency.

This Lenovo Press publication provides an overview of the Lenovo Client Virtualization solution, which is based on Citrix XenDesktop that is running on System x rack and blade servers. It highlights key components, architecture, and benefits of this solution. It also provides planning and deployment considerations, and step-by-step instructions about how to perform specific tasks.

This book is intended for IT professionals who are involved in the planning, design, deployment, and management of the Lenovo Client Virtualization that is built on System x family of servers that are running Citrix XenDesktop.

Authors

This book was produced by in collaboration with Lenovo Press by a team of subject matter experts from around the world.

Ilya Krutov is a Project Leader at Lenovo Press. He manages and produces pre-sale and post-sale technical publications for various IT topics, including x86 rack and blade servers, server operating systems and software, virtualization and cloud, and data center networking. Ilya has more than 15 years of experience in the IT industry, performing various roles, including Team Leader, Portfolio Manager, Brand Manager, IT Specialist, and Certified Instructor. He has written more than 200 books, papers, and other technical documents. He has a Bachelor's degree in Computer Engineering from the Moscow Engineering and Physics Institute (Technical University).

David Blair has 20 years of experience in the IT industry. As a Product Field Engineer for 16 years, he worked on new technologies, including the early implementations of Microsoft's Cluster Server on Windows NT, through to the latest version on 2012 R2. In that time, he gained much experience with Storage technologies, including Fibre Channel, iSCSI, and more recently Fibre Channel over Ethernet. For the last two years, he worked as a pre-sales engineer on PureFlex System and Flex System™, particularly in the virtualization arena on server workloads and more recently on client workload provision.

Reza Fanaei Aghdam is a Senior IT Specialist working in Zurich, Switzerland. He has 19 years of professional experience with x86-based hardware, storage technologies, and systems management, with more than 12 years at IBM. He instructs Business Partners and customers about how to configure and install System x, BladeCenter®, Systems Director, Storage, VMware, and Hyper-V. He is an IBM Certified Systems Expert (System x, BladeCenter), Midrange Storage Technical Support, and VMware Certified Professional.

Gica Livada is a Certified IT Specialist with more than 20 years experience in the IT field. He joined IBM in 2006 and has held several positions, including System Administrator, Customer Technical Leader, and Virtualization Specialist. He is a member of the VMware Center of Excellence team and preparing to become an IT Architect. He is passionate about virtualization and cloud technologies, and he has multiple certifications from Citrix, Microsoft, and NetApp.

Thanks to the authors of the first edition of this book:

- ▶ Ilya Krutov
- ▶ Andreas Groth
- ▶ Gica Livada
- ▶ Diego Pereira
- ▶ Jean-Baptiste Valette
- ▶ Brad Wasson

Thanks to the following people for their contributions to this project:

Amy Freeman
Michael Perks
David Watts
Lenovo

Karen Lawrence
IBM ITSO

Part 1

Introduction to Virtual Desktop Infrastructure

In this part, we introduce virtual desktop infrastructure (VDI) and provide an overview of its components and building blocks. This part includes the following chapters:

- ▶ Chapter 1, “Lenovo Client Virtualization overview” on page 1
- ▶ Chapter 2, “Components of the virtual desktop infrastructure” on page 9

Lenovo Client Virtualization overview

This chapter introduces Lenovo Client Virtualization and describes one of its solutions with Citrix XenDesktop.

This chapter includes the following topics:

- ▶ 1.1, “Virtual desktop infrastructure overview” on page 2
- ▶ 1.2, “Lenovo Client Virtualization” on page 2
- ▶ 1.3, “Citrix XenDesktop” on page 5

1.1 Virtual desktop infrastructure overview

Today, businesses are looking for ways to securely bring in new ways for people to communicate at work without having to limit them to an office. Personal tablets, smartphones, and netbooks now dominate a landscape that was owned by the personal computer. Delivering the same business applications securely to these new devices drives the adoption of the virtual desktop infrastructure (VDI).

VDI is based on a desktop-centric model to provide an environment to the remote networked-based user. The user accesses the desktop by using a remote display protocol on their device in a secure manner. The resources are centralized and users can move between locations while accessing the applications and data. By using this feature, administrators have better control over the management of the desktop and tighter security.

One of the most important aspects of deploying a virtual desktop solution is to control costs while providing familiar user experience and functionality. The other important aspect is the ability to scale to the demanding needs of the user. Too many times, businesses are excited by a solution but soon outgrow the initial deployment and find it difficult to add the next 100 users or 100 TB of storage. Therefore, careful planning and analysis must be done to ensure the successful implementation of VDI projects.

Lenovo VDI solutions are consolidated under the Lenovo Client Virtualization umbrella.

1.2 Lenovo Client Virtualization

Lenovo Client Virtualization offers robust, cost-effective, and manageable virtual desktop solutions for a wide range of clients, user types, and industry segments. These solutions can help to increase business flexibility and staff productivity, reduce IT complexity, and simplify security and compliance. Based on a reference architecture approach, this infrastructure supports various hardware, software, and hypervisor platforms.

The Lenovo Client Virtualization solution with Citrix XenDesktop that is running on Lenovo x86 rack and blade servers offers tailored solutions for every business, from the affordable all-in-one Citrix VDI-in-a-Box to the enterprise-wide Citrix XenDesktop. XenDesktop is a comprehensive desktop virtualization solution with multiple delivery models that is optimized for flexibility and cost-efficiency.

The hosted virtual desktop (HVD) approach, combined with the hosted applications, is the most common form of implementing a virtualized user desktop environment. With HVDs, all applications and data that the user interacts with are stored centrally and securely in the data center. These applications never leave the data center boundaries. This setup makes management and administration much easier and gives users access to data and applications from anywhere and at anytime.

The following drivers are key for virtual desktops in today's business climate:

- ▶ Data security and compliance concerns
- ▶ Complexity and costs of managing existing desktop environments
- ▶ An increasingly mobile workforce
- ▶ The changing ownership of endpoint devices with bring-your-own-device (BYOD) programs
- ▶ The need for rapid recovery from theft, failure, and disasters

Lenovo Client Virtualization offers the following benefits:

- ▶ Lowers the total cost of ownership (TCO) over an extended period compared to traditional PCs
- ▶ Simplifies desktop administration, support, and management
- ▶ Enhances security and compliance management
- ▶ Improves availability and reliability
- ▶ Enables users to work anytime, anywhere quickly and easily, regardless of location or device
- ▶ Supports growth initiatives for mobility and flexible work locations better

The Lenovo Client Virtualization solution with Citrix XenDesktop that is running on Lenovo x86 servers includes the following components:

- ▶ Virtual infrastructure software: Citrix XenDesktop
- ▶ Hardware platform:
 - System x
 - Flex System
 - IBM Storwize family from Lenovo

Figure 1-1 shows the functional components of the Lenovo Client Virtualization solution.

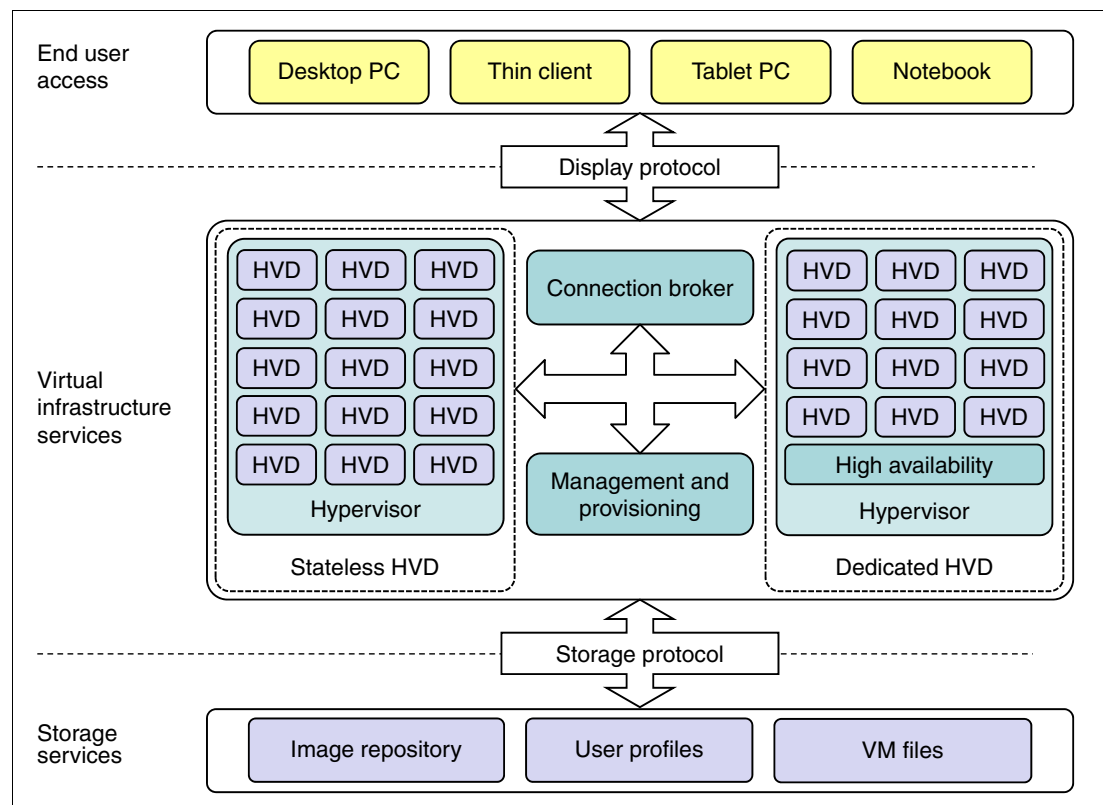


Figure 1-1 Lenovo Client Virtualization functional components

The Lenovo Client Virtualization solution consists of the following functional layers:

- ▶ User access layer

The user access layer is a user entry point into the virtual infrastructure. Devices that are supported at this layer include traditional desktop PCs, thin clients, notebooks, and handheld mobile devices.

- ▶ Virtual infrastructure services layer

The virtual infrastructure services layer provides the secure, compliant, and highly available desktop environment to the user. The user access layer interacts with the virtual infrastructure layer through display protocols. The Remote Desktop Protocol (RDP), half-duplex (HDX), and Independent Channel Architecture (ICA) display protocols are available in Citrix XenDesktop.

- ▶ Storage services layer

The storage services layer stores user persona, profiles, gold master images, and actual virtual desktop images. The storage protocol is an interface between virtual infrastructure services and storage services. The storage protocols supported by Citrix XenDesktop include Network File System (NFS), Common Internet File System (CIFS), iSCSI, and Fibre Channel.

The virtual infrastructure services layer includes the following key functional components:

- ▶ Hypervisor

The hypervisor provides a virtualized environment for running virtual machines (VMs) with the desktop operating systems in them. These VMs are called *hosted virtual desktops*.

- ▶ Hosted virtual desktops

An HVD is a VM that runs a user desktop operating system and applications.

- ▶ Connection broker

This broker is the point of contact for the client access devices that request the virtual desktops. The connection broker manages the authentication function and ensures that only valid users can access to the infrastructure. When authenticated, it directs the clients to their assigned desktops. If the virtual desktop is unavailable, the connection broker works with the management and provisioning services to have the VM ready and available.

- ▶ Management and provisioning services

The management and provisioning services allow the centralized management of the virtual infrastructure, which provides a single console to manage multiple tasks. They provide image management, lifecycle management, and monitoring for hosted VMs.

- ▶ High availability services

High availability (HA) services ensure that the VM is up and running, even if a critical software or hardware failure occurs. HA can be a part of connection broker function for stateless HVDs or a separate failover service for dedicated HVDs.

The following types of assignment models are available for the user HVDs:

- ▶ Persistent

A *persistent* (also known as stateful or dedicated) HVD is assigned permanently to the specific user (similar to a traditional desktop PC). Users log in to the same virtual desktop image when they connect. All changes that they make and each application that they install are saved when the user logs off. The dedicated desktop model is best for users who need the ability to install more applications, store data locally, and retain the ability to work offline.

- ▶ Non-persistent

A *non-persistent* (also known as pooled or stateless) HVD is allocated temporarily to the user. After the user logs off, changes to the image are discarded (reset). Then, the desktop becomes available for the next user, or a desktop is created for the next user session. A persistent user experience (the ability to personalize the desktop and save data) is achieved through user profile management, folder redirection, and similar approaches. Specific individual applications can be provided to nonpersistent desktops by using application virtualization technologies, if required.

Functional layers and components are supported by a hardware infrastructure platform that provides the following features:

- ▶ Sufficient computing power to support demanding workloads
- ▶ Scalability to satisfy future growth requirements
- ▶ Reliability to support business continuity and 24x7 operations
- ▶ High-speed, low-latency networking for a better user experience
- ▶ Cost-efficient storage to handle large amounts of VM and user data
- ▶ Centralized management of combined physical and virtual infrastructure from a single user interface to simplify and automate deployment, maintenance, and support tasks

System x rack and blade servers in a Lenovo Client Virtualization solution can help to achieve the following advantages:

- ▶ Better VM density because of support for top Intel Xeon processors and large memory and I/O capacity
- ▶ Better virtual desktop performance and better use of VDI server resources with flexible local SSD support
- ▶ Transparent support for high-performance remote graphics with GPU adapters installed
- ▶ Simplified deployment and management of physical and virtual infrastructures because of System x management capabilities

1.3 Citrix XenDesktop

Lenovo Client Virtualization with Citrix XenDesktop can help to transform Microsoft Windows desktops, applications, and data into a cloud-type service that is accessible on virtually any device, anywhere. Citrix offers tailored solutions that range from the affordable, all-in-one Citrix VDI-in-a-Box for simple IT organizations to the enterprise-wide Citrix XenDesktop. XenDesktop is a comprehensive desktop virtualization solution for every user with multiple delivery models that are optimized for flexibility and cost efficiency. Both solution types deliver a rich, high-definition user experience across any network that uses Citrix HDX technologies.

By using the open architecture of Citrix XenDesktop, clients can adopt desktop virtualization quickly and easily with any hypervisor, storage, or management infrastructure.

The following Citrix XenDesktop features provide a familiar experience for the user:

- ▶ Multiple monitor support
- ▶ 3D graphics business application support
- ▶ Multimedia support
- ▶ Printing from a virtual desktop
- ▶ Accessing USB devices and other peripheral devices
- ▶ Roaming user profiles

Citrix XenDesktop offers several levels of security, including the following features:

- ▶ Multi-factor authentication
- ▶ Traffic encryption
- ▶ Built-in password management
- ▶ Secure Sockets Layer (SSL) tunneling to ensure that all connections are encrypted

The following Citrix XenDesktop features provide centralized administration and management:

- ▶ Microsoft Active Directory
- ▶ Web-based administrative console
- ▶ Automated desktop provisioning and storage optimization

Citrix XenDesktop includes the following scalability, integration, and optimization features:

- ▶ VMware vSphere, Microsoft Hyper-V, and XenServer hypervisor support
- ▶ Integration with VMware vCenter to achieve cost-effective densities, high levels of availability, and advanced resource allocation control for virtual desktops
- ▶ Automated provisioning of desktop images that share virtual disks with a master image

Citrix XenDesktop software components are shown in Figure 1-2.

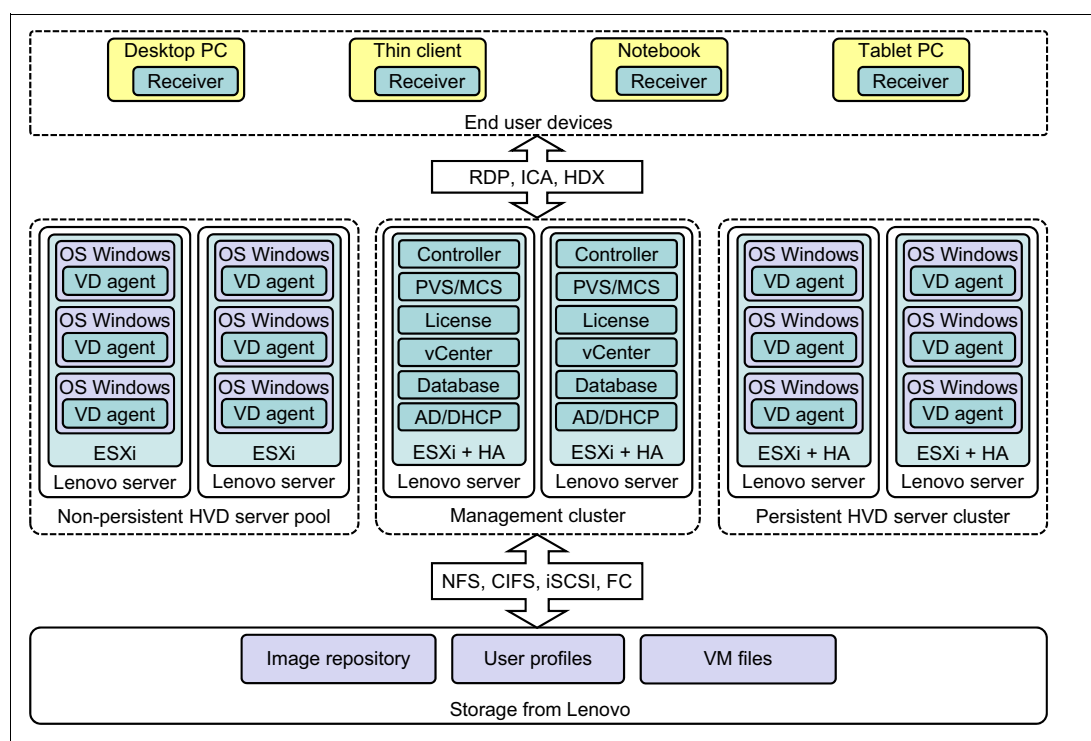


Figure 1-2 Citrix XenDesktop software components

The Citrix XenDesktop core services have the following software components:

- ▶ Citrix Receiver
Citrix Receiver is a client software for accessing virtual desktops by using the Independent Channel Architecture (ICA) protocol. The client software can run on different types of user access devices, including desktop PCs, notebooks, thin clients, and others.

- ▶ Citrix Virtual Desktop Agent
Citrix Virtual Desktop Agent is installed on virtual desktops and supports Citrix Receiver direct connections through the ICA.
- ▶ Citrix XenDesktop Controller
Citrix XenDesktop Controller is a software service that is responsible for connection brokering, authenticating users, and starting virtual desktops and user persona management, if required. Authentication of users is performed through Windows Active Directory.
- ▶ Citrix Provisioning Services or Machine Creation Services
Citrix Provisioning Services and Machine Creation Services create and provision virtual desktops from desktop images. Provisioning Services support stateless HVD pools, and Machine Creation Services can support stateless and dedicated HVD pools.
- ▶ Citrix License Server
Citrix License Server manages licenses for all XenDesktop components.
- ▶ Citrix Data Store
Citrix Data Store is a database that stores configuration information for the XenDesktop environment.
- ▶ VMware ESXi
VMware ESXi is a hypervisor that is used to host VMs.
- ▶ VMware vCenter
The VMware vCenter service acts as a central administrator for VMware ESX and ESXi servers that are connected on a network. vCenter Server provides a central point for configuring, provisioning, and managing VMs in the data center.

Components of the virtual desktop infrastructure

This chapter introduces Lenovo VDI solution components (servers, networking, storage, and management) to consider during the design of the virtual desktop infrastructure (VDI). It also provides guidelines about how to use the components to optimize the solution.

This chapter includes the following topics:

- ▶ 2.1, “Planning for Lenovo VDI components” on page 10
- ▶ 2.2, “VDI servers” on page 11
- ▶ 2.3, “Networking components” on page 18
- ▶ 2.4, “Storage components” on page 29
- ▶ 2.5, “Management components” on page 39

2.1 Planning for Lenovo VDI components

To design your Lenovo VDI infrastructure, you must determine the resources that are needed by your infrastructure servers and persistent and non-persistent desktops.

Each category of user operates a specific software platform with a specific workload, which involves different hardware resources. Consumption assessment on resource usage must be performed for each category of user for the following resources:

- ▶ Processor
- ▶ Memory
- ▶ I/O characteristics: size, percentage of reads and writes, and type of access (random or sequential)
- ▶ Size of user data and user profile
- ▶ Graphic usage profile

Then, for each category of user or workload profile, you can translate the assessed requirements into compute node resources. Processor, memory, and graphic requirements must be considered for VDI server design. Requirements for I/O and storage for data determine the network and storage design.

Consider the following points when you are sizing your VDI servers:

- ▶ Do not overcommit memory because disk swapping deteriorates the performance.
- ▶ Do not overcommit processors. If too many virtual machines (VMs) are used, the response time deteriorates quickly.
- ▶ Plan for failover. If one or more compute nodes fail, the user VMs that are hosted on the failed compute nodes must be reallocated over the remaining compute nodes. Consider allowing for overhead of 20% in memory and processor to support these extra VMs without reaching the compute node resource boundaries.
- ▶ The hypervisor often uses 3 GB - 6 GB of server memory and one processor core.

To define the storage solution, consider the subject in the following parts:

- ▶ Storage for the infrastructure servers. A shared storage is the best solution.
- ▶ Storage for the persistent VMs. Privilege is also a shared storage.
- ▶ Storage for the non-persistent VMs. Consider the use of local storage or shared storage with high I/O performance.
- ▶ Storage connectivity. Consider a separate Fibre Channel SAN to achieve potentially better performance, availability, scalability, and security with moderate to heavy storage workloads. Consider converged FCoE or iSCSI or unified NAS storage to achieve potentially better total cost of ownership (TCO) with light-to-medium storage workloads.

To achieve higher availability, consider redundancy for the I/O modules: Ethernet network switches and storage SAN switches.

2.2 VDI servers

This section describes Lenovo VDI server offerings and includes the following topics:

- ▶ 2.2.1, “Flex System Enterprise Chassis”
- ▶ 2.2.2, “Flex System compute nodes” on page 12
- ▶ 2.2.3, “System x3550 M4” on page 16
- ▶ 2.2.4, “System x3650 M4” on page 17

2.2.1 Flex System Enterprise Chassis

The Flex System Enterprise Chassis with its flexible design is a 10U integrated infrastructure platform with integrated chassis management that supports a mix of compute, storage, and networking resources to meet the IT demands. It is designed for a simple deployment and can scale up to meet future needs. It also meets the needs of varying workloads with scalable IT resource pools for higher usage and lower cost per workload.

Although increased security and resiliency protect vital information and promote maximum uptime, the integrated, easy-to-use management system can reduce setup time and complexity, which provides a quicker path to return on investment.

The Flex System chassis is shown in Figure 2-1.



Figure 2-1 Flex System chassis

The Flex System Enterprise Chassis has 14 node bays that support up to 14 half-width, one-bay compute nodes, or up to seven full-width two-bay x86 compute nodes. You can use one-bay and two-bay compute nodes to meet your specific hardware needs. Also, the rear of the chassis has four high-speed networking switches bays. The compute nodes share common resources, such as power, cooling, management, and I/O resources in the chassis.

The chassis' I/O architecture with flexibility in fabric and speed and the ability to use Ethernet, InfiniBand, Fibre Channel, FCoE, and iSCSI can meet the growing and future I/O needs of large and small businesses.

The Flex System Enterprise Chassis includes the following key features:

- Flexibility and efficiency

The 14 bays in the chassis allow the installation of compute or management nodes, with networking modules in the rear. A single chassis or a group of chassis can be fully customized to the specific needs of the computing environment. IT can meet the needs of the business by using a single system for multiple operating environments.

- Easily scalable with simple administration

Because the Flex System Enterprise Chassis is an all-in-one solution, it is designed for growth from a single chassis to many. Adding compute or networking capability is as simple as adding nodes, modules, or chassis. The simple, highly integrated management system allows you to use the Chassis Management Modules that are integrated into each chassis to administer a single chassis, or Flex System Manager that controls up to 16 chassis from a single panel.

- Designed for multiple generations of technology

The Flex System Enterprise Chassis is designed to be the foundation of your IT infrastructure now and into the future. Compute performance requirements are always on the rise and networking demands continue to grow with rising bandwidth needs and a shrinking tolerance for latency. The chassis is designed to scale to meet the needs of your future workloads and offer the flexibility to support current and future innovations in compute, storage, and networking technology.

2.2.2 Flex System compute nodes

The following choices of compute nodes are wide, and depend on the computing requirements for the VMs hosted:

- Flex System x222 is designed for virtualization, dense cloud deployments, and hosted clients. It is a good choice for the clients that want to virtualize their general-purpose user applications while maximizing the density of their computing resources.
- Flex System x240 is a good choice for VDI workloads that require more memory and I/O bandwidth.
- For resource demanding VMs, Flex System x440 brings massive compute power and memory resources. A high VM density on a compute node can be reached. The effect on the users in a compute node failure is proportional.

Table 2-1 lists key features of the compute nodes.

Table 2-1 x222, x240, and x440 compute node feature comparison

Feature	x222 (one half)	x240	x440
Processor	E5-2400	E5-2600 v2	E5-4600
Number of sockets	2	2	4
Memory (max)	384 GB	768 GB	1.5 TB
Local storage (max)	1 TB	3.2 TB	3.2 TB
I/O ports (max)	4	8	16

The following sections describe these compute nodes:

- ▶ “Flex System x222 Compute Node”
- ▶ “Flex System x240 Compute Node” on page 14
- ▶ “Flex System x440 Compute Node” on page 15

Flex System x222 Compute Node

The Flex System x222 Compute Node is a high-density dual-server that is designed to maximize the computing power that is available in the data center. With a balance between cost and system features, the x222 is an ideal platform for dense workloads, such as virtualization, cloud deployments, and hosted clients.

The x222 has two independent servers in one mechanical package, which means that the x222 has a double-density design that allows up to 28 servers to be housed in a single 10U Flex System Enterprise Chassis.

The x222 is the ideal platform for clients that want to virtualize their workloads while maximizing the density of their computing resources.

A Flex System x222 Compute Node is shown in Figure 2-2.



Figure 2-2 Flex System x222 Compute Node

This half-wide high-density server offers the following key features for VDI:

- ▶ **Processor:** The Intel Xeon Processor E5-2400 with up to 8-core processors and up to 2.4 GHz core speeds, depending on the CPU's number of cores, up to 20 MB of L3 cache, and QPI interconnect links of up to 8 Gigaticks per second (GTps) is available. The x222 supports up to 32 cores with up to four processors in a standard (half-width) Flex System form factor.

Note: The two servers are independent and cannot be combined to form a single four-socket system.

- ▶ **Memory:** Up to 24 DIMM sockets in a standard (half-width) Flex System form factor is available. Each server provides up to 12 DIMM sockets DDR3 ECC memory. RDIMMs provide speeds up to 1600 MHz and a memory capacity of up to 384 GB. Load-reduced DIMMs (LRDIMMs) are supported by a maximum capacity of 768 GB.
- ▶ **Network:** Up to eight virtual I/O ports per each server (up to 16 per one node) with integrated 10 GbE ports (for more information, see “Virtual Fabric adapters” on page 27), which offers the choice of Ethernet, Fibre Channel, iSCSI, or FCoE connectivity.
- ▶ **Disk:** Each half-height server has one 2.5-inch simple-swap SATA drive bay that supports SATA drives and SSDs. Optional solid-state drive (SSD) mounting kit to convert a 2.5-inch simple-swap bay into two 1.8-inch hot-swap SSD bays also is available.

Flex System x240 Compute Node

The Flex System x240 Compute Node is a high-performance Intel Xeon processor-based server that offers outstanding performance for virtualization with new levels of processor performance and memory capacity, and high networking bandwidth.

The x240 Compute Node is an efficient server that runs a broad range of workloads. Armed with advanced management capabilities, by using this Compute Node, you can manage your physical and virtual IT resources from a single pane of glass.

The x240 Compute Node is a high-availability, scalable compute node that is optimized to support the next-generation microprocessor technology and is ideally suited for medium and large businesses. A Flex System x240 Compute Node is shown in Figure 2-3.



Figure 2-3 Flex System x240 Compute Node

This half-wide server offers the following key features for VDI:

- ▶ **Processor:** The Intel Xeon Processor E5-2600 v2 with up to 12-core processors and up to 3.5 GHz core speeds is available depending on the CPU's number of cores, up to 30 MB of L3 cache, and QPI interconnect links of up to two 8 GTps. Up to 2 processors, 24 cores, and 48 threads maximize the concurrent running of multi-threaded applications.
- ▶ **Memory:** Up to 24 DDR3 ECC memory RDIMMs provide speeds up to 1866 MHz and a memory capacity of up to 384 GB. Load-reduced DIMMs (LRDIMMs) are supported by a maximum capacity of 768 GB.
- ▶ **Network:** Up to 16 virtual I/O ports per compute node with integrated 10 Gb Ethernet ports (for more information, see "Virtual Fabric adapters" on page 27) is available, which offers the choice of Ethernet, Fibre Channel, iSCSI, or FCoE connectivity.
- ▶ **Disk:** Two 2.5-inch hot-swap SAS/SATA drive bays support SAS, SATA, and SSDs. Support for up to eight 1.8-inch SSDs is available.

The x240 compute node can also be equipped with the Flex System PCIe Expansion Node, which is used to attach extra PCI Express cards, such as next-generation graphics processing units (GPUs), to it. This capability is ideal for many desktop applications that require hardware acceleration that use a PCI Express GPU card.

Flex System PCIe Expansion Node

For VDI, you can use the Flex System PCIe Expansion Node to attach next-generation GPUs to x240 compute nodes. The PCIe Expansion Node supports up to four PCIe adapters and two other Flex System I/O expansion adapters.

Figure 2-4 shows the PCIe Expansion Node that is attached to a compute node.



Figure 2-4 Flex System PCIe Expansion Node attached to a compute node

The PCIe Expansion Node includes the following features:

- ▶ Support for up to four standard PCIe 2.0 adapters:
 - Two PCIe 2.0 x16 slots that support full-length, full-height adapters (1x, 2x, 4x, 8x, and 16x adapters supported)
 - Two PCIe 2.0 x8 slots that support low-profile adapters (1x, 2x, 4x, and 8x adapters supported)

- ▶ Support for PCIe 3.0 adapters by operating them in PCIe 2.0 mode
- ▶ Support for one full-length, full-height double-wide adapter (by using the space of the two full-length, full-height adapter slots)
- ▶ Support for PCIe cards with higher power requirements

The Expansion Node provides two auxiliary power connections, up to 75 W each for a total of 150 W of more power by using standard 2x3, +12 V six-pin power connectors. These connectors are placed on the base system board so that they both can provide power to a single adapter (up to 225 W), or to two adapters (up to 150 W each). Power cables are used to connect from these connectors to the PCIe adapters and are included with the PCIe Expansion Node.

- ▶ Two Flex System I/O expansion connectors

These I/O connectors expand the I/O capability of the compute node.

The following PCIe GPU adapters can be used in the VDI solutions with the PCIe Expansion Node:

- ▶ NVIDIA GRID K1 for Flex System PCIe Expansion Node
- ▶ NVIDIA GRID K2 for Flex System PCIe Expansion Node

NVIDIA GRID GPU adapters are designed for VDI applications in fields including seismic processing; computational biology and chemistry; weather and climate modeling; image, video and signal processing; computational finance, computational physics; CAE and CFD; and data analytics. NVIDIA GRID cards can be shared between multiple concurrent users to support heavy 3D applications and simulations.

Flex System x440 Compute Node

The Flex System x440 Compute Node is a four-socket Intel Xeon processor-based server that is optimized for high-end virtualization, mainstream database deployments, and memory-intensive high performance environments.

Compared to the x240 compute node, it provides double the amount of memory capacity and processor sockets, and a high networking bandwidth. A Flex System x440 Compute Node is shown in Figure 2-5.



Figure 2-5 Flex System x440 Compute Node

This full-wide server offers the following key features for VDI:

- ▶ Processor: The Intel Xeon processor E5-4600 with 8-core processors and up to 2.9 GHz core speeds, up to 20 MB of L3 cache, and up to two 8 GTps QPI interconnect links is available. Up to 4 processors, 32 cores, and 64 threads maximize the concurrent running of multithreaded applications.
- ▶ Memory: Up to 48 DDR3 ECC memory RDIMMs provide speeds up to 1600 MHz and a memory capacity of up to 768 GB. Load-reduced DIMMs (LRDIMMs) are supported by a maximum capacity of 1.5 TB of memory.
- ▶ Network: Up to 32 virtual I/O ports per compute node with integrated 10 Gb Ethernet ports offer the choice of Ethernet, Fibre Channel, iSCSI, or FCoE connectivity. Optionally, you can have up to 64 virtual I/O ports by installing four CN4054 10Gb Virtual Fabric adapters.
- ▶ Disk: Two 2.5-inch hot-swap SAS/SATA drive bays support SAS, SATA, and SSD drives.

2.2.3 System x3550 M4

The x3550 M4 is a cost- and density-balanced 1U, 2-socket business-critical server that offers improved performance and pay-as-you grow flexibility with new server management features. Its energy-efficient design supports more cores, memory, and data capacity in a scalable 1U package that is easy to service and manage. The powerful system is designed for your most important business applications, such as VDI solutions and cloud deployments.

Combining balanced performance and flexibility, the x3550 M4 is a great choice for VDI solutions. It can provide outstanding uptime to keep VDI solutions and cloud deployments running safely. Ease-of-use and comprehensive management tools make it easy to deploy.

System x3550 M4 is shown in Figure 2-6.



Figure 2-6 System x3550 M4

The following x3550 M4 server components are key for VDI:

- ▶ **Processor:** Intel Xeon processor E5-2600 v2 with 12-core processors and up to 3.5 GHz core speeds, up to 30 MB of L3 cache, and up to two 8 GTps QPI interconnect links are available. Up to 2 processors, 24 cores, and 48 threads maximize the concurrent running of multi-threaded applications.
- ▶ **Memory:** Supports up to 24 Load Reduced DIMMs (LRDIMMs) of 1866 MHz DDR3 ECC memory that provides speed, high availability, and a memory capacity of up to 768 GB.
- ▶ **Video:** The NVIDIA Quadro K600 GPU adapter is supported by E5-2600 v2 processors and available via CTO only.
- ▶ **Network:** Features four integrated Gigabit Ethernet 1000BASE-T ports (RJ-45); two embedded 10 Gb Ethernet ports (10GBASE-T RJ-45 or 10GBASE-SR SFP+ based) on optional 10 Gb Ethernet mezzanine card (does not use PCIe slot).
- ▶ **Disks:** Up to eight 2.5-inch hot-swap SAS/SATA HDDs, or up to three 3.5-inch hot-swap SAS/SATA HDDs, or up to three 3.5-inch Simple Swap SATA HDDs are supported.

2.2.4 System x3650 M4

The System x3650 M4 server provides great performance on a flexible and scalable design. Its energy-efficient design supports more cores, memory, and data capacity in a scalable 2U package that is easy to service and manage.

The x3650 M4 is an outstanding 2U 2-socket business-critical server that offers improved performance and pay-as-you grow flexibility along with new features that improve server management capability. This powerful system is designed for your most important business applications and cloud deployments.

It completes the VDI infrastructure by providing a solution to support graphics-intensive virtual desktops that run 3D or CAD applications. The x3650 M4 is shown in Figure 2-7.



Figure 2-7 System x3650 M4

The following components are key for VDI:

- ▶ **Processor:** An Intel Xeon processor E5-2600 v2 with 12-core processors and up to 3.5 GHz core speeds, up to 30 MB of L3 cache, and up to two 8 GTps QPI interconnect links is featured. Up to 2 processors, 24 cores, and 48 threads maximize the concurrent running of multi-threaded applications.
- ▶ **Memory:** Supports up to 24 Load Reduced DIMMs (LRDIMMs) of 1866 MHz DDR3 ECC memory that provides speed, high availability, and a memory capacity of up to 768 GB.

- ▶ Video: The following GPUs are supported for VDI:
 - NVIDIA Quadro K600
 - NVIDIA Quadro K2000
 - NVIDIA Quadro K5000
- ▶ Network: Four integrated Gigabit Ethernet 1000BASE-T ports (RJ-45); two embedded 10 Gb Ethernet ports (10GBASE-T RJ-45 or 10GBASE-SR SFP+ based) on optional 10 Gb Ethernet mezzanine card (does not use PCIe slot).
- ▶ Disk: Up to 32 1.8-inch SSD bays, or 16 2.5-inch hot-swap SAS/SATA bays, or up to six 3.5-inch hot-swap SAS/SATA bays, or up to eight 2.5-inch Simple Swap SATA bays, or up to six 3.5-inch Simple Swap SATA bays.

2.3 Networking components

This section describes the Lenovo Ethernet networking options that can be used in the VDI environments and includes the following topics:

- ▶ 2.3.1, “Flex System networking I/O modules”
- ▶ 2.3.2, “RackSwitch offerings” on page 22
- ▶ 2.3.3, “Virtual Fabric adapters” on page 27

2.3.1 Flex System networking I/O modules

This section describes the following Flex System Ethernet I/O modules:

- ▶ “Flex System Fabric EN4093R 10Gb Scalable Switch” on page 19
- ▶ “Flex System Fabric CN4093 10Gb Converged Scalable Switch” on page 20
- ▶ “Flex System Fabric SI4093 System Interconnect Module” on page 21
- ▶ “Flex System EN4091 10Gb Ethernet Pass-thru Module” on page 22

Flex System Fabric EN4093R 10Gb Scalable Switch

The Flex System Fabric EN4093R 10Gb Scalable Switch provides unmatched scalability, port flexibility, and performance, while also delivering innovations to help address several networking concerns today and providing capabilities that help you prepare for the future.

This switch can support up to 64 10 Gb Ethernet connections while offering Layer 2/3 switching, in addition to OpenFlow and “easy connect” modes. This switch can help clients migrate to a 10 Gb or 40 Gb Ethernet infrastructure and offers cloud ready virtualization features, such as Virtual Fabric and VMready® and is Software Defined Network (SDN) ready.

Flex System Fabric EN4093 10Gb Scalable Switch is shown in Figure 2-8.



Figure 2-8 Flex System Fabric EN4093R 10Gb Scalable Switch

The EN4093R switch is initially licensed for 24x 10 GbE ports. More ports can be enabled with Upgrade 1 and Upgrade 2 license options. Upgrade 1 must be applied before Upgrade 2 can be applied.

By using flexible port mapping for the EN4093R switch, you can buy only the ports that you need, when you need them.

The switches offer the following key features and benefits for VDI:

- Optimized network virtualization with virtual NICs

Virtual Fabric provides a way for companies to carve up 10 Gb ports into virtual NICs. For large-scale virtualization, the Flex System solution can support up to 32 vNICs by using a pair of CN4054 10Gb Virtual Fabric adapters in each compute node and four EN4093R 10Gb Scalable Switches in the chassis.

The EN4093R switch offers next-generation vNIC™ - Unified Fabric Port (UFP). UFP is an advanced solution that provides a flexible way for clients to allocate, reallocate, and adjust bandwidth to meet their requirements.

- Increased performance

The EN4093R is the embedded 10 GbE switch for a server chassis to support submicrosecond latency and up to 1.28 Tbps, while also delivering full line rate performance, which makes it ideal for managing dynamic workloads across the network. This switch also provides a rich Layer 2 and Layer 3 feature set that is ideal for many of today's data centers and it offers industry-leading uplink bandwidth by being the first integrated switch to support 40 Gb uplinks.

- VM-aware networking

The EN4093R switch simplifies management and automates VM mobility by making the network VM aware with VMready, which works with all the major hypervisors.

- Transparent networking capability

With a simple configuration change to Easy Connect mode, the EN4093R switch becomes a transparent network device. By emulating a host NIC to the data center core, it accelerates the provisioning of VMs by eliminating the need to configure the typical access switch parameters.

Flex System Fabric CN4093 10Gb Converged Scalable Switch

The Flex System Fabric CN4093 10Gb Converged Scalable Switch provides unmatched scalability, performance, convergence, and network virtualization. The switch offers full Layer 2/3 switching and FCoE Full Fabric and Fibre Channel NPV Gateway operations to deliver a truly converged integrated solution. The switch can help clients migrate to a 10 Gb or 40 Gb converged Ethernet infrastructure and offers virtualization features, such as Virtual Fabric and VMready.

Flex System Fabric CN4093 10Gb Converged Scalable Switch is shown in Figure 2-9.



Figure 2-9 Flex System CN4093 10Gb Converged Switch

The CN4093 has flexible port licensing. The base switch configuration includes 14 10 GbE connections to the node bays, two 10 GbE SFP+ ports, and six Omni Ports™ with SFP+ connectors. The client then has the flexibility of turning on more 10 GbE connections to the internal node bays, and more Omni Ports and 40 GbE QSFP+ uplink ports (or 4x 10 GbE SFP+ DAC uplinks on each QSFP+ port) when needed. The client turns them on by using Features on Demand (FoD) licensing capabilities that provide “pay as you grow” scalability without the need for more hardware.

The switches offer the following key features and benefits for VDI:

- Optimized network virtualization with virtual NICs
Virtual Fabric provides a way for companies to divide 10 Gb ports into virtual NICs. For large-scale virtualization, the Flex System solution can support up to 32 vNICs by using a pair of CN4054 10Gb Virtual Fabric adapters in each compute node.
The CN4093 switch offers next-generation vNIC - Unified Fabric Port (UFP). UFP is an advanced solution that provides a flexible way for clients to allocate, reallocate, and adjust bandwidth to meet their requirements.
- Increased performance
The CN4093 is the embedded 10 Gb switch for a server chassis to support aggregated throughput of 1.28 Tbps, while also delivering full line rate performance on Ethernet ports, which makes it ideal for managing dynamic workloads across the network. It also offers industry-leading uplink bandwidth by being the integrated switch to support 40 Gb uplinks.
- VM-aware networking
Flex System CN4093 simplifies management and automates VM mobility by making the network VM aware with VMready, which works with all the major hypervisors.

Flex System Fabric SI4093 System Interconnect Module

The Flex System Fabric SI4093 System Interconnect Module enables simplified integration of Flex System into your networking infrastructure.

The SI4093 System Interconnect Module requires no management for most data center environments, which eliminates the need to configure each networking device or individual port. As a result, the number of management points is reduced. It provides a low latency, loop-free interface that does not rely upon spanning tree protocols, which removes one of the greatest deployment and management complexities of a traditional switch.

The SI4093 System Interconnect Module offers administrators a simplified deployment experience while maintaining the performance of intra-chassis connectivity.

The SI4093 System Interconnect Module is shown in Figure 2-10.



Figure 2-10 Flex System Fabric SI4093 System Interconnect Module

The SI4093 System Interconnect Module is initially licensed for 14 enabled 10 Gb internal ports and 10 enabled 10 Gb external uplink ports. More ports can be enabled, including 14 internal ports and two 40 Gb external uplink ports by using the FoD licensing mode.

The switch offers the following key features and benefits for VDI:

- ▶ **Transparent (or VLAN-agnostic) mode**
The interconnect module provides traffic consolidation in the chassis to minimize TOR port usage. It also enables server-to-server communication for optimum performance; for example, vMotion.
- ▶ **Optimized network virtualization with virtual NICs**
Virtual Fabric provides a way for companies to divide 10 Gb ports into virtual NICs. For large-scale virtualization, the Flex System solution can support up to 32 vNICs by using a pair of CN4054 10Gb Virtual Fabric adapters in each compute node.
- ▶ **VM-aware networking**
The SI4093 simplifies management and automates VM mobility by making the network VM aware with VMready, which works with all the major hypervisors. Network policies migrate automatically along with VMs to ensure that security, performance, and access remain intact as VMs move from server to server.
- ▶ **Increased performance**
The SI4093 is the embedded 10 Gb interconnect Module for a server chassis to support aggregated throughput of 1.28 Tbps, while also delivering full line rate performance on Ethernet ports, which makes it ideal for managing dynamic workloads across the network. It offers industry-leading uplink bandwidth by being the integrated switch to support 40 Gb uplinks.

The SI4093 also offers increased security and performance advantage when it is configured in VLAN-aware mode; it does not force communications upstream into the network, which reduces latency and generates less network traffic.

- ▶ **Transparent networking**

The SI4093 is a transparent network device that is not apparent to the upstream network. By emulating a host NIC to the data center core, it accelerates the provisioning of VMs by eliminating the need to configure the typical access switch parameters.

Flex System EN4091 10Gb Ethernet Pass-thru Module

The Flex System EN4091 10Gb Ethernet Pass-thru Module offers easy connectivity of the Flex System Enterprise Chassis to any external network infrastructure. This unmanaged device enables direct Ethernet connectivity of the compute node in the chassis to an external TOR data center switch. This module can function at 1 Gb and 10 Gb Ethernet speeds. It has 14 internal 1 Gb or 10 Gb links, and 14 external 1 Gb or 10 Gb SFP+ uplinks.

Flex System EN4091 10Gb Ethernet Pass-thru Module is shown in Figure 2-11.



Figure 2-11 Flex System EN4091 10Gb Ethernet Pass-thru Module

The Flex System EN4091 offers the following key features:

- ▶ Intelligent workload deployment and management for maximum business agility
- ▶ High-speed performance, complete with integrated servers, storage, and networking
- ▶ Independently scalable IT resource pools for higher usage and lower cost per workload

2.3.2 RackSwitch offerings

Lenovo Ethernet switch family (which is a RackSwitch™ family) is designed to bring speed and intelligence to the edge of your network where it is closer to your business, users, and innovations. Lenovo top of rack products are lossless, low latency, and low power.

The following RackSwitches are described in this section:

- ▶ “RackSwitch G8124E”
- ▶ “RackSwitch G8264” on page 24
- ▶ “RackSwitch G8264CS” on page 25

RackSwitch G8124E

As shown in Figure 2-12 on page 23, the RackSwitch G8124E is designed with top performance in mind. This low-latency switch provides line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data.




Figure 2-12 RackSwitch G8124 TOR switch

The RackSwitch G8124E offers the following feature benefits regarding VDI environments:

- ▶ High performance

The 10G Low Latency (as low as 570 nanoseconds) switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management.

- ▶ Lower power and better cooling

The G8124E uses as little power as two 60 W light bulbs, which is a fraction of the power usage of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the G8124E rear-to-front cooling design reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power usage.

- ▶ Virtual Fabric support

Virtual Fabric can help customers address I/O requirements for multiple NICs while also helping reduce cost and complexity. Virtual Fabric allows for dividing a physical NIC into multiple virtual NICs (2 - 8 vNICs) and creates a virtual pipe between the adapter and the switch for improved performance, availability, and security while reducing cost and complexity.

- ▶ VM-aware networking

VMready software on the switch helps reduce configuration complexity while significantly improving security levels in virtualized environments. VMready automatically detects VM movement from one physical server to another. It also instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with all leading VDI VM providers, such as VMware and Microsoft.

- ▶ Layer 3 functionality

The switch includes Layer 3 functionality, which provides security and performance benefits as inter-VLAN traffic stays within the chassis. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) for enterprise customers.

- ▶ Seamless Interoperability

RackSwitches perform seamlessly with other vendors' upstream switches.

- ▶ Fault tolerance

These switches learn alternative routes automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, advanced VLAN-based failover, VRRP, Hot Links, Uplink Failure Detection (UFD), IGMP V3 snooping, and OSPF.

- Converged fabric

The switch supports CEE/DCB and connectivity to FCoE gateways. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable highly efficient block storage over Ethernet for consolidating server network connectivity.

RackSwitch G8264

The RackSwitch G8264 is a 10 Gb/40 Gb Top-of-Rack switch that is for applications that require the highest performance at low latency. It combines 1.28 Tbps throughput with up to 64 10 Gb SFP+ ports in an ultra-dense 1U form factor.

The front view of the RackSwitch G8264 is shown in Figure 2-13.




Figure 2-13 RackSwitch G8264 front view

The RackSwitch G8264 offers the following benefits regarding VDI environments:

- High performance

The 10 Gb/40 Gb switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management. It has a throughput of 1.28 Tbps.

- Stacking

With the G8264, a single switch image and configuration file can be used for up to eight switches, sharing only one IP address and one management interface.

- Lower power and better cooling

The RackSwitch G8264 uses as little as 330 W of power, which is a fraction of the power usage of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the front-to-rear or rear-to-front cooling design of the G8264 reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power usage.

- Virtual Fabric

The G8264 can help customers address I/O requirements for multiple NICs while reducing cost and complexity. By using Virtual Fabric, you can divide a physical NIC into multiple virtual NICs (2 - 8 vNIC) and create a virtual pipe between the adapter and the switch for improved performance, availability, and security.

- VM-aware networking

VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement between physical servers and instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with all leading VM providers, such as VMware and Microsoft.

- Layer 3 functionality

The G8264 includes Layer 3 functionality, which provides security and performance benefits, as inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as OSPF and BGP for enterprise customers.

- Seamless interoperability

RackSwitches perform seamlessly with other vendors' upstream switches.

- Fault tolerance

The G8264 switch learns alternative routes automatically and performs faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, advanced VLAN-based failover, VRRP, and Hot Links.

- Converged fabric

The G8264 switch supports CEE and connectivity to FCoE gateways. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable highly efficient block storage over Ethernet for consolidating server network connectivity.

- Transparent networking capability

With a simple configuration change to Easy Connect Mode, the RackSwitch G8264 becomes a transparent network device that is not apparent to the core, which eliminates network administration concerns of Spanning Tree Protocol configuration/interoperability, VLAN assignments, and avoids any possible loops.

By emulating a host NIC to the data center core, it accelerates the provisioning of VMs by eliminating the need to configure the typical access switch parameters.

RackSwitch G8264CS

The RackSwitch G8264CS is an enterprise-class switch that offers high-bandwidth performance with 36 1/10 Gb SFP+ connections, 12 Omni Ports that can be used for 10 Gb SFP+ connections, 4 Gb or 8 Gb Fibre Channel connections or both, and four 40 Gb QSFP+ connections.

It simplifies the deployment with its innovative Omni Port technology and offers the flexibility to choose 10 Gb Ethernet, 4 Gb or 8 Gb Fibre Channel, or both for upstream connections. In FC mode, Omni Ports provide convenient access to FC storage.

The G8264CS provides 100% line rate performance with low latency and 1.28 Tbps non-blocking switching throughput (full duplex) on Ethernet ports, which makes it an optimal choice for managing dynamic workloads across the network. It provides a rich Layer 2 and Layer 3 feature set that is ideal for many of today's data centers.

Also, its Omni Port technology helps in consolidating the enterprise storage, networking, data, and management into a simple to manage single fabric. It also reduces costs that are associated with energy and cooling, management and maintenance, and capital costs.

The RackSwitch G8264CS is shown in Figure 2-14.



Figure 2-14 RackSwitch G8264CS

The RackSwitch G8264 offers the following benefits regarding VDI environments:

- ▶ **High performance**
The 10-Gb/40-Gb switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management. It has a throughput of up to 1.28 Tbps.
- ▶ **Lower power and better cooling**
The G8264CS uses as little as 330 W of power, which is a fraction of the power usage of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the front-to-rear or rear-to-front cooling design of the G8264CS switch reduces the costs of data center air conditioning by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power usage.
- ▶ **Support for Virtual Fabric**
The G8264CS can help customers address I/O requirements for multiple NICs while reducing cost and complexity. By using Virtual Fabric, you can divide a physical dual-port NIC into multiple vNICs (2 - 8 vNICs) to create a virtual pipe between the adapter and the switch for improved performance, availability, and security. With support for FCoE or iSCSI, two vNICs on a dual-port adapter can be configured as CNAs to allow for more cost savings through convergence.
- ▶ **VM-aware networking**
VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects VM movement between physical servers and instantly reconfigures the network policies of each VM across VLANs to keep the network up and running without interrupting traffic or affecting performance. VMready works with leading VDI VM providers, such as VMware and Microsoft.
- ▶ **Layer 3 functionality**
The G8264CS includes Layer 3 functionality, which provides security and performance benefits because inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as OSPF and BGP for enterprise customers.
- ▶ **Seamless interoperability**
RackSwitch perform seamlessly with other vendors' upstream switches.

- ▶ Fault tolerance

The G8264CS switches learn alternative routes automatically and perform faster convergence if there is a link, switch, or power failure. The switch uses proven technologies, such as L2 trunk failover, advanced VLAN-based failover, VRRP, and Hot Links.

- ▶ Converged fabric

The RackSwitch G8264CS supports CEE and full fabric FCoE connectivity. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable highly efficient block storage over Ethernet for consolidating server network connectivity.

2.3.3 Virtual Fabric adapters

The family of Virtual Fabric adapters for System x rack servers and Flex System compute nodes helps provide flexible, scalable, and efficient network connectivity for the servers and storage in VDI solutions.

The following Virtual Fabric adapter (VFA) choices are available:

- ▶ Dual-port embedded adapters

Embedded VFAs are built into the system board (LAN-on-motherboard - LOM) on selected compute nodes or available as optional mezzanine cards on selected System x rack servers:

- Emulex Dual Port 10GbE SFP+ Embedded adapter
- Emulex Dual Port 10GbE SFP+ Embedded VFA IIIr

- ▶ Dual-port PCIe VFAs for System x rack servers:

- Emulex Dual Port 10GbE SFP+ VFA III
- Emulex Dual Port 10GbE SFP+ VFA IIIr

- ▶ Quad-port mezzanine VFAs for Flex System compute nodes:

- Flex System CN4054 10Gb Virtual Fabric adapter
- Flex System CN4054R 10Gb Virtual Fabric adapter

VFAs offer the following operational mode choices:

- ▶ pNIC mode (multichannel disabled)

The adapter operates as a standard 10 Gbps Ethernet adapter (dual-port of four-port depending on the adapter), and it functions with any 10 GbE switch.

- ▶ vNIC mode (multichannel enabled)

This mode enables up to four virtual NIC interfaces per 10 Gb physical port (a total of eight for dual-port VFAs and 16 for quad-port VFAs). It uses the IEEE 802.1Q VLAN tag, which is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch.

You can also use the following vNIC linking options:

- Virtual Fabric mode works with EN4093R, CN4093, G8124E, G8264, and G8264CS switches. In this mode, the adapter communicates with the switch module to obtain vNIC parameters (by using DCBX). Also, a special tag within each data packet is added and later removed by the NIC and switch for each vNIC group to maintain separation of the virtual channels.

vNIC bandwidth allocation and metering are performed by the switch and the adapter. In such a case, a bidirectional virtual channel of an assigned bandwidth is established between them for every defined vNIC.

- Switch Independent Mode works with any switch, and the vNIC bandwidth metering and control are performed on the adapter side only, which forms unidirectional virtual channel (server-to-switch). This mode extends the client's VLANs to the virtual NIC interfaces.

vNIC bandwidth allocation and metering are performed only by the adapter. In such a case, a unidirectional virtual channel is established in which the bandwidth management is performed only for the outgoing traffic on a network adapter side (server-to-switch). The incoming traffic (switch-to-server) uses the all of the available physical port bandwidth because there is no metering that is performed on a switch side.

- Unified Fabric Port (UFP) mode is the current direction of NIC virtualization, and it provides a more feature-rich solution that is compared to the original vNIC Virtual Fabric mode. UFP mode is supported by EN4093R, CN4093, G8124E, G8264, and G8264CS switches and SI4093 interconnect modules.

As with Virtual Fabric mode vNIC, UFP allows dividing a single 10 Gb port into four virtual NICs (called vPorts in UFP). UFP supports the following modes:

- Tunnel mode
Provides Q-in-Q mode, where the vPort is customer VLAN-independent (similar to vNIC Virtual Fabric Dedicated Uplink Mode).
- Trunk mode
Provides a traditional 802.1Q trunk mode (multi-VLAN trunk link) to the virtual NIC (vPort) interface; that is, permits host side tagging.
- Access mode
Provides a traditional access mode (single untagged VLAN) to the virtual NIC (vPort) interface that is similar to a physical port in access mode.
- FCoE mode
Provides FCoE functionality to the vPort.
- Auto-VLAN mode
Auto VLAN creation for Qbg and VMready environments.

Consider configuring the VFA adapters in UFP mode with the supported switches to distribute the 10 GbE network bandwidth flexibly to the VLANs that are used within the VDI infrastructure.

If you choose to implement an FCoE or iSCSI converged network for storage and network connectivity, an optional Advanced Upgrade can be activated on VFAs to enable FCoE or iSCSI processing.

2.4 Storage components

Some of the storage options to consider for the VDI storage design are described in this section, which includes the following topics:

- ▶ 2.4.1, “Fibre Channel connectivity”
- ▶ 2.4.2, “Converged fabrics” on page 31
- ▶ 2.4.3, “Solid-state drives in the VDI solution” on page 32
- ▶ 2.4.4, “RAID considerations” on page 33
- ▶ 2.4.5, “Flex System Storage Expansion Node” on page 33
- ▶ 2.4.6, “IBM Storwize V7000” on page 34
- ▶ 2.4.7, “IBM Storwize V3700” on page 36

2.4.1 Fibre Channel connectivity

Fibre Channel is well-established in the open systems environment as the underlying architecture of the SAN. Fibre Channel is a technology standard that allows data to be transferred from one network node to another at high speeds.

Fibre Channel is ideal for moving large volumes of data across long distances quickly and reliably. In current implementations, the Fibre Channel standard speed is generally available 2 Gbps - 16 Gbps; however, older 2 Gbps and 4 Gbps equipment is being replaced by faster connections.

In this section, we focus on the following Fibre Channel components that are available on Flex System:

- ▶ “FC5022 16Gb SAN Scalable Switch”
- ▶ “Fibre Channel adapters” on page 30

FC5022 16Gb SAN Scalable Switch

The Flex System FC5022 16Gb SAN Scalable Switch is a high-density, 48-port 16 Gbps Fibre Channel switch that is used in the Flex System chassis. The switch provides 28 internal ports to compute nodes by way of the midplane, and 20 external SFP+ ports. These SAN switch modules deliver an embedded option for Flex System users that are deploying storage area networks in their enterprise. The modules offer end-to-end 16 Gb and 8 Gb connectivity.

The N-Port Virtualization mode streamlines the infrastructure by reducing the number of domains to manage while enabling the ability to add or move servers without affecting the SAN. Monitoring is simplified via an integrated management appliance or clients that are using end-to-end Brocade SAN that can use the Brocade management tools.

The FC5022 switch is shown in Figure 2-15.



Figure 2-15 Flex System FC5022 16Gb Scalable Switch

The FC5022 16Gb Switch supports multi-tenancy in cloud environments through VM-aware end-to-end visibility and monitoring, quality of service (QoS), and fabric-based advanced zoning features.

FC3171 8Gb SAN Switch and Pass-thru

The Flex System FC3171 8Gb SAN Switch is a full-fabric Fibre Channel component with expanded functionality. The SAN switch supports high-speed traffic processing for Flex System configurations. It also offers scalability in external SAN size and complexity and enhanced systems management capabilities.

The Flex System FC3171 8Gb Pass-thru supports a fully interoperable solution for seamless integration of the Fibre Channel initiators to a fabric. The pass-thru module uses industry-standard N_Port ID virtualization (NPIV) technology to provide a cost-effective connectivity solution for the Flex System chassis.

FC3171 is shown in Figure 2-16.



Figure 2-16 Flex System FC3171 8Gb SAN Switch

Fibre Channel adapters

If you decided to implement Fibre Channel connectivity for your VDI storage, the following adapters are available:

- ▶ Flex System FC3172 2-port and FC3052 2-port 8Gb FC adapters
- ▶ Flex System FC5022 2-port and FC5054 4-port 16Gb 16Gb FC adapters
- ▶ “Flex System FC5024D 4-port 16Gb FC adapter” on page 31
- ▶ “Flex System FC5172 2-port 16Gb FC adapter” on page 31

Flex System FC3172 2-port and FC3052 2-port 8Gb FC adapters

The Flex System FC3172 2-port and FC3052 2-port 8Gb FC adapters enable high-speed access for Flex System compute nodes to connect to a Fibre Channel SAN. The adapters connect to the midplane directly, without having to use cables or small form-factor pluggable (SFP) modules. By eliminating these components for up to 14 servers, the resulting savings can cover the investment in the chassis. Both adapters also offer comprehensive virtualization capabilities with support for NPIV and virtual fabric.

Flex System FC5022 2-port and FC5054 4-port 16Gb 16Gb FC adapters

The Flex System FC5022 2-port and FC5054 4-port 16Gb FC adapters enable high-speed access for compute nodes to an external SAN. These adapters are based on Brocade architecture and offer end-to-end 16 Gb connectivity to SAN. The adapters also offer enhanced features, such as N_Port trunking, NPIV, and boot-from-the-SAN with automatic LUN discovery and end-to-end server application optimization. Having 16 Gb adapters and switches also offers future investment protection by enabling the density of VMs to be increased on a compute node.

The FC5022 2-port and FC5024 4-port 16Gb FC adapters have the following features:

- ▶ Direct I/O enables native (direct) I/O performance by allowing VMs to bypass the hypervisor and communicate directly with the adapter.
- ▶ Uses 16 Gbps bandwidth to eliminate internal oversubscription.
- ▶ Over 500,000 IOPS per port, which maximizes transaction performance and density of VMs per compute node.

Flex System FC5024D 4-port 16Gb FC adapter

The Flex System FC5024D 4-port 16Gb FC adapter is a quad-port mid-mezzanine card for the Flex System x222 Compute Node. The FC5024D provides Fibre Channel connectivity to both servers in the x222, with two ports that are routed to each server. This adapter offers end-to-end 16 Gb connectivity to SAN.

The Flex System FC5024D 4-port 16Gb FC adapter has the following enhanced features:

- ▶ Direct I/O enables native (direct) I/O performance by allowing VMs to bypass the hypervisor and communicate directly with the adapter.
- ▶ Uses 16 Gbps bandwidth to eliminate internal oversubscription.
- ▶ Over 500,000 IOPS per port, which maximizes transaction performance and density of VMs per compute node.

Flex System FC5172 2-port 16Gb FC adapter

The Flex System FC5172 2-port 16Gb FC adapter from QLogic enables high-speed access for Flex System Enterprise Chassis compute nodes to connect to a Fibre Channel SAN. It works with the 8 Gb or 16 Gb Flex System Fibre Channel switch modules

2.4.2 Converged fabrics

As the name implies, converged fabrics are all about taking a set of protocols and data that is designed to run on top of one type of physical medium, and allowing them to be carried on top of a different physical medium. This configuration provides a number of cost benefits, such as reducing the number of physical cabling plants that are required, removing the need for separate physical NICs and HBAs, and potentially reducing power and cooling. From an OpEx perspective, it can reduce the cost that is associated with the management of separate physical infrastructures. In the data center world, two of the most common forms of converged fabrics are FCoE and iSCSI.

FCoE allows a host to use its 10 Gb Ethernet connections to access Fibre Channel attached storage, as though it were physically Fibre Channel that is attached to the host. In fact, the FC traffic is encapsulated into FCoE frames and carried to the remote storage via an Ethernet network.

iSCSI takes a protocol that was originally designed for hosts to talk to relatively close physical storage over physical SCSI cables and converts it to use IP and run over an Ethernet network; therefore, it can access storage way beyond the limitations of a physical SCSI-based solution.

iSCSI can be used in existing (lossy) and new (lossless) Ethernet infrastructures, with different performance characteristics. However, FCoE requires a lossless converged enhanced Ethernet network, and it relies on extra functionality that is known from Fibre Channel (for example, nameserver, zoning).

Fibre Channel over Ethernet

FCoE assumes the existence of a lossless Ethernet, such as one that implements the Data Center Bridging (DCB) extensions to Ethernet. The basic notion of FCoE is that the upper layers of FC are mapped onto Ethernet. The upper layer protocols and services of FC remain the same in an FCoE deployment. Zoning, fabric services, and similar services still exist with FCoE.

The difference is that the lower layers of FC are replaced by lossless Ethernet, which also implies that FC concepts, such as port types and lower-layer initialization protocols, must be replaced by new constructs in FCoE. Such mappings are defined by the FC-BB-5 standard.

The EN4093R, CN4093, G8264, and G8264CS switches and SI4093 interconnect modules support FCoE. The G8264, EN4093R, and SI4093 functions as an FCoE transit switch while the CN4093 and G8264CS have Omni Ports that can be set to function as FC ports or Ethernet ports as specified in the switch configuration.

iSCSI

The iSCSI protocol allows for longer distances between a server and its storage when compared to the traditionally restrictive parallel SCSI solutions or the newer serial-attached SCSI (SAS). iSCSI technology can use a hardware initiator, such as an HBA, or a software initiator to issue requests to target devices. Within iSCSI storage terminology, the initiator is typically known as a *client*, and the target is the storage device. The iSCSI protocol encapsulates SCSI commands into protocol data units (PDUs) within the TCP/IP protocol and then transports them over the network to the target device.

iSCSI provides block-level access to storage (as does Fibre Channel) but uses TCP/IP over Ethernet instead of Fibre Channel protocol. Therefore, iSCSI is attractive for its relative simplicity and usage of widely available Ethernet skills. Its chief limitations often are the relatively lower speeds of Ethernet compared to Fibre Channel and the extra TCP/IP encapsulation that is required. With lossless 10 Gb Ethernet now available, the attractiveness of iSCSI is expected to grow rapidly. TCP/IP encapsulation are still used, but 10 Gbps Ethernet speeds dramatically increase the appeal of iSCSI.

2.4.3 Solid-state drives in the VDI solution

A hard disk drive (HDD) is a proven technology with excellent reliability and performance, given the physical limitations of its spinning platters and moving arms. A solid-state drive (SSD) uses non-volatile flash memory rather than spinning magnetic media to store data. The main advantage of SSDs for VDI is the lower access latency that is 10 times faster than in an HDD. All of the System x and Flex System servers for VDI and the Storwize V7000 support SSD disks within the internal drive bay.

Consider SSD disks for the following reasons:

- ▶ Provide the best performance for the non-persistent VDI hosts by installing two SSDs, which are configured as RAID 0
- ▶ Implement the Easy Tier function on the Storwize V7000 to increase its IOPS performance on the most frequently accessed data
- ▶ Accelerating Virtualized Applications:
 - VM resiliency (ultra-fast check-pointing)
 - Decreasing write latency
 - IO Read Caching (VMware, KVM, and so on)
 - Eliminating I/O contention
 - Faster response time

2.4.4 RAID considerations

The RAID configuration affects only the performance for write operations. Read operations are not affected. The write penalty is the consequence of the RAID data protection technique, which requires multiple disk IOPS requests for each user write IOPS. RAID penalty is used to determine the functional IOPS of an array. The following formulas are used:

- ▶ Raw IOPS = Disk Speed IOPS x Number of disks
- ▶ Functional IOPS = (Raw IOPS x Write % / RAID Penalty) + (RAW IOPS x Read %)

Table 2-2 provides the write penalty for RAID configuration.

Table 2-2 RAID penalty

RAID	Write penalty
0	1
1	2
5	4
6	6
10	2

In scalable implementations, hosted virtual desktops can generate substantial IOPS workload on the storage part of the VDI infrastructure. Therefore, select the appropriate RAID level to match the following workload:

- ▶ For a read-intensive workload, use RAID 0, RAID 1, RAID 5, and RAID 10 levels that spread read operations across multiple disks simultaneously. If the volume of data is important, you can also use a RAID level that optimizes disk usability.
- ▶ For a write-intensive workload, use a RAID level that offers a low write penalty, such as RAID 0 and RAID 10.
- ▶ To store the PVS write cache, redundant RAID configuration is not needed because of the non-persistence of the environment.

2.4.5 Flex System Storage Expansion Node

The Flex System Storage Expansion Node (SEN) is a storage enclosure that attaches to a single half-wide compute node to provide that compute node with more direct-attach local storage. The SEN adds 12 hot-swap 6.35-cm (2.5-inch) drive bays and an LSI RAID controller and connects to the compute node via its PCIe expansion connector. A SEN that is attached to a x240 compute node is shown in Figure 2-17.



Figure 2-17 SEN that is attached to an x240 compute node

The x240 compute node with the SEN can be used as an entry-level, NAS-only, or unified server storage in VDI deployments.

The following features were retained for VDI:

- ▶ Support for 6 Gbps SAS and SATA drives; HDDs and SSDs
- ▶ Support for RAID 0, 1, 5, 10, and 50 as standard
- ▶ Support for logical unit number (LUN) sizes up to 64 TB
- ▶ Optional support for SSD performance acceleration and SSD caching with FoD upgrades

2.4.6 IBM Storwize V7000

As members of the IBM Storwize family, IBM Storwize V7000 is a virtualized, enterprise-class storage system that provides the foundation for implementing an effective storage infrastructure and transforming the economics of data storage. With industry-first hardware accelerated Real-time Compression, they can reduce the cost of storage by up to half while maintaining application performance.

The Storwize V7000 family includes the capability to virtualize its own internal storage and external SAN-attached storage in the same manner as the System Storage SAN Volume Controller does. Also, the V7000 uses the advanced functions of the IBM System Storage DS8000 family for its RAID configurations of the internal disks, and the highly flexible graphical user interface (GUI) of the IBM XIV Storage Subsystem for management.

One key feature is the IBM System Storage Easy Tier. The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, which places such data in a faster tier of storage.

The Storwize V7000 is shown in Figure 2-18.




Figure 2-18 Storwize V7000 with 2.5-inch drives

The following sections provide an overview of the hardware and software.

Hardware overview

The Storwize V7000 solution provides a choice of up to 480 x 3.5-inch or 960 x 2.5-inch serial-attached SCSI (SAS) drives for the internal storage in a clustered system and uses SAS cables and connectors to attach to the optional expansion enclosures. In a clustered system, the V7000 can provide 1.92 PiB capacity.

When virtualizing external storage arrays, a Storwize V7000 system can provide up to 32 PiB of usable capacity. A Storwize V7000 system supports a range of external disk systems, similar to what the SAN Volume Controller supports today.

The Storwize V7000 solution consists of one to four control enclosures and optionally, up to 36 expansion enclosures (and supports the intermixing of the different expansion enclosures). Within each enclosure are two canisters. Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

The Storwize V7000 family provides several configuration options that are aimed at simplifying the implementation process. It also includes automated instruction steps, called Directed Maintenance Procedures (DMP) to help resolve any events that might occur. Storwize V7000 is a clusterable, scalable, storage system, and an external virtualization device.

Software overview

The Storwize V7000 provides thin provisioning, automated tiering for automated SSD optimization, internal and external virtualization, clustering, replication, multiprotocol support, and a next-generation graphical user interface (GUI).

The Storwize V7000 software performs the following functions:

- ▶ Creates a single pool of storage
- ▶ Provides logical unit virtualization
- ▶ Manages logical volumes
- ▶ Manages physical resources, including drives

The Storwize V7000 system also provides the following functions:

- ▶ Large scalable cache
- ▶ Thin provisioning
- ▶ Volume mirroring
- ▶ FlashCopy:
 - Full and Incremental copy
 - Multi-target FlashCopy
 - Cascaded FlashCopy
 - Reverse FlashCopy
 - FlashCopy nocopy with thin provisioning
 - Consistency groups
- ▶ Remote Copy feature:
 - Metro Mirror (synchronous copy)
 - Global Mirror (asynchronous cop
- ▶ Data Migration
- ▶ System Storage Easy Tier

It provides a mechanism to seamlessly migrate hot spots to the most appropriate tier within the Storwize V7000 system. This migration can be to internal drives within the Storwize V7000 system or to external storage systems that are virtualized by the Storwize V7000 system.

- ▶ Real-time Compression

Provides for data compression that uses the IBM Random-Access Compression Engine (RACE), which can be performed on a per volume basis in real time on active primary workloads. Real-time Compression can provide as much as a 50% compression rate for data that is not already compressed. It can help with reducing the amount of capacity needed for storage, which can help with delaying further growth purchases. Real-time Compression supports all storage that is attached to the Storwize V7000 system whether internal, external, or external virtualized storage.

- External Storage Virtualization

With this feature, an external storage subsystem can be attached through the Fibre Channel or by FCoE to the Storwize V7000 system. After the storage from the external system is integrated into Storwize V7000 and added to a storage pool, it is available to be virtualized and used by any of the features and functions of the Storwize V7000 system.

2.4.7 IBM Storwize V3700

IBM Storwize V3700 Storage System is a member of the Storwize family of disk systems. By using IBM Storwize V7000 Storage System and IBM SAN Volume Controller functions, interoperability, and management tools, Storwize V3700 delivers innovation and new levels of storage efficiency with ease of use in an entry disk system to enable organizations to overcome their storage challenges.

Storwize V3700 controller unit models include six 6 Gb SAS and four 1 Gb Ethernet ports standard for SAS and iSCSI connectivity. They can be optionally configured with eight 8 Gb Fibre Channel (FC) ports, four 10 Gb Ethernet (iSCSI/FCoE) ports, or extra 6 Gb SAS or 1 Gb Ethernet ports.

The Storwize V3700 is shown in Figure 2-19.




Figure 2-19 IBM Storwize V3700

Storwize V3700 delivers powerful and intuitive storage with the following features:

- Dual-active intelligent array node canisters with up to 8 GB cache per canister
- 6 Gb SAS and 1 Gb iSCSI connectivity standard with optional 8 Gb Fibre Channel (FC) or 10 Gb iSCSI or Fibre Channel over Ethernet (FCoE) connectivity
- Support for 12 3.5-inch large form factor or 24 2.5-inch small form factor drives
- Scalable up to 240 drives per system with the attachment of Storwize V3700 expansion units
- Rich set of standard functions, including virtualized internal storage, thin provisioning, data migration, and data replication
- Optional licensed functions, including Turbo performance, Easy Tier, and remote mirroring
- Innovative, intuitive, web-based GUI for easy system setup and management

Storwize V3700 supports the complete range of data storage requirements, from highly used applications to high-capacity, low usage applications.

The following 3.5-inch drives are supported:

- High-performance, enterprise class disk drives:
 - 300 GB and 600 GB 15,000 rpm
 - 900 GB and 1.2 TB 10,000 rpm

- ▶ High-capacity, archival-class nearline disk drives of 2 TB, 3 TB, and 4 TB 7 200 rpm

The following 2.5-inch drives are supported:

- ▶ Flash drives in 200 GB, 400 GB, and 800 GB
- ▶ High-performance, enterprise class disk drives:
 - 146 GB, 300 GB, and 600 GB 15,000 rpm
 - 600 GB, 900 GB, and 1.2 TB 10,000 rpm
- ▶ High-capacity, archival-class nearline disk drives: 1 TB 7 200 rpm

All drives are dual-port and hot-swappable. Drives of the same form factor can be intermixed within the appropriate enclosure, which provides the flexibility to address performance and capacity needs within a single enclosure.

Up to nine Storwize V3700 expansion units are supported by a single Storwize V3700 controller unit. You can intermix 3.5-inch and 2.5-inch expansion units behind a 3.5-inch or 2.5-inch controller unit. This configuration delivers the added flexibility to mix 3.5-inch and 2.5-inch drives within a single system.

More drives and expansion units are designed to be dynamically added with virtually no downtime, which helps to quickly and seamlessly respond to ever-growing capacity demands.

The Storwize V3700 Storage System is designed to offer high system and data availability with the following features:

- ▶ Dual-active, intelligent node canisters with mirrored cache
- ▶ Dual-port disk drives with automatic disk drive failure detection and RAID rebuild with global hot spares
- ▶ Redundant hardware, including power supplies and fans
- ▶ Hot-swappable and customer replaceable components
- ▶ Automated path failover support for the data path between the server and the drives

The following functions are included with every Storwize V3700:

- ▶ RAID levels 0, 1, 5, 6, and 10: Provide the flexibility to choose the level of data protection required.
- ▶ Virtualization of internal storage: Enables rapid, flexible provisioning and simple configuration changes.
- ▶ Thin provisioning: Optimizes efficiency by allocating disk storage space in a flexible manner among multiple users, based on the minimum space that is required by each user at a specific time. With thin provisioning, applications use only the space they are using, not the total space that was allocated to them.
- ▶ Data migration: Enables easy and nondisruptive moves of volumes from another storage system onto the Storwize V3700 Storage System by using FC or SAS connectivity.
- ▶ FlashCopy: Enables the creation of copies of data for backup, parallel processing, testing, and development, and have the copies available almost immediately. Storwize V3700 supports up to 64 FlashCopy targets per system.

Storwize V3700 capabilities can be expanded with optional licensed functions. Each function is licensed to a Storwize V3700 controller unit and covers the entire Storwize V3700 Storage System (controller unit and all attached expansion units).

To help evaluate the benefits of these new capabilities, licensed functions (except for FlashCopy upgrade) can be enabled at no charge for a 90-day trial period. Trials are started from the Storwize management GUI and do not require any Lenovo intervention. Upon expiration of the trial, the function is automatically disabled unless a license key for that function was installed onto the machine.

The following optional licensed functions are available:

- ▶ **Turbo performance:** Turbo performance increases the maximum IOPS and throughput of a Storwize V3700 Storage System. Configurations with greater than 80 disk drives or more than five SSDs are ideal candidates to benefit from the increased IOPS that is offered with Turbo performance. Configurations with greater than 30 disk drives are suited to benefit from the throughput increase offered with Turbo performance.
- ▶ **FlashCopy upgrade:** FlashCopy allows the creation of copies of data for backup, parallel processing, testing, and development, and have the copies available almost immediately. All Storwize V3700 Storage Systems support up to 64 targets per system at no charge. The FlashCopy upgrade option increases this support to 2,040 FlashCopy targets per system.
- ▶ **Easy Tier:** Storage tiering helps optimize storage use with data location to improve system performance, reduce costs, and simplify management. Easy Tier automatically and dynamically moves frequently accessed data to flash (solid state) drives in the system, which results in flash drive performance without manually creating and managing storage tier policies. Easy Tier makes it easy and economical to deploy flash drives in the environment.

- ▶ **Remote mirroring:** This feature provides storage system-based data replication that uses synchronous or asynchronous data transfers over IP, FC, or FCoE communication links. Metro Mirror maintains a fully synchronized copy at metropolitan distances (up to 300 km). Global Mirror operates asynchronously and helps maintain a copy at much greater distances (up to 8000 km). Both functions support VMware Site Recovery Manager to help speed disaster recovery.

For ultimate flexibility, Storwize V3700 remote mirroring is designed to interoperate with any other IBM Storwize family system, including Storwize V7000, Storwize V5000, and SAN Volume Controller.

The remote mirroring option must be acquired (or licensed) for the primary (local) and secondary (remote) systems. If Storwize V3700 is mirrored to a system other than Storwize V3700, the other system must have the appropriate and applicable license for remote mirroring.

Each Storwize V3700 Storage System includes a simple and intuitive GUI that is designed to allow storage to be quickly deployed and efficiently managed. The GUI runs on the Storwize V3700 system, so there is no need for a separate console. You need only to point your web browser to the system. It is based on the Storwize V7000 management GUI and has a similar look and feel.

2.5 Management components

This section describes the following management components and features:

- ▶ 2.5.1, “Integrated Management Module II”
- ▶ 2.5.2, “Chassis Management Module”
- ▶ 2.5.3, “Flex System Manager” on page 40
- ▶ 2.5.4, “Introduction to Upward Integration” on page 41

2.5.1 Integrated Management Module II

The Integrated Management Module II (IMM2) is the next generation of the integrated service processors for the System x server family. The IMM2 enhancements include a more responsive user interface, faster power-on capability, and increased remote presence performance.

The IMM2 provides the following major features as standard:

- ▶ Intelligent Peripheral Management Interface (IPMI) V2.0 compliance
- ▶ Remote configuration of IMM2 and Unified Extensible Firmware Interface (UEFI) settings without the need to power on the server
- ▶ Remote access to system fan, voltage, and temperature values
- ▶ Remote IMM and UEFI update
- ▶ UEFI update when the server is powered off
- ▶ Remote console by way of a serial over LAN
- ▶ Remote access to the system event log
- ▶ Predictive failure analysis and integrated alerting features; for example, by using Simple Network Management Protocol (SNMP)
- ▶ Remote presence, including remote control of server by using a Java or Active x client
- ▶ Operating system failure window (blue screen) capture and display through the web interface
- ▶ Virtual media that allow the attachment of a diskette drive, CD/DVD drive, USB flash drive, or disk image to a server
- ▶ Syslog alerting mechanism that provides an alternative to email and SNMP traps
- ▶ Support for FoD enablement of server functions, option card features, and System x solutions and applications

2.5.2 Chassis Management Module

The CMM provides single-chassis management, and it is used to communicate with the management controller in each compute node. It provides system monitoring, event recording, and alerts and manages the chassis, its devices, and the compute nodes.

The chassis supports up to two CMMs. If one CMM fails, the second CMM can detect its inactivity, activate, and take control of the system without any disruption. The CMM is central to the management of the chassis, and is required in the Enterprise Chassis.

Through an embedded firmware stack, the CMM implements functions to monitor, control, and provide external user interfaces to manage all chassis resources. By using the CMM, you can perform the following tasks:

- ▶ Define login IDs and passwords
- ▶ Configure security settings, such as data encryption and user account security
- ▶ Select recipients for alert notification of specific events
- ▶ Monitor the status of the compute nodes and other components
- ▶ Find chassis component information
- ▶ Discover other chassis in the network and enable access to them
- ▶ Control the chassis, compute nodes, and other components
- ▶ Access the I/O modules to configure them
- ▶ Change the start sequence in a compute node
- ▶ Set the date and time
- ▶ Use a remote console for the compute nodes
- ▶ Enable multichassis monitoring
- ▶ Set power policies and view power consumption history for chassis components
- ▶ Perform diagnostic tests for the chassis, I/O options, and compute nodes
- ▶ Initialize chassis and compute nodes
- ▶ Perform resource discovery and inventory management
- ▶ Perform resource alerts and monitoring management
- ▶ Perform chassis and compute nodes power management

2.5.3 Flex System Manager

Flex System Manager is a high-performance scalable systems management appliance with a preinstalled software stack. It is designed to optimize the physical and virtual resources of the Flex System infrastructure while simplifying and automating repetitive tasks. Flex System Manager provides easy system setup procedures with wizards and built-in expertise and consolidated monitoring for all of your resources, including compute, storage, networking, and virtualization resources.

Flex System Manager has full, built-in virtualization support of servers, storage, and networking to speed provisioning and increase resiliency. In addition, it supports open industry standards, such as operating systems, networking and storage fabrics, virtualization, and system management protocols to easily fit within existing and future data center environments.

Flex System Manager provides the following advantages:

- ▶ Reduce the number of interfaces, steps, and clicks that it takes to manage IT resources
- ▶ Allows IT staff to intelligently manage and deploy workloads that are based on resource availability and predefined policies.
- ▶ Provides IT staff with the tools to manage events and alerts to increase system availability and to reduce downtime.
- ▶ Reduces operational costs by increasing overall efficiency of your operational teams.

The Flex System Manager management appliance is shown in Figure 2-20.



Figure 2-20 Flex System Manager management appliance

As an appliance, Flex System Manager is delivered preinstalled onto a dedicated compute node platform, which is designed to provide a specific purpose. It is intended to configure, monitor, and manage Flex System resources in up to 16 Flex System Enterprise Chassis, which optimizes time-to-value. Flex System Manager provides an instant resource-oriented view of the Enterprise Chassis and its components, which provides vital information for real-time monitoring.

The Flex System Manager offers the following functions:

- ▶ Support for up to 16 managed chassis
- ▶ Support for up to 5,000 managed elements
- ▶ Auto-discovery of managed elements
- ▶ Overall health status
- ▶ Monitoring and availability
- ▶ Hardware management
- ▶ Security management
- ▶ Administration
- ▶ Network management (Network Control)
- ▶ Storage management (Storage Control)
- ▶ VM lifecycle management (VMControl Express)
- ▶ I/O address management (Fabric Manager)

2.5.4 Introduction to Upward Integration

For VMware vSphere and Microsoft Hyper-V virtualized environments, Lenovo offers powerful extensions that are called Upward Integration Modules (UIMs). These UIMs integrate hardware management features, such as status monitoring, firmware upgrades, and predictive failure alerts (PFA) into a management application (VMware vCenter and Microsoft System Center).

Upward integration for VMware vSphere

System x UIMs for VMware vSphere provides IT administrators with the ability to integrate the management features of the System x offerings with VMware vCenter.

UIM for VMware expands the virtualization management capabilities of VMware vCenter with System x hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration. It provides the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

When combined with the management features of System x offerings, VMware vCenter enhances and extends VMware's virtualization technologies and hardware service management to help you dramatically reduce complexity and cost.

The following key features are included:

- ▶ Overview of the host or cluster status, including information summary and health messages of the managed entities.
- ▶ Collects and analyzes system information to help diagnose system problems.
- ▶ Acquires and applies the latest UpdateXpress System Packs™ and individual firmware updates to your ESXi system.
- ▶ Provides nondisruptive system updates, which automates the update process of the hosts in a cluster environment without any workload interruption.
- ▶ Monitors and provides a summary of power usage, thermal history, and fan speed and a trend chart of the managed host. Enable or disabled the Power Metric function on a host and set the power capping for a power-capping capable host to limit the server power usage. Support power throttling and provide notification if the server power usage exceeds the specific value.
- ▶ Manage the current system settings on the host including IMM, uEFI, and boot order settings for the host.
- ▶ Monitor the server hardware status and automatically evacuate VMs in response to predictive failure alerts to protect your workloads.

For more information, see the following System x Upwards Integration Modules for VMware vSphere website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?Indocid=migr-vmware>

Upward integration for Microsoft System Center

The System x UIM for Microsoft System Center v5.5 is a new offering that provides IT administrators with the ability to integrate the management features of the System x, BladeCenter, and Flex System servers with Microsoft System Center.

UIM expands Microsoft System Center server management capabilities by integrating System x hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration. It provides the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

The System x UIM for Microsoft System Center provides the following features:

- ▶ Integrated end-to-end management of System x hardware with monitoring of physical and virtual server health
- ▶ Operating system deployment with the latest System x firmware and driver update management
- ▶ Automated VM migration that is based on server health or power usage
- ▶ Capability to perform hardware configuration and firmware and driver updates and to check for the latest updates from the support website
- ▶ Capability to collect specific hardware inventory of System x rack, tower, and blade servers
- ▶ Capability to power on and off blades via Microsoft System Center Console

- ▶ Capability to author configuration packs to perform compliance checking on System x or BladeCenter x86 blade servers
- ▶ Capability to manage servers remotely independent of operating system state
- ▶ One year of software service and maintenance (three years available as an option)

System x UIM for Microsoft System Center can be purchased as a 1-year or 3-year software service and maintenance license.

For more information, see the following Upward Integration for Microsoft System Center bundle website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?lnocid=migr-5087849>

Part 2

VDI design considerations

In this part, we describe design considerations for VDI solutions that are based on System x offerings.

This part includes the following chapters:

- ▶ Chapter 3, “VMware vSphere design considerations” on page 47
- ▶ Chapter 4, “Microsoft Hyper-V design considerations” on page 59
- ▶ Chapter 5, “Citrix XenDesktop design considerations” on page 67

VMware vSphere design considerations

VMware vSphere management infrastructure and ESXi hypervisor are the preferred choice for virtualization components that are used in the XenDesktop on the Flex System solution. One of the key features of ESXi hypervisor is its ability to be embedded into a server, which provides “bare metal” virtualization capabilities without a need to perform extra installation and configuration tasks.

The technical and commercial features of VMware vSphere steadily evolved from version to version. Throughout this progression, two main components continued to define the fundamental virtualization platform: the ESXi hypervisor and the vCenter management layer.

In addition to these two significant components, it is important to note that a complete virtualization platform includes two other components: storage and networking.

This chapter provides an overview of VMware vSphere 5.x and design guidelines that are adapted for System x hardware and Citrix XenDesktop.

This chapter includes the following topics:

- ▶ 3.1, “ESXi and vSphere features” on page 48
- ▶ 3.2, “Networking considerations” on page 54
- ▶ 3.3, “Storage considerations” on page 56

3.1 ESXi and vSphere features

This section describes the main features of VMware vSphere 5.x, including the ESXi Hypervisor, the VMware vCenter server, vMotion, the Distributed Resource Scheduler (DRS), and vSphere High Availability.

The VMware vCenter server is one of the core management components for virtual desktop infrastructure (VDI). It is widely used to manage the full virtual machine (VM) lifecycle and to monitor the virtual environment. It also provides advanced functionality for VMs, such as high availability, live migration, and workload allocation.

In a VDI environment, advanced functions are used to provide required levels of availability for the server management components and for persistent virtual desktops. For non-persistent desktops, these advanced functions are not required because the availability and workload management functions are performed by the XenDesktop Controller.

3.1.1 ESXi hypervisor

In the vSphere infrastructure, the hypervisor is VMware vSphere Hypervisor (ESXi). ESXi provides a virtualization layer and creates an abstraction layer for processor, memory, storage, and networking resources of the hosted client.

Evolved from ESX, ESXi has a small disk footprint that allows it to be stored on internal flash memory, such as a USB key that is plugged into the system board of a supported System x rack or blade server.

The System x customized version of ESXi provides more drivers and Common Information Model (CIM) modules that are specific to System x hardware to provide online platform management, including updating and configuring firmware, platform diagnostics, and enhanced hardware alerts. The System x ESXi option is delivered on a USB flash drive.

Choosing the ESXi embedded USB option on the servers that make up the VDI infrastructure provides the following benefits:

- ▶ Reduce server deployment time.
- ▶ Use a diskless compute node, which reduces cost and security exposures.
- ▶ Compute node's local disks are available for hosting non-persistent virtual desktops.

Starting with vSphere Version 5.0, VMware offers only ESXi as the hypervisor.

In the VDI environment, ESXi supports management clusters and VDI compute clusters, as described in Chapter 5, “Citrix XenDesktop design considerations” on page 67.

Consider the following design suggestions:

- ▶ Consult the VMware HCL when the server model is selected.
- ▶ Use the latest stable version of ESXi that is compatible with all other products that are used in the solution.
- ▶ Select hosts with a higher CPU core count per CPU socket to minimize VMware licensing costs.
- ▶ Consider the use of fewer, larger hosts in large environments and more, smaller hosts in smaller environments.

- ▶ Typically, memory over commitment is not used, or it is used only for non-critical environments. If the ESXi host does not run into memory contention issues, ballooning or swapping do not occur.
- ▶ ESXi hosts must have fully redundant hardware components, including redundant network cards, redundant host bus adapters (HBAs) for SAN access, and redundant power supplies.

3.1.2 VMware vCenter Server

VMware vCenter Server is a mandatory component to provide a centralized and extensible platform for managing virtual infrastructure. Many advanced features such as high availability HA, DRS, vMotion, and dvSwitches are available through vCenter Server only.

VMware vCenter Server is also a critical component in a XenDesktop environment because of its central role of managing all communication between XenDesktop and vSphere. Each VMware cluster relies on vCenter to perform cluster management and other hosting infrastructure tasks. The delivery of desktops might be affected if vCenter becomes slow or unresponsive under high stress conditions, such as in a large XenDesktop environment with many morning logons or rapid shift changes.

The VMware recommendation is to use vCenter as a VM, which allows for protection of vCenter with HA. Achieving HA for the VMware vCenter Server is also recommended by Citrix for XenDesktop deployments.

Starting with Version 5.1, the vCenter architecture was changed by decoupling components, such as inventory services or by introducing new components, such as single sign-on (SSO), which can be installed on separate servers. This configuration allows more flexibility in sizing and designing. Version 5.5 improves scalability and performance, especially for vCenter appliance (vCSA) and makes it a valid alternative to Windows version of vCenter.

In addition to the classic vSphere client, VMware introduced a new web client. All operations are possible now via the web client and certain operations can be performed only via the web client.

An important consideration is the communication between vCenter Server and XenDesktop Desktop Delivery Controller (DDC); a third-party or self-signed certificate must be installed on the vCenter server and the DDCs that are in the environment. Although a self-signed certificate can be used in non-production environments, Citrix suggests the use of a certificate that is provided by a third-party certificate authority (CA) or an internal enterprise CA for production use.

3.1.3 vMotion

By using live migration or vMotion technology, you can move running VMs from one physical server to another with no downtime. This ability enables companies to perform hardware maintenance without disrupting business operations.

vMotion relies on the following mechanisms:

- ▶ Encapsulation of the VM state in files that are stored on shared storage
- ▶ Transfer of the active memory of a VM over a network
- ▶ Virtualized network that is used by the VM to ensure that the network identity and network connections are preserved

vMotion preserves the execution state, network identity, and active network connections with no disruption to users.

Storage vMotion enables moving VM disks from one physical storage location to another without an outage in the guest operating system and applications. Storage vMotion is used by system administrators to relocate VMs when changes must be implemented in the physical infrastructure, or when the VM must grow its storage and there is not enough available space in the current physical container.

Before vSphere 5.1, vMotion required shared storage between hosts. Storage vMotion required a host to have access to the source and destination data stores.

vSphere 5.1 or newer removes this requirement and allows combining vMotion and storage vMotion into one process. This combined migration copies the VM memory and its disk over the network to the destination host. After all memory and disk data are sent, the destination VM resumes and the source VM is powered off (see Figure 3-1).




Figure 3-1 vMotion

In the VDI environment, vMotion is used to provide live migration capabilities for management servers VMs and persistent virtual desktops.

The following designs and preferred practices are suggested:

- ▶ VMs use the latest virtual hardware (Version 9 for vSphere 5.5).
- ▶ Separate the vMotion network from management and VM networks.
- ▶ If possible, leave some CPU resources for vMotion operations. To ensure the ability to use full network bandwidth, ESXi reserves CPU resources on the source and destination hosts.

3.1.4 Distributed Resource Scheduler

vSphere DRS works with vMotion (see Figure 3-2) to provide automated resource optimization and VM placement. DRS uses vMotion to balance the workload across all hosts in a cluster based on CPU and memory activity.

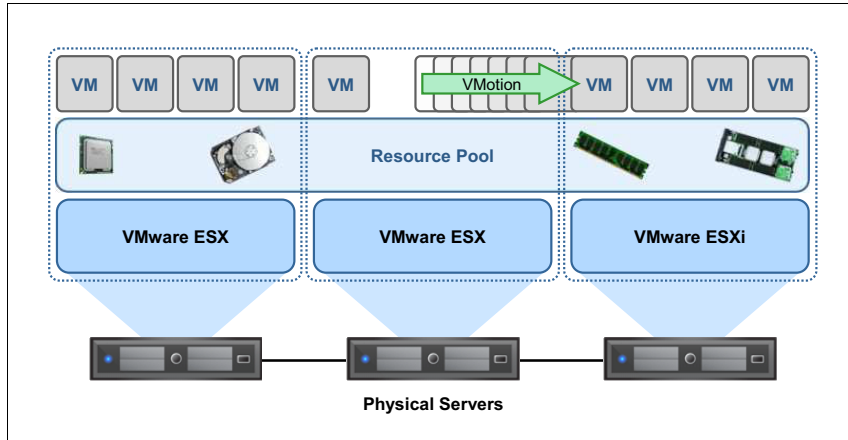


Figure 3-2 Distributed Resource Scheduler

DRS enhances the consolidation ratio by deciding how the resources can be optimized in terms of workload placement. It enables performance management and capacity planning savings, and incident management savings. It is also used to automate workload distribution when physical hosts are placed in maintenance mode during changes.

With DRS enabled, you can create resource pools that span all hosts in the cluster and apply cluster-level resource allocation policies.

DRS also can perform the following functions:

- Initial placement

When a VM is powered on, DRS places it on an appropriate host or generates a recommendation, depending on the automation level.

- Load balancing

DRS distributes VM workloads across the vSphere hosts inside the cluster. DRS continuously monitors the workload and the available resources and performs or suggests VM migrations to maximize workload performance.

- Power management

Distributed Power Management (DPM) can place vSphere hosts in standby mode or power them back on as capacity needs. DPM can also be set to issue recommendations for power on/off operations.

- Constraint correction

DRS redistributes VMs across vSphere hosts as needed to adhere to user-defined affinity and anti-affinity rules following host failures or hosts that are placed in maintenance.

The following designs and preferred practices are suggested:

- ▶ Enable DRS on the entire cluster in fully automated mode, unless there are specific constraints.
- ▶ If needed, you can change the default DRS settings on specific VMs.
- ▶ Configure affinity and anti-affinity rules and DRS groups only when necessary (if certain VMs must run on certain hosts, run the VMs on separate hosts or the same host). A use case for these rules is vCenter, which must run on 1 - 2 hosts to locate it faster for troubleshooting purposes.

3.1.5 High Availability

vSphere HA provides an automated process for restarting VMs when a physical host becomes unavailable (see Figure 3-3). VMs are automatically registered and restarted on the remaining hosts in the cluster.

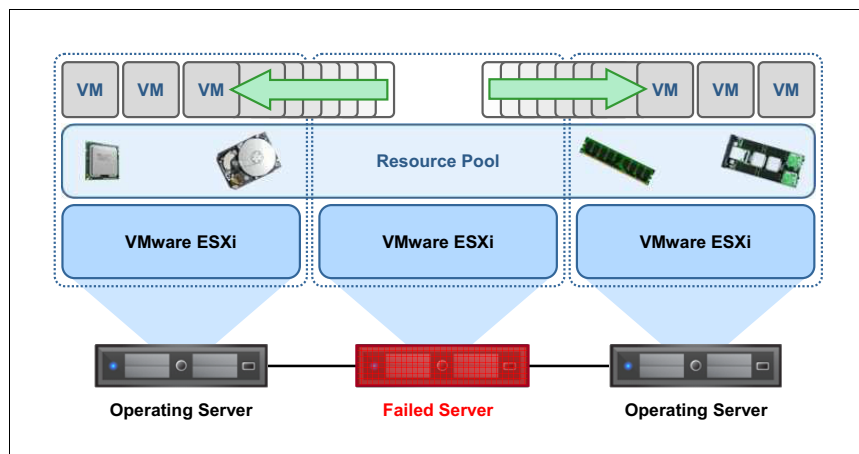


Figure 3-3 High Availability

HA helps to meet service level agreements (SLAs) and to manage the risk of having aggressive consolidation ratios on physical hosts by reducing the potential for long outages. When hardware failures occur, HA helps to reduce labor by providing recovery automation.

When vSphere HA is enabled for a cluster, all active hosts choose the cluster's master host. Only one master host exists per cluster and all other hosts are slave hosts. A new election is held if the master host fails, is shut down, or is removed from the cluster.

The master host in a cluster has the following responsibilities:

- ▶ Monitoring the state of slave hosts. If a slave host fails or becomes unreachable, the master host identifies the VMs that must be restarted.
- ▶ If VM monitoring is enabled, the master host monitors the power state of all protected VMs. If one VM fails, the master host ensures that it is restarted.
- ▶ Managing the lists of cluster hosts and protected VMs.
- ▶ Acting as the vCenter Server management interface to the cluster and reporting the cluster health state.

The slave hosts primarily run VMs, which monitor their runtime states and report state updates to the master host. A master host can also run and monitor VMs. Slave hosts and master hosts implement the VM and Application Monitoring features.

In the VDI environment, VMware HA provides HA for management services VMs and persistent virtual desktops, if required.

The following designs and preferred practices are suggested:

- ▶ HA is enabled on all clusters with strict admission control, with one exception: the cluster for non-persistent desktops, which have HA disabled.
- ▶ Clusters that have 12 or fewer hosts must allow for the loss of at least one physical host. Clusters with more than 12 hosts must allow for the loss of at least two physical hosts.
- ▶ Configure the Percentage of Cluster Resources Reserved policy and reserve failover capacity for at least one host. Use the Host Failures Cluster Tolerates policy if VM reservations are not used and you do not need granular control of reserved failover capacity. Use the percentage policy if you have a cluster of only two hosts. There might be a requirement for desktop groups to offer varying levels of redundancy. For example, a desktop group might require $N+100\%$ redundancy while another one might require $N+10\%$ only.
- ▶ HA works even if vCenter is down; however, vCenter is needed to initially configure HA.

3.1.6 vSphere licensing considerations

vSphere 5.x is licensed on a per-processor basis. Each physical processor (CPU) in a server must have at least one vSphere 5.1 or 5.5 processor license key that is assigned to run vSphere. vRAM entitlement that was introduced in vSphere 5.0 was ended with vSphere 5.1.

vCenter Server is licensed per instance. One instance is required in a vSphere deployment to enable the centralized management and deployment of core vSphere features, such as vMotion, Distributed Resource Scheduler, and others.

vSphere is available in three editions that provide basic features in the Standard edition to the full range of features in the Enterprise Plus edition.

The vSphere features and the required vSphere edition are listed in Table 3-1.

Table 3-1 vSphere features and editions

Feature	vSphere edition
Thin provisioning	Standard
vMotion	Standard
High Availability	Standard
Hot-Add RAM and CPU	Enterprise
Fault Tolerance	Enterprise
DRS	Enterprise
Storage multipathing	Enterprise
Storage vMotion	Enterprise
Host profiles	Enterprise Plus
Storage DRS	Enterprise Plus
Storage I/O Control	Enterprise Plus
Network I/O Control	Enterprise Plus

Feature	vSphere edition
Distributed Switches	Enterprise Plus

3.1.7 System x integration with VMware

The Flex System Manager accelerates the provisioning of compute node, networking, and storage resources to the VMware ESXi software layer and supporting Citrix XenDesktop components. These capabilities decrease deployment time significantly.

VMware integration offers the following features:

- ▶ Deploying hardware patterns from the Flex System Manager to new compute nodes, which ensures that standard hardware adapter interfaces are logically assigned to the compute resources as suitable for a new vCloud Suite compute node.
- ▶ Installing Lenovo customized ESXi 5.x images to the new compute nodes from inside the Flex System Manager interface.
- ▶ Providing VMware environment visibility and ESXi resource inventory and topology views from within the Flex System Manager interface, including the ability to deploy new VM images.
- ▶ Providing extensibility from the native vCenter server to the System x and Flex System hardware by using the specialized Upward Integration Module (UIM). Capabilities of the UIM include monitoring power and thermals of the hardware components, viewing and updating firmware and software levels for various components, and modifying settings for predictive failure alerts.

3.2 Networking considerations

Networking can be seen as physical network infrastructure and VMware vSphere virtual network infrastructure.

From a physical network perspective, the host's networking resources are shared by the virtual desktops that it supports. If there is no sufficient bandwidth, users experience a degraded level of performance. It is suggested to use fast network cards and Flex System compute nodes to address this issue by using 10 Gb network cards.

Also, performance might be improved by separating different types of network traffic. For example, management, VM, storage, provisioning, and backup traffic can all be isolated from each other. For more information about network design, see 5.5, "Network configuration" on page 79.

The VMware virtual network consists of various subcomponents, such as virtual switches (standard and distributed), ports, port groups, virtual Ethernet adapters, and uplink ports. These components build the communication channel between the VMs and the external or physical network.

3.2.1 Virtual switch

Virtual switches (vSwitches) are a software-based switch that is in the VMkernel and provides traffic management for VMs. There are two types of virtual switches in vSphere: standard switch (vSS) and distributed switch (vDS).

Although standard virtual switches are defined at the host level, distributed virtual switches are defined at the data center level, which means that virtual switch configuration is then pushed consistently to all hosts that are within the same data center. A vSphere Distributed Switch is abbreviated as VDS and is also called a dvSwitch.

In addition, distributed virtual switches enable advanced features, such as Rx traffic shaping, improved monitoring through port mirroring (dvMirror), consistent network statistic monitoring or Link Layer Discovery Protocol (LLDP), which is a vendor-neutral standard equivalent of Cisco Discovery Protocol (CDP).

The vDS requires an Enterprise license. Use distributed switches to enable the benefits.

Note: When dvSwitches are used, they often can be controlled only from your vCenter server. If your vCenter Server becomes unavailable, networking continues to function, but you cannot make any modifications until the vCenter Server is back online.

3.2.2 Ports and port groups

A *port* or *port group* is a logical object on a vSwitch that provides specialized services for the VMkernel or VMs. A virtual switch can contain a VMkernel port or a VM port group. On a vSphere Distributed Switch, these groups are called *dvPort groups*.

A *VMkernel port* is a specialized virtual switch port type that is configured with an IP address to allow vMotion, iSCSI storage access, network-attached storage (NAS), or Network File System (NFS) access, or vSphere Fault Tolerance (FT) logging. Because vSphere 5.x includes ESXi hosts only, a VMkernel port also provides management connectivity for managing the host. A VMkernel port is also referred to as a *vmknics*.

A *VM port group* is a group of virtual switch ports that share a common configuration and allow VMs to access other VMs or the physical network.

3.2.3 Uplink ports

Uplink ports are ports that are associated with physical adapters, which provide a connection between a virtual network and a physical network.

Distributed Virtual Uplinks (dvUplinks) are a new concept that was introduced with vDS. dvUplinks provide a level of abstraction for the physical NICs (vmnics) on each host. NIC teaming, load balancing, and failover policies on the vDS and DV Port Groups are applied to the dvUplinks and not the vmnics on individual hosts. Each vmnic on each host is mapped to a dvUplink, which permits teaming and failover consistency regardless of vmnic assignments.

The following designs and preferred practices are suggested:

- ▶ Ensure redundancy by using a single dvSwitch with more uplinks on all of the hosts in the cluster.
- ▶ Create separate, highly available port groups for each management and vMotion traffic. The Flex System platform positions Ethernet switch hardware inside the chassis, which provides inherent network performance improvement for activities that use network bandwidth (such as VMware vMotion) from traditional top-of-rack (TOR) network switching.

Note: Internal Layer 2 switches provide a more effective approach for communication between co-resident servers by using an east-west approach. Communication between nodes uses an internal, active Layer 2 switch to pass traffic to one another. By containing network traffic within the Flex System chassis, latency is improved compared to a north-south approach. In the north-south approach, all of the traffic is routed to the top-of-rack (TOR) switch and the flow goes up to the TOR switch and down to the co-located server.

- ▶ Improve the network performance by using the TCP offload engine (TOE) capabilities of integrated network adapters on the Flex System compute nodes, by enabling stateless offload of the following tunables:
 - Checksum offload
 - TCP segmentation offload (TSO)
 - Jumbo frames (JF)
 - Large receive offload (LRO)

3.3 Storage considerations

Storage has a significant effect on the performance, scalability, and availability of the Citrix XenDesktop implementation.

3.3.1 Local or shared storage

Virtual deployments often use shared storage in preference to local storage. Shared storage is required to support vMotion, DRS, and HA. Although these features are less critical when hosting non-persistent virtual desktops, they are important for management server workloads and persistent desktops.

3.3.2 Tiered storage

A one-size-fits-all storage solution is unlikely to meet the requirements of most virtual desktop implementations. The use of tiered storage, where storage technologies, such as solid-state drives (SSDs) and network-attached and Fibre Channel-attached storage systems, and drive access technologies, such as SAS and SATA, are grouped into storage tiers, provides an effective mechanism for offering a range of storage options that are differentiated by performance, scalability, redundancy, and cost. In this way, different virtual workloads with similar storage requirements can be grouped and a similar cost model applied.

3.3.3 Redundancy

vSphere Datastores must be designed to meet the redundancy requirements of the components that they support, such as RAID levels, storage adapters, and the back-end storage configuration. The preferred practice for shared storage is to configure two NICs or HBAs in a bonded or multipath setup.

VMware vSphere uses a default storage multipath policy of *Fixed (VMware)*, which means that the same storage path is always used to access that specific logical unit number (LUN). If you have a configuration where you have multiple access paths to your storage LUNs, this policy is not the optimal multipath policy because it does not make the most of your redundant hardware.

Selecting Round Robin (VMware) often is a good choice. It means that at any time, the LUN is accessed over a single path, but that path changes the next time that it is accessed.

Microsoft Hyper-V design considerations

Microsoft Windows Hyper-V technology offers a basis for the virtualization layer that is underlying Citrix XenDesktop on System x servers.

The capabilities of Hyper-V technology evolved from the initial offering of Hyper-V 2005 as an add-on to the Windows Server 2003 operating system. This evolution includes the base hypervisor and the development of the management infrastructure to the level offered by the current Microsoft System Center product. In addition, Microsoft added software that defined storage and software defined networking.

This chapter provides the design considerations for Microsoft Hyper-V when used with Citrix XenDesktop on System x platform.

This chapter includes the following topics:

- ▶ 4.1, “Hyper-V virtualization and management features” on page 60
- ▶ 4.2, “Networking considerations” on page 63
- ▶ 4.3, “Storage Considerations” on page 64

4.1 Hyper-V virtualization and management features

This section describes the features of Microsoft Hyper-V, including Hyper-V Server, which is a free download from Microsoft. It also describes the use of the Microsoft System Center Virtual Machine Manager (SCVMM) to manage virtualization hosts and guest virtual machines (VMs).

In a virtual desktop infrastructure (VDI) environment, advanced functions are used to provide required levels of availability for the server management components, and for persistent virtual desktops. For non-persistent desktops, these advanced functions are not required because the availability and workload management functions are performed by the XenDesktop Controller.

4.1.1 Hyper-V overview

The Hyper-V hypervisor can take the form of a role in the full Windows operating system (OS) whether Core or GUI-based. It can also take the form of the Hyper-V Server, which is a free download. In either form, it provides the virtualization layer and creates an abstraction layer to provide processor, memory, storage, and networking resources for the guest OS.

Hyper-V includes the following advanced features:

- ▶ Awareness of Non-Uniform Memory Architecture (NUMA) hosts.
- ▶ Support of NUMA guests to match NUMA hosts.
- ▶ High availability (HA) by using Hyper-V clusters.
- ▶ Physical to virtual (P2V) host to guest transitions, which are scriptable.
- ▶ Virtual to Virtual (V2V) guest transition (for example, VMware to Hyper-V).
- ▶ Bare metal deployment of new servers to Hyper-V clusters.
- ▶ Intelligent placement of VMs on hosts, including Anti-Affinity of HA VMs to avoid a single point of failure of one host.
- ▶ Dynamic Optimization of VMs and Power Optimization by using this technology.
- ▶ Shared storage hosts that use cluster volumes or SMB V3 shares, and shared nothing hosts that use network connectivity only.
- ▶ Live migration of storage.

In the VDI environment, Hyper-V supports management clusters and VDI compute farms, as explained in Chapter 5, “Citrix XenDesktop design considerations” on page 67.

Consider the following design suggestions:

- ▶ Consult the Microsoft HCL and ServerProven® when the server model is selected.
- ▶ Review the Microsoft licensing model, which is based on a license that supports two sockets.
- ▶ Note the unlimited VM license model of the Microsoft Datacenter offering.
- ▶ Virtualization hosts ideally have fully redundant hardware components, including network ports, host bus adapters for storage access, and redundant power supplies. System x rack servers and Flex System compute nodes match well with these requirements.

4.1.2 Hyper-V management

Hyper-V can be managed at multiple levels. On a hierarchical basis, these levels start with command-line management through the OS command-line interface (CLI), the Windows Management Interface (WMI) that uses the WMI command line (WMIC) and Microsoft Powershell. The next level is the Microsoft management console (MMC), which can be used from the full OS installation in GUI mode or from the equivalent console in a client OS by using the plug-in that is available for download from the Microsoft website.

Note: Only Windows 8 and above can manage Microsoft Windows Server 2012 or 2012 R2. Similarly, Windows Server 2008 can be managed only from these operating systems or client systems at the Vista or Windows 7 level. The only overlap is that Windows Server 2008 R2 can be managed by either environment.

Alternatively, Hyper-V can be managed from the Microsoft SCVMM, which manages individual Hyper-V hosts and Hyper-V clusters. SCVMM enables the advanced features that are described in section 4.1.1, “Hyper-V overview” on page 60. Consider embedding the management server in a VM to use the capabilities of the infrastructure to provide HA, which is also considered by Citrix for XenDesktop environments.

Citrix XenDesktop uses the Microsoft Structured Query Language (SQL) server for its Site Configuration Database.

Note: If Citrix policy details are stored in Active Directory (AD) instead of the Site Configuration Database, Microsoft Group Policy Management Console is required.

Microsoft AD is used in a Citrix XenDesktop environment to ensure secure communications between its individual elements. The AD schema is extended beyond the standard tree that Microsoft provides to support XenDesktop. For more information, see this website:

<http://support.citrix.com/proddocs/topic/xenapp-xendesktop/cds-xenapp-xendesktop-75-landing.html>

4.1.3 Lenovo management devices and tools for Hyper-V

The Flex System Manager accelerates the provisioning of compute node, networking, and storage resources for the installation of Hyper-V and supporting Citrix XenDesktop components. These capabilities decrease deployment time significantly.

The Flex System Manager offers the following key features:

- ▶ Configuration patterns, which configure Flex System nodes and chassis with the wanted hardware settings to ease the installation of Hyper-V. These patterns also can be applied to the placeholder chassis before the hardware is installed. The placeholder configurations can then be deployed as the installation progresses, which ensures a homogeneous installation base.
- ▶ Network configuration of chassis switches by using templates to ensure that they are configured to match the patterns on the nodes.

In addition to the Flex System Manager, Upwards Integration Modules (UIMs) for SCVMM and other modules for Systems Center are available to assist with managing the hardware (System x rack servers and Flex System compute nodes) that is underlying the Hyper-V installation. Capabilities of the UIM include monitoring power and thermal values, viewing and updating firmware and software levels, and enabling SCVMM to react to predictive failure alerts of the servers.

4.1.4 VM and Storage Migration

Live migration on Hyper-V can take the following forms:

- ▶ Clustered Hyper-V hosts on Windows Server 2012 and 2012 R2 can auto migrate between clustered nodes that are based on Cluster Shared Volumes as in previous versions of Hyper-V.
- ▶ Migrate VMs by using file-based access on an SMB version 3 share.
- ▶ Manual migration on a shared-nothing basis. This migration can be between clusters, stand-alone nodes, or clustered and stand-alone nodes in either direction.

Hyper-V can migrate running VMs between heterogeneous storage. The following components can be migrated:

- ▶ Virtual hard disk drives (VHD)
- ▶ Configuration files
- ▶ Snapshots

A snapshot is a checkpoint in time of a VM that can be taken by SCVMM, even when a VM is running. No downtime is needed for this process to occur. Hyper-V can also replicate entire hosts configurations or guest VMs between nodes and clusters by using heterogeneous storage.

In a VDI environment, live and storage migration can be used for management server VMs and persistent desktops. Consider separating migration networks from management and data networks for guest VMs.

4.1.5 VM placement

SCVMM implements automated resource optimization and VM placement. It balances the workload across all hosts in individual clusters based on processor and memory activity. It enhances the consolidation ratio by deciding how the resources can be optimized in terms of workload placement. It also enables performance management and capacity planning savings and incident management savings. It is also used to automate workload distribution when physical hosts are placed in maintenance mode.

SCVMM can perform the following tasks:

- ▶ Initial placement
When a VM is powered on, SCVMM offers an automated or manual placement of the workload that is based on pre-configured preferences.
- ▶ Load balancing
By using Dynamic Optimization, SCVMM dynamically distributes workloads across clusters, which monitor the physical hosts to ensure that changes in individual guest workloads are balanced across the cluster.

- Power management

Dynamic optimization ensures that as workload needs decline, physical hosts are cleared of any workload and powered off. Out-of-band management is used to affect this control power and to power on sufficient hosts to react to any increase in guest workloads.

- VM Distribution

SCVMM redistributes guest VMs according to user predefined rules to achieve affinity, where VMs should be on the same host. It also uses the same mechanism for ant-affinity for example where multiple nodes of the same guest cluster should not be on the same physical host.

4.1.6 High Availability with Hyper-V clusters

SCVMM provides HA by using automated processes to migrate VMs when machines are placed into a maintenance mode or when they suffer any unexpected downtime. This process is achieved when sufficient capacity exists on the other physical hosts of the cluster. To aid this process, SCVMM can shut down low priority VMs based on pre-existing user rules. In addition, any low priority VMs from the previous host are not started if they are marked as not auto relocatable. By using these techniques, SCVMM optimizes the availability of critical VMs and allows more aggressive consolidation ratios on physical hosts with its ability to prioritize workloads.

Arbitration of failover clustering in Windows Server 2012 R2 depends on a quorum vote where the ownership of cluster resources is decided by most of the nodes. SCVMM monitors the health of the hosts in the cluster and migrates VMs on any system crash. In addition, the individual hosts monitor guest VMs by using Integrated Components (IC), including monitoring applications and services to determine whether they should be restarted or migrated to another host. ICs are built in to Windows Server 2012 and 2012 R2.

In the VDI environment, HA provides availability for management VMs and persistent desktops. Up to 64 hosts can be configured in one Hyper-V cluster. Preferred practices suggest that when such a cluster is designed, sufficient free resources are available to absorb the workload of at least two missing hosts, with a suggestion that the loss of up to one host in 10 can be supported.

4.2 Networking considerations

From a physical network perspective, the hosts' networking resources are shared by the virtual desktops that they support. If there is insufficient bandwidth, users might have a degraded desktop experience. For that reason, the preferred practice is to use the fastest network hardware possible. System x platform addresses this issue by offering 10 GbE network infrastructure from the compute servers outwards.

Also, performance can be affected if there is contention between differing network activity. For example, management, storage, provisioning, and backup network activity can each have peaks of activity that contend for network bandwidth. Consider isolating these networks from each other by using virtual NICs in separate VLANs as offered by the Virtual Fabric Adapters (integrated LOMs or PCIe cards) for System x servers.

The Hyper-V virtual network infrastructure consists of extensible switches, software-defined networks for multi-tenant isolation, virtual network adapters, and extensions for capturing, filtering, and forwarding functionality within the virtual network. These components build the communication channels between the VMs and the internal virtual and external physical networks.

4.2.1 Hyper-V Network Virtualization

Hyper-V Network Virtualization (HNV) was developed to address the need to transform networking in a virtualized environment. The aim is to deliver networking as part of an automated infrastructure with multi-tenant isolation capable of expansion without disruption to capability while reducing operational complexity. HNV uses Software Defined Networking (SDN) to abstract the physical network with virtual networks, span policies across both types of networks, and to control data center traffic flow.

HNV achieves multi-tenant isolation by encapsulating traffic by using Network Virtualization Generic Routing Encapsulation (NVGRE). This process uses a Virtual Subnet ID (VSID) as a key to differentiate between traffic for isolated virtual networks, which provides a customer address space that is within the provider address space. This customer address space can be extended by using a Virtual Private Network to allow separate tenants to connect in isolation from each other while retaining their own IP address space. Windows Server 2012 R2's implementation of NVGRE supports NIC teaming and encapsulated task offload to improve network throughput.

This design of extensible software defined switch also provides the capability for third-party extensions to be embedded within the virtual switch for traffic filtering or diagnostic purposes. The virtual switch also was designed to learn dynamic IP addresses on VMs in the virtual network, which enables the use of DHCP within the user address space even when routed across the provider network.

4.3 Storage Considerations

Storage has a significant influence on the performance, scalability, and availability of the Citrix XenDesktop implementation.

4.3.1 Local or shared storage

Virtual deployments often use shared storage to enable migration technologies for HA and dynamic placement that is supported by SCVMM. This feature is critical for management workloads and persistent desktops but are of less relevance for non-persistent desktop provisioning.

4.3.2 Tiered storage

Different workloads impose differing requirements on storage provisioning. Persistent desktops and management VMs (which by their nature do not need frequent reboots) perform to an acceptable level with normal shared storage. Non-persistent desktops (which by their nature feature peaks of activity at the start and end of the working day) benefit from tiered storage. Some dense implementations might even need flash storage to handle peaks of activity.

Microsoft offers the possibility of tiered storage with their Storage Spaces technology, which uses SSD flash drives as cache on top of SAS disks. This implementation uses software-defined arrays, not hardware RAID. These storage spaces can be used in a cluster environment. The Storwize family offers Easy Tier where a storage volume is constructed of a hybrid array of SSD and SAS disks. The storage controller assigns hot data to the SSDs to ensure that it is served in as fast a manner as possible.

4.3.3 Usage of flash based storage

In a VDI environment, some elements of the solution benefit from flash-only storage. This storage can take the form of direct attached solid-state disk (SSD) drives in the systems, shared storage with SSD arrays (for example a Storwize V7000), or a dedicated flash storage unit, such as the FlashSystem 840.

4.3.4 Redundancy and load balancing

Hyper-V disk resources must be designed to meet the redundancy requirements of the components that they support. To meet this requirement, appropriate RAID levels should be selected and redundant storage adapters should be used for shared storage.

The preferred practice is to use Microsoft Multi-Path I/O (MPIO) that is specific to the storage controllers in use. The architecture of the Microsoft multipath driver is to use a standard MPIO to front end a Disk Specific Module (DSM), which is designed for the storage system in use. The module for V7000 is called the Subsystem Device Driver DSM or SDDDSM.

Citrix XenDesktop design considerations

This chapter describes the Citrix XenDesktop design process that is based on Lenovo Client Virtualization Reference Architecture (RA) for Citrix XenDesktop. A Lenovo validated reference design, the RA for XenDesktop offers support for different hypervisors, including a VMware ESXi back-end infrastructure that is managed by a vCenter Server and Microsoft Hyper-V managed by System Center.

The Lenovo RA is updated regularly to include new features and components. The most recent Lenovo RA for Citrix XenDesktop is available at this website:

<http://lenovopress.com/tips1278>

This chapter includes the following topics:

- ▶ 5.1, “Citrix XenDesktop components” on page 68
- ▶ 5.2, “Desktop and application delivery” on page 70
- ▶ 5.3, “Citrix XenDesktop provisioning” on page 71
- ▶ 5.4, “Storage configuration” on page 76
- ▶ 5.5, “Network configuration” on page 79
- ▶ 5.6, “Operational model and sizing guidelines” on page 80

5.1 Citrix XenDesktop components

Figure 5-1 shows the components of a Citrix XenDesktop architecture that supports several models for desktop delivery: Hosted Virtual Desktop (HVD), streamed HVD, and hosted and streamed applications.

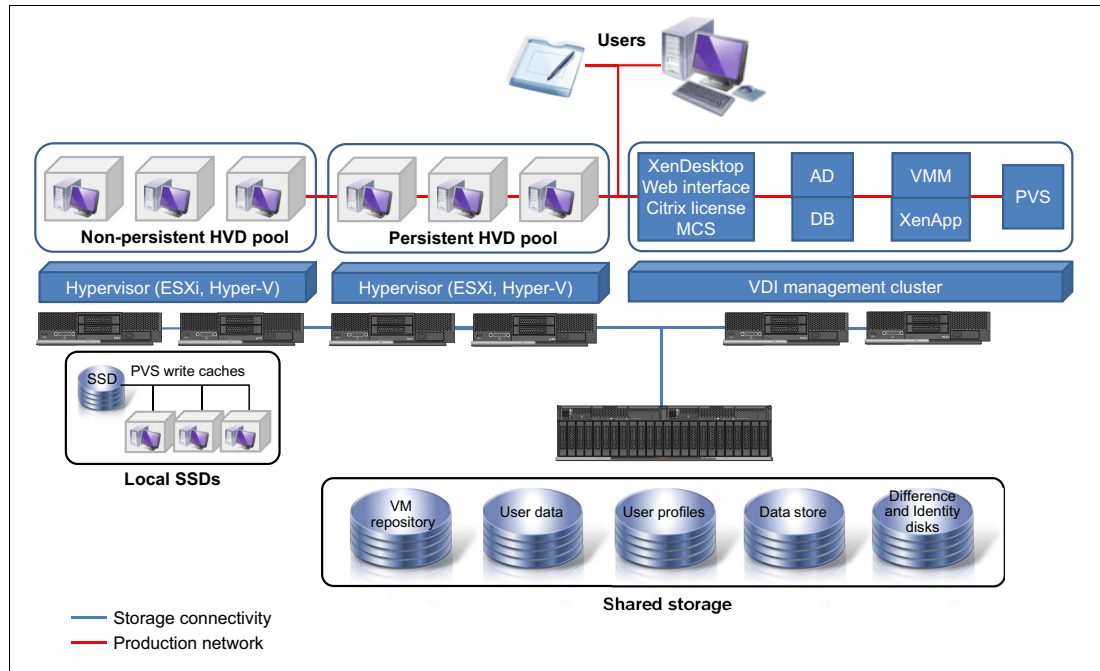


Figure 5-1 Citrix XenDesktop components

The Citrix XenDesktop architecture includes the following main components:

- ▶ **Compute clusters (persistent or non-persistent):**
 - Hosts the virtual desktop workloads.
 - Consists of multiple compute servers.
 - Quantities of compute servers per cluster vary with building block type (for more information, see 5.6, “Operational model and sizing guidelines” on page 80).
 - Must not be used to host workloads other than virtual desktops.
- ▶ **Management cluster:**
 - Hosts the hypervisor and Citrix XenDesktop management components.
 - Can be hosted on an existing or new hypervisor environment.
 - Cluster contains VMware vCenter or Microsoft infrastructure, Citrix XenDesktop, Machine Creation Services (MCS), Provisioning Services (PVS), License Server, Web Interface, database server, and other optional components.
 - Can host more infrastructure services (Active Directory (AD), Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and so on) if they are not already in the environment.

The following components are included:

- Web Interface

The Web Interface provides the user interface to the XenDesktop environment. Web Interface brokers user authentication, enumerates the available desktops and upon start, delivers a .ica file to the Citrix Receiver on the user's local device to start a connection. Because the Web Interface is a critical component, redundant servers must be available to provide fault tolerance.

- Domain Controller

The Domain Controller hosts the following services:

- AD: Provides a common namespace and secure method of communication between all the servers and desktops in the environment.
- DHCP: DHCP is used by the virtual desktops to request and obtain IP addresses. DHCP uses Option 66 and 67 to specify the bootstrap file location and file name to a virtual desktop. The DHCP service receives requests on UDP port 67 and sends data to UDP port 68 on a virtual desktop. Citrix Provisioning Services then streams the operating system over the network to the virtual desktops.
- DNS: The DNS server provides IP host name resolution for the core XenDesktop infrastructure components.

- Desktop delivery controller

The desktop delivery controllers (DDCs) are responsible for maintaining the proper level of idle desktops to allow for instantaneous connections, monitoring the state of online and connected desktops, and shutting down desktops, as needed.

A XenDesktop farm is a larger grouping of virtual machine (VM) servers. The primary DDC is configured as the XenDesktop farm master server. The master focuses on farm management while another DDC acts as a dedicated XML server. The XML server is responsible for brokering user authentication, resource enumeration, and starting the desktop. Because a failure in the XML service results in users being unable to start their desktops, it is suggested that you configure multiple controllers per farm.

- Provisioning Services (PVS) or Machine Creation Services (MCS)

PVS can be used to provision non-persistent VMs. MCS can be used to provision persistent and non-persistent VMs (for more information, see 5.3, "Citrix XenDesktop provisioning" on page 71 for details).

- Virtual machine management (VMM) infrastructure

vCenter Server is the managing server for VMware ESXi hypervisor. By using a single console, it provides centralized management of the VMs.

Redundancy for vCenter Server is achieved through VMware high availability (HA). The vCenter Server also includes a licensing server for VMware ESXi.

- License Server

The Citrix License Server is responsible for managing the licenses for all XenDesktop components. XenDesktop has a 30-day grace period that allows the system to function normally for 30 days if the license server becomes unavailable. This grace period offsets the complexity of otherwise building redundancy into the license server.

- ▶ XenDesktop SQL Server

Each Citrix XenDesktop farm requires an SQL Server database that is called the *data store*, which is used to centralize farm configuration information and transaction logs. The data store maintains all static information about the XenDesktop environment. Because the XenDesktop SQL server is a critical component, redundant servers must be available to provide fault tolerance.

- ▶ Virtual Desktop Agent (VDA)

Each VM needs a Citrix VDA to capture VM data and send it to the Receiver in the client device. The VDA also emulates the keyboard and gestures that are sent from the Receiver.

Note: The VDA is different for HDX 3D Pro because it must capture data from a graphics processing unit (GPU) that is rendering a 3D scene. Independent Channel Architecture (ICA) is the Citrix display protocol for 2D and 3D virtual desktop infrastructure (VDI).

- ▶ Client devices

XenDesktop supports a broad set of devices, including PCs, Mac OS devices, tablets, smartphones, and thin clients, along with all major device operating platforms, including Apple iOS, Google Android, and Google Chrome OS. XenDesktop enables a rich, native experience on each device, including support for gestures and multi-touch features, which customizes the experience based on the type of device. Each client device has a Citrix Receiver, which acts as the agent to communicate with the virtual desktop by using the ICA/HDX protocol.

- ▶ Hypervisor

XenDesktop has an open architecture that supports the use of XenServer, Microsoft Hyper-V, and VMware ESX or vSphere hypervisors. VMware vSphere 5.x and Microsoft Hyper-V are covered in the Lenovo Reference Architecture for Citrix XenDesktop.

- ▶ Citrix XenApp

Citrix XenApp allows most Windows applications to be instantly delivered as a service to users anywhere on any device. It can be used to deliver virtualized applications and virtualized desktops. In the Hosted VDI model, XenApp often is used for on-demand access to streamed and hosted applications.

- ▶ Shared storage

Shared storage is used to store user profiles and user data files. Depending on the provisioning model that is used, different data is stored for VM images. Shared storage also holds the redirected vSwap files.

5.2 Desktop and application delivery

Citrix XenDesktop offers FlexCast delivery technology that provides flexible desktop and application delivery that ranges from hosted shared desktops and applications to hosted virtual desktops by using *published*, *installed*, or *streamed* deployment models.

The choice of specific delivery model or a combination of the delivery models depends on user and application compatibility and customization requirements, as shown in Table 5-1 on page 71.

Note: The terms Low, Medium, and High that are used in Table 5-1 are relative indicators for comparison purposes and do not represent any meaning in terms of absolute values. For example, values in the Relative user density row mean that non-persistent desktops have better user density than persistent desktops; hosted applications have better user density than non-persistent desktops.

Table 5-1 Application and desktop delivery model comparison

Feature or requirement	Hosted virtual desktops		Hosted applications	
	Persistent	Non-persistent	Published	Streamed
Application compatibility with desktop OS	Yes	Yes		
Application compatibility with server OS			Yes	Yes
User customization	Yes	Yes		
Application customization	Yes			
Standard application installer	Yes	Yes	Yes	Yes
Custom application installer	Yes	Yes	Yes	
Multi-user aware application	Yes	Yes	Yes	Yes
Single-user application	Yes	Yes		
Provisioning model	MCS	MCS or PVS	XenApp	XenApp
Management	Complex	Simplified	Simplified	Simplified
Relative storage IOPS	High	Low	Low	Low
Relative user density	Low	Medium	High	High
Relative cost	High	Medium	Low	Low

If the application can work in a multi-user environment, requires no user customization, and is compatible with the server operating system, the most cost-efficient way to deploy VDI is to use hosted applications with a published or streamed delivery model.

For highly customized user application environments, hosted virtual desktops provide an efficient way to deploy a centralized desktop infrastructure, with a non-persistent model that is cost optimized and a persistent model that is application customization optimized.

5.3 Citrix XenDesktop provisioning

Citrix XenDesktop supports the following primary provisioning models:

- ▶ MCS
- ▶ PVS

MCS is a part of the XenDesktop Studio management console, but is limited to hosted virtual desktops only (pooled or dedicated). Organizations that want use a streamed VHD model must use PVS. However, PVS requires a separate server and potentially multiple servers within the infrastructure.

Another consideration is the requirement for dedicated private desktops. By using *private desktops*, users can control their virtual desktops. With private desktops, the initial delivery of the desktop is identical. After it is deployed, each desktop becomes unique as changes persist across reboots. Within the hosted VDI desktop FlexCast model, this level of personalization can be achieved with installed images, MCS images, and PVS images.

By using built-in technology to provide each desktop with a unique identity, MCS thin-provisions each desktop from a master image. Only changes that are made to the desktop use more disk space. PVS also uses built-in technology to provide each desktop with a unique identity, but it uses a complete copy of the base desktop image in read/write mode. Each copy uses disk space that expands as the user adds items to the desktop image.

When dedicated desktops are required, most organizations use MCS images. Most organizations use PVS for pooled desktop configurations because PVS requires fewer IOPS and offers faster patching and image updates. A single XenDesktop environment can host any mix of PVS and MCS desktops that an organization needs to meet its design goals.

Tip: When managing a VDI farm, pooled VDI desktops provide better total cost of ownership (TCO) and reduced administrative overhead. Although most organizations typically require a few dedicated desktops, it is better to limit their use when possible.

For more information about choosing the appropriate image delivery option, see “XenDesktop Planning Guide: Desktop Image Delivery”, which is available at this website:

<http://support.citrix.com/article/CTX128643>

For more information about PVS, see 5.3.1, “Provisioning Services solution” on page 72. For more information about MCS, see 5.3.2, “Machine Creation Services” on page 74.

5.3.1 Provisioning Services solution

Hosted VDI desktops can be deployed with or without Citrix PVS. The advantage of the use of PVS is that you can stream a single desktop image to create multiple virtual desktops on one or more servers in a data center.

Figure 5-2 shows the sequence of operations that are run by XenDesktop to deliver a Hosted VDI virtual desktop to the user.

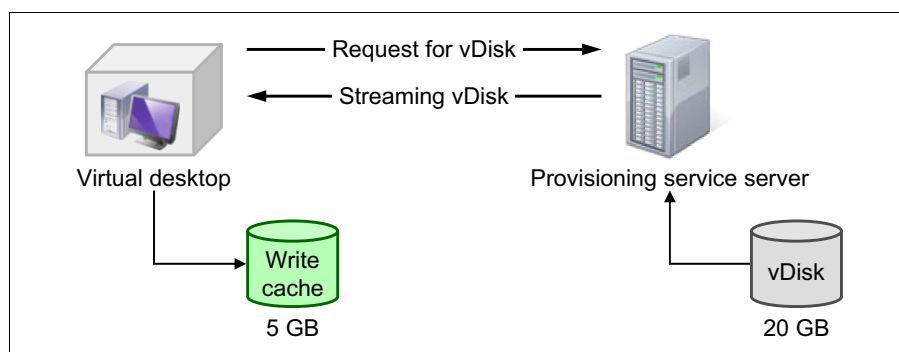


Figure 5-2 PVS provisioning steps

When PVS is used, the administrator performs the following steps:

1. Prepare a master target device to be imaged by installing an operating system and software.

2. Create a virtual disk (vDisk) image from the hard disk drive (HDD) of the master target device and save it to the PVS server.

The PVS server streams vDisk contents to the target device on demand in real time by using software-streaming technology.

After the vDisk image is available from the network, the VM on a target device no longer needs its local HDD to operate; instead, it boots directly from the network and behaves as though it is running from a local drive on the target device. For this reason, PVS is suggested for stateless virtual desktops. PVS is not used for dedicated virtual desktops because the write cache is not stored on shared storage.

PVS is also used with Microsoft Roaming Profiles (MSRPs) so that the user's profile information can be separated out and reused. Profile data is available from shared storage.

For more information about PVS streaming configuration, see 9.3.1, "Configuring streamed desktops" on page 198.

Write cache options

PVS supports several write cache destination options. The write cache destination for a vDisk is selected on the General tab on the vDisk File Properties dialog.

The following write cache destinations are valid:

- **Cache on Device Hard Drive**

The write cache can exist as a file in New Technology File System (NTFS) format that is on the target device's HDD. This write cache option frees up the PVS server because it does not have to process write requests and does not have the finite limitation of RAM.

Note: The write cache file is temporary unless the vDisk mode is set to Difference Disk Image mode.

- **Cache in Device RAM**

In this case, the write cache exists as a temporary file in the target device's RAM. This configuration provides the fastest method of disk access because memory access is always faster than hard disk access.

- **Cache on a Server**

The write cache can exist as a temporary file on a PVS server, which can increase disk I/O and network traffic.

For extra security, the PVS server can be configured to encrypt write cache files. Because the write cache file is on the HDD between reboots, the data is encrypted if the HDD is stolen.

- **Cache on Server Persistent**

This cache option allows for the saving of changes between reboots. By using this option, a target device after rebooting can retrieve changes that were made from previous sessions that differ from the read-only vDisk image. If a vDisk is set to Cache on Server Persistent, each target device that accesses the vDisk automatically has a device-specific, writable disk file created. Any changes that are made to the vDisk image are written to that file, which is not automatically deleted upon shutdown.

The drawback of the use of this cache option is that the cache file is available only while the file remains valid. Any changes that are made to the vDisk force the cache file to be marked invalid. For example, if the vDisk is set to Private Image Mode, all associated cache files are marked invalid.

Write cache sizing

The size of the cache file for each VM depends on several factors, including the types of applications, user workloads, and reboot frequency. A general estimate is 300 MB - 500 MB for the cache size of a provisioned workstation. If a workstation is not rebooted often, or uses applications that are virtualized by using Microsoft App-V or similar programs, cache size can grow much larger.

Because application workloads can vary for each environment, perform a detailed analysis to determine expected cache file sizes for your environment.

For more information about write cache sizing, see this website:

<http://blogs.citrix.com/2011/10/06/pvs-write-cache-sizing-considerations/>

Communication Ports

The following User Datagram Protocol (UDP) ports often are defined:

- ▶ PVS server to target device communication
Each PVS server must be configured to use the same UDP ports to communicate with target devices (by using the StreamProcess). The port range is configured by using the Console's Network tab on the Server Properties dialog. Default ports are UDP ports 6910 - 6930.
- ▶ Login server communication
Each PVS server that is used as a login server must be configured on the Stream Servers Boot List dialog when the administrator runs the configuration wizard. The default port for login servers is UDP 6910.

For more information about the best practices for PVS sizing and configuration, see the *Provisioning Services 5.x and 6.x Best Practices* article that is available at this website:

<http://support.citrix.com/article/ctx127549>

5.3.2 Machine Creation Services

Unlike PVS, MCS does not require extra servers. Instead, it uses integrated functionality in XenDesktop Studio and communicates through the respective APIs with VMware vSphere or MS SCVMM. Each desktop has one difference disk and one identity disk, as shown in Figure 5-3 on page 75.

The *difference disk* is used to capture any changes that are made to the master image. The *identity disk* is used to store information, such as machine name and password.

The following types of image assignment models are available for MCS:

- ▶ Pooled-Random
Desktops are assigned randomly. When they log off, the desktop is free for another user. When rebooted, any changes that are made are destroyed.

- Pooled-Static

Desktops are permanently assigned to a single user. When a user logs off, only that user can use the desktop, regardless of whether the desktop is rebooted. During reboots, any changes that are made are destroyed.

- Dedicated

Desktops are permanently assigned to a single user. When a user logs off, only that user can use the desktop, regardless of whether the desktop is rebooted. During reboots, any changes that are made persist across subsequent restarts.

Figure 5-3 shows MCS provisioning.

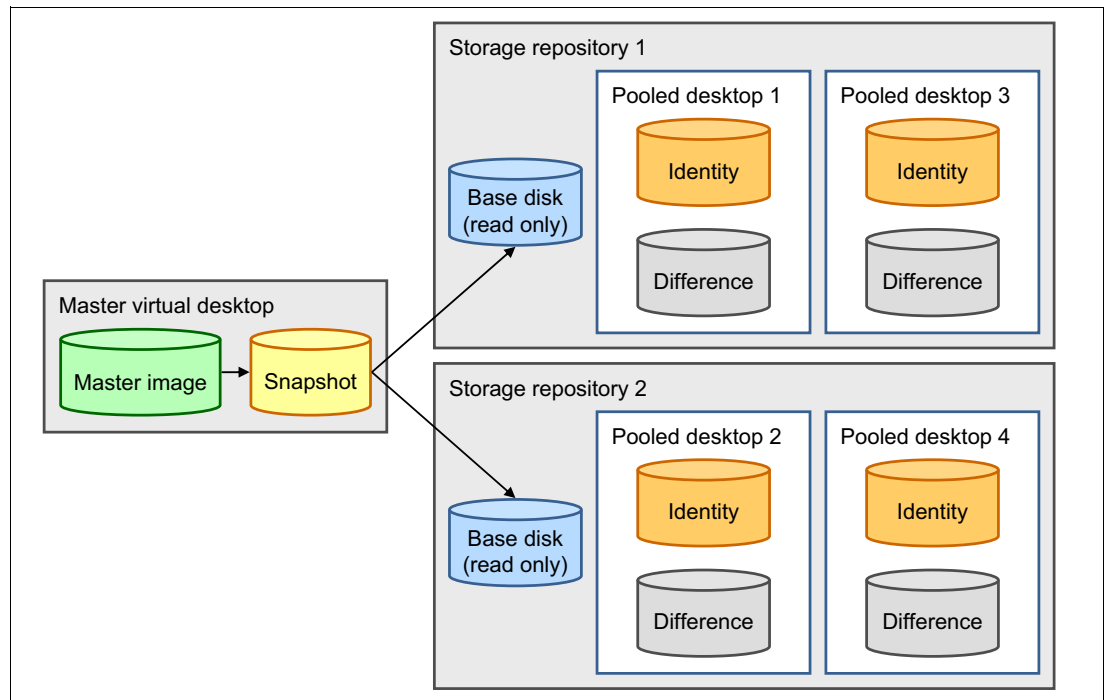


Figure 5-3 MCS provisioning

For more information about MCS, see this website:

<http://bit.ly/lacsLuK>

5.3.3 Personal vDisk

The personal vDisk feature in XenDesktop provides single image management for the administrator. At the same time, it provides users with complete personalization.

Personal vDisk technology enables the single-image management of pooled and streamed desktops while enabling users to install applications and change desktop settings as they do in a dedicated user-to-image model. In a traditional VDI deployment with pooled desktops, users lose customizations and personal applications when an administrator alters the base virtual machine (VM).

In contrast, XenDesktop deployments that use personal vDisks can retain those changes across reboots and base image updates. Therefore, administrators can easily and centrally manage base VMs while providing users with a customized and personalized desktop experience.

Personal vDisks provide this separation by redirecting all changes that are made on the user's VM to a separate disk (the personal vDisk) that is attached to the user's VM. The content of the personal vDisk is blended at run time with the content from the base VM to provide a unified experience. In this way, users can still access applications that were provisioned by their administrator in the base VM.

This user's specific Virtual Hard Disk file (a .vhd or .vmdk file) contains all of the user's customizations, such as applications that are installed in the C:\Program Files directory. Physically, a personal vDisk does not need to be stored on the same storage with the base VM but can be on other data stores.

Personal vDisk provides the best option, which is single image management with complete user personalization and customization. For more information about whether personal vDisks are the correct approach for your environment, see the following documents:

- ▶ *Citrix Personal vDisk Technology Planning Guide:*
<http://support.citrix.com/article/CTX133227>
- ▶ *Personal vDisk FAQs:*
<http://support.citrix.com/article/CTX131553>

5.3.4 Image assignment models

In this section, we describe the following primary user-to-image assignment models:

- ▶ Non-persistent (preferred)
- ▶ Persistent (if needed)

Citrix has various image delivery options with which you can achieve either of the two assignment models that use PVS and MCS. In reality, they can be combined with various profile management options and the Personal vDisk feature. By using the Personal vDisk feature, you can simulate a persistent user experience (for example, installing personal applications) even when non-persistent images are used, which provides a “simulated” or “hybrid” model.

In this book, we describe the following image assignment scenarios:

- ▶ Non-persistent: Streamed PVS image with Microsoft Roaming Profiles (MSRP)
- ▶ Persistent: Pooled image that is delivered through Machine Creation Services with MSRP

5.4 Storage configuration

The architecture assumes that all hypervisor data stores are hosted on a supported shared storage that uses one of the storage protocols: Fibre Channel, Fibre Channel over Ethernet (FCoE), Internet Small Computer System Interface (iSCSI), SMB v3 in the case of Hyper-V, or Network File System (NFS) in the case of VMware.

The following sections describe the required storage-related components and configurations for each model.

Non-persistent model (PVS)

The following local storage components make up the PVS model:

- ▶ Hypervisor

For VMware, each compute node is running the ESXi custom image that is uploaded from an internal USB key.

For Hyper-V, the OS must be installed locally on the HDD. Because of the highly dense nature of Flex System compute nodes, this OS must use the available drive slots and the stateless virtual desktops must be stored on shared storage. The preferred practice is to use shared SSD-based storage for this purpose.

- ▶ Local storage for each compute node

For VMware, two local SSD disks are configured in an RAID-0 configuration to store the PVS Write Cache (delta files).

For Hyper-V, two local HDDs are configured in a RAID 1 to store the OS and two local SSD disks are configured as for VMware. Note the restriction for Flex System compute nodes that is described in the previous bullet.

Because of the stateless nature of the architecture, there is little added value in configuring reliable SSD drives in more redundant RAID configurations. Redundancy is not achieved on a host level, but achieved inherently through the ability of a user to connect to virtual desktops that are hosted on any of the surviving nodes if there is an individual node failure.

The following shared storage components are valid for the PVS model:

- ▶ Data stores

For VMware PVS-delivered stateless virtual desktops, shared storage is used to host only the redirected vswap files for the hosts. Any similar Hyper-V based desktops might need SSD-based shared storage if Flex nodes are used.

- ▶ User profiles (if Roaming Profiles is used):

Common Internet File System (CIFS): A CIFS-based file share to host the user profile data (for Microsoft Roaming Profiles) is required.

- ▶ User data on network drives

CIFS: Specifically for the stateless user model, it is essential to redirect persistent user data (documents, other file repositories, and so on) to user-specific file shares (CIFS-based) or network drives. The detailed designs of these network shares (for example, aspects of redundancy and performance) is not within the scope of this book.

Persistent model (MCS)

The following local storage components make up the persistent model:

- ▶ Hypervisor

Each compute node is running a hypervisor that is booted locally from an HDD or an internal USB memory key.

- ▶ Local storage for each compute node

For dedicated hosts, no local storage is configured for use by the desktop.

The following shared storage components are valid for the persistent model:

- Shared storage data stores

Hypervisor data stores are required to host all VM-associated data: MCS-provisioned desktops (including base disks, identity, and difference disks), the MCS Master Image, and the vswap files.

- User profiles (for example, if Roaming Profiles are used)

CIFS: The user profiles are typically hosted on a CIFS-based file share. There are many ways to manage user profiles from native Microsoft profile management over Citrix's profile management solution to third-party solutions. Our environment assumes the use of Microsoft Roaming Profiles.

- User data and network drives

CIFS-based file shares are primarily used for the stateless user model to redirect persistent user data (documents, other file repositories, and so on) to user-specific file shares or network drives. However, they can also be used to complement the dedicated model.

Figure 5-4 shows the required storage tiers. It represents a XenDesktop hybrid environment that consists of a non-persistent PVS-delivered model and a dedicated MCS pool that is connected to the same storage system.

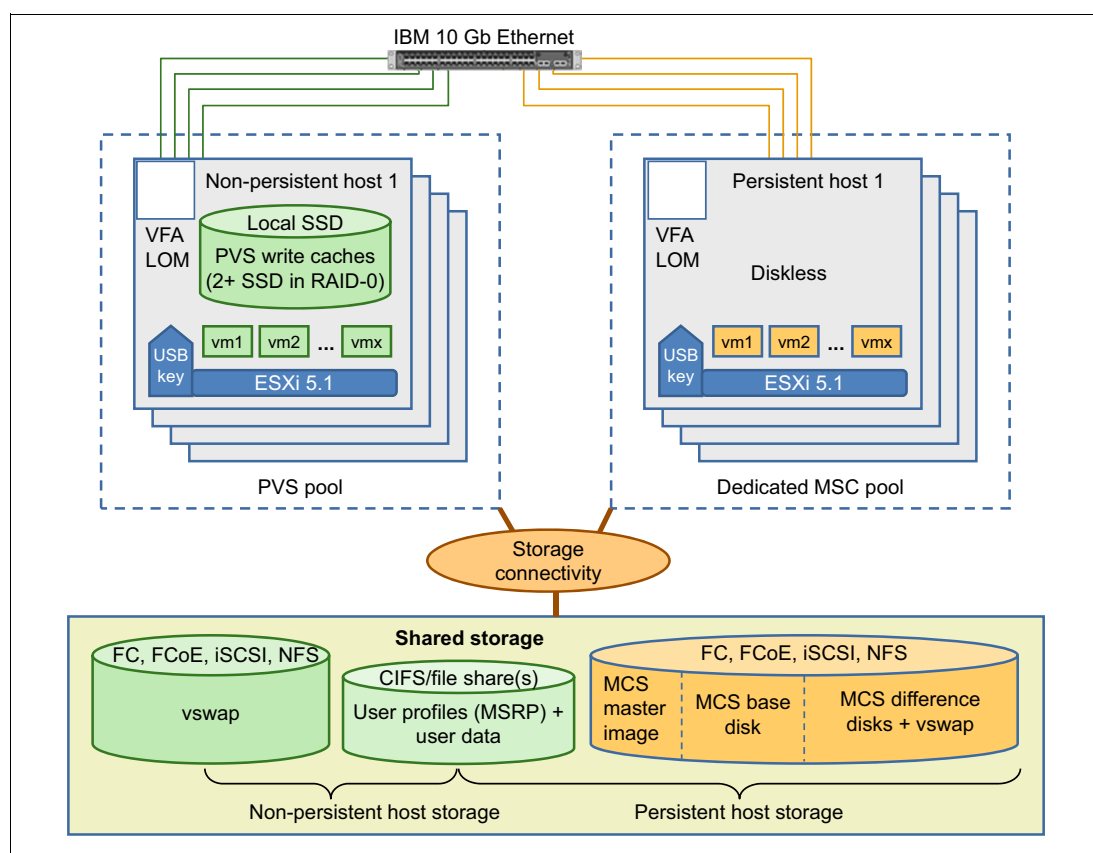


Figure 5-4 Storage layout for Citrix XenDesktop (VMware example): Persistent versus non-persistent

The vDisk that is used to stream the images to the individual targets is on the PVS server on local or shared storage. In our environment, we assume that the PVS server is in a VM that is hosted on the dedicated VDI management cluster (typically on shared storage for HA purposes).

5.5 Network configuration

A redundant 10 Gb Ethernet network infrastructure is used to provide the network connectivity between all components of the Citrix XenDesktop architecture, including storage.

The following VLANs are commonly deployed:

- ▶ Storage VLAN to provide storage connectivity unless FC is the only storage protocol.
- ▶ VM data VLAN for production (user) access and PVS image streaming (if there is no separate VLAN for PVS image streaming)
- ▶ Management VLAN for dedicated access to the management interface of systems
- ▶ Dedicated PVS VLAN for desktop image streaming in highly scalable deployments
- ▶ VM control traffic VLAN for inter-VM communications, such as vMotion or live migration between Hyper-V hosts.

On the server side, the minimum network provision is provided by a single dual-port 10 GbE Virtual Fabric LAN-on-motherboard (LOM) or Network Interface Card (NIC). Each physical 10 Gbps port can be divided into four virtual ports with bandwidth allocation in 100 Mbps increments to the maximum 10 Gbps per physical port.

Note: The actual VLAN configuration and bandwidth allocation depends on your individual requirements; ensure that you have adequate bandwidth available for each traffic type.

Figure 5-5 shows logical network separation.

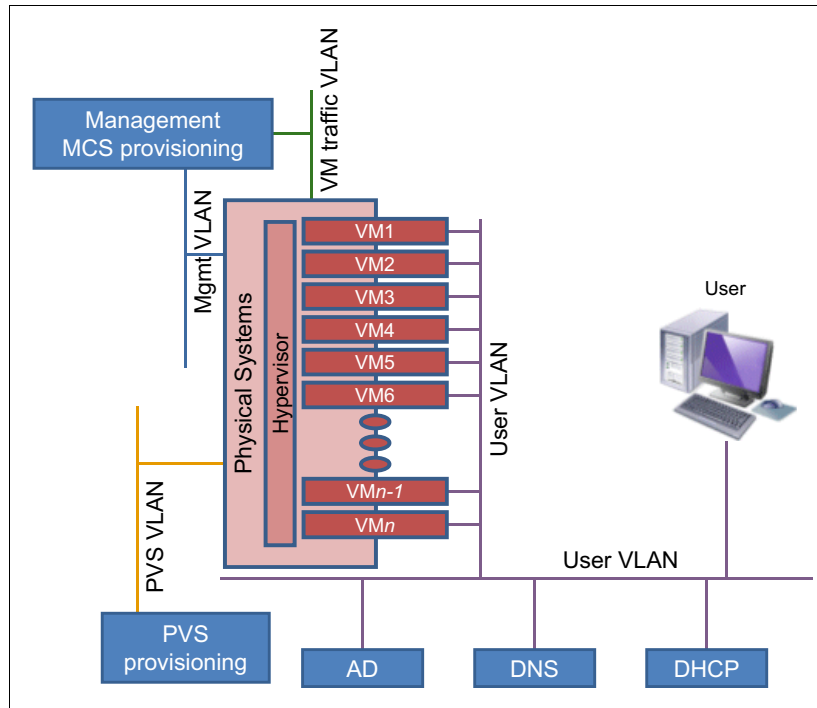


Figure 5-5 Logical network separation

Consider the following bandwidth allocation requirements as a starting point:

- ▶ Management: 0.5 Gbps
- ▶ VM control traffic: 1 Gbps

- ▶ VM data: 1 - 2 Gbps
- ▶ PVS image streaming (if separate VLAN): 1 Gbps
- ▶ Storage (if used): 1-2 Gbps

5.6 Operational model and sizing guidelines

Two separate main operational models are described in this book to cover non-persistent and persistent image models. In some client environments, non-persistent and persistent image models might be required; therefore, a mixed operational model is required. To show the operational model for different client's environments and size needs, four configurations are described to support 600, 1,500, 4,500, and 10,000 users. Because the operational model for 10,000 users is roughly seven times larger than the model for 1,500 users, you can estimate the needs for intermediate numbers of users by using different multiples of the 1500-user model.

5.6.1 VDI server configuration

The VDI server is the base system unit that makes up the compute clusters. The compute clusters can consist of any System x or Flex System servers that are listed in 2.2, "VDI servers" on page 11.

VDI servers run the VMware ESXi or Microsoft Hyper-V hypervisor and host Citrix XenDesktop user VMs. For non-persistent users, the typical range of memory that is required for each desktop VM is 1.5 GB - 4 GB. For persistent users, the typical range of memory for each desktop VM is 2 GB - 6 GB. High-end computer-aided design (CAD) users that need 3D VDI technology might require 8 GB - 16 GB of memory per desktop. In general, power users that require larger memory sizes also require more virtual processors. The virtual desktop memory must be large enough so that swapping is not needed and vSwap can be disabled.

As a part of the Reference Architecture validation, System x servers with Intel Xeon processor E5-2400 and E5-2600 v2 product family that are running VMs with different memory sizes of 1.5 GB, 2 GB, and 3 GB were tested¹. The results are listed in Table 5-2.

Table 5-2 Number of virtual desktops per compute node

Feature (per node)	VM memory size		
	1.5 GB	2 GB	3 GB
x222 compute nodes with Intel Xeon E5-2400 processors			
Processor	2x E5-2470	2x E5-2470	2x E5-2470
System memory	384 GB (2 x 192 GB)	384 GB (2 x 192 GB)	384 GB (2 x 192 GB)
Desktop VMs	204 (2 x 102)	156 (2 x 78)	104 (2 x 52)
Desktop VMs (failover)	250 (2 x 125)	188 (2 x 94)	126 (2 x 63)
System x servers and Flex System compute nodes with Intel Xeon E5-2600 v2 processors			
Processor	2x E5-2650 v2	2x E5-2650 v2	2x E5-2690 v2

¹ Lenovo Client Virtualization Reference Architecture for Citrix XenDesktop: <http://lenovopress.com/tips1278>

Feature (per node)	VM memory size		
	1.5 GB	2 GB	3 GB
System memory	256 GB	256 GB	384 GB
Desktop VMs	125	105	105
Desktop VMs (failover)	150	126	126

If a server goes down, users on that server must be transferred to the remaining servers. For the degraded failover case, it is typical to keep 25% headroom on servers to cope with possible failover scenarios.

VDI hosts can be used in the following sample configurations:

► Non-persistent host:

- Processor: Dual socket (8-core Intel Xeon processor E5-2650 v2, E5-2690 v2, or E5-2470)
- Memory: 256 GB (16 x 16 GB) or 384 GB (24 x 16 GB)
- Disks: SATA HS HDDs (to install Microsoft Hyper-V, if used) and MLC HS SSDs (to store VMs files)

Consideration: If you use a Flex System-based, non-persistent VDI environment, VM files should be placed onto external shared storage.

- Hypervisor: USB Memory Key for VMware ESXi (with ESXi Custom Image) or Microsoft Hyper-V that is installed on the local SATA HDDs that are configured in a RAID 1 array
- Network adapter: Dual Port 10 GbE Virtual Fabric Adapter

► Persistent host:

- Processor: Dual socket (8-core Intel Xeon processor E5-2650 v2, E5-2690 v2 or E5-2470)
- Memory: 256 GB (16 x 16 GB) or 384 GB (24 x 16 GB)
- Disks: SATA HDDs (to install Microsoft Hyper-V if used); no drives needed if VMware ESXi USB memory key is used
- Hypervisor: USB Memory Key for VMware ESXi (with ESXi Custom Image) or Microsoft Hyper-V installed on the local SATA HDDs configured in a RAID 1 array
- Network adapter: Dual Port 10 GbE Virtual Fabric Adapter

If you intend to use the host in a dedicated user model (MCS in “dedicated” mode), you can remove the local SSD drives because all data other than the hypervisor is on shared external storage.

Table 5-3 on page 82, Table 5-4 on page 82, and Table 5-5 on page 82 show the number of compute nodes that are needed for each user size (based on desktop VM quantity per server from Table 5-2 on page 80 for different VM sizes).

Table 5-3 Compute nodes needed for VM size of 1.5 GB

Description	600 users	1,500 users	4,500 users	10,000 users
x222 compute nodes with Intel Xeon E5-2400 processors				
Compute nodes @ 204 users	4	8	22	49
Compute nodes @ 250 users (failover)	3	6	18	40
Failover ratio	3:1	3:1	4.5:1	4.5:1
System x servers with Intel Xeon E5-2600 v2 processors				
Compute nodes @ 125 users	5	12	36	80
Compute nodes @ 150 users (failover)	4	10	30	68
Failover ratio	4:1	5:1	5:1	7:1

Table 5-4 Compute nodes needed for different numbers of users for VM size of 2 GB

Description (VM size of 2 GB)	600 users	1,500 users	4,500 users	10,000 users
x222 compute nodes with Intel Xeon E5-2400 processors				
Compute nodes @ 156 users	5	10	30	65
Compute nodes @ 188 users (failover)	4	8	24	54
Failover ratio	4:1	4:1	4:1	5:1
System x servers with Intel Xeon E5-2600 v2 processors				
Compute nodes @ 105 users	6	14	42	96
Compute nodes @ 126 users (failover)	5	12	36	80
Failover ratio	5:1	6:1	6:1	5:1

Table 5-5 Compute nodes needed for different numbers of users for VM size of 3 GB

Description (VM size of 3 GB)	600 users	1,500 users	4,500 users	10,000 users
x222 compute node (dual-server; each server with 2x E5-2470 processors)				
Compute nodes @ 104 users	6	15	45	96
Compute nodes @ 126 users (failover)	5	12	36	80
Failover ratio	5:1	4:1	4:1	5:1
System x server with 2x E5-2650 v2 processors				
Compute nodes @ 105 users	6	14	42	96
Compute nodes @ 126 users (failover)	5	12	36	80
Failover ratio	5:1	6:1	6:1	5:1

Management services

A typical Citrix XenDesktop environment requires several management components. It is suggested that you install the management components on a separate management environment (for example, on a virtual management cluster instance). However, to separate desktop and server workloads for organizational, licensing, and workload attribute reasons, install management components on a cluster other than the one that is used for VDI compute nodes.

In practice, a management cluster can be built on a vSphere environment with spare capacity or you can use more Lenovo systems to create a management cluster with the hypervisor of your choice. For larger scale implementations, it makes sense to have a separate vCenter instance that is dedicated to the management components of all building blocks.

When Provisioning Services is used, it is suggested to keep the PVS server close to the compute nodes that are running the target VMs to optimize network traffic.

The following example shows the VMs that are required to host the management components that are on the management cluster (the Reference Architecture assumes that you run each of these components in VMs):

- ▶ VM1: Citrix Provisioning Server
- ▶ VM2: Citrix XenDesktop Controller
- ▶ VM3: SQL Server
- ▶ VM4: License Server
- ▶ VM5: Web Interface Server
- ▶ VM6: VMware vCenter Server

Depending on your existing environment, you can also host more infrastructure servers on this cluster for other VMs (such as Active Directory and associated services) or the XenApp Controllers for Application Delivery.

Management servers have the same hardware specification as VDI compute nodes (for more information, see 5.6.1, “VDI server configuration” on page 80) so they can be used interchangeably in a worst-case scenario. The management servers are also hypervisor based, but have management VMs instead of user VMs. Table 5-6 lists the VM requirements and performance characteristics of each management service.

Table 5-6 VM requirements for management services

Management service	Virtual processors	Memory	Storage	Windows OS	HA needed	Performance characteristic
vCenter server	4	4 GB	15 GB	2008 R2	No	Up to 2,000 desktops
vCenter SQL server	4	4 GB	15 GB	2008 R2	Yes	Double the virtual processors and memory for more than 2,500 users
DDC	4	4 GB	15 GB	2008 R2	Yes	5,000 user connections
Web server	4	4 GB	15 GB	2008 R2	Yes	30,000 connections per hour
Licensing server	2	4 GB	15 GB	2008 R2	No	170 licenses per second
XenDesktop SQL server	2	4 GB	15 GB	2008 R2	Yes	Double the virtual processor and memory for more than 2,500 users
PVS server	4	32 GB	40 GB	2008 R2	Yes	Up to 1,000 desktops, memory must be a minimum of 2 GB, plus 1.5 GB per image served

Table 5-7 lists the number of management VMs for each size of users following the HA and performance characteristics that are listed. The number of vCenter servers is half of the number of vCenter clusters that is shown in Table 5-3 on page 82. This difference is the result of the fact that each vCenter server can handle two clusters of up to 1,000 desktop VMs and each cluster is on two vCenter servers.

Table 5-7 Management VMs needed

Management service	600 users	1,500 users	4,500 users	10,000 users
vCenter server	1	1	3	7
vCenter SQL server	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)
XenDesktop SQL server	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)
Web server	N/A	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)
Controller: ► Includes Licensing server ► Includes Web server	2 (1 + 1) Yes Yes	2 (1 + 1) No No	2 (1 + 1) No No	4 (3 + 1) No No
Licensing server	N/A	1	1	1
PVS server	2 (1 + 1)	4 (2 + 2)	8 (6 + 2)	14 (10 + 4)

It is assumed that common services, such as Microsoft Active Directory, Dynamic Host Configuration Protocol (DHCP), DNS server, and Microsoft licensing servers, are in the client's environment.

Based on the number and type of VMs, Table 5-8 lists the appropriate number of physical management servers. In all cases, there is redundancy in the management servers and the management VMs.

Table 5-8 Physical management servers needed

Delivery model	600 users	1,500 users	4,500 users	10,000 users
Persistent	2	2	2	4
Non-persistent	2	2	4	7

5.6.2 Shared storage

VDI workloads, such as virtual desktop provisioning, VM loading across the network, and access to user profiles and data files, place huge demands on network shared storage. In this book, we describe the performance requirements of non-persistent and persistent virtual desktops and then show the storage configuration that meets those requirements.

Experimentation with VDI infrastructures shows that the input/output operation (IOP) performance takes precedence over storage capacity. This precedence means that more of the slower speed drives are needed to get the required performance than higher speed drives. Even with the fastest drives that are available as of this writing (15,000 rpm), there still can be an excess capacity in the storage system.

The large rate of IOPs (and therefore large number of drives that are needed for dedicated virtual desktops) can be ameliorated to some extent by implementing SSD storage combined with Easy Tier functionality in the V7000 storage systems.

The storage configurations are based on the peak performance requirement, which usually occurs during a so-called “login storm”. This state occurs when all workers at a company arrive at the same time in the morning and try to start their virtual desktops at the same time. The storage configurations that are described in this section have conservative assumptions about the VM size, changes to the VM, and user data sizes to ensure that the configurations can cope with the most demanding user scenarios.

The storage configurations tend to have more storage than is strictly required to meet the performance objectives for IOPs. In our experience, this “extra” storage is more than sufficient for the other types of data that is needed for VDI, such as SQL databases and transaction logs.

The storage configurations do not include facilities for data replication, data compression, or data deduplication. Although these features might not be required, they can affect the storage configuration. The storage configurations, where possible, include flash memory as a means to cache frequently used data.

Non-persistent virtual desktops

Non-persistent virtual desktops that use Citrix XenDesktop are provisioned from shared storage by using PVS. The PVS write cache is maintained on a local SSD. Table 5-9 lists the peak IOPs and shared disk space requirements for stateless virtual desktops on a per-user basis.

Table 5-9 Shared storage performance requirements: Non-persistent VHD

Data type	Protocol	Size	IOPS	% write
User data files	CIFS or NFS	5 GB	1	75%
User profiles (through MSRP)	CIFS	100 MB	0.8	75%

Persistent virtual desktops

Table 5-10 lists the peak IOPs and disk space requirements for persistent virtual desktops on a per-user basis. The last two rows are the same as used for non-persistent desktops.

Table 5-10 Shared storage performance requirements: Persistent VHD

Data type	Protocol	Size	IOPS	Percentage write
Master image	FC, FCoE, iSCSI, NFS	30 GB	18	85%
Difference disks		10 GB		
User “AppData” folder				
User data files	CIFS	5 GB	1	75%
User profiles (MSRP)	CIFS	100 MB	0.8	75%

The sizes and IOPS for user data files and user profiles that are listed in Table 5-9 and Table 5-10 can vary depending on the client’s environment. For example, power users might require 10 GB and 5 IOPS for user files because of the applications they use. It is assumed that 100% of the users at peak load times require concurrent access to user data files and profiles.

Storage capacity estimation

For our example, we assume that each user has 5 GB for shared folders and profile data and uses an average of 2 IOPS to access those files. Reviewing the performance shows that 600 GB 10,000 rpm drives in a RAID 10 array give the best ratio of I/O operation performance-to-disk space. We found that 300 GB 15,000 rpm drives have the required performance, but extra drives are needed even when configured as RAID 5. Therefore, it is suggested to use a mixture of drives for persistent desktops, shared folders, and profile data.

If users need more than 5 GB, the 900 GB 10,000 rpm drives can be used instead of 600 GB. If less capacity is needed, the 300 GB 15,000 rpm drives can be used for shared folders and profile data.

Depending on the number of master images, one or more RAID 1 arrays of SSDs can be used to store the VM master images. This configuration helps with the performance of provisioning virtual desktops, which is a “boot storm”. Each master image requires at least double the space. The actual number of SSDs in the array depends on the number and size of images. In general, more users require more images.

Table 5-11 shows an example scenario of calculating storage capacity for VM images.

Table 5-11 Storage capacity for storing VM images

Description	600 users	1,500 users	4,500 users	10,000 users
Image size	30 GB	30 GB	30 GB	30 GB
Number of master images	2	4	8	16
Required disk space (doubled)	120 GB	240 GB	480 GB	960 GB
400GB SSD configuration	RAID 1 (2)	RAID 1 (2)	2 x RAID 1 (4)	4 x RAID 1 (8)

For stateless desktops, the Storwize storage configuration is listed in Table 5-12.

Table 5-12 Storwize configuration for stateless desktops

Stateless desktops	600 users	1,500 users	4,500 users	10,000 users
400 GB SSDs in a RAID 1 for master images	2 (1 x RAID 1)	2 (1 x RAID 1)	4 (2 x RAID 1)	8 (4 x RAID 1)
Hot spare SSDs	2	2	4	4
600 GB 10K rpm HDDs in a RAID 10 for users	12	28	80	168
Hot spare 600 GB HDDs	2	2	4	12
Storwize Control Enclosure	1	1	1	1
Storwize Expansion Enclosure	0	1	3	7

Consideration: The Storwize V3700 can support up to 7,000 stateless users based on the requirements that are listed in Table 5-12.

For persistent desktops or stateless desktops that use shared storage to store VM files, the Storwize storage configuration is listed in Table 5-13.

Table 5-13 Storwize configuration for persistent or stateless desktops

Persistent or stateless desktops (shared storage)	600 users	1,500 users	4,500 users	10,000 users
400 GB SSDs in a RAID 1 for master images	2 (1 x RAID 1)	2 (1 x RAID 1)	4 (2 x RAID 1)	8 (4 x RAID 1)
Hot spare SSDs	2	2	4	8
600 GB 10K rpm HDDs in a RAID 10 for users	12	28	80	168
Hot spare 600 GB HDDs	2	2	4	12
300 GB 15 K rpm HDDs in RAID 10 for persistent desktops	40	104	304	672
Hot spare 300 GB drives	2	4	4	12
400 GB SSDs for Easy Tier	4	12	32	64
Storwize Control Enclosure	1	1	2	4
Storwize Expansion Enclosure	2	6	16 (2 x 8)	36 (4 x 9)

Consideration: The Storwize V3700 can support up to 1,200 persistent or stateless users based on the requirements that are listed in Table 5-13.

It is typical to cluster multiple Storwize V7000 storage systems by using a separate control enclosure for every 2,500 dedicated desktops.

If CIFS or NFS services do not exist, they can be enabled in the VDI environment with Windows Storage Server. In this case, two more physical management nodes are added to the solution, and Windows Storage Server is deployed on them in a highly available cluster.

Part 3

VDI deployment and management

In this part, we provide step-by-step instructions about how to perform specific tasks as a part of the deployment and management process for a Citrix XenDesktop VDI solution on System x servers.

This part includes the following chapters:

- ▶ Chapter 6, “Citrix XenDesktop lab environment” on page 91
- ▶ Chapter 7, “Deploying Flex System” on page 101
- ▶ Chapter 8, “Deploying Citrix XenDesktop” on page 147
- ▶ Chapter 9, “Operating Citrix XenDesktop” on page 179
- ▶ Chapter 10, “Managing System x and Flex System hardware in a VDI environment” on page 243

Citrix XenDesktop lab environment

This chapter describes the structure of the environment and the implementation plan for Citrix XenDesktop in the lab scenario.

The lab setup that is described in this chapter shows the main infrastructure patterns that were applied to the production virtual desktop infrastructure (VDI) environments.

This chapter includes the following topics:

- ▶ 6.1, “Lab environment” on page 92
- ▶ 6.2, “Use case for the lab environment” on page 93
- ▶ 6.3, “Component model” on page 95
- ▶ 6.4, “Operational model” on page 95
- ▶ 6.5, “Logical design” on page 97

6.1 Lab environment

The physical environment and the software components that were used in the implementation of the Citrix XenDesktop landscape are described in this section.

The physical environment consists of the following components:

- ▶ Flex System Enterprise chassis
- ▶ Flex System x240 compute nodes
- ▶ Flex System V7000 that is used as shared storage
- ▶ Flex System Fabric EN4093 10Gb Ethernet Scalable Switch for Ethernet and iSCSI storage connectivity
- ▶ Flex System FC3171 8Gb SAN Switch for FC storage connectivity

Table 6-1 lists the software components that were used in the landscape and their roles.

Table 6-1 Software solution

Software component	Description
Microsoft Hyper-V 2012 R2 or VMware ESXi 5.1	Hypervisor officially that is supported by Lenovo and used to virtualize the XenDesktop landscape.
Microsoft System Center Virtual Machine Manager 2012 R2 or VMware vCenter 5.1	This component manages the hypervisor environment.
Windows 2012 R2	Operating system for the Citrix products and SCCM 2012 R2.
Windows 2008 R2	Operating system for the database server and vCenter
Windows 7	Operating system for virtual desktops.
SQL Server 2008 R2	Database server to store Citrix configuration databases of Provisioning Services, XenApp, and XenDesktop.
Citrix StoreFront 2.5	This component provides users with access to their virtualized desktops on XenDesktop and virtualized applications on XenApp.
Citrix Provisioning Services Version 7.1	This component enables a standardized desktop image to be streamed to all desktops while centralizing the administrative efforts.
Citrix Licensing Server Version 11.11	This component manages the Citrix licenses.
Citrix XenApp Version 7.5	This component virtualizes the application to deliver it integrated with XenDesktop.
Citrix XenDesktop Version 7.5	This component virtualizes the desktops.

6.2 Use case for the lab environment

Upward Integration Modules for VMware vSphere and Microsoft System Center are used to monitor the hardware components in the vSphere and manage the hardware components in Hyper-V based VDI environments.

The VDI is distributed across three different clusters within the same data center. This distribution allows for segmentation of the resource usage and to align with a standard pattern that is deployed in production environments in which each cluster has a specific purpose. The following clusters are available:

- Management cluster

The management cluster concentrates all infrastructure server components, such as Active Directory, database server, and Citrix infrastructure servers.

Note: The provisioning services server has a critical role in the Virtual Desktop environment because this server is responsible for managing the target devices and for streaming the standard image for these desktops. In scenarios where it is necessary to have many different images (for example, with specific requirements for financial areas and industry areas), it is common to have a separate computer node for provisioning to avoid bottlenecks that affect the environment's usage.

- Persistent desktop cluster

This cluster is responsible for processing the persistent desktops. The nodes are separate from the audience for persistent desktops because, in general, persistent desktops are more sensitive than non-persistent desktops and require more aggressive service-level agreements (SLAs).

- Non-persistent desktop cluster

The non-persistent desktops are separated from the persistent desktop cluster because, in general, they do not need high availability (HA) enabled. If a failure occurs at the computer nodes, desktops are restarted during the normal recovery process by using the standard image that is configured at the Provisioning Services.

For management and persistent desktop clusters, consider enabling HA in a production environment.

The storage that is used by the clusters is a shared storage is provided by the Flex System V7000 via Fibre Channel or iSCSI. The following volumes are presented to physical hosts:

- One volume to be used as a data store for the management cluster
- One volume to be used as a data store for the persistent and non-persistent clusters

The network traffic is split on different VLANs and managed by the Flex System Fabric EN4093R 10Gb Ethernet Scalable Switch. It also provides external connectivity for client device connection and, if iSCSI is used, manages the storage flows via iSCSI protocol.

When Fibre Channel (FC) is used, the storage flows are managed by the Flex System FC3171 8Gb SAN Switch. All nodes are on the same storage zone.

Table 6-2 lists the position of the software components on the infrastructure servers when Microsoft Hyper-V is used as the hypervisor.

Table 6-2 Software components that are installed on the servers: Hyper-V environment

Server name	Component installed	Other use
SCCM ^a	<ul style="list-style-type: none"> ▶ Windows 2012 R2 ▶ System Center Operations Manager (SCOM) 2012 R2 ▶ System Center Virtual Machine Manager (SCVMM) 2012 R2 	Microsoft Management Consoles for infrastructure components: DNS, DHCP, AD users and computers, Hyper-V, and so on
vCenter ^b	<ul style="list-style-type: none"> ▶ Windows 2012 R2 ▶ VMware vCenter 	N/A
File Server	<ul style="list-style-type: none"> ▶ Windows 2012 R2 	Host roaming profiles, folder redirection
SQL Server	<ul style="list-style-type: none"> ▶ Windows 2008 R2 ▶ Microsoft SQL Server 2008 R2 	N/A
Domain Controller	<ul style="list-style-type: none"> ▶ Windows 2012 R2 	DNS server
Citrix StoreFront	<ul style="list-style-type: none"> ▶ Windows 2008 R2 ▶ Citrix StoreFront 	N/A
Provisioning Services	<ul style="list-style-type: none"> ▶ Windows 2012 R2 ▶ Citrix Provisioning Services 	Dynamic Host Configuration Protocol (DHCP), Trivial File Transfer Protocol (TFTP), and Preboot Execution Environment (PXE)
License Server	<ul style="list-style-type: none"> ▶ Windows 2012 R2 ▶ Citrix Licensing Server 	N/A
XenDesktop Controller	<ul style="list-style-type: none"> ▶ Windows 2012 R2 ▶ Citrix XenDesktop 	N/A
XenApp	<ul style="list-style-type: none"> ▶ Windows 2012 R2 ▶ Citrix XenApp 	N/A

a. System Center Virtual Machine Manager is used with Hyper-V

b. vCenter is necessary with VMware ESXi

Table 6-3 lists the position of the software components on the infrastructure servers when VMware ESXi is used as the hypervisor.

Table 6-3 Software components that are installed on the servers: vSphere environment

Server name	Component installed	Other use
vCenter	<ul style="list-style-type: none"> ▶ Windows 2008 R2 ▶ VMware vCenter 	N/A
File Server	<ul style="list-style-type: none"> ▶ Windows 2008 R2 	Host roaming profiles, folder redirection
SQL Server	<ul style="list-style-type: none"> ▶ Windows 2008 R2 ▶ SQL Server 2008 R2 	N/A
Domain Controller	<ul style="list-style-type: none"> ▶ Windows 2008 R2 	DNS server
Citrix StoreFront	<ul style="list-style-type: none"> ▶ Windows 2008 R2 ▶ Citrix StoreFront 	N/A
Provisioning Services	<ul style="list-style-type: none"> ▶ Windows 2008 R2 ▶ Citrix Provisioning Services 	Dynamic Host Configuration Protocol (DHCP), Trivial File Transfer Protocol (TFTP), and Preboot Execution Environment (PXE)

Server name	Component installed	Other use
License Server	<ul style="list-style-type: none"> Windows 2008 R2 Citrix Licensing Server 	N/A
XenDesktop Controller	<ul style="list-style-type: none"> Windows 2008 R2 Citrix XenDesktop 	N/A
XenApp	<ul style="list-style-type: none"> Windows 2008 R2 Citrix XenApp 	N/A

6.3 Component model

The component model of the VDI is shown in Figure 6-1.

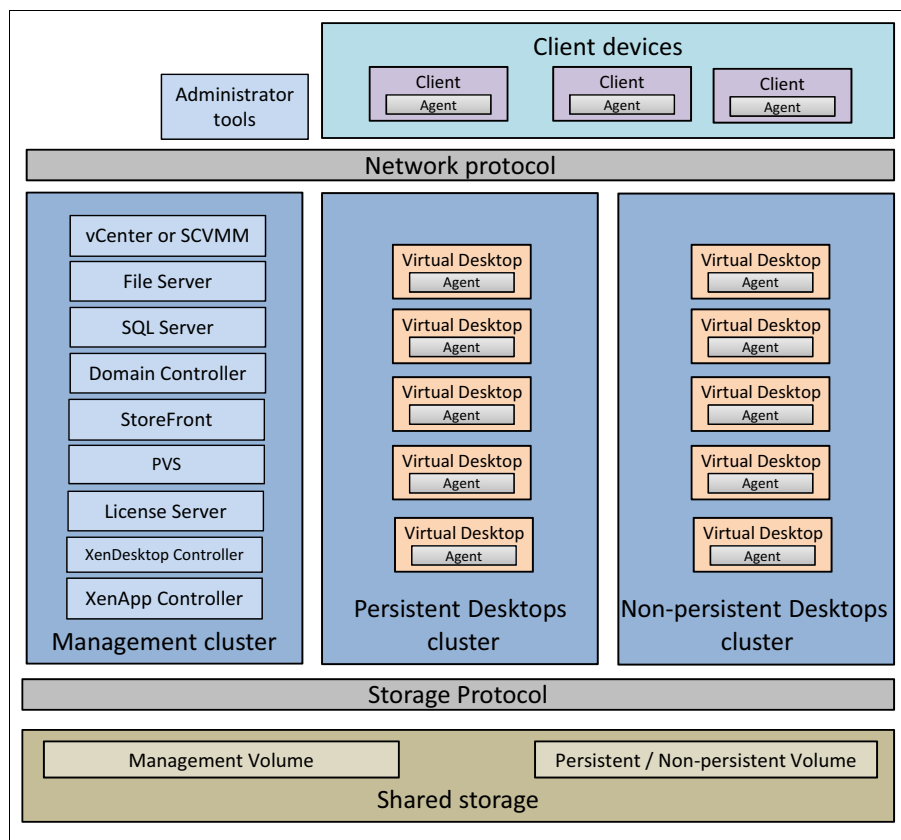


Figure 6-1 Component model

6.4 Operational model

The components that are listed in Table 6-1 on page 92 are installed on virtual machines (VMs) that are running on the compute nodes in the Flex System Enterprise Chassis. It integrates compute nodes, storage, Ethernet switches, and SAN switches in a single machine.

The rear of the chassis shows common management modules (CMMs) and Ethernet and SAN switches, which manage internal communication flows within the chassis and communication flows with external infrastructure components.

Figure 6-2 shows the physical view of the network between components when FC is used.

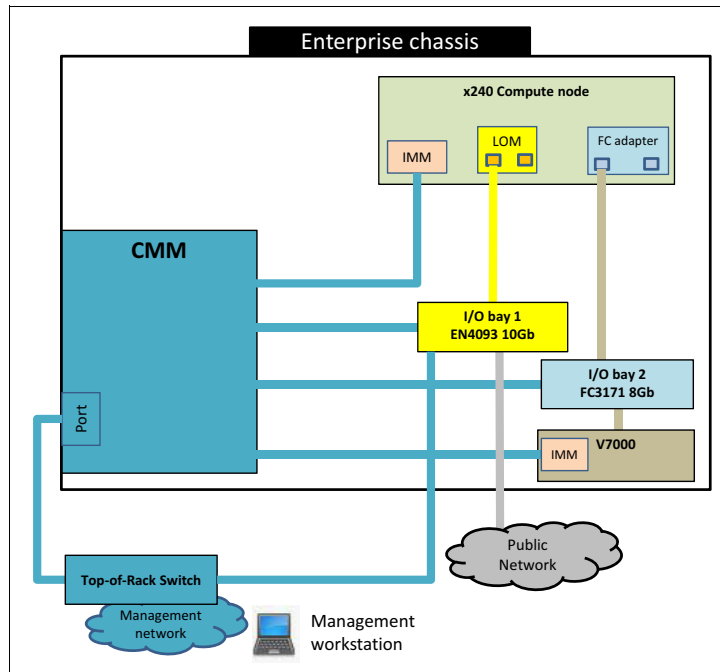


Figure 6-2 Physical view of the network: FC use case

Figure 6-3 shows the physical view of network between components when iSCSI is used.

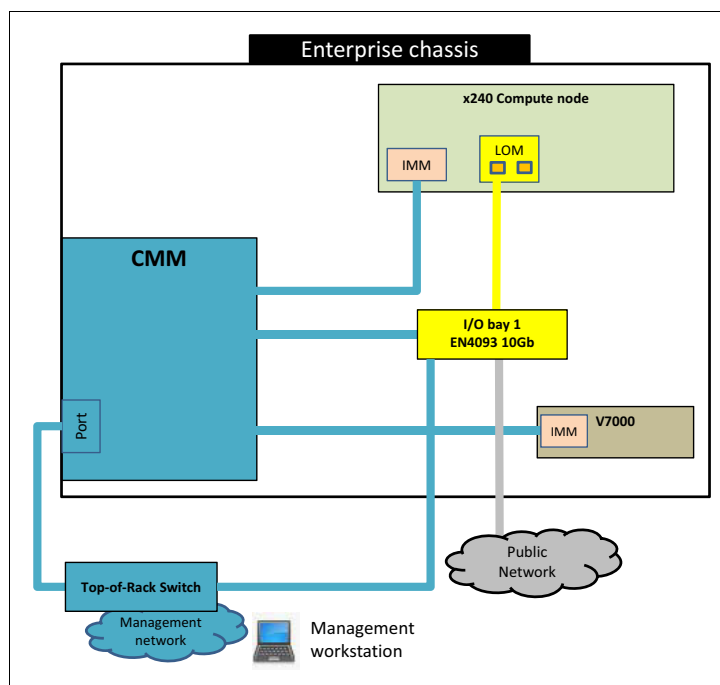


Figure 6-3 Physical view of the network: iSCSI use case

6.5 Logical design

The different network flows (internal and external) between the Flex System components and software components, administrators, and users are described in this section.

Note: The landscape that was created in the lab is for demonstration purposes and it does not cover cluster elements for HA and fault tolerance (FT). You must consider HA and FT in a production environment.

6.5.1 Ethernet segment

The Ethernet segment is configured to split the traffic according to the software component requirements by using the following perspective:

- ▶ Management VLAN (VLAN42)

The management VLAN allows the technical support team to connect to the environment for management purposes. It connects all Flex System components (compute nodes, storage, and switches) and VDI hosts. For security reasons, management traffic is not shared with the user access segment (Public/Access).

- ▶ Live Migration/vMotion VLAN (VLAN10)

When FC is used to access the shared storage, the Live Migration/vMotion virtual LAN (VLAN) is used for Microsoft Hyper-V or VMware ESXi operation. This VLAN is responsible for allowing the VMs to be transferred from one physical node to another in case of maintenance or a hardware failure.

- ▶ iSCSI

When iSCSI is used to access the shared storage, VLAN 10 is dedicated to iSCSI and configured accordingly.

- ▶ Public/Access Network (VLAN20)

This VLAN network segment is used to access the desktops. It is also available for software component communication needs; for example, Active Directory authentication, database access, and other Citrix traffic (with the licensing server, XenApp, and so on).

- ▶ PVS Network VLAN (VLAN30)

This segment was created to isolate the streaming traffic from Provisioning Services to desktops. The traffic segment is a good way to avoid the network conflicts or bottlenecks that can negatively affect the desktop deployment.

Figure 6-4 shows the logical view of the networks between components.

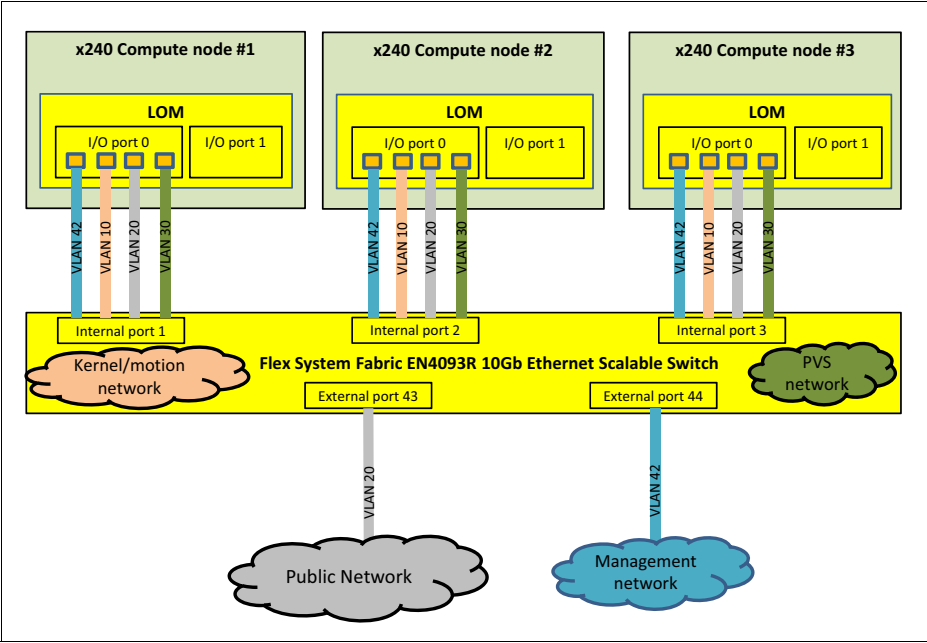


Figure 6-4 Network logical view

Bandwidth allocation is a process step to virtualize the network adapters, where virtual network interface cards (vNICs) are created to be presented to the hosts as traditional adapters that are configured at your own VLAN.

Table 6-5 shows how bandwidth is allocated for each VLAN when FC is used.

Table 6-4 Ethernet adapter bandwidth allocation: FC use case

VLAN	Bandwidth
VLAN42	10%
VLAN10	10%
VLAN20	40%
VLAN30	40%

Table 6-5 shows how bandwidth is allocated for each VLAN/iSCSI when iSCSI is used.

Table 6-5 Ethernet adapter bandwidth allocation: iSCSI use case

VLAN/iSCSI	Bandwidth
VLAN42	10%
VLAN20	15%
VLAN30	25%
iSCSI (VLAN 10)	50%

6.5.2 Storage disk and host mapping

The MDisk is created on Flex System V7000. It is divided into two volumes with Thin Provision preset.

For external storage access, the following volumes are created:

- Management: One volume that is used to store all infrastructure servers and Citrix components.
- Desktops: One volume that is dedicated to store the desktops' write cache disks and persistent vDisks.

Figure 6-5 shows how the volumes are mapped to the hosts.




Figure 6-5 Volume mapping

If we use VMware vSphere, two data stores (formatted with VMware Virtual Machine Filesystem 5) are created from these volumes.

If we use Microsoft Hyper-V, two Cluster Shared Volumes are created from these volumes.

Deploying Flex System

This chapter describes the initial setup and configuration tasks that must be performed on Flex System for virtual desktop infrastructure (VDI) deployment.

This chapter includes the following topics:

- ▶ 7.1, “Initial configuration of the Chassis Management Module” on page 102
- ▶ 7.2, “Firmware updates and basic hardware configuration” on page 113
- ▶ 7.3, “Configuring Active Directory Integration for CMM” on page 121
- ▶ 7.4, “Configuring the EN4093 10Gb Ethernet Switch” on page 122
- ▶ 7.5, “Enabling UFP on the x240 compute node” on page 130
- ▶ 7.6, “Configuring iSCSI on the x240 compute node” on page 131
- ▶ 7.7, “V7000 configuration” on page 132

7.1 Initial configuration of the Chassis Management Module

This section describes how to initially configure the Chassis Management Module (CMM) to enable chassis management tasks.

The following tasks are described:

- ▶ 7.1.1, “Connecting to the Chassis Management Module” on page 102
- ▶ 7.1.2, “Using the initial setup wizard” on page 104
- ▶ 7.1.3, “Configuring IP addresses for the chassis components” on page 112

7.1.1 Connecting to the Chassis Management Module

You can cable the CMM to support a management connection that best matches your site configuration. You must connect a client system to the CMM to configure and manage the operation of the Flex System Enterprise Chassis.

By default, the CMM does not have a fixed static IPv6 IP address. For initial access to the CMM in an IPv6 environment, you can use the IPv4 IP address or the IPv6 link-local address.

By default, the CMM is configured to respond to Dynamic Host Configuration Protocol (DHCP) first before it uses its static IP address.

The HTTP connection is not available when the CMM security policy is set to Secure (which is the manufacturing default setting). When the security policy is set to Secure, Ethernet connections must be made by using HTTPS.

To connect to the CMM, complete the following steps:

1. Ensure that the subnet of the client computer is set to the same value in the CMM (the default CMM subnet is 255.255.255.0). The IP address of the CMM must also be in the same local domain as the client computer. To connect to the CMM for the first time, you might have to change the Internet Protocol properties on the client computer.
2. Open a web browser on the client computer and browse to the CMM IP address. For the first connection to the CMM, use the default IP address of the CMM, as shown in Figure 7-1.




Figure 7-1 Log in to the CMM by using the default IP address

Clarification: The CMM has the following default settings:

- ▶ Subnet: 255.255.255.0
- ▶ User ID: USERID (all capital letters)
- ▶ Password: PASSWORD (note the number zero, not the letter O, in PASSWORD)
- ▶ IP address: 192.168.70.100

3. In the CMM window that is shown in Figure 7-2, log in to the CMM by using the default credentials: USERID/PASSWORD. Click **Log In**.

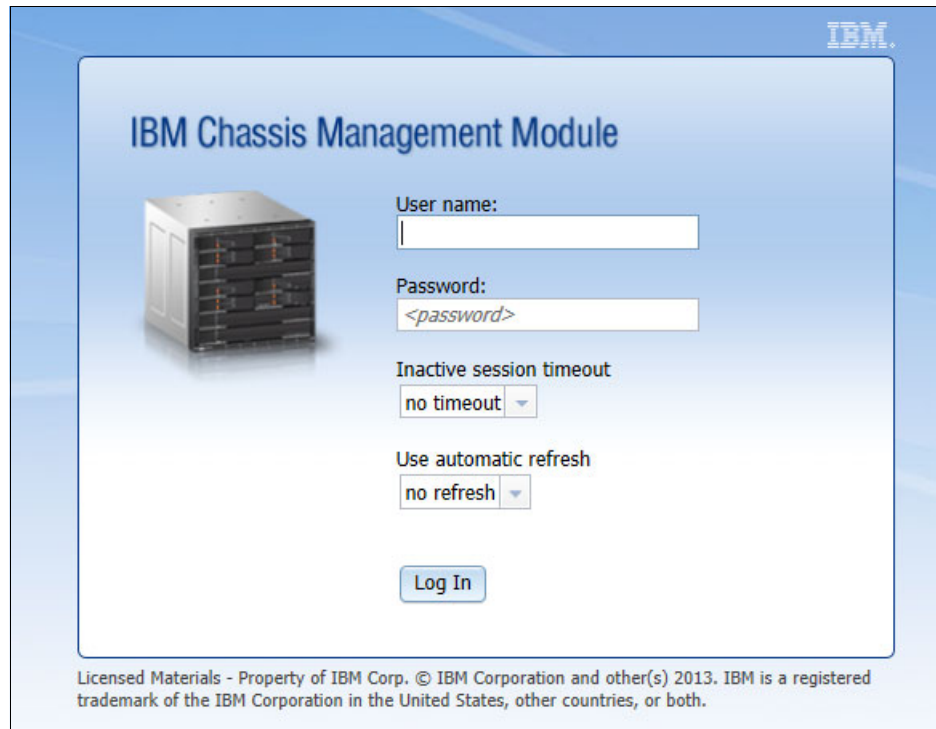


Figure 7-2 CMM login

The CMM main window opens, as shown in Figure 7-3.

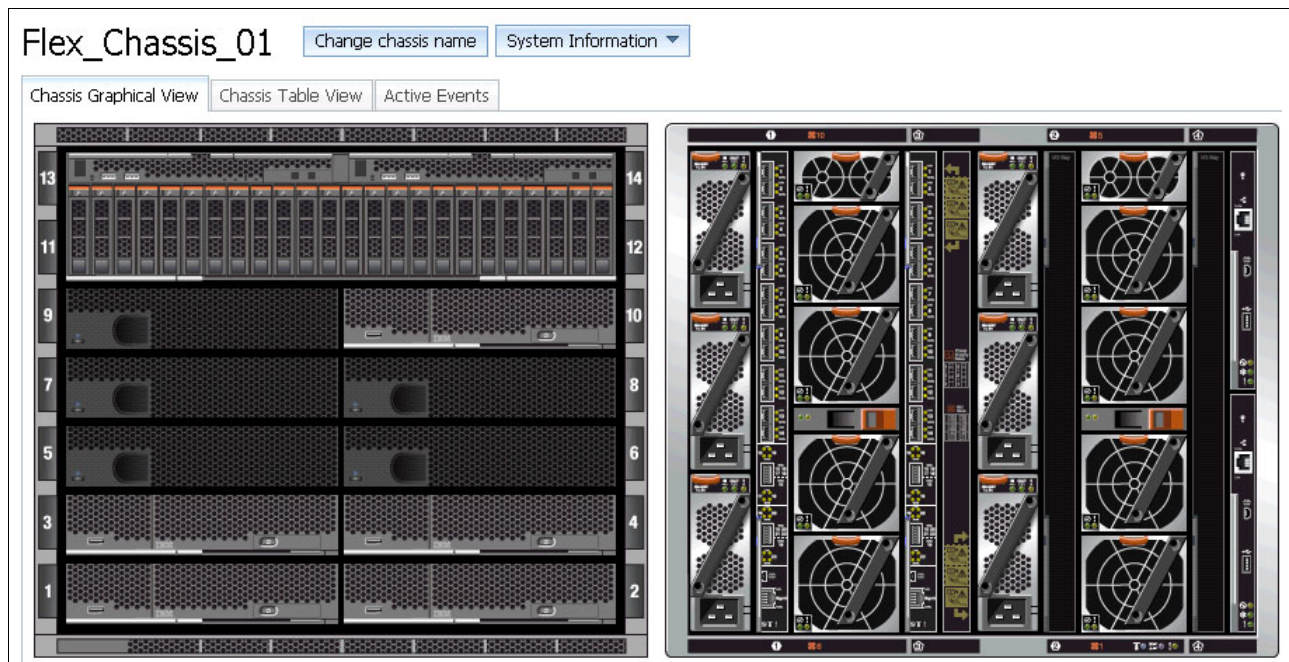


Figure 7-3 CMM main window

7.1.2 Using the initial setup wizard

The next step is the initial configuration of the CMM. The initial setup wizard can help you configure the CMM by using a web interface. The wizard starts automatically when you first access the web interface of a new CMM or a CMM that was reset to its default settings.

Complete the following steps to manually start the initial setup wizard and perform the initial configuration:

1. From the CMM web interface home window, click **Mgt Module Management** → **Configuration**, as shown in Figure 7-4.

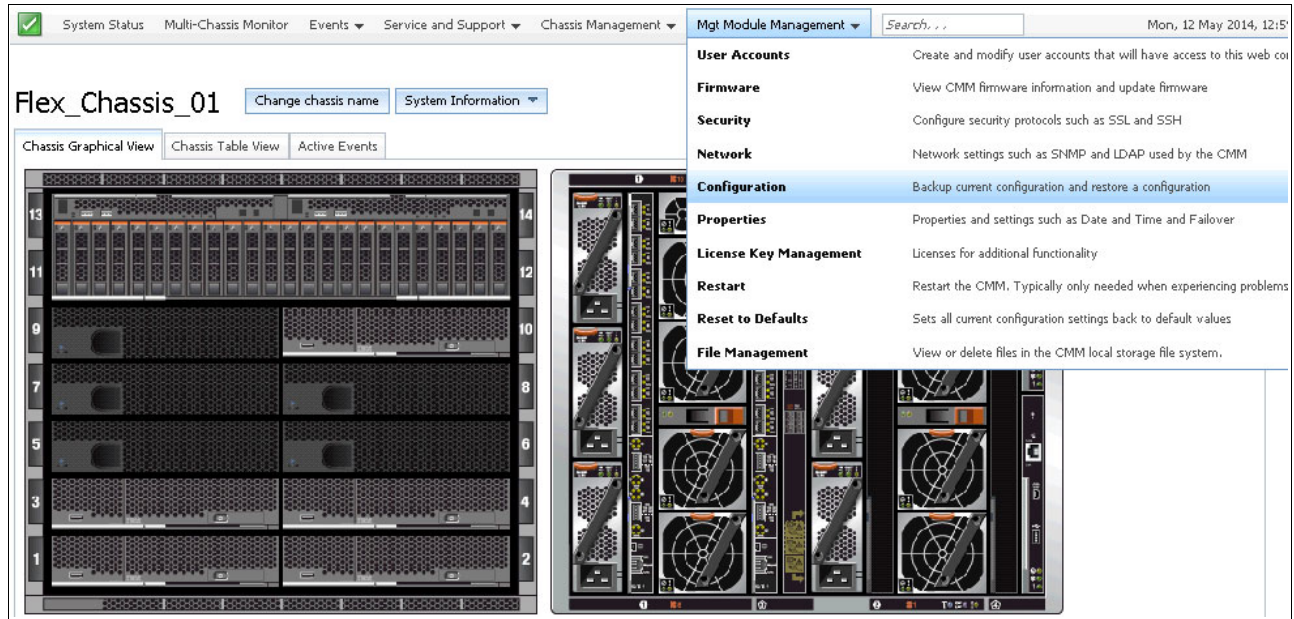


Figure 7-4 CMM main window: Mgt Module Management

The initial setup wizard is included in the Configuration menu. Several options are displayed for managing the CMM configuration.

2. For the first-time connection, click **Initial Setup Wizard**, as shown in Figure 7-5.




Figure 7-5 Manage Configuration window

3. When the wizard starts, the first window displays on the left side of the window the steps to be performed. The basic description of the steps is displayed in the main portion of the window.

Figure 7-6 shows the Welcome window of the setup wizard. Navigation buttons for the wizard are in the lower-left corner of each window. Click **Next**.

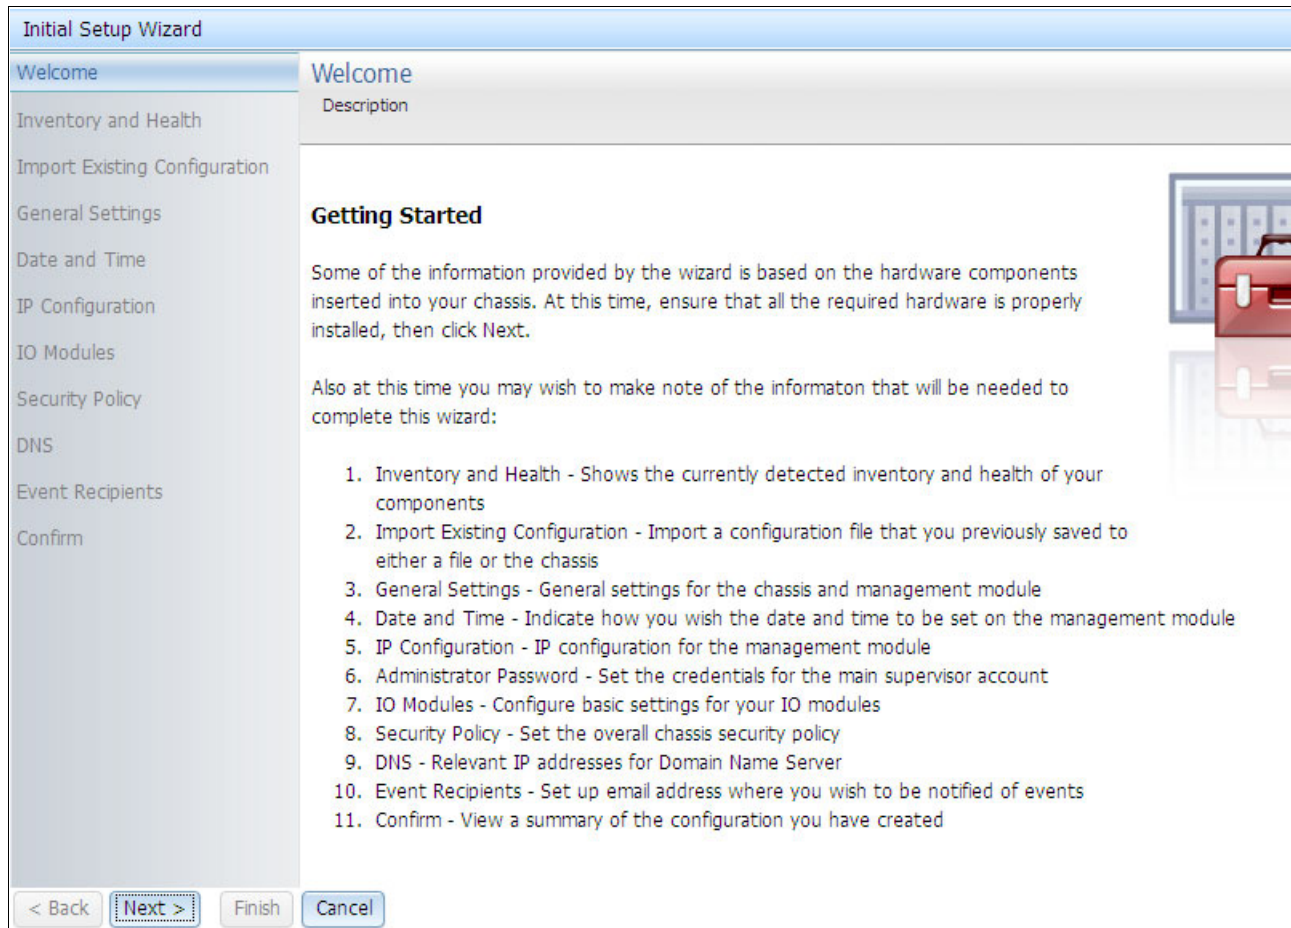


Figure 7-6 Welcome window

4. Select the **Health status** tab on the Inventory and Health window to view the detected components in Chassis and their current health status, as shown in Figure 7-7. Click **Next**.

Initial Setup Wizard

☒ Welcome

Inventory and Health
Shows the currently detected inventory and health of your components

Import Existing Configuration

General Settings

Date and Time

IP Configuration

IO Modules

Security Policy

DNS

Event Recipients

Confirm

Examine the list of your components below and confirm that all components are present

Health status | Active events

Device Name	Device Type	Health Status	Bay	Machine T
SN#Y030BG1CL001	Management Module	✓ Normal	1	
node01	Blade	✓ Normal	1	
node02	Blade	✓ Normal	2	
node03	Blade	✓ Normal	3	
node04	Blade	✓ Normal	4	
node05	Blade	✓ Normal	5	
node06	Blade	✓ Normal	6	
node08	Blade	✓ Normal	8	
node09	Blade	✓ Normal	9	
node10	Blade	✓ Normal	10	
Power Module 1	Power Module	✓ Normal	1	
Power Module 2	Power Module	✓ Normal	2	

Figure 7-7 Inventory and Health window

5. If you saved a configuration file, you can select the file that you created by using the Import Existing Configuration window. The appropriate values are automatically entered in the fields of the wizard, as shown in Figure 7-8. Click **Next**.

Initial Setup Wizard

☒ Welcome

☒ Inventory and Health

Import Existing Configuration

General Settings

Date and Time

IP Configuration

IO Modules

Security Policy

DNS

Event Recipients

Confirm

To facilitate your task of setting up the management module, you can import a configuration file that chassis. Importing a configuration will automatically fill in the fields of this wizard with the appropriate

If this is your first time setting up a chassis, you will not have a configuration file to import your management module settings, or for configuring multiple chassis. To create a config console under Mgt Module Management -> Configuration.

i Some restore operations may cause a temporary loss of web connectivity. Under the confirmation popup and restore log may not be available. If web connectivity is lost and restart your session. At this point, check the event log for messages related to the c

Passphrase:

Confirm pass:

Upload configuration file:
backup_20140508_211846.bkp

Figure 7-8 Import Existing Configuration window

- The General Settings window prompts you to enter some descriptive information about Chassis, including location and contact person, as shown in Figure 7-9. Click **Next**.

The screenshot shows the 'Initial Setup Wizard' window. On the left is a sidebar with a list of steps: Welcome, Inventory and Health, Import Existing Configuration, General Settings (highlighted), Date and Time, IP Configuration, IO Modules, Security Policy, DNS, Event Recipients, and Confirm. The main area is titled 'General Settings' with the subtitle 'General settings for the chassis and management module'. It contains several input fields: 'Management module name' (SN#Y0308G1CL001), 'Chassis description' (empty), 'Contact person' (No Contact Configured), 'Chassis location' (No Location Configured), 'Room ID' (empty), 'Rack ID' (empty), 'Lowest U-position' (0), and 'Unit height of chassis' (10).

Figure 7-9 General Settings window

- Set the date and time for the CMM in the Date and Time window, as shown in Figure 7-10. There are two options to sync the time: by using Network Time Protocol (NTP) or setting it manually. Click **Next**.

The screenshot shows the 'Initial Setup Wizard' window with the 'Date and Time' step selected in the sidebar. The main area is titled 'Date and Time' with the subtitle 'Date and time settings for the management module'. It includes instructions: 'Indicate how you wish the date and time to be set on the management module. The management module will use this time for the event log, for example.' The 'Select method' dropdown is set to 'Synchronize with an NTP server'. Other fields include 'NTP server host name and/or IP address' (9.42.170.223), 'Synchronization frequency (minutes)' (20), and a checked box for 'Enable NTP v3 Authentication'. Below this, 'NTP v3 Authentication key index' is 2 and 'NTP v3 Authentication key (M - MD5)' is BBB8F9C3. A status message says 'NTP last updated the clock on 05/09/2014 16:18:53 by 0 s.' The 'GMT Offset' dropdown is set to '-5:00 - Eastern Standard Time (Eastern USA, Ontario, Quebec)'. A note states: 'Unable to automatically determine the daylight saving time to use. Please provide the DST'. The 'Selected GMT offset' is the same as the GMT Offset. The 'Available schemes' dropdown is set to 'USA and Canada'. Finally, the 'Automatically adjust for daylight savings time (DST)' checkbox is checked.

Figure 7-10 Date and Time window

8. Each CMM is configured with the same static IP address. Use the IP Configuration window that is shown in Figure 7-11 to create a unique static IP address for each CMM. If DHCP is not used, only one CMM at a time can be added onto the network for discovery. Adding more than one CMM to the network without a unique IP address assignment for each CMM results in IP address conflicts. Click **Next**.

The screenshot shows the 'Initial Setup Wizard' window with the 'IP Configuration' tab selected. The left sidebar lists various setup steps: Welcome, Inventory and Health, Import Existing Configuration, General Settings, Date and Time, IP Configuration (highlighted), IO Modules, Security Policy, DNS, Event Recipients, and Confirm. The main area is titled 'IP Configuration' and 'IP configuration for the management module'. It contains fields for 'Host name' (set to 'CMM') and 'Domain name' (empty). A checkbox for 'Register this interface with DNS' is unchecked. Below these are two tabs: 'IPv4' (selected) and 'IPv6'. The 'IPv4' tab displays 'Currently assigned IPv4 address information' with 'IP address: 9.42.170.215', 'Subnet mask: 255.255.254.0', and 'Default gateway: 9.42.170.1'. It also shows 'IP address assignment methods: Use static IP address'. At the bottom, 'Static IP Address Settings' include a note '*Changing settings requires a CMM restart.' and fields for 'Static address: 9.42.170.215', 'Subnet mask: 255.255.254.0', and 'Default gateway: 9.42.170.1'.

Figure 7-11 IPv4 tab configuration window

9. If you need to set up IPv6, select the **IPv6** tab, as shown in Figure 7-12. Click **Next**.

The screenshot shows the 'Initial Setup Wizard' window with the 'IP Configuration' tab selected. The left sidebar is the same as in Figure 7-11. The main area is titled 'IP Configuration' and 'IP configuration for the management module'. It contains fields for 'Host name' (set to 'CMM') and 'Domain name' (empty). A checkbox for 'Register this interface with DNS' is unchecked. Below these are two tabs: 'IPv4' and 'IPv6' (selected). The 'IPv6' tab displays 'Enable IPv6' with a checked checkbox. It shows 'Link local address: fe80::5ef3:fcff:feff:73d8', 'Stateless address: None assigned', 'Default gateway: 0::0', and 'Stateful address:'. Under 'IP address assignment methods', three options are listed: 'Use stateless address autoconfiguration' (checked), 'Use stateful address configuration (DHCPv6)' (checked), and 'Use statically assigned IP address' (unchecked).

Figure 7-12 IPv6 configuration window

10. You can view the status and configure the options for the I/O modules that are connected to the CMM, as shown in Figure 7-13. Click **Next**.

Device Name	Health Status	Enable external ports	Enable external manag. over all ports	Private IP
IO Module 1	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
IO Module 2	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	

Figure 7-13 I/O Modules window

11. Select the security policy for your CMM, as shown in Figure 7-14. Click **Next**.

Policy Setting: Secure

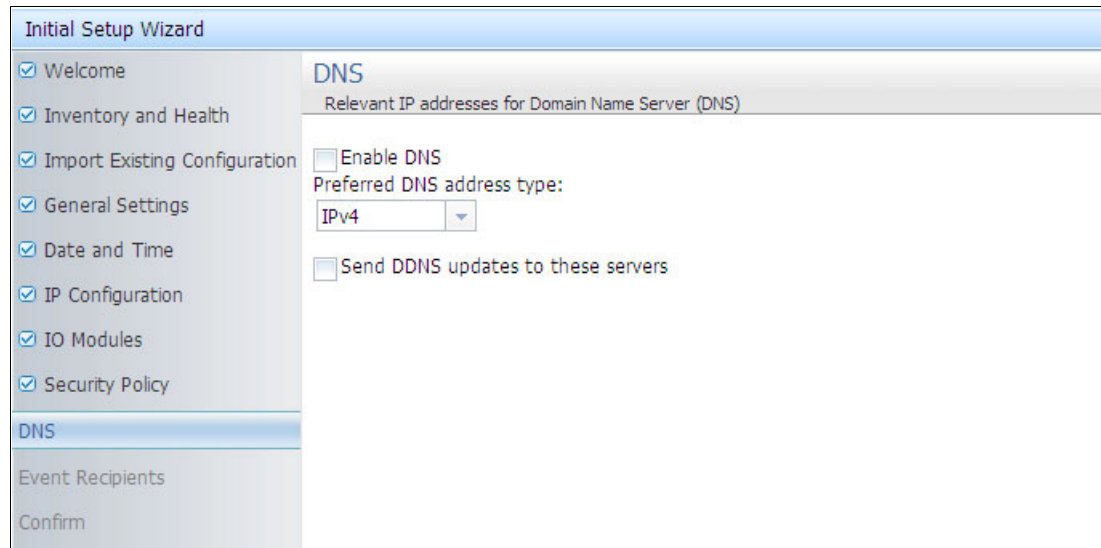
This is the default security setting that establishes a secure chassis infrastructure with a level of user control over the chassis configuration. The Secure level of security policy provides the following settings:

- Password policies are automatically checked and required to be strong
- Well-known passwords for network login are automatically required to be changed after initial setup
- Only secure communication protocols may be enabled
- Certificates for establishing secure and trusted connections to applications running on management processors are automatically generated and managed by the system

Figure 7-14 Security Policy window

Restriction: When the CMM is set to *Secure* security mode, only the secure file transfer methods HTTPS and Secure File Transfer Program (SFTP) can be used for firmware updates and other tasks that involve file transfers. These other tasks include transferring a backup configuration file to restore a configuration. The insecure file transfer protocols, HTTP, FTP, and Trivial File Transfer Protocol (TFTP) are disabled when security is set to the *Secure* mode.

12. Select the appropriate Domain Name Server (DNS) options for your CMM, as shown in Figure 7-15. Click **Next**.



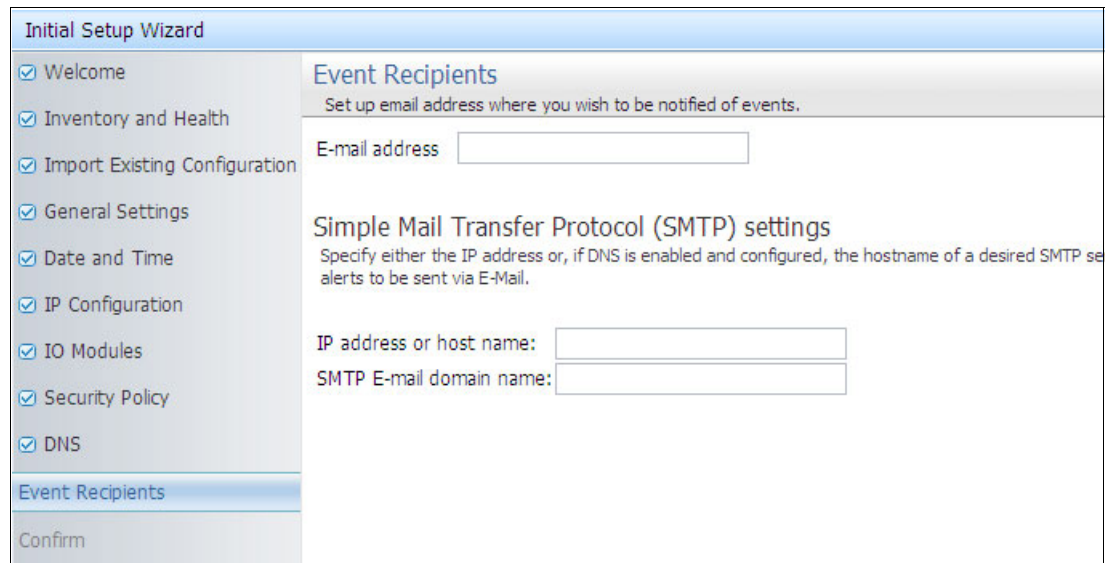
The screenshot shows the 'Initial Setup Wizard' window with the 'DNS' step selected in the left sidebar. The main area is titled 'DNS' and contains the following options:

- ☒ Enable DNS
- Preferred DNS address type:
- ☐ Send DDNS updates to these servers

The left sidebar lists the following steps: Welcome, Inventory and Health, Import Existing Configuration, General Settings, Date and Time, IP Configuration, IO Modules, Security Policy, DNS (selected), Event Recipients, and Confirm.

Figure 7-15 DNS setup window

13. Enter the email addresses to which notifications are sent as CMM events occur, as shown in Figure 7-16. Click **Next**.



The screenshot shows the 'Initial Setup Wizard' window with the 'Event Recipients' step selected in the left sidebar. The main area is titled 'Event Recipients' and contains the following options:

- Set up email address where you wish to be notified of events.
- E-mail address:
- Simple Mail Transfer Protocol (SMTP) settings
- Specify either the IP address or, if DNS is enabled and configured, the hostname of a desired SMTP server to be sent via E-Mail.
- IP address or host name:
- SMTP E-mail domain name:

The left sidebar lists the following steps: Welcome, Inventory and Health, Import Existing Configuration, General Settings, Date and Time, IP Configuration, IO Modules, Security Policy, DNS, Event Recipients (selected), and Confirm.

Figure 7-16 Event Recipients window

14. Confirm all of the information that was entered in the setup wizard, as shown in Figure 7-17. Click **Finish**.

Initial Setup Wizard	
<input checked="" type="checkbox"/> Welcome	<p>You have completed entry of all the information necessary to get your chassis running and communicating with network.</p> <p>Step 4 - General Settings</p> <p>Management module name: SN#Y011BG25302F</p> <p>Chassis description: Flex_Chassis_01</p> <p>Contact person: No Contact Configured</p> <p>Location: No Location Configured</p> <p>Room ID:</p> <p>Rack ID:</p> <p>Lowest U-position: 0</p> <p>Unit height of chassis: 10</p> <p>Step 5 - Date and Time</p> <p>Select method: Synchronize with an NTP server</p> <p>Date: Fri May 9 00:00:00 EDT 2014</p> <p>Time: 12:19 PM</p> <p>GMT Offset: -5:00 - Eastern Standard Time (Eastern USA, Ontario, Quebec)</p> <p>Automatically adjust for daylight savings time (DST): Enabled</p> <p>NTP server host name and/or IP address: 9.42.170.223</p> <p>Synchronization frequency (minutes): 20</p> <p>Enable NTP v3 Authentication: Enabled</p> <p>NTP v3 Authentication key index: 2</p> <p>NTP v3 Authentication key (M - MD5): BBB8F9C3</p>
<input checked="" type="checkbox"/> Inventory and Health	
<input checked="" type="checkbox"/> Import Existing Configuration	
<input checked="" type="checkbox"/> General Settings	
<input checked="" type="checkbox"/> Date and Time	
<input checked="" type="checkbox"/> IP Configuration	
<input checked="" type="checkbox"/> IO Modules	
<input checked="" type="checkbox"/> Security Policy	
<input checked="" type="checkbox"/> DNS	
<input checked="" type="checkbox"/> Event Recipients	
Confirm	
<p>< Back Next > Finish Cancel</p>	

Figure 7-17 Confirm window

7.1.3 Configuring IP addresses for the chassis components

By using the Component IP Configuration menu, you can set the IP parameters on I/O modules and compute nodes, as shown in Figure 7-18.

Component IP Configuration

Configure IPv4 and IPv6 address information for the components below.

I/O Modules

Bay	Device Name	IPv4 Enabled	IP Address
1	IO Module 1	Yes	View
2	IO Module 2	Yes	View

Compute Nodes

Bay	Device Name	IPv4 Enabled	IP Address
1	Node 01 (x240_01)	Yes	View
2	Node 02 (x240_02)	Yes	View
3	Node 03 (x240_03)	Yes	View
4	Node 04 (x240_04)	Yes	View
10	Node 10 (FSM)	Yes	View

Storage Nodes

Bay	Device Name	IPv4 Enabled	IP Address
11-14:1	Node 11 - 01	Yes	View
11-14:2	Node 11 - 02	Yes	View

Figure 7-18 Component IP Configuration window

Click the I/O module or compute node link to open its IP properties window, as shown in Figure 7-19.

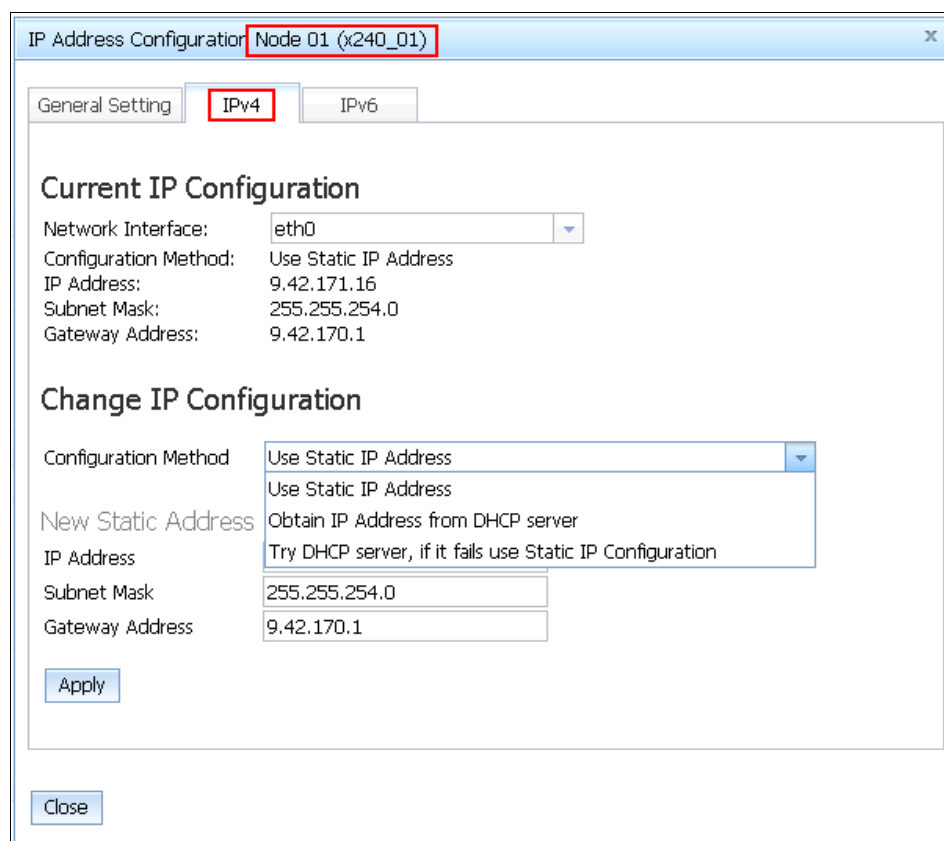


Figure 7-19 IP Address Configuration node01 window

7.2 Firmware updates and basic hardware configuration

In our lab, we are using FastSetup tool to perform firmware updates and basic hardware configuration of the hardware environment.

FastSetup is a no-cost software tool that helps simplify the maintenance and deployment of select x86 servers, including System x rack servers, BladeCenter blade servers, and Flex System compute nodes. The intuitive GUI starts all phases of server setup, including discovery, updating, and configuration. Features include templates that enable replication of settings across many servers and automation that reduces hands-on time and user errors. Wizards and other default settings enable flexible customization capabilities. The low-touch, set-once and walk-away feature reduces the hands-on server setup time from days to minutes, particularly for larger deployments.

For more information about and to download FastSetup, see this website:

<https://www-947.ibm.com/support/entry/myportal/docdisplay?Indocid=T00L-FASTSET>

Complete the following steps to use the FastSetup tool to configure Flex System:

1. Start the FastSetup tool. A welcome window opens. Click **Next**.
2. In the Initial Configuration window, configure the proxy server (if needed) and select the Network interface that is used to communicate with Flex System, as shown in Figure 7-20. Click **Next**.

Network Access

Tell IBM FastSetup how your local workstation is connected to the Internet and connected to the LAN. A connection to Internet is required to download firmware updates from ibm.com. The LAN is used to access the resources to be managed. The proxy configuration will be saved when Internet Explorer is allowed to save the cookie from IBM FastSetup.

Proxy Settings(optional)

If your local workstation requires a proxy server to connect to the Internet, enter the information below. The proxy configuration will be saved when Internet Explorer is allowed to save the cookie from IBM FastSetup.

IP/host name:

Port:

User name:

Password:

LAN Access

IBM FastSetup has detected the following network adapters. Select the adapter corresponding to the network IBM FastSetup should use to access the resources you want to manage.


	Ethernet Adapter	Description	IP Address
	Local Area Connection 5	Broadcom BCM5709S NetXtreme II GigE (NDIS VBD Client) #4	9.42.171.21

Figure 7-20 FastSetup: Network Access

3. In the Resource Selection window, select **Flex System (CMM, x86 Compute Nodes and I/O Modules)**. Click **Next**.
4. In the Task Selection window, select a task that is based on your needs. In our example, we use Full Setup. Click **Next**.
5. In the System Discovery page, you can allow Fast Setup to scan your network and find all accessible CMMs or you can add your CMM manually. After the CMM is discovered (as shown in Figure 7-21 on page 115), click **Next**.

System Discovery






Select a method for discovering a system in your environment. You can choose to automatically discover a system, manually enter a system IP address to discover, or select a system from a list of previously discovered systems.

 **EIZUI1111V**
 Run successfully.
[Details...](#)

☐ Automatically discover systems in this subnet
☒ Manually enter a system IP address or host name to discover

System IP address or host name:

IP Addresses/Host Name to Discover:

 |     Actions ▼

	IP Address/Host Name
<input checked="" type="checkbox"/>	9.42.170.215

☐ Select from a list of previously discovered systems

Select the system that you would like to configure:


	Name	Model	Machine Type	URL/IP Address	Status
<input checked="" type="checkbox"/>	LAB_Chassi	Flex Chassis	8721HC1	9.42.170.215	 Valid

Figure 7-21 Discovering CMM in FastSetup

6. It takes some time to discover the details about your Flex System. You can monitor discovery status on the Inventory and Health page. After the discovery is finished, click **Next**.

7. In the Device Selection window, select the component that you want to configure and update it by using Fast Setup, as shown in Figure 7-22. Click **Next**.

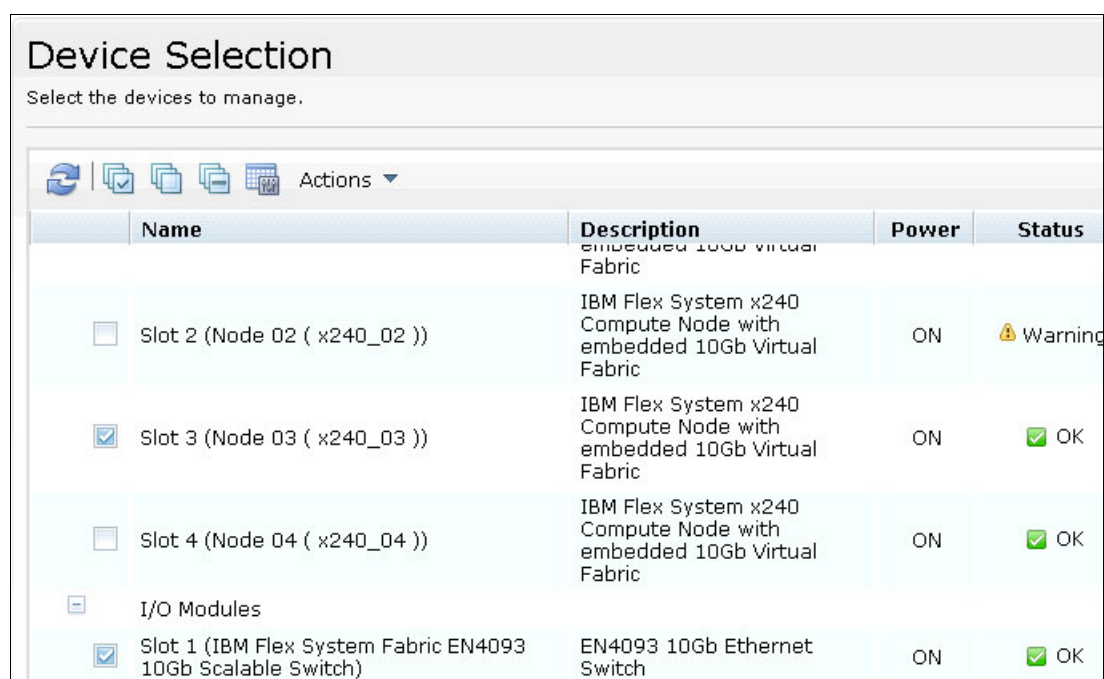


Figure 7-22 Select components

8. In the Temporary IP Settings window, enter the temporary IP for the Flex System compute node, as shown in Figure 7-23. This IP address is used to boot the compute node from the FastSetup machine. Ensure that this IP address has full network access to a machine from where Fast Setup is run. Click **Next**.

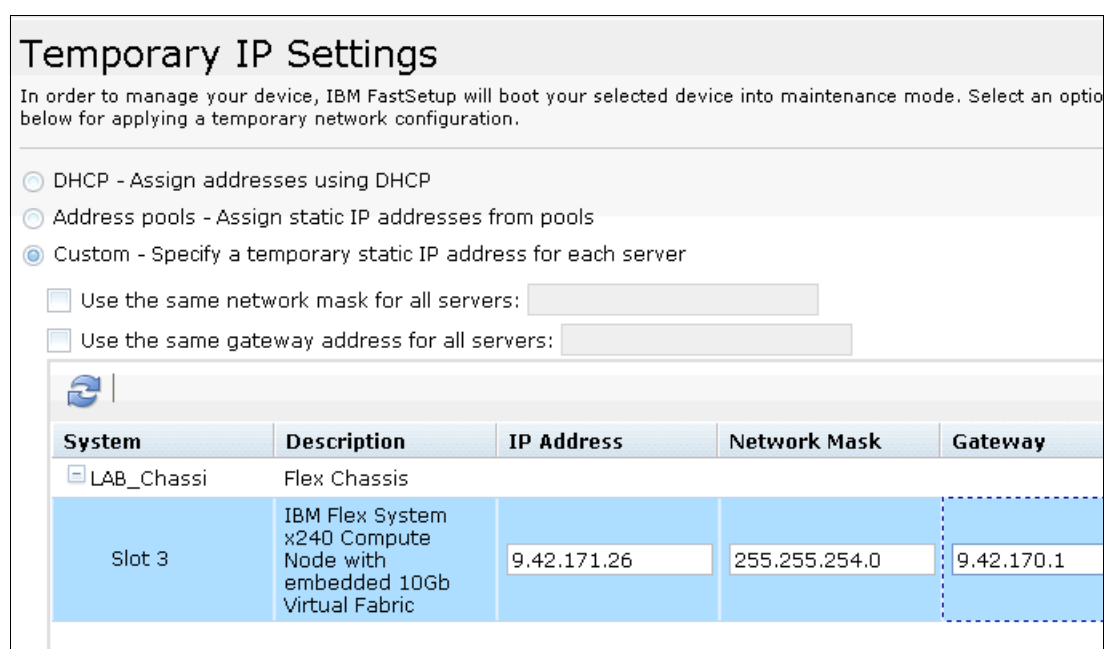


Figure 7-23 Configure temporary IP address for the compute node

- In the Adapter Port Settings window, select which NIC Fast Setup must use for communication, as shown in Figure 7-24. Click **Next**.

Adapter Port Settings

For each server, specify the adapter port connected to the data network.

☐ All servers use the same adapter port: NIC 1 (I/O Bay 1(A3))

Actions ▼

System	Description	IP Address	Adapter Port - MAC Address (I/O Bay(Internal Port))
LAB_Chassi			
Slot 3	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric	9.42.171.26	NIC 1 - 34:40:b5:be:7d:00 (I/O Bay 1(A3))

Figure 7-24 Select NIC to use

- The compute node must be rebooted after you click **Next** in the previous step. A pop-up window opens in which you confirm your choice.
- FastSetup is collecting all possible information about the compute node. After the process completes (as shown in Figure 7-25), click **Next**.

Device Inventory

A detailed inventory is being collected on the selected devices to obtain the current firmware levels.



Actions ▼

Device Name	Description	Build ID	Release Date	Version	Status
LAB_Chassi	Flex Chassis				✓ Finished
Servers					
slot 3	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric				✓ Finished



Figure 7-25 Gathering details about the compute node

Management Module Configuration

Select the Management Module you want to configure and select Configure Settings.

Configure Settings

	Chassis	Bay	IP Address	Host Name	Status
	LAB_Chassi	2	9.42.170.215	flexchassis1	 100%


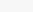
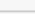
13. In the Management Module Update window (as shown in Figure 7-27), you can apply Firmware updates to CMM. Select the applicable updates and then click **Next**. If no updates are needed, click **Next** only.

Management Module Updates

Select the type of update you want to apply, select the Management Modules to which it will be applied, and click Apply Updates.

☒ Update using the latest available Management Module firmware.

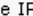

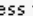
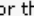
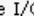
☐ Select from a list of all available Management Module firmware levels.

	Chassis	Description	Installed Version	Pending Version	Status
					
	LAB_Chassi		2PET12I	<div>unavailable</div>	N/A

14. In the I/O Module Configuration window (as shown in Figure 7-28), configure the network parameters for all selected I/O modules. Click **Next**.

I/O Module Configuration

Configure the IP address for the I/O Modules in each chassis.






 Actions ▼

Chassis	Description	State	Configuration Type	IP Address	Network Mask	Gateway
<input type="checkbox"/> LAB_Chassis1	Flex Chassis					
Bay 1	IBM Flex System Fabric EN4093 10Gb Scalable Switch	Enabled ▼	Static ▼	9.42.171.8	255.255.254.0	9.42.170.1

118 Implementing Lenovo Client Virtualization with Citrix XenDesktop

15. In the I/O Module Updates window (as shown in Figure 7-29), you can update the firmware of your I/O module. If there are applicable updates, select those updates and click **Next**. If no updates are needed, click **Next**.

I/O Module Updates

Select the type of update you want to apply, select the I/O Modules to which the update will be applied, and click Apply Updates.

☒ Update using the latest available I/O Module firmware
 ☐ Select from a list of all available I/O Module firmware levels

	System	Description	Installed Version	Pending Version	Status
<div> <div></div> <div>LAB_Chassi</div> </div>					
	Slot 1	IBM Flex System Fabric EN4093 10Gb Scalable Switch	Boot ROM: 7.5.3.0 Main Application 1: 7.5.3.0 Main Application 2: 7.2.2.2		N/A

Figure 7-29 I/O modules firmware update

16. Compute node firmware updates can be applied on the Server Updates page of the FastSetup tool, as shown in Figure 7-30. If there are applicable updates, select those updates and click **Next**. If no updates are need, click **Next**.

Server Updates

Select the update type you want to apply. Then select the servers or components to which it will be applied and click Apply Updates.

☐ Update using the UpdateXpress System Pack (UXSP) - server level only
 ☐ Update using the latest available component firmware
 ☒ Select from a list of all available component firmware levels

	System	Description	Installed Version	Pending Version	Status
<div> <div></div> <div>LAB_Chassi</div> </div>					
	Slot 3	IBM Flex System x240 Compute Node with embedded 10Gb Virtual Fabric			Loaded
<input checked="" type="checkbox"/>		Emulex UCNA Adapter Firmware Update	Unrecognized (N/A)	4.6.281.21-1(12/03/2013)	
<input checked="" type="checkbox"/>		IBM Online SAS/SATA Hard Disk Drive Update Program	Unrecognized (N/A)	1.13.04(02/27/2014)	

Figure 7-30 Update firmware on the compute node

17. Complete the following steps in the Server Configuration window:

- a. In the IMM Configuration window, configure the network settings for the IMM on the compute node, as shown in Figure 7-31.

Integrated Management Module Configuration
Select the Management Module you want to configure, then select Configure Settings.

Configure Settings

	Chassis	Bay	IP Address	Host Name	Status
<input checked="" type="radio"/>	LAB_Chassi	3	9.42.171.18	IMM2-3440b5bf4d71	100%

Figure 7-31 Network configuration of the IMM

- b. In the RAID Configuration window, configure the RAID level that is to be used on your compute node, as shown in Figure 7-32.

Configure RAID Array

☒ Select the desired RAID level and the drives to be included in the array.
☐ Create a RAID0 array using all available drives

Server: Node 03 (x240_03) **RAID Controller:** SAS2004

Desired RAID level: RAID0 **Volume Size(2048-1907738):** 2,048 MB

Minimum drives: 2 **Added number of drives:** 0

Available Drives					Added Drives	
<input type="checkbox"/>	953	Drive 1	SAS_H DD	Yes		
<input type="checkbox"/>	953	Drive 2	SAS_H DD	Yes		

>>

Figure 7-32 Configure RAID on the compute node

- c. In the UEFI Settings page (as shown in Figure 7-33), configure the Boot Order for UEFI or reset the order to the default settings.

Configure Basic Settings

☒ Set all UEFI settings to default values
☐ Specify boot order and set all other UEFI settings to default values
☐ Specify boot order only

Apply Cancel

Figure 7-33 Configure UEFI settings in FastSetup

18. You can perform the following tasks on the Summary page:

- Export settings for future use.
- Export Firmware repository. This task can be useful if you must run Fast Setup on a machine that does not have an Internet connection.

19. In the Completion State window, you can select the action that you want to perform on the compute nodes after all of the actions are performed. You can shut down all of the servers after all of the tasks are complete or reboot them. Select one of the options and click **Next**.

7.3 Configuring Active Directory Integration for CMM

The CMM offers the possibility to integrate user security into your centralized user management system by using LDAP or Microsoft Active Directory integration. In this section, we describe how to integrate it with Active Directory by using basic settings.

Note: Make sure that you configured DNS settings and your domain is discoverable via DNS_SRV records.

Complete the following steps to configure LDAP authentication on the CMM:

1. Log in to the CMM web console by using your local account.
2. Click **Mgt Module Management** → **Network** → **LDAP Client**.
3. In the LDAP Authentication section, select **Use LDAP Servers for Authentication Only (with local authorization)**.
4. In the LDAP Servers field section, select **Use DNS to find LDAP Servers**.
5. In the Miscellaneous Settings Finding Method section, select **w/ Login credentials**.
6. You can leave the others fields empty and click **Apply**, as shown in Figure 7-34.

Lightweight Directory Access Protocol (LDAP) Client

The CMM contains a LDAP client that can be configured to provide user authentication through one or more LDAP servers. The client can be configured to be discovered dynamically or manually pre-configured. Use the dropdown list to select which of these two methods you want to use.

LDAP Authentication: Use LDAP Servers for Authentication Only (with local authorization) ▼

LDAP Servers: Use DNS to find LDAP Servers ▼

Active Directory Forest Name:

Domain Name:

Active Directory Settings

Use Mgt Module Management > User accounts for user configuration

Miscellaneous Settings

Root DN:

UID search attribute:

Binding method: w/ Login credentials ▼

Figure 7-34 LDAP configuration

Complete the following steps to pair Active Directory Groups to CMM roles:

1. Click **Mgt Module Management** → **User Accounts** → **Group Profiles**, as shown in Figure 7-35.

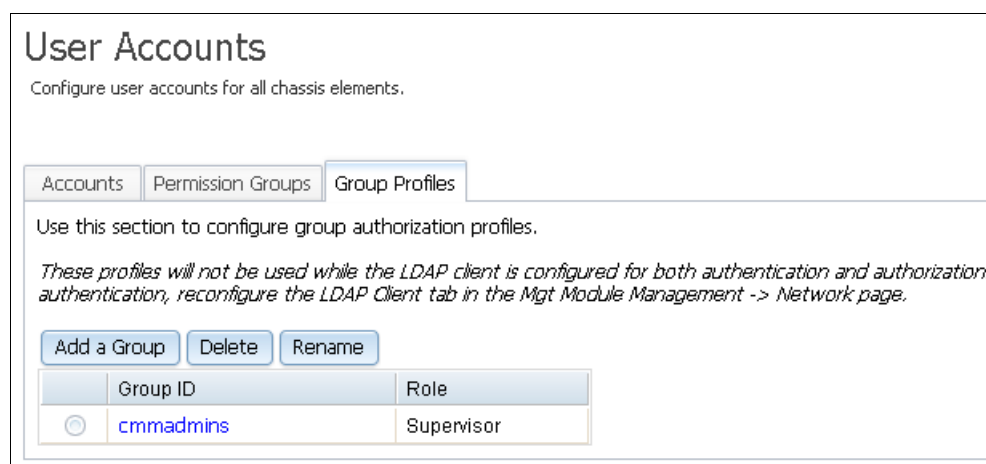


Figure 7-35 Active Directory Group-Role Mapping

2. Click **Add a Group**. Complete the following steps:
 - a. In the Group Profile Name tab, enter Group Name (the Active Directory group name). Click **Next**.
 - b. Define the Role and click **Next**.
 - c. Define the Authority and click **Next**.
 - d. Define the Access Scope and click **Finish**.
3. Log out and log in by using a domain account.

7.4 Configuring the EN4093 10Gb Ethernet Switch

In the example that is described in this section, we use the Switch Center tool to create VLANs and assign those VLANs to network ports in the network switch, which is installed into the Flex System chassis. We also configure Unified Fabric Port (UFP) on the 10 GbE ports of the compute node to create separated network subinterfaces for different types of traffic.

Switch Center provides remote monitoring and management of Ethernet and converged switches from Lenovo. It is designed to simplify and centralize the management of your BladeCenter, Flex System, and RackSwitch Ethernet and converged switches.

The Switch Center offers the following features:

- ▶ Improve network visibility and drive availability, reliability, and performance
- ▶ Simplify management of large groups of switches by using automatic discovery
- ▶ Automate and integrate management, deployment, and monitoring
- ▶ Provide simple network management protocol (SNMP) based configuration and management
- ▶ Support network policies for virtualization
- ▶ Authentication and authorization

- ▶ Fault and performance management
- ▶ Integration with VMware Virtual Center and vSphere clients

For more information about Switch Center, see this website:

<http://www.ibm.com/systems/networking/software/snsc/index.html>

Any third-party management platforms that support SNMP also can be used to configure and manage the modules.

The following tasks are described next:

- ▶ “Adding a switch to Switch Center”
- ▶ “Configuring VLANs on the switch with Switch Center” on page 126
- ▶ “Enabling UFP on the EN4093R by using Switch Center” on page 128
- ▶ “Enabling UFP on the x240 compute node” on page 130

Adding a switch to Switch Center

Complete the following steps to add a Flex System I/O module to Switch Center:

1. Run Switch Center on the machine where it is installed. A login window opens. If this log in is the first time that you logged in to Switch Center, use the default credentials. Consider changing the password on your first login.

Note: Switch Center uses the following default credentials:

- ▶ User name: admin
- ▶ Password: admin

2. In the main window of Switch Center, click **Device List Page**, as shown in Figure 7-36.

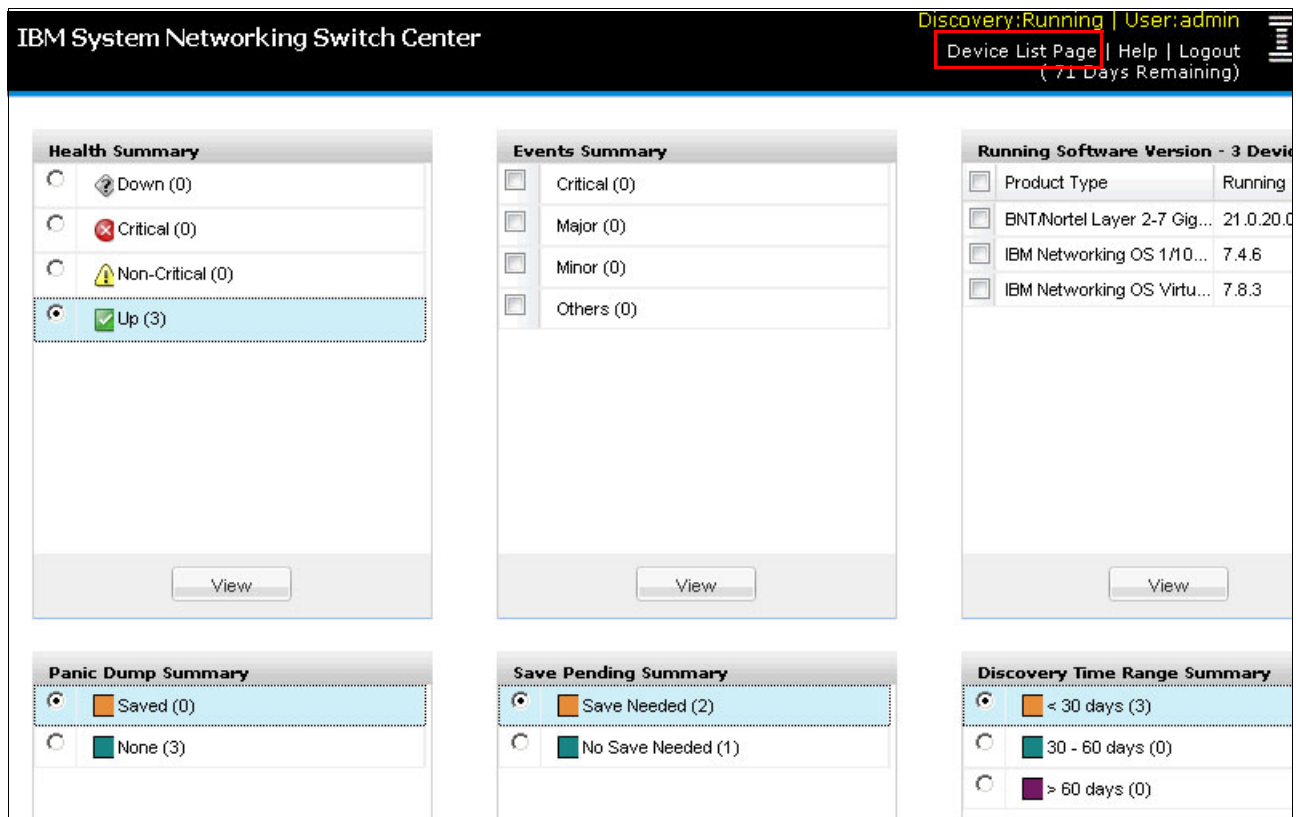


Figure 7-36 Switch Center main window

3. In the Device List window, browse to the Flex System category and click **Add a Switch**, as shown in Figure 7-37.




Figure 7-37 Add a switch in Switch Center

4. In the Add a Switch window, enter the switch connection details and then, click **Open**, as shown in Figure 7-38.

Figure 7-38 Connection details for a network switch

5. After the switch is added, it is shown in the list of available switches (the BladeCenter and Flex System I/O modules are listed in our example). Click the highlighted IP of the switch to see the Device Console, as shown in Figure 7-39.

	Product Name	IP Address
BladeCenter (3 devices)		
<input type="checkbox"/>	IBM Networking OS 1/10Gb Uplink Ethernet Swi...	9.42.171.84
<input type="checkbox"/>	IBM Networking OS Virtual Fabric 10Gb Switch ...	9.42.171.85
<input type="checkbox"/>	BNT/Nortel Layer 2-7 Gigabit Ethernet Switch M...	9.42.171.86
Flex System (1 device)		
<input type="checkbox"/>	IBM Flex System Fabric EN4093 10Gb Scalable ...	9.42.171.8

Figure 7-39 Selecting the switch to configure

Configuring VLANs on the switch with Switch Center

Complete the following steps to configure VLANs on the switch by using Switch Center:

1. To add VLANs to the switch, in Device Console, click the **Configure** tab, select the **Layer 2** folder and then, click **Virtual LANs**. Click **Insert** to insert a new VLAN, as shown in Figure 7-40.

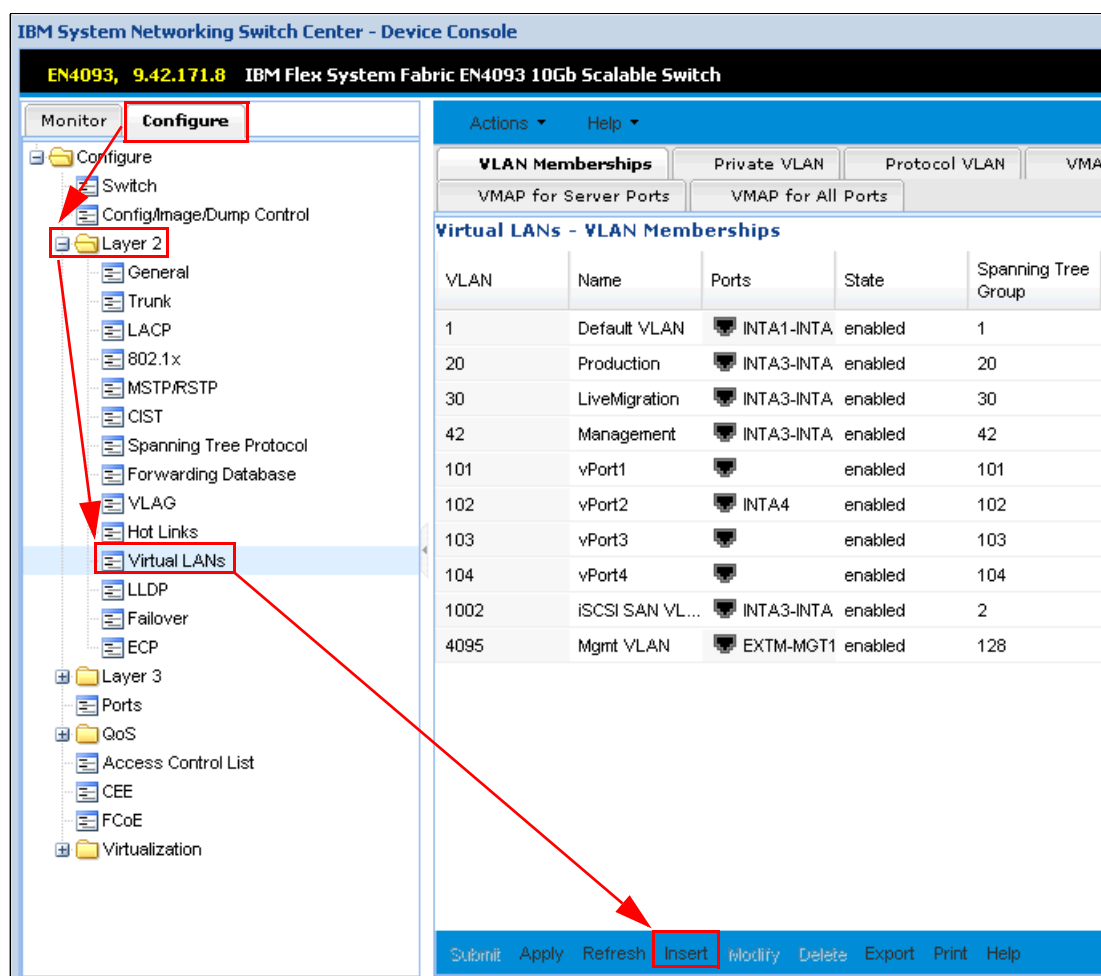


Figure 7-40 Virtual LANs menu

2. Enter the required details for the VLAN. In our example, we add VLAN 80 and assign that VLAN to ports INTA3 and EXT1, as shown in Figure 7-41. Click **OK**.

Virtual LANs - VLAN Memberships - Insert Form

VLAN: 80 1..4095

Name: Management_VLAN_80 0..32 characters

Ports: INTA3;EXT1 Browse...

State: ☒ enabled ☐ disabled

Spanning Tree Group: 80 0..127 [1-127 for 802.1d; 1 for RSTP; 0-32 for MSTP]

Management State: ☒ enabled ☐ disabled

Virtual Ports: Browse...

OK Cancel

Figure 7-41 VLAN details

3. To apply and save the configuration, click **Submit** at the bottom center of the window. Click **Apply** (which is next to Submit). To preserve this change across the reboot of the switch, click **Actions** → **Save**, as shown in Figure 7-42.

EN4093, 9.42.171.8 IBM Flex System Fabric EN4093 10Gb Scalable Switch

Monitor **Configure**

Actions Help

Apply **Save** Diff Config Diff Flash Config Dump Syslog Dump Revert Revert Apply Reboot Switch Clear Panic Dump Exit

Private VLAN Protocol VLAN VMAP

VMAP for All Ports

Memberships

Ports	State	Spanning Tree Group	M S
INTA1-INTA	enabled	1	c
INTA3-INTA	enabled	20	c
INTA3-INTA	enabled	30	c
INTA3-INTA	enabled	42	c
INTA3;EXT1	enabled	80	e
...	enabled	101	c
INTA4	enabled	102	c
...	enabled	103	c
104	vPort4	104	c
1002	iSCSI SAN VL...	2	c
4095	Mgmt VLAN	128	e

Figure 7-42 Selecting Actions → Save

4. Complete the following steps:
 - a. Click **Ports** on the left side pane.
 - b. Find the port to which you assigned a created VLAN.
 - c. Double-click **VLAN Tag State** and change it to tagged.
 - d. Double-click **Default VLAN** in front of it and select the VLAN that you want to use as the default, as shown in Figure 7-43.

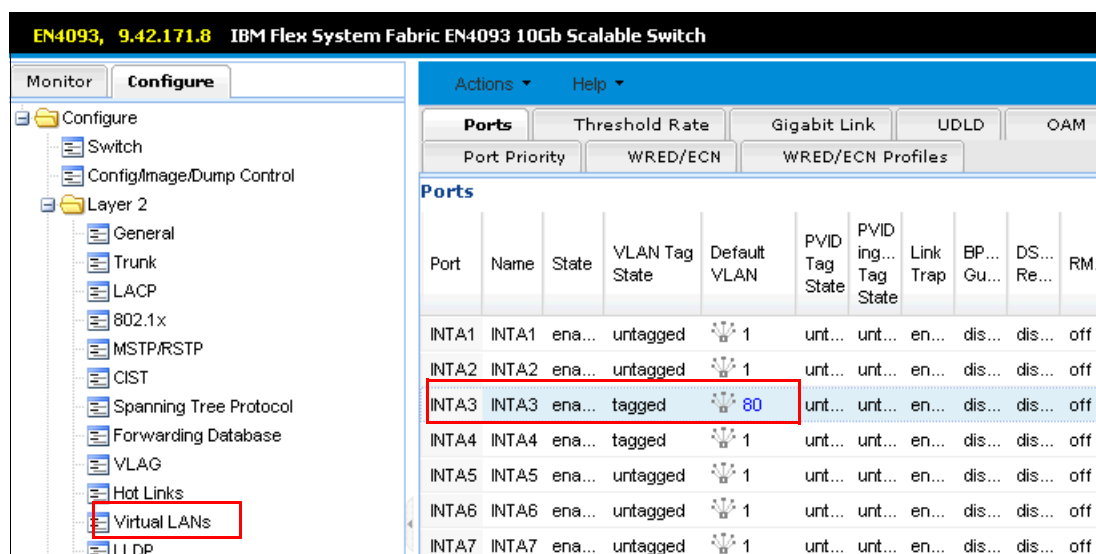


Figure 7-43 Configure port

- e. Click **Submit** at the bottom center of the page.
- f. Click **Apply** (which is next to Submit).
- g. To preserve this change across reboot of the switch, click **Actions** → **Save**.

The switch is now configured to accept traffic from VLAN 80.

Repeat these steps to configure other VLANs on the switch.

Enabling UFP on the EN4093R by using Switch Center

UFP is an approach to NIC virtualization. It is similar to Virtual Fabric vNIC, but with enhanced flexibility and should be considered the direction for future development in the virtual NIC area for Lenovo switching solutions. With Flex System, UFP is supported today on the EN4093R 10Gb Scalable Switch, CN4093 10Gb Converged Scalable Switch, and SI4093 System Interconnect Module.

For more information about UFP and other NIC virtualization choices, see *NIC Virtualization in Flex System Fabric Solutions*, SG24-8223, which is available at this website:

<http://lenovopress.com/sg248223>

Complete the following steps to enable UFP on the switch and on the port:

1. In Switch Center, open Device Console, then click **Configure** → **Virtualization** → **UFP**. In the General tab, select **enabled** and then click **Apply**, as shown in Figure 7-44.

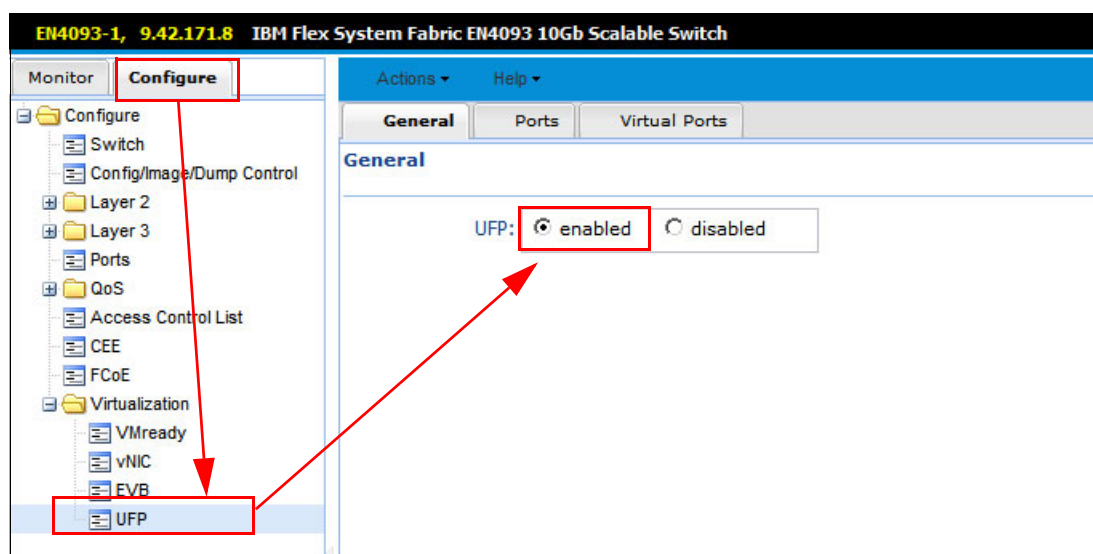


Figure 7-44 Enable UFP

2. Click the **Ports** tab, enable UFP for the wanted port by selecting **enable** in the state field and then click **Apply**, as shown in Figure 7-45.

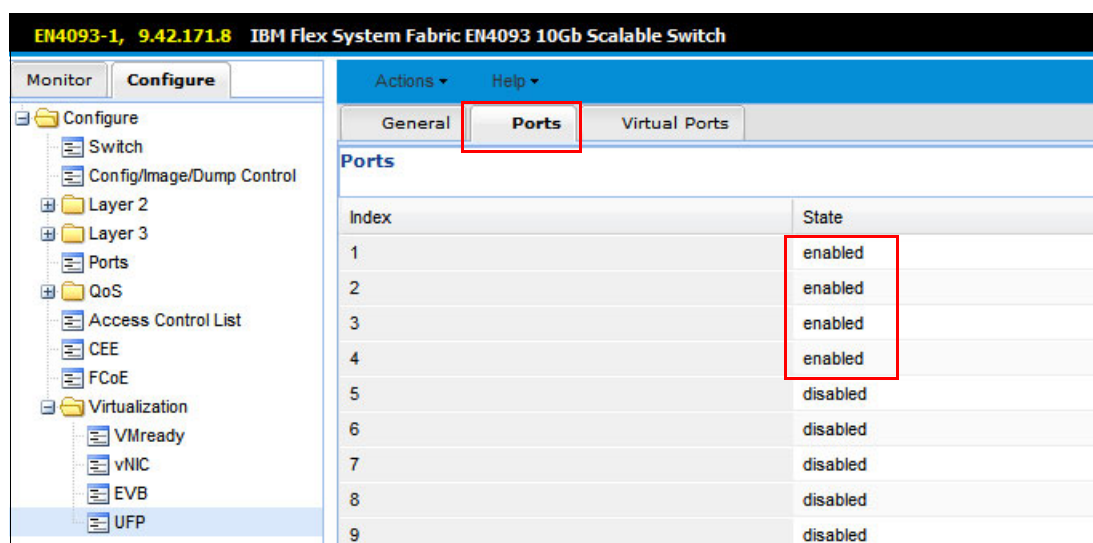


Figure 7-45 Enable UFP for Port

3. In the Virtual Ports tab, you can configure virtual ports, their VLANs membership, Network mode, and bandwidth, as shown in Figure 7-46.

Port Index	vPort Index	State	Network Mode	Network Default VLAN	Network Default Tag
1	1	enabled	access	32	disabled
1	2	enabled	access	1002	disabled
1	3	enabled	access	30	disabled
1	4	disabled	tunnel	0	disabled
2	1	enabled	access	32	disabled
2	2	enabled	access	1002	disabled
2	3	enabled	access	30	disabled
2	4	disabled	tunnel	0	disabled

Figure 7-46 Virtual Ports configuration

7.5 Enabling UFP on the x240 compute node

To enable the UFP function, you must configure Flex System compute node. Complete the following steps:

1. Connect to the console of your compute node and enter the UEFI configuration by pressing F1 during start.
2. Enable Multichannel Mode in UFP Mode personality for the network adapter by clicking **System Settings** → **Network** in UEFI. Complete the following steps:
 - h. From the Network Device List, open the first adapter.
 - i. Press Enter to enter configuration mode.
 - j. Change Multichannel Mode to Unified Fabric Protocol Mode, as shown in Figure 7-47.

Emulex NIC Selection	
Emulex OC111102-F-X Virtual Fabric Adapter 2-port 10Gb LOM	
Firmware Version	: 4.6.281.26
Bus:Device:Function	: C:0:0
Link Speed	: 10 Gbps
Advanced Mode	<Enable>
Personality	<iSCSI>
Multichannel Mode	<IBM Unified Fabric Protocol Mode>
Controller Configuration	
Port Management	

Multichannel Configuration. Use of IBM Virtual Fabric Mode or IBM Unified Fabric Protocol Mode requires the switch to also support the functionality...
SYSTEM RESET REQUIRED

Figure 7-47 Emulex NIC configuration utility

- k. Exit and apply to the second adapter, if needed.
3. Exit and save UEFI configuration.
4. Reboot the server and your UFP configuration is finished.

After you configure network adapters in UFP mode, you are configuring virtual ports settings only in the network switch side; no other configuration in UEFI is needed.

7.6 Configuring iSCSI on the x240 compute node

The Converged Network Adapter that is included with Flex System can work as an iSCSI initiator or as a FCoE HBA. In our example, we describe the scenario of configuring it for use with an iSCSI SAN. Complete the following steps to configure your iSCSI initiator:

1. Connect to the console of your compute node and enter the System Setup utility.
2. Enable iSCSI personality for network adapter by clicking **System Settings** → **Network**. Complete the following steps:
 - a. From Network Device List, open first Adapter.
 - b. Press Enter to begin the configuration mode, as shown in Figure 7-47 on page 130.
 - c. Change Personality to iSCSI.
 - d. Exit and configure the second NIC, if needed (the second port on the same NIC is configured with the first port).
 - e. Exit and save the UEFI configuration.
 - f. Reboot the system and reenter the System Setup utility.
3. Click **System Settings** → **Storage**. Complete the following steps:
 - a. Enter the Emulex iSCSI Utility for the particular adapter, as shown in Figure 7-48.

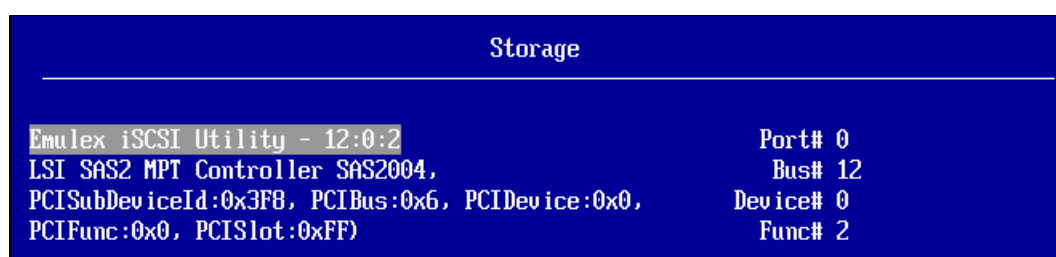


Figure 7-48 UEFI: Emulex iSCSI Utility

- b. Enter the iSCSI Initiator Name and save the changes, as shown in Figure 7-49.



Figure 7-49 uEFI iSCSI Controller Configuration menu

- c. Enter the Network Configuration.
 - d. Enter the iSCSI Target Configuration.
 - e. Exit the configuration window.
 - f. Repeat steps a - e for the second network adapter, if needed.
4. Reboot the computer.

7.7 V7000 configuration

In this section, we describe how to configure the V7000 Storage System. The following topics are included:

- ▶ 7.7.1, “V7000 initial configuration”
- ▶ 7.7.2, “V7000 Storage Node setup wizard” on page 134
- ▶ 7.7.3, “Configuring storage volumes” on page 140
- ▶ 7.7.4, “Configuring hosts” on page 144

7.7.1 V7000 initial configuration

Complete the following steps to configure the V7000 Storage System when the Flex System Manager web user interface is used:

1. Open a web browser and browse to the IP address of the management interface of one of the V7000 nodes. A welcome window from the V7000 GUI opens. You are prompted to create a system (cluster) or add to a system, as shown in Figure 7-50 on page 133. Select **Create a new system**. Click **Next**.



Figure 7-50 V7000 first-time set up welcome window

2. In the window that is shown in Figure 7-51, select whether you are using an IPv4 or IPv6 management IP address and enter the IP address (you can use DHCP or the static address that was assigned). The subnet mask and gateway show the defaults, which you can edit.

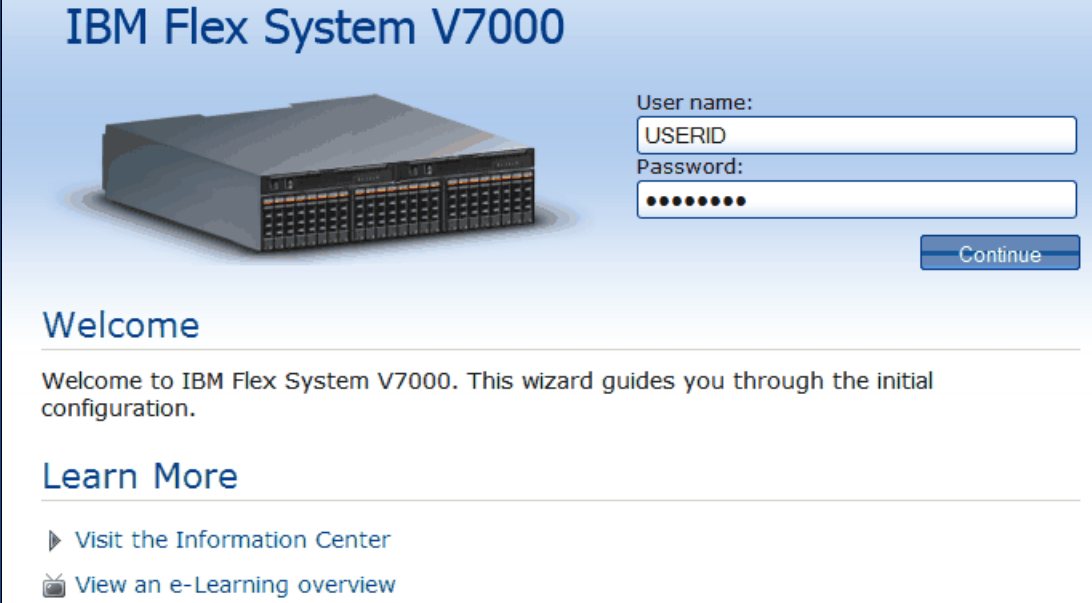
Figure 7-51 Create a new storage cluster

3. Click **Finish** to set the management IP address for the system. System initialization begins and might take several minutes to complete.

When system initialization is complete, system setup is started automatically. The setup wizard takes you through the steps to configure basic system settings, such as time and date, system name, and hardware detection and verification.

7.7.2 V7000 Storage Node setup wizard

After the initial configuration process that is described in 7.7.1, “V7000 initial configuration” on page 132 is complete, the V7000 Storage Welcome window opens, as shown in Figure 7-52.



The image shows the IBM Flex System V7000 Storage Welcome window. At the top, the title "IBM Flex System V7000" is displayed in blue. Below the title is a 3D rendering of the V7000 storage node. To the right of the rendering are two input fields: "User name:" with the text "USERID" and "Password:" with a masked password of eight dots. A blue "Continue" button is located to the right of the password field. Below the rendering and login fields, the word "Welcome" is displayed in blue. Underneath, a message reads: "Welcome to IBM Flex System V7000. This wizard guides you through the initial configuration." Below this message, the text "Learn More" is displayed in blue. At the bottom, there are two links: "Visit the Information Center" with a right-pointing triangle icon and "View an e-Learning overview" with a television icon.

Figure 7-52 V7000 Storage Welcome window

Tip: During the initial setup of the Flex System V7000, the installation wizard prompts you for various information that you must have available during the installation process. If you do not have this information or choose not to configure some of the items now, you can configure them later through the GUI.

Complete the following steps:

1. Read and accept the license agreement, as shown in Figure 7-53. Click **Next**.

License Agreement (Step 1 of 7)

Read the license agreement carefully.

License IBM Notices Java Notices Non-IBM Licenses Additional Licenses and Notices

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON AN "ACCEPT" BUTTON, OR OTHERWISE USING THE PROGRAM, LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF LICENSEE, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND LICENSEE TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

* DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, CLICK ON AN "ACCEPT" BUTTON, OR USE THE PROGRAM; AND

* PROMPTLY RETURN THE UNUSED MEDIA, DOCUMENTATION, AND PROOF OF ENTITLEMENT TO THE PARTY FROM WHOM IT WAS OBTAINED FOR A REFUND OF THE AMOUNT PAID. IF THE PROGRAM WAS DOWNLOADED, DESTROY ALL COPIES OF THE PROGRAM.

1. Definitions

"Authorized Use" - the specified level at which Licensee is authorized to execute or run the Program. That level may be measured by number of users, millions of service units ("MSUs"), Processor Value Units ("PVUs"), or other level of use specified by IBM.

"IBM" - International Business Machines Corporation or one of its subsidiaries.

☒ I agree with the terms in the license agreement.

☐ I do not agree with the terms in the license agreement.

Next >

Figure 7-53 Setup wizard: License Agreement

2. Enter a System Name and Superuser Password, as shown in Figure 7-54. Click **Next**.

System Name and Superuser Password (Step 2 of 7)

System Name

* Name: FlexSystem_V7000

Superuser Password

New Superuser Password

Verify New Superuser Password

Next >

Figure 7-54 Setup wizard: Set system name and superuser password

3. Set up the system date and time, as shown in Figure 7-55. Click **Next**

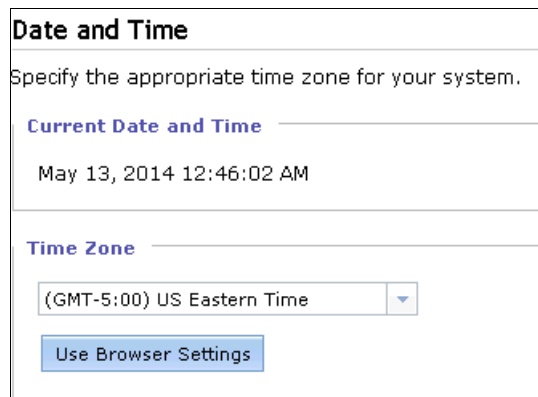


Figure 7-55 Set up wizard: Set Date and Time

4. Optionally, you can enter system licenses (as shown in Figure 7-56) and click **Next**. The system licenses include External Virtualization Limit, Remote-Copy Limit, and IBM Real-time Compression Limit. The virtualization license for all directly attached expansion enclosures is included in the system license and is not added here.

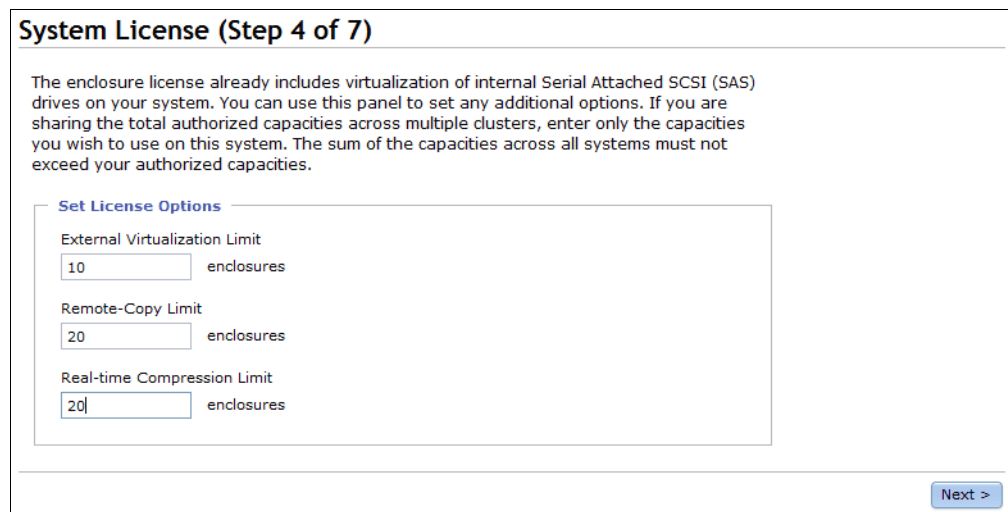


Figure 7-56 System license window

5. Configure the support notifications, as shown in Figure 7-57. Click **Next**.

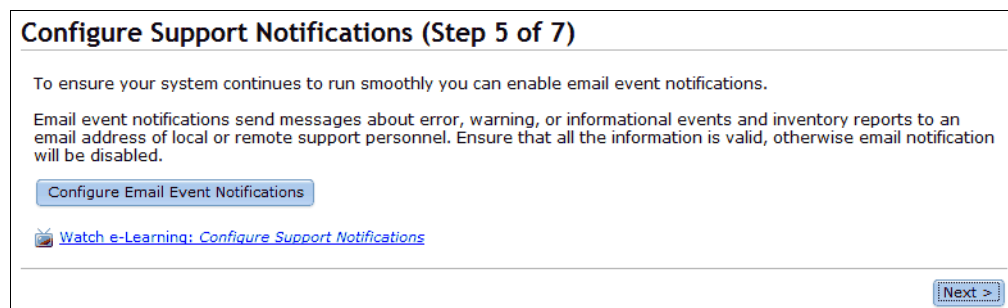


Figure 7-57 Configure Support Notifications window

6. Define any company contact information, as shown in Figure 7-58. Click **Next**.

Configure Support Notifications Step 1 of 4

Define Company Contact

Support personnel can contact this person to assist with problem resolution. Ensure that all contact information is valid.

Email Contact

* Contact Name	* Email Reply Address	
<input type="text"/>	<input type="text"/>	
* Machine Location	* Telephone (Primary)	Telephone (Alternate)
<input type="text" value="305"/>	<input type="text" value="9091234567"/>	<input type="text"/>

* Required

Next > Cancel

Figure 7-58 Define Company Contact window

7. Verify that all hardware was correctly detected by the system, as shown in Figure 7-59. Click **Next**.

Hardware (Step 6 of 7)

Actions

- FlexSystem_V7000
 - Enclosure 1
 - Drive Slots
 - Canisters**
 - Canister 1
 - Canister 2

Verify that all the installed hardware has been detected by the system. If the enclosure is not displayed, ensure it has been cabled correctly and is powered on.

Figure 7-59 Verify hardware

Note: Do not select Yes to automatically configure internal storage now because a customized storage layout is created if Yes is selected.

8. Click **Finish** to complete the setup wizard task and log in to V7000, as shown in Figure 7-60. You log in as a Superuser with your newly defined password. If you did not change the password, the default is passw0rd.



Figure 7-60 Setup wizard task complete

After a successful login, the V700 Home Overview window looks similar to the window that is shown in Figure 7-61.



Figure 7-61 V7000 Home Overview window

The V7000 initial configuration is complete and the cluster is up and running, as shown in Figure 7-62.

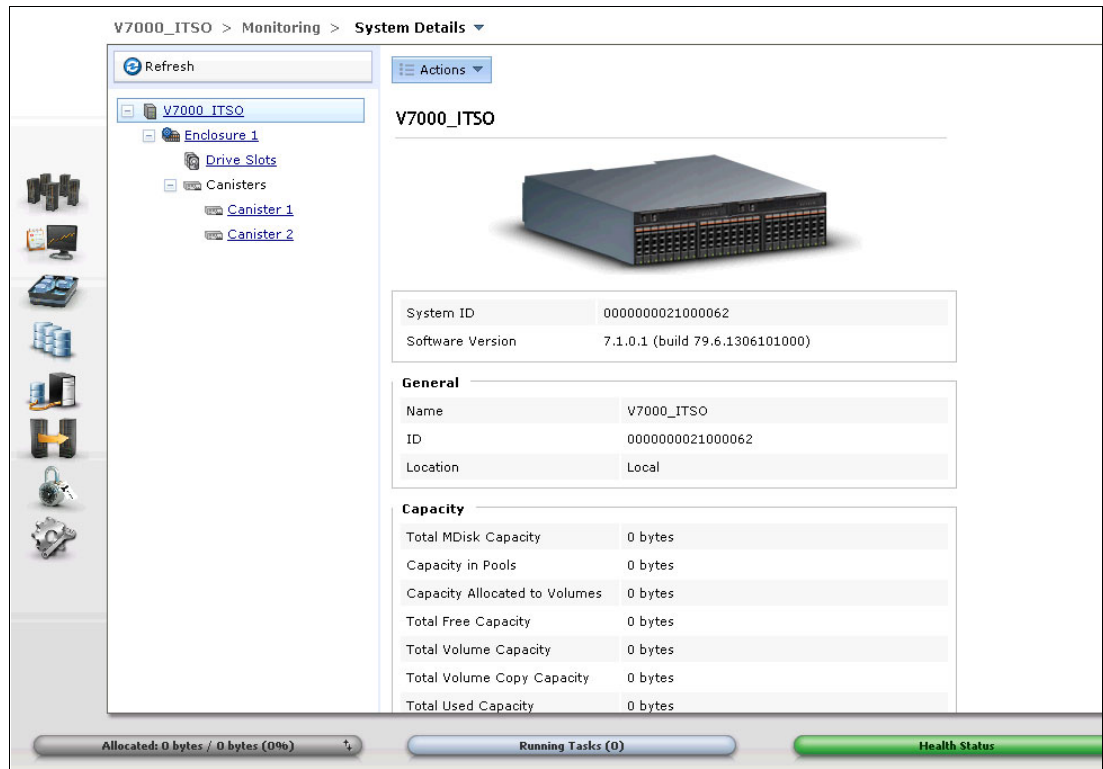


Figure 7-62 System Details view in the management GUI

9. You can continue to configure more functions and features for your environment to meet your implementation requirements.

7.7.3 Configuring storage volumes

Complete the following steps to configure the MDisk storage volumes:

1. Go back to the Overview window, as shown in Figure 7-63. Click the **Pools** icon. In the Pools menu, select **Internal Storage**.

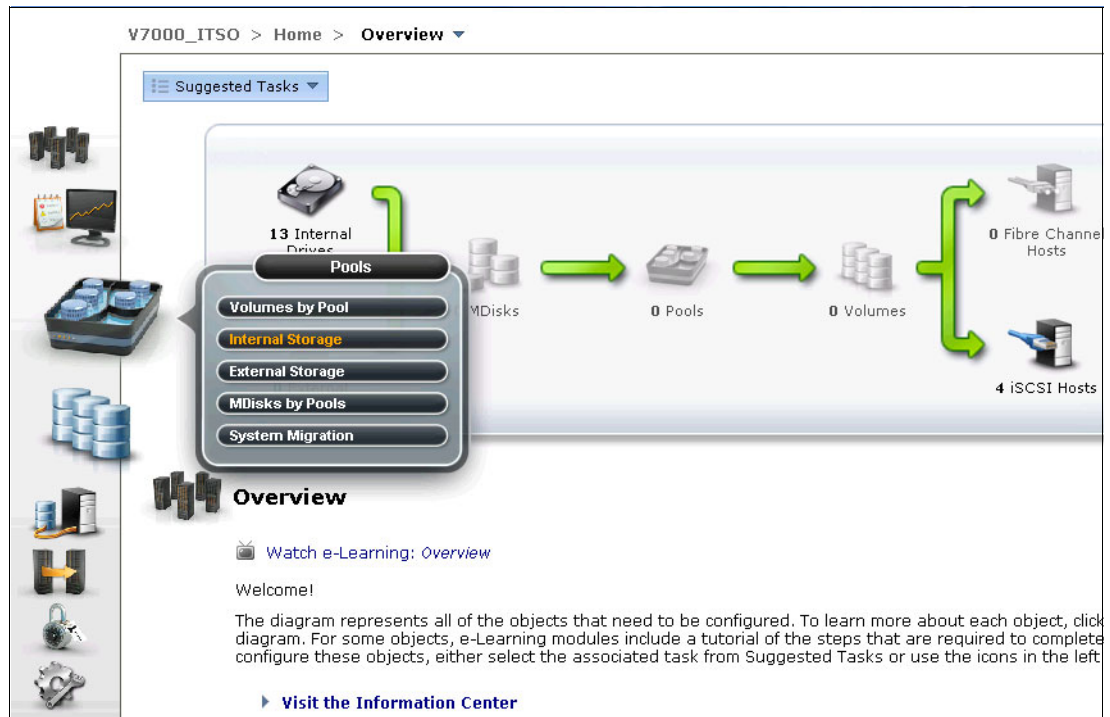


Figure 7-63 Flex System V7000: Overview

2. Click **Configure Storage**, as shown in Figure 7-64.

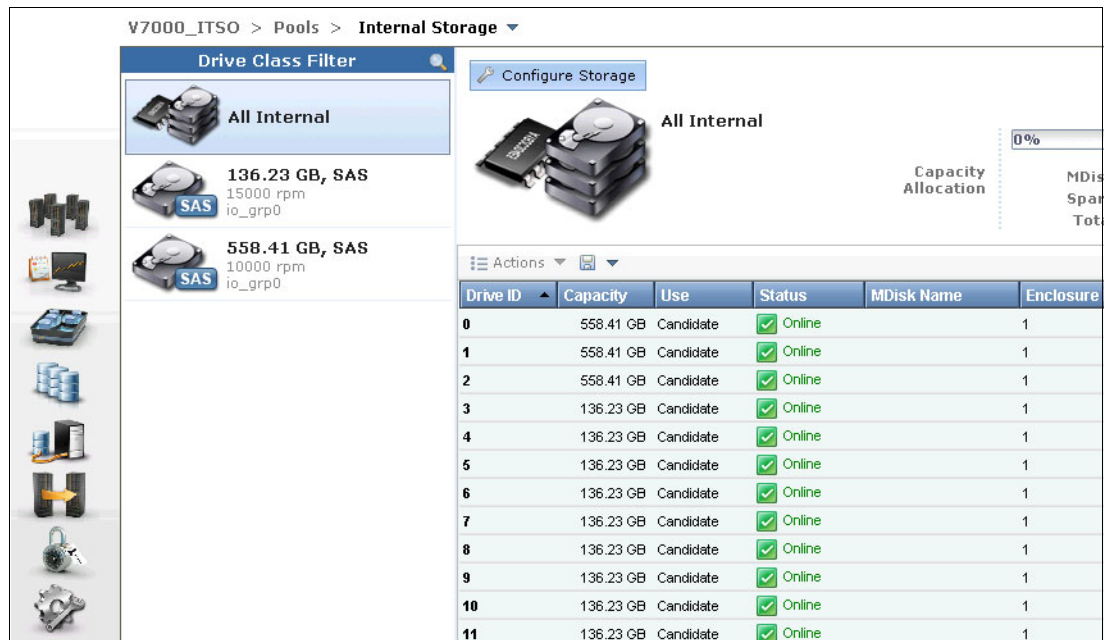


Figure 7-64 Flex System V7000: Internal Storage

- Choose **Select a different configuration**. In the Preset drop-down list, select **Basic RAID-5**. In the Number of drives to provision field, enter the number of drives (10, in our example). As shown in Figure 7-65, the RAID-5 is constituted by nine drives and one drive is a Hot Spare. Click **Next**.



Use this wizard to allocate RAID arrays to storage pools. After this configuration wizard completes, you can create volumes from these storage pools.

Storage Found:
(10 drives) 136.23 GB, SAS, 15000 rpm, io_grp0

☐ Use the recommended configuration: **Basic RAID-5**
Select this option to configure all available drives based on recommended values for the RAID level and drive class. The recommended configuration uses all the drives to build arrays that are protected with the appropriate amount of spare drives.

☒ **Select a different configuration**

Preset: Basic RAID-5

☒ **Automatically configure spares**

☐ Optimize for Performance

☐ Optimize for Capacity

10 Number of drives to provision

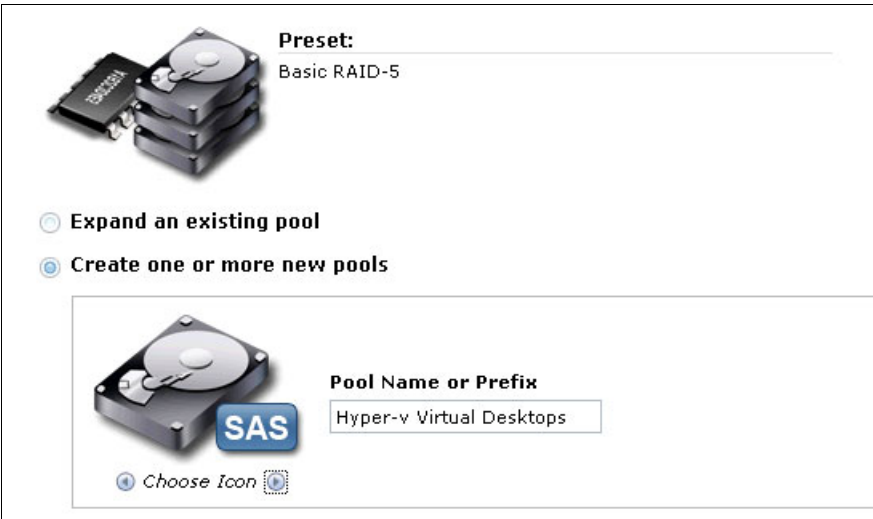
Configuration Summary:

2 x Basic RAID-5 (136.23 GB, SAS, 15000 rpm, io_grp0):

- 5, 4 drives
- 1 Hot Spares
- 0 Unconfigured Drives

Figure 7-65 Configure Internal Storage: RAID configuration

- Select **Create one or more new pools** and enter a Pool name, as shown in Figure 7-66. Click **Finish**.



Preset:
Basic RAID-5

☐ Expand an existing pool

☒ **Create one or more new pools**

Pool Name or Prefix
Hyper-v Virtual Desktops

Choose Icon

Figure 7-66 Configure Internal Storage: Creating a pool

5. When the task is completed, click **Close**, as shown in Figure 7-67.




Figure 7-67 Create RAID Arrays task window

Complete the following steps to configure the volumes:

1. Click the **Volumes** icon. In the Volumes menu, select **Volumes**, as shown in Figure 7-68.



Figure 7-68 Volumes creation

2. Select **New Volume**, as shown in Figure 7-69.




Figure 7-69 New Volume

3. Select **Thin-Provision**, as shown in Figure 7-70.




Figure 7-70 Preset selection

4. Select the pool that was created earlier that is named **Virtual Desktops**, as shown in Figure 7-71.

Select a Preset

Generic **Thin-Provision** Mirror Thin Mirror Compressed

Select a Pool

Name	Status	Free Capacity	Capacity
Hyper-v MGMT	Online	1.08 TB	1.08 TB
Hyper-v Virtual Desktops	Online	947.00 GB	947.00 GB

Figure 7-71 Pool selection

5. Create a volume that is named Shared for Management Volume and set the size to 500. Select **GB**, as shown in Figure 7-72. Click **Create**.

Select a Preset

Generic **Thin-Provision** Mirror Thin Mirror Compressed

Select a Pool

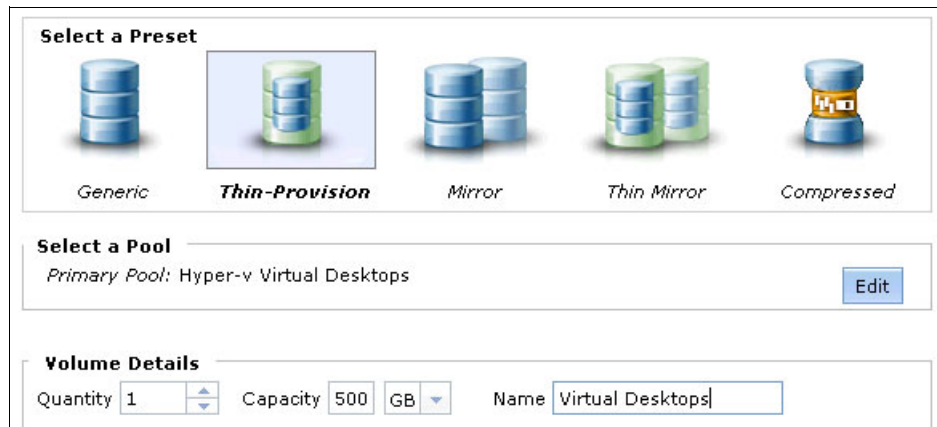
Primary Pool: Hyper-v MGMT Edit

Volume Details

Quantity Capacity Name

Figure 7-72 Select Names and Sizes

6. Click **Close** in the Create Volumes window when the task is completed.
7. Create a volume that is named Virtual desktops Volume and set the size to 500. Select **GB**, as shown in Figure 7-73 on page 144. Click **Create**.



Select a Preset

Generic **Thin-Provision** Mirror Thin Mirror Compressed

Select a Pool

Primary Pool: Hyper-v Virtual Desktops Edit

Volume Details

Quantity Capacity GB Name

Figure 7-73 Virtual desktops Volume creation

7.7.4 Configuring hosts

Use the following steps to configure the hosts:

1. Click the **Hosts** icon. In the menu, select **Hosts**, as shown in Figure 7-74.

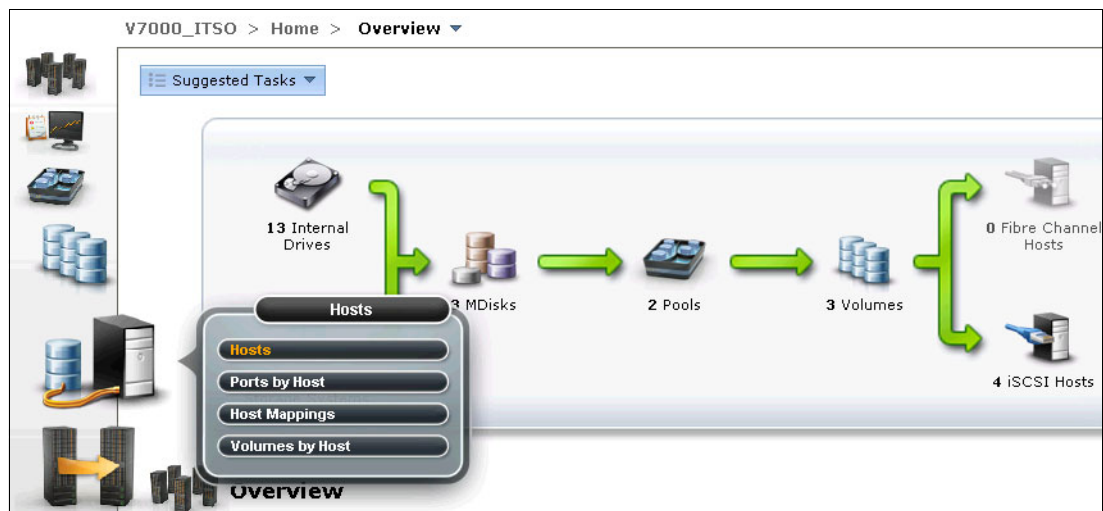


Figure 7-74 Hosts

2. Click **New Host**, as shown in Figure 7-75.




Figure 7-75 New Host

3. Choose Fibre Channel Host or iSCSI Host, depending on your storage connectivity protocol, as shown in Figure 7-76. In our example, we select **iSCSI Host**.



Figure 7-76 Choose the Host Type

The process of adding an FC host is similar to adding an iSCSI host.

4. Specify the Host Name. We entered hyperv03. Copy the iSCSI Qualified Name (IQN) of your VDI host and paste it in the iSCSI Ports field and click **Add Port to List**. Click **Create Host**, as shown in Figure 7-77. Repeat this step for all hosts.

Figure 7-77 Create Host

- To modify the host mappings, select a host, right-click to display the menu, and select **Modify Mappings**, as shown in Figure 7-78.

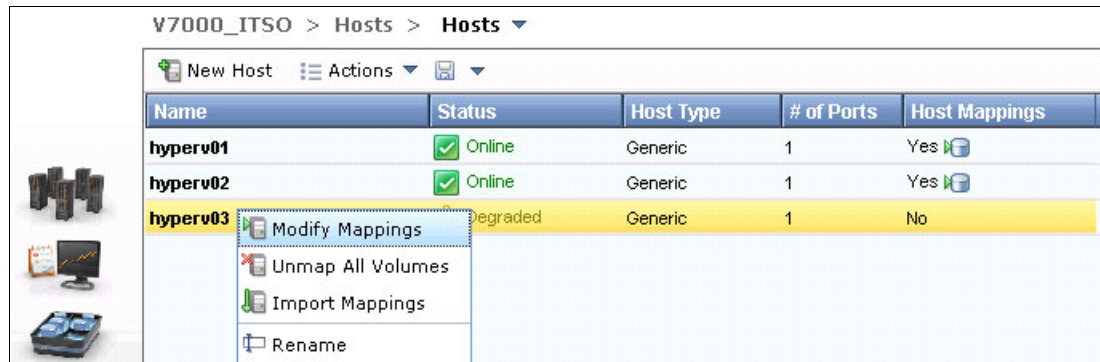


Figure 7-78 Modify Mappings

- Assign the needed volumes to each host, and click **Apply**, as shown in Figure 7-79.




Figure 7-79 Modify Host Mappings

Deploying Citrix XenDesktop

In this chapter, we describe how to provision virtual machines (VMs) for Citrix XenDesktop components and how to install the Citrix XenDesktop components.

This chapter includes the following topics:

- ▶ 8.1, “Configuring utility services” on page 148
- ▶ 8.2, “Provisioning VMs for Citrix XenDesktop components” on page 149
- ▶ 8.3, “Installing Citrix XenDesktop Controller” on page 153
- ▶ 8.4, “Installing Citrix XenApp” on page 163
- ▶ 8.5, “Installing Citrix StoreFront” on page 167
- ▶ 8.6, “Installing Citrix Provisioning Services” on page 171

8.1 Configuring utility services

To save time, we used a template that was created with System Center Virtual Machine Manager (SCVMM) or vCenter to deploy multiple VMs. A *template* is a master copy of a VM that can be used to create and provision VMs. In this scenario, we performed a classic build for the Microsoft Windows 2012 R2 server in a VM named W2K12R2 and converted it in the template.

Notes: For more information about templates in Hyper-V environments, see the following Microsoft documentation:

<http://bit.ly/1gn5mRf>

For more information about templates in VMware vSphere environments, see the following VMware documentation:

<http://bit.ly/16MDNGQ>

In this scenario, the template W2K12R2 is used to build the necessary VMs for installing network and utility services.

A summary of VMs and their characteristics is listed in Table 8-1.

Table 8-1 Network and utility services VMs

VM name	vCPU (number)	RAM (GB)	VMDK (GB)	Network (VLAN)	Purpose
AD	2	4	30	20	Domain Controller and Remote Desktop Licensing Manager
FS	2	4	30 + 10	20	File server
SQL	4	4	15	20	MS SQL 2008 R2
SCCM	4	4	15	42, 20	SCVMM
Win7	1	2	30	42, 20	User desktop

Notes: Consider the following points:

- ▶ Installing and configuring Windows Server roles, such as Domain Controller, Internet Information Services (IIS), or File Server, are beyond the scope of this book and are not documented.
- ▶ Microsoft System Center Virtual Machine Manager and vCenter installation is not documented in this book.
- ▶ For more information about installing Microsoft System Center Virtual Machine Manager, see this website:
<http://bit.ly/1o9QCb7>
- ▶ For more information about installing vCenter Server 5.x, see this website:
<http://bit.ly/19ut0gp>

8.2 Provisioning VMs for Citrix XenDesktop components

A summary of VMs for installing Citrix components is listed in Table 8-2.

Table 8-2 Citrix XenDesktop VMs

VM name	vCPU	RAM (GB)	VMDK (GB)	VLAN	Purpose
XLIC	2	4	15	20	License server
XDC	4	4	15	20	XenDesktop Controller
XAP	2	4	30	20	XenApp
XFS	4	4	15	20	FrontStore
PVS	4	32	40	20, 30	Provisioning services

8.2.1 Installing the Citrix License Server

In this section, we describe how we installed the Citrix License Server on a separate VM.

Note: In this scenario, License Server Version 11.11 is installed. For more information about Citrix documentation for License Server, see this website:

<http://bit.ly/1hcpws8>

Complete the following steps to install License Server Version 11.11:

1. Mount the remote media on the XLIC VM. After mounting the Citrix XenDesktop installation media, the AutoRun window opens. Select **Start** from XenDesktop section, then click **Get Started** from Delivery Controller.
2. Accept the licensing agreement and click **Next**.
3. Select **License Server** and clear the other options, as shown in Figure 8-1. Click **Next**.

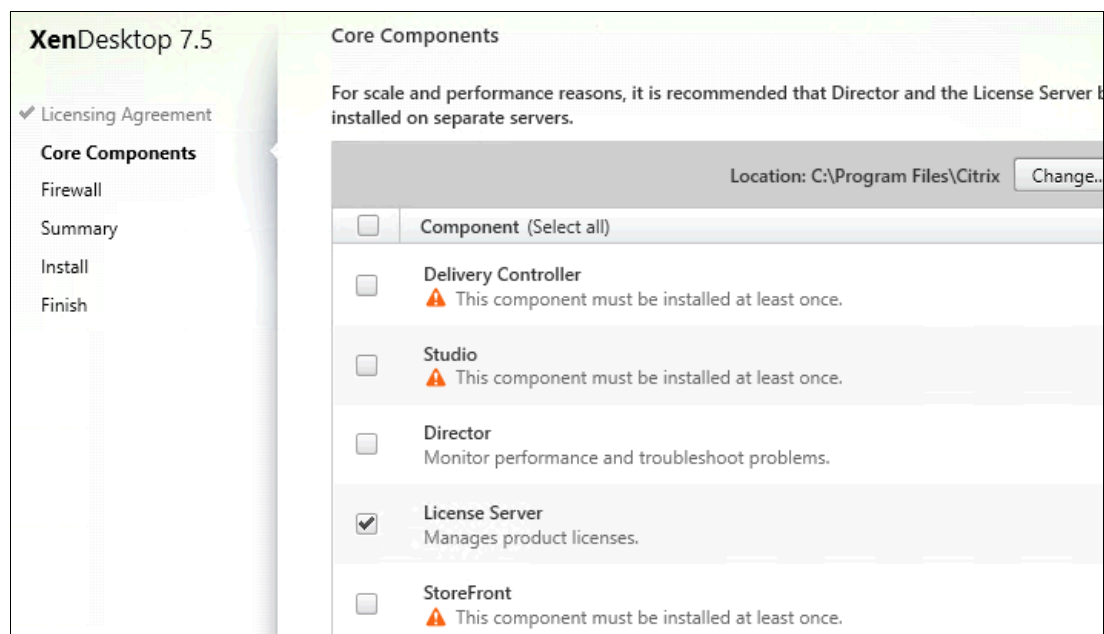


Figure 8-1 Selecting components to install

4. Select **Automatically** configure firewall rules to allow TCP Ports 27000, 7279, 8083, and 8082 to be used for License Server connections, as shown in Figure 8-2. Click **Next**.

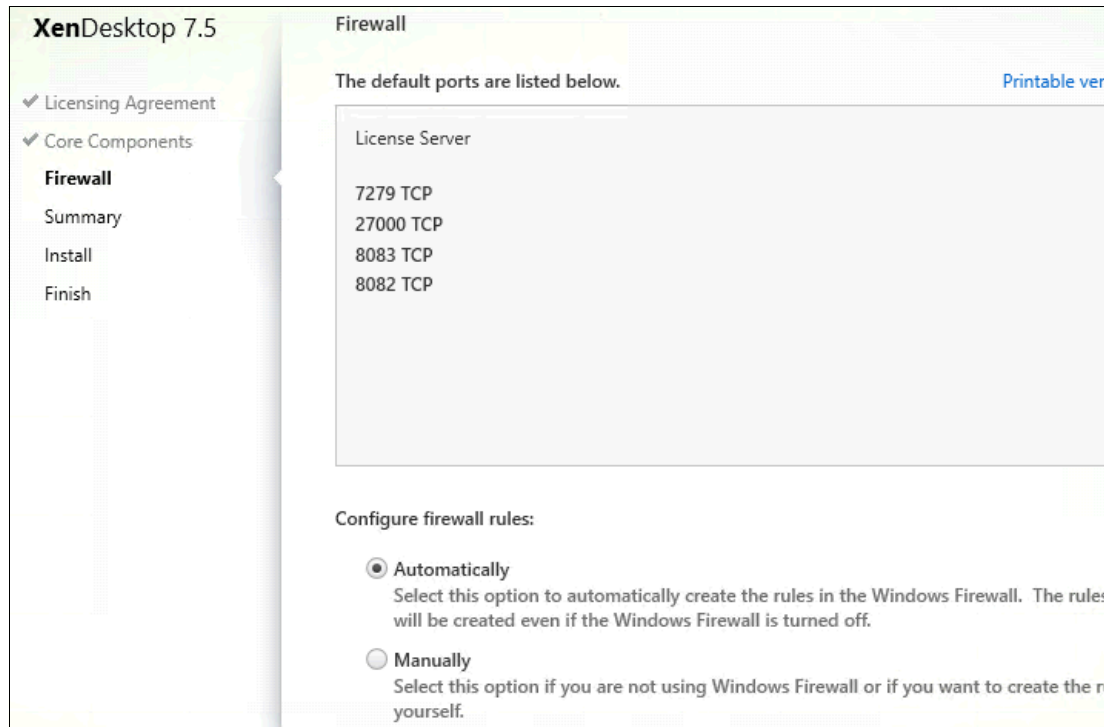


Figure 8-2 Firewall Configuration window

5. Review the summary for the installation, as shown in Figure 8-3. Click **Install**.



Figure 8-3 Summary window

6. Allow the Setup Wizard to complete the installation. After the installation is complete, a final summary is displayed, as shown in Figure 8-4. Click **Finish**.



Figure 8-4 Installation successfully completed

8.2.2 Configuring the licenses

After the License Server is installed, the licenses must be configured. The setup process creates a shortcut to the License Administrator Console and opens a browser to `https://localhost:8082`, as shown in Figure 8-5.

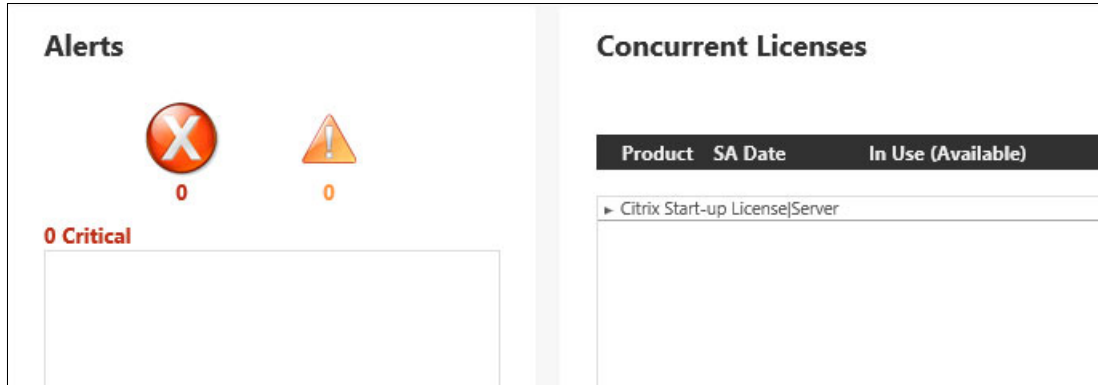


Figure 8-5 Citrix Licensing Console window

Complete the following steps to configure the licenses:

1. Select **Administration**.
2. After entering the credentials that are used for installation, select **Vendor Daemon Configuration**. Vendor Daemon Configuration window opens, as shown in Figure 8-6.

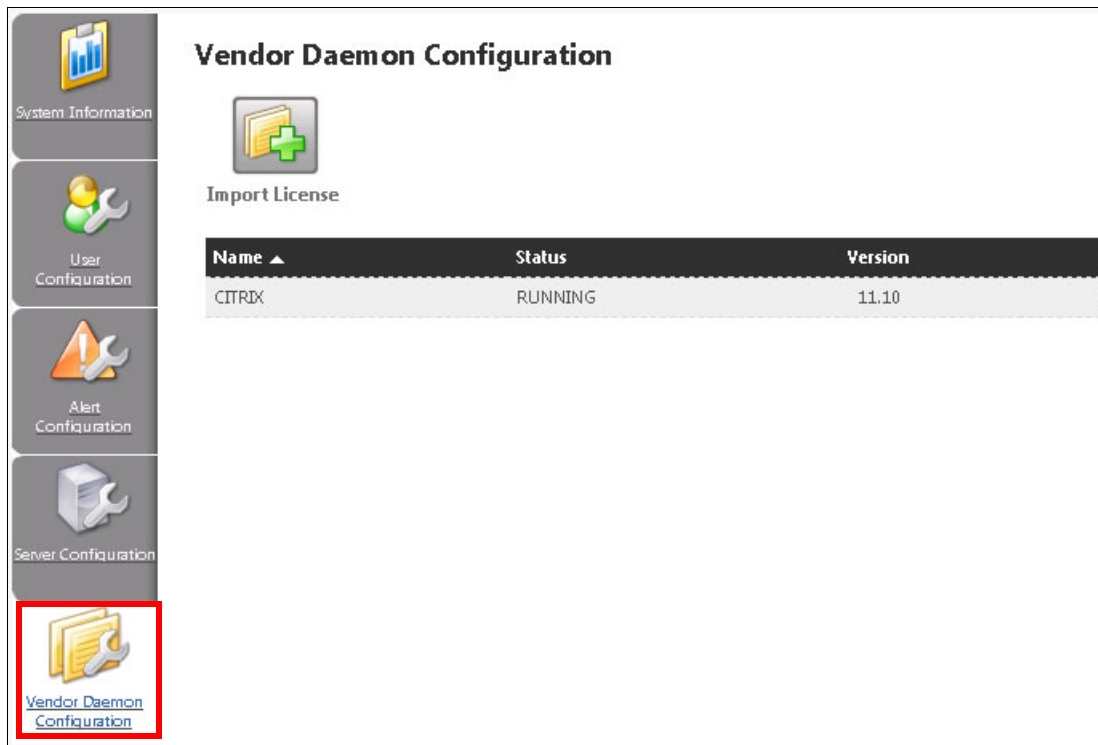


Figure 8-6 License Server Vendor Daemon Configuration window

3. Import the license file by selecting **Import License File**. Click **Browse** to locate the correct file and then select **Import License** (see Figure 8-7).

Import License File

License File from Your Local Machine:
 C:\Users\svcd75\Desktop\FID__398b9677_145ee7dec Browse...

☐ Overwrite License File on License Server

Import License **Cancel**

1. Allocate your license from [My Citrix](#).
2. Type the path or use the Browse button to locate the license file that you copied. If the file has the same name as an existing one, or if you copied the file directly to the MyFiles directory, select the Overwrite License File on License Server check box.
3. Click Import License, then OK. The License Administration Console copies the file from its existing location into the MyFiles directory where it can be read by the license server.
4. Click the Administer link in the Citrix vendor daemon line.
5. Click Reread License Files to allow the license server to recognize the new file.

Figure 8-7 Import License File window

Note: It is not in the scope of this book to describe the process of obtaining the license file from the Citrix website. However, it is important to remember that the name of the server on which the licensing components are installed is encoded in the license file and is case-sensitive.

4. Restart the vendor daemon service and select **Reread License Files** (see Figure 8-8).

Vendor Daemon: CITRIX

Vendor Daemon Port in Use: 7279

Vendor Daemon Actions:

Stop **Reread License Files**

Report Log Name: **Rotate Report Logs**

General Configuration

* License File or Directory:
 C:\Program Files (x86)\Citrix\Licensing\MyFiles
 C:\Program Files (x86)\Citrix\Licensing\MyFiles\FID__398b9677_145ee7dec_2017.lic

Figure 8-8 Reread License Files window

5. The message “The license file was successfully reread” displays, as shown in Figure 8-9.

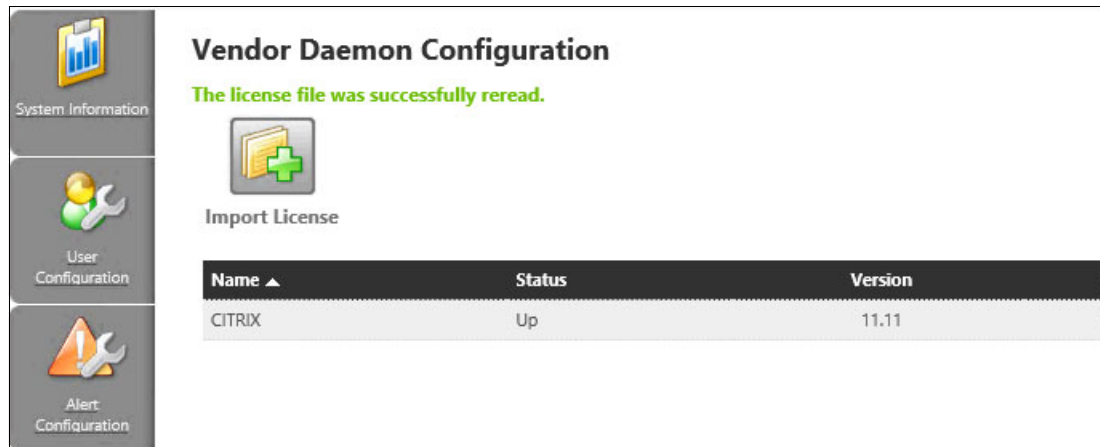


Figure 8-9 License Server with license

8.3 Installing Citrix XenDesktop Controller

The Desktop Delivery Controller brokers the connections between the user and the virtual desktop, and the creation and management of virtual desktops on the provisioning and hypervisor infrastructures. The controllers enumerate resources for the users and direct user launch requests to the appropriate virtual desktop.

Note: Ensure that you install this component on a separate VM.

If multiple XenDesktop Controllers are planned, complete the next procedure for each controller.

8.3.1 Installing the XenDesktop Controller

Complete the following steps to install the XenDesktop Controller:

1. Mount the remote media on the XDC VM. After mounting the Citrix XenDesktop 7.5 installation media, the AutoRun window opens. Select **Get Started** from Delivery Controller section.
2. Accept the license agreement and click **Next**.
3. Select the components that you want to install and clear the other option selections (see Figure 8-10). Click **Next**.

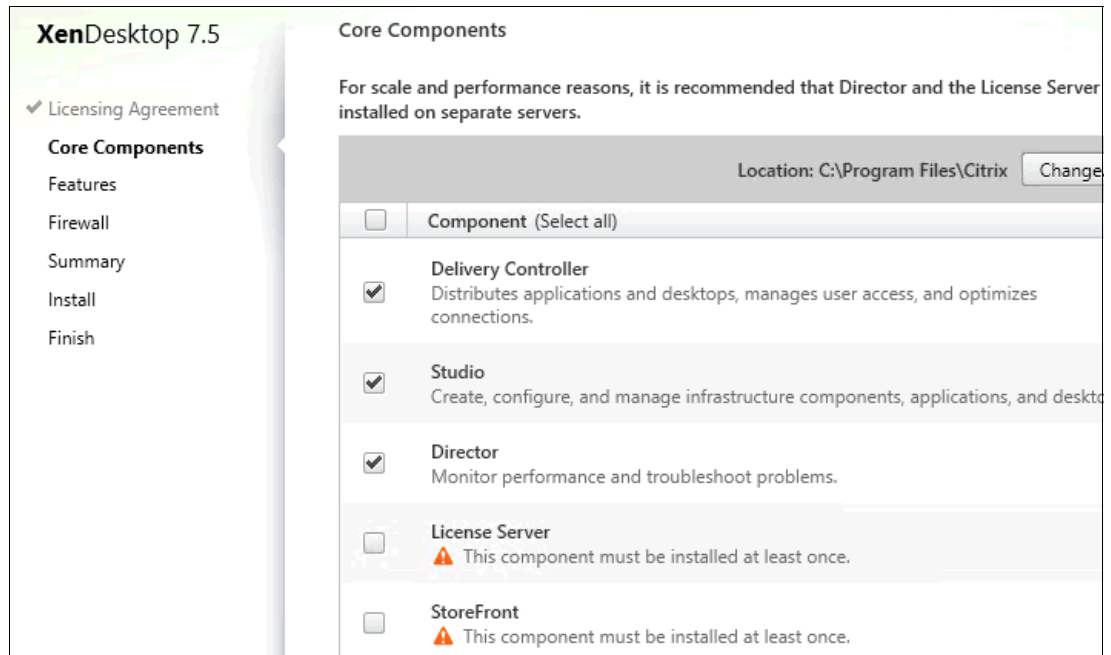


Figure 8-10 Selecting components to install

4. The Feature window opens, as shown in Figure 8-11. In our example, we have a dedicated SQL instance; therefore, clear **Install Microsoft SQL Server 2012 Express**. Click **Next**.

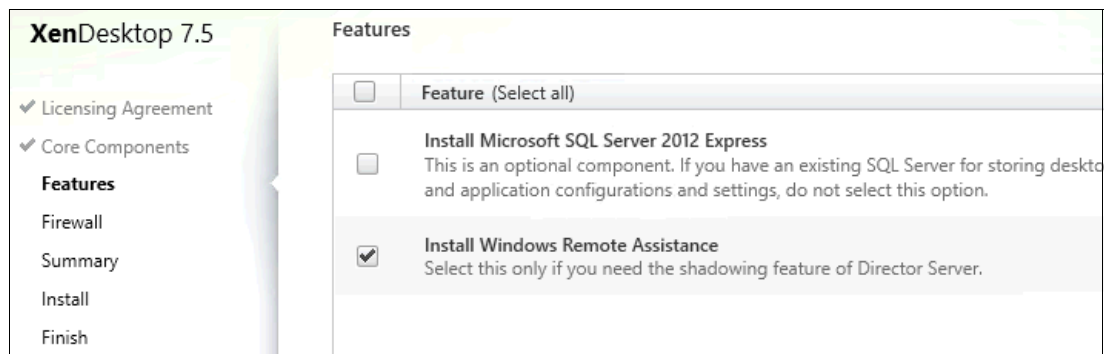


Figure 8-11 XenDesktop Features window

5. The Firewall Configuration window opens, as shown in Figure 8-12. Click **Next**.

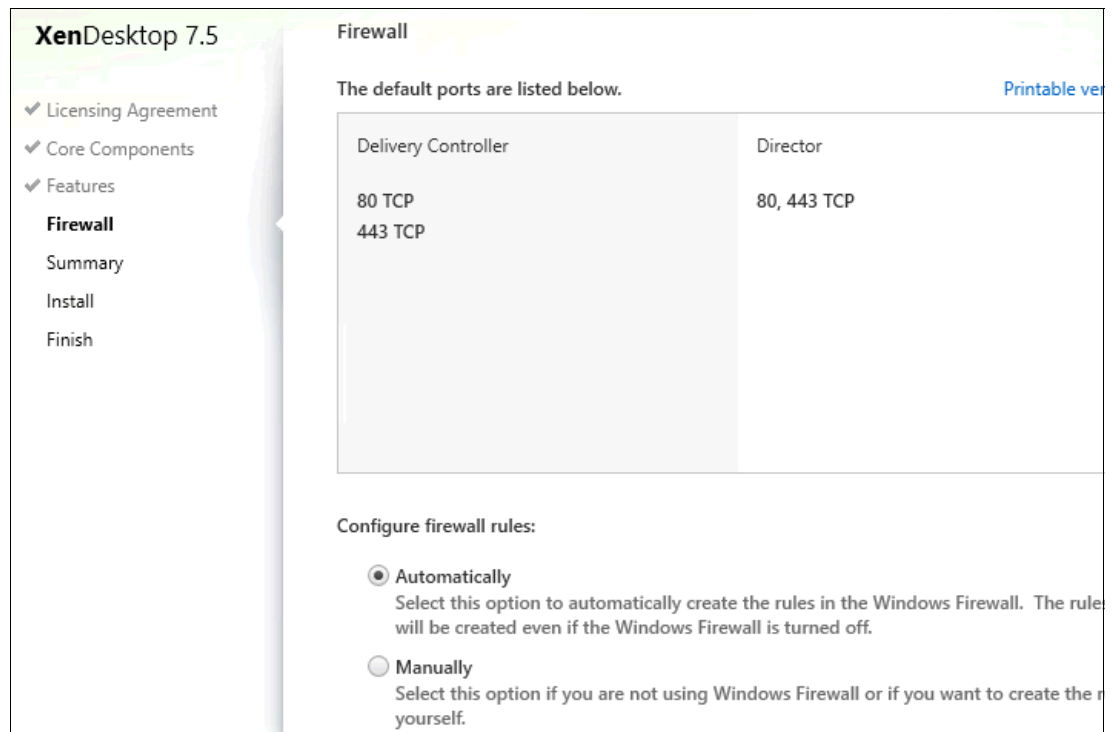


Figure 8-12 Controller Firewall Configuration window

6. Review the summary for the installation, as shown in Figure 8-13. Click **Install**.

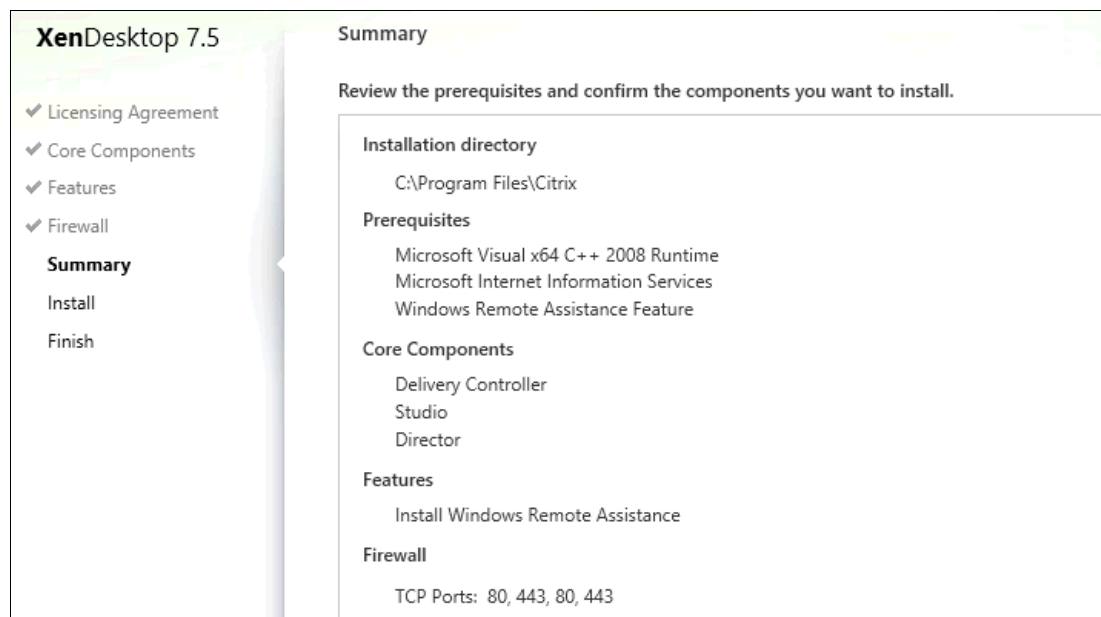


Figure 8-13 Controller Summary window

7. Allow the installation process to finish. If you want to configure Machine Provisioning Services, select **Configure XenDesktop after closing**, as shown in Figure 8-14. Click **Close**.

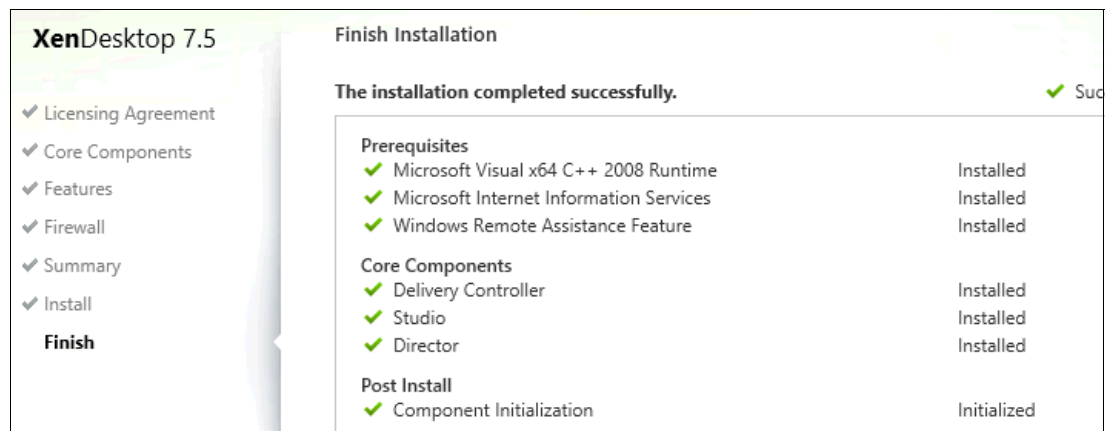


Figure 8-14 Controller Installation Successful window

8. If you chose to configure XenDesktop, the Citrix Desktop Studio console is automatically started. Alternatively, you can select **Start** → **Citrix** → **Desktop Studio**. Select **Desktop deployment**, as shown in Figure 8-15.

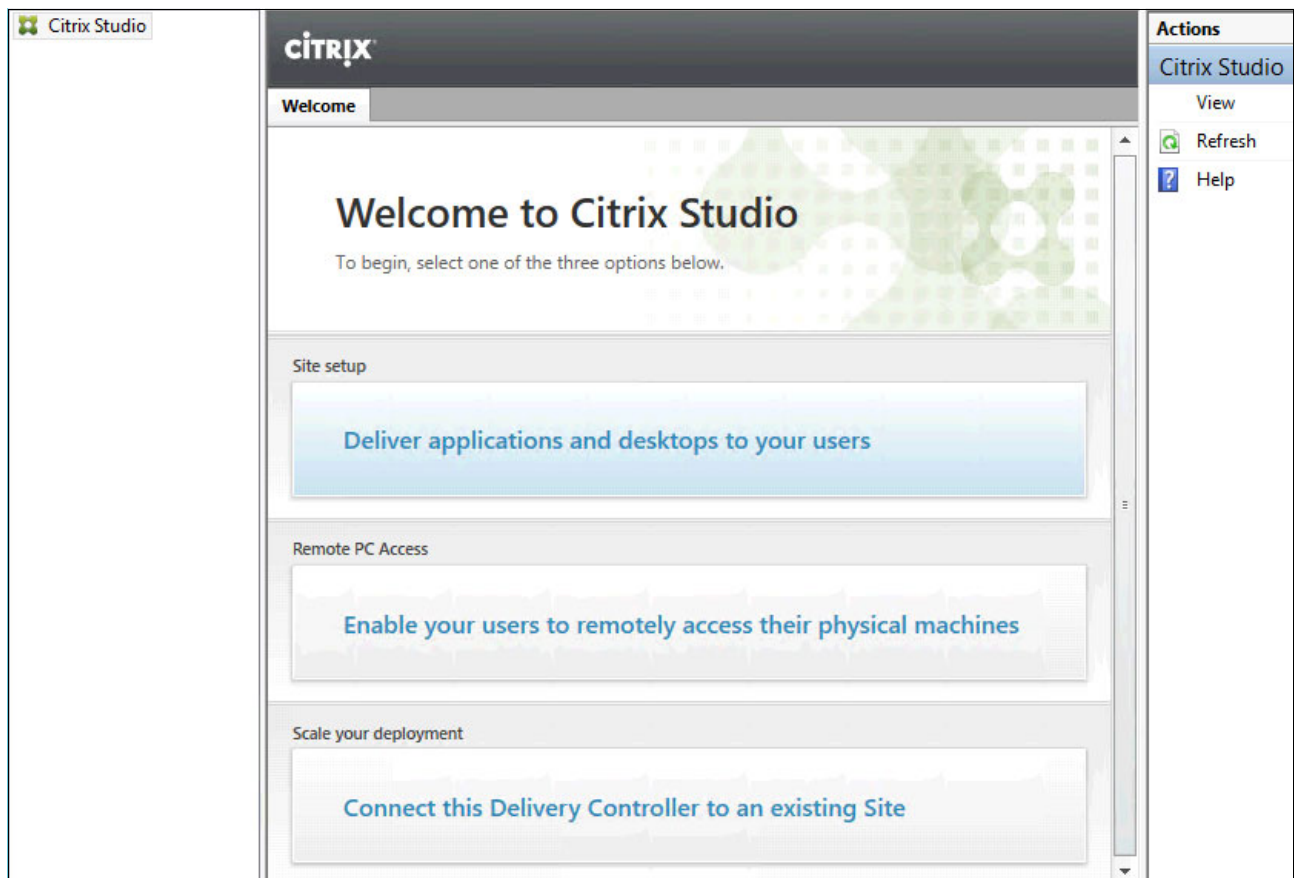


Figure 8-15 Citrix Desktop Studio's Welcome window

9. The initial configuration starts. Specify the site name, as shown in Figure 8-16. Click **Next**.

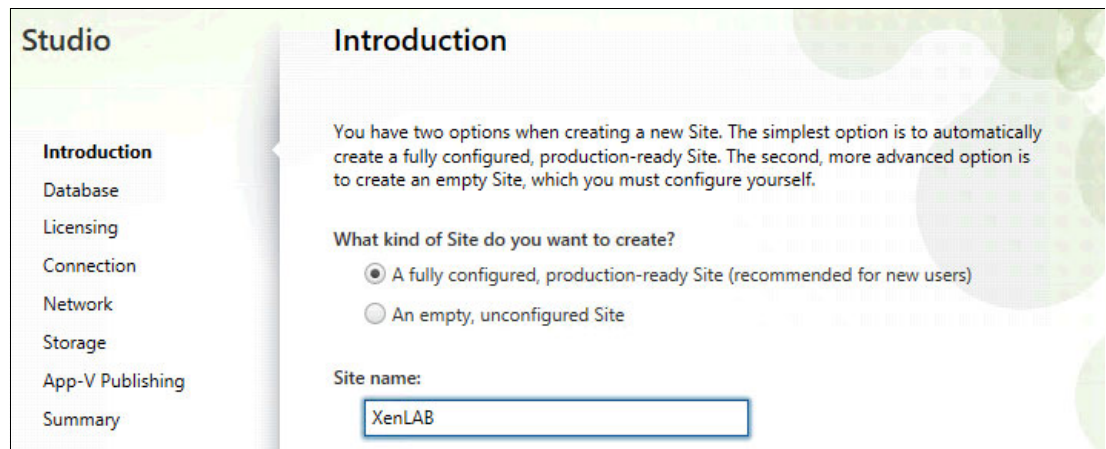
The screenshot shows the 'Studio' application window with the 'Introduction' tab selected in the left sidebar. The main area is titled 'Introduction' and contains text explaining the two options for creating a new Site: a fully configured, production-ready Site (recommended) or an empty, unconfigured Site. Below this, there are two radio buttons. The first radio button is selected and is labeled 'A fully configured, production-ready Site (recommended for new users)'. The second radio button is labeled 'An empty, unconfigured Site'. Below the radio buttons, there is a 'Site name:' label followed by a text input field containing the text 'XenLAB'.

Figure 8-16 Controller site setup Introduction configuration

10. On the database page, complete the Database server location and Database name fields. Select **Test connection** and enter the credentials that are required for database connection, as shown in Figure 8-17. Click **OK**.

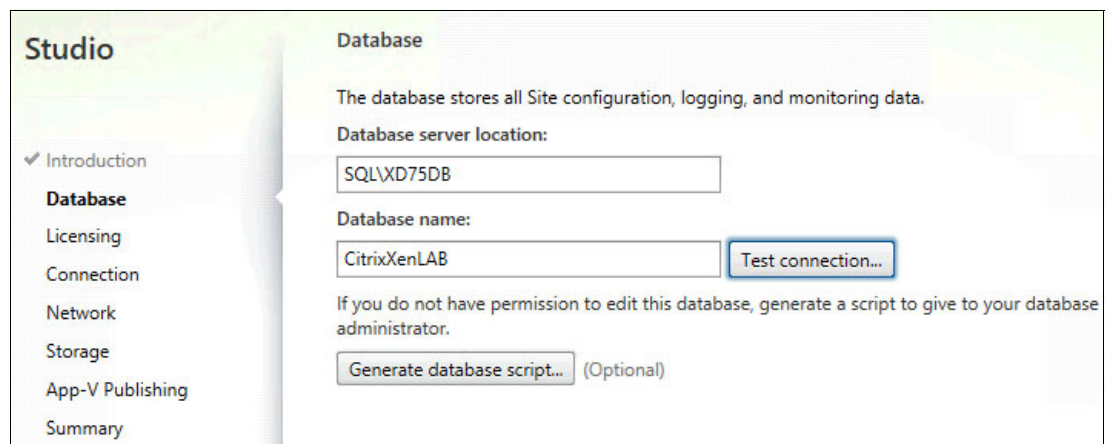
The screenshot shows the 'Studio' application window with the 'Database' tab selected in the left sidebar. The main area is titled 'Database' and contains text explaining that the database stores all Site configuration, logging, and monitoring data. Below this, there are two text input fields. The first is labeled 'Database server location:' and contains the text 'SQL\XD75DB'. The second is labeled 'Database name:' and contains the text 'CitrixXenLAB'. To the right of the 'Database name:' field is a button labeled 'Test connection...'. Below these fields, there is a paragraph of text: 'If you do not have permission to edit this database, generate a script to give to your database administrator.' Below this text is a button labeled 'Generate database script...' followed by the text '(Optional)'.

Figure 8-17 Database connection window

11. If the test is successful, the message “All database connection tests passed” is displayed. Click **Close**, and then click **Next**.

12. Configure the license server. Set your XenDesktop edition and licensing model (as shown in Figure 8-18) or you can add license files and edit your licensing model later by using XenDesktop Studio. Click **Next**.

Figure 8-18 License server configuration window

The wizard shows the connection window. Here, we can choose the connection type depending on technology that is used for virtualization, as shown in Figure 8-19.

Figure 8-19 Controller Connection Type window

Note: The next steps are similar for VMware ESXi and Microsoft Hyper-V. Our scenario is based on Hyper-V though.

13. Complete the following steps to configure the type and connection details for Microsoft Hyper-V virtualization layer (see Figure 8-20 on page 159):
 - a. Choose **Microsoft System Center Virtual Machine Manager** as the connection type.
 - b. Enter the Fully Qualified Domain Name (FQDN) in the Address field.
 - c. Enter a user name and password that is used for connection to SCVMM.
 - d. Enter a connection name.

- e. Select **Studio Tools (Machine Creation Services)**, as shown in Figure 8-20.
- f. Click **Next**.

Studio

✓ Introduction

✓ Database

✓ Licensing

Connection

Resources

Storage

App-V Publishing

Summary

Connection

Select a Connection type. If machine management is not used (for example when using physical hardware), select 'No machine management.'

Connection type:

Microsoft® System Center Virtual Machine Mana...

Connection address:

scmm.xenlab.local

User name:

xenlab\VMMAdmin

Password:

.....

Connection name:

SCVMM Connection

The Connection name appears in Studio; it helps administrators identify the Connection.

Create virtual machines using:

☒ Studio tools (Machine Creation Services)

☐ Other tools

Figure 8-20 Controller connection to Hyper-V virtualization layer

14. In the Resource section, enter the resource name, then select the cluster that is used for desktops (in our example, Non-persistent Desktops), as shown in Figure 8-21. Select the network for the VMs to use (in our example, **Virtual Switch VLAN 20**) and click **Next**.

Studio

✓ Introduction

✓ Database

✓ Licensing

✓ Connection

Resources

Storage

App-V Publishing

Summary

Name for these resources:

HyperV Cluster

Cluster

Select a cluster for the new virtual machines.

HYPERV_Cluster

Browse...

Select one or more networks for the virtual machines to use:

☒ Virtual Switch VLAN20

☐ Virtual Switch VLAN30

☐ Virtual Switch VLAN42

Figure 8-21 Selecting the cluster and network for desktops

15. Continue with selecting data stores that are used for desktops (see Figure 8-22). Use the default setting in the Personal vDisk storage section. Click **Next**.

Studio

- ✓ Introduction
- ✓ Database
- ✓ Licensing
- ✓ Connection
- ✓ Resources
- Storage**
- App-V Publishing
- Summary

Storage

Select one or more storage devices for the new virtual machines:

Shared

Name
<input checked="" type="checkbox"/> CSVVolume2 on HYPERVCluster.xenlab.local
<input type="checkbox"/> CSVVolume1 on HYPERVCluster.xenlab.local

Personal vDisk storage (Desktop OS only): [Learn more](#)

☒ Use same storage for virtual machines and Personal vDisks
☐ Use different storage for Personal vDisks

Select storage... (None selected)

Figure 8-22 Selecting data stores

16. Leave the information in the App-V Publishing window unchanged. Click **Next**. Review the summary and click **Finish** (see Figure 8-23).

Studio

- ✓ Introduction
- ✓ Database
- ✓ Licensing
- ✓ Connection
- ✓ Resources
- ✓ Storage
- ✓ App-V Publishing
- Summary**

Summary

Site name:	XenLAB
Database server:	SQL\XD75DB
Database name:	CitrixXenLAB
License server:	xlic.xenlab.local
Connection type:	Microsoft® System Center Virtual Machine Manager
Connection address:	sccm.xenlab.local
Connection name:	SCVMM Connection
Create virtual machines with:	Studio tools (Machine Creation Services)
Networks:	Virtual Switch VLAN20
Virtual Machine storage:	CSVVolume2 on HYPERV_Cluster.xenlab.local
Personal vDisk storage:	Use same storage as Virtual Machines
App-V:	Not configured

Figure 8-23 Initial configuration summary window

The Desktop Studio displays the results that are shown in Figure 8-24.

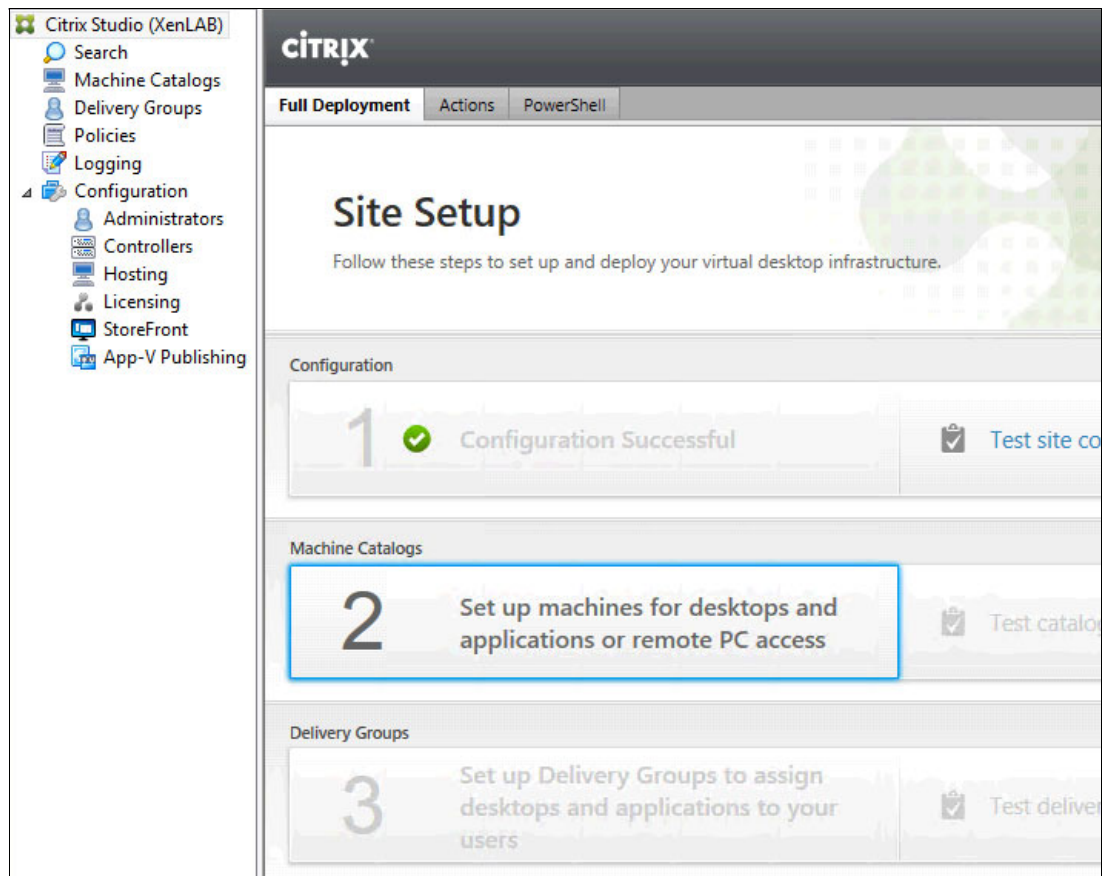


Figure 8-24 Desktop Studio window

8.3.2 Advanced settings

To finish the initial configuration, the following advanced settings are needed:

- Store Controller information in Active Directory (AD)

This setting stores the FarmGUID and connector information in AD so that when the virtual desktop agent is loaded, the FarmGUID can be pulled from AD and is not entered manually.

To perform AD-based controller discovery, run the PowerShell script `Set-ADControllerDiscovery.ps1` that is installed on each controller in the directory `$Env:ProgramFiles\Citrix\Broker\Service\Setup Scripts` (see Figure 8-25 on page 162). The script must be run on a controller in the site by a user who is a full administrator of the controller and who has the appropriate permissions to make changes in the relevant organizational unit (OU) in AD.

- Domain Name System (DNS) server resolution of desktops

This setting is an enabled or disabled setting that is disabled by default and helps with desktop discovery. Run the **`Set-BrokerSite -DnsResolutionEnabled 1`** command in PowerShell SDK (see Figure 8-25 on page 162).

- Trust XML service requests

This setting is an enabled or disabled setting that is disabled by default. Enter the **Set-BrokerSite -TrustRequestsSentToTheXmlServicePort \$true** command in PowerShell SDK (see Figure 8-25).

```
PS C:\Program Files\Citrix\Broker\Service\Setup Scripts> .\Set-ADControllerDiscovery.ps1 -on -existingOUDn "OU=XenLAB, OU=Resources, DC=xenlab, DC=local"
0
1
Important: You must restart the Citrix Broker Service on all DDCs to complete the configuration change
Important: You must make sure that your UDAs are using the correct OU (e.g. set HKLM\SOFTWARE\Citrix\VirtualDesktopAgent\FarmGUID to "0c0180cb-e412-4e5a-900c-6b47d76d7db9")
PS C:\Program Files\Citrix\Broker\Service\Setup Scripts> _
```

Figure 8-25 Controller advanced settings in PowerShell SDK

- Change the default port of 80 for XML

Port 80 is used by other services and processes and that can sometimes lead to conflicts. It is a good idea to change this port; Citrix suggests the use of 8082. Enter the **BrokerService.exe -wiport 8082** command in a command prompt that is opened with administrative-level rights (see Figure 8-26).

```
PS C:\Windows\system32> cd 'C:\Program Files\Citrix\Broker\Service'
PS C:\Program Files\Citrix\Broker\Service> .\BrokerService.exe /show
SDK Port: 80
UDA Port: 80
WI Port: 80
WI SSL Port: 443
Log File:
PS C:\Program Files\Citrix\Broker\Service> .\BrokerService.exe -wiport 8082
Stopping service: CitrixBrokerService
Starting service: CitrixBrokerService
Command completed successfully
PS C:\Program Files\Citrix\Broker\Service> .\BrokerService.exe /show
SDK Port: 80
UDA Port: 80
WI Port: 8082
WI SSL Port: 443
Log File:
PS C:\Program Files\Citrix\Broker\Service> _
```

Figure 8-26 Controller advanced settings in a command prompt

Note: For more information about the Active Directory OU configuration for XenDesktop, see the following eDoc:

<http://bit.ly/ljxMBdf>

For more information about enabling the DNS server resolution for XenDesktop, see this website:

<http://bit.ly/1k89iW7>

For more information about how to change communication ports in XenDesktop 7.x, see this website:

<http://bit.ly/SnCQUX>

8.4 Installing Citrix XenApp

In version 7.5, XenApp and XenDesktop share a unified architecture and management. XenDesktop and XenApp achieve true abstraction of the operating system, applications, and user data.

Note: The availability of some features depends on product edition and licenses. For more information about differences between XenApp 7.5 and previous versions, see this website:

<http://bit.ly/lqXXZ7i>

Complete the following steps to install Citrix XenApp on a separate VM:

1. Connect to the server that is selected to be the XenApp Server, browse to the location of the XenDesktop Installer Media, and start the installer media by using AutoPlay. Select **Start** in the XenApp Deliver Applications section and then select **Delivery Controllers**.
2. After you accept the license agreement, select the components that are to be installed, as shown in Figure 8-27.

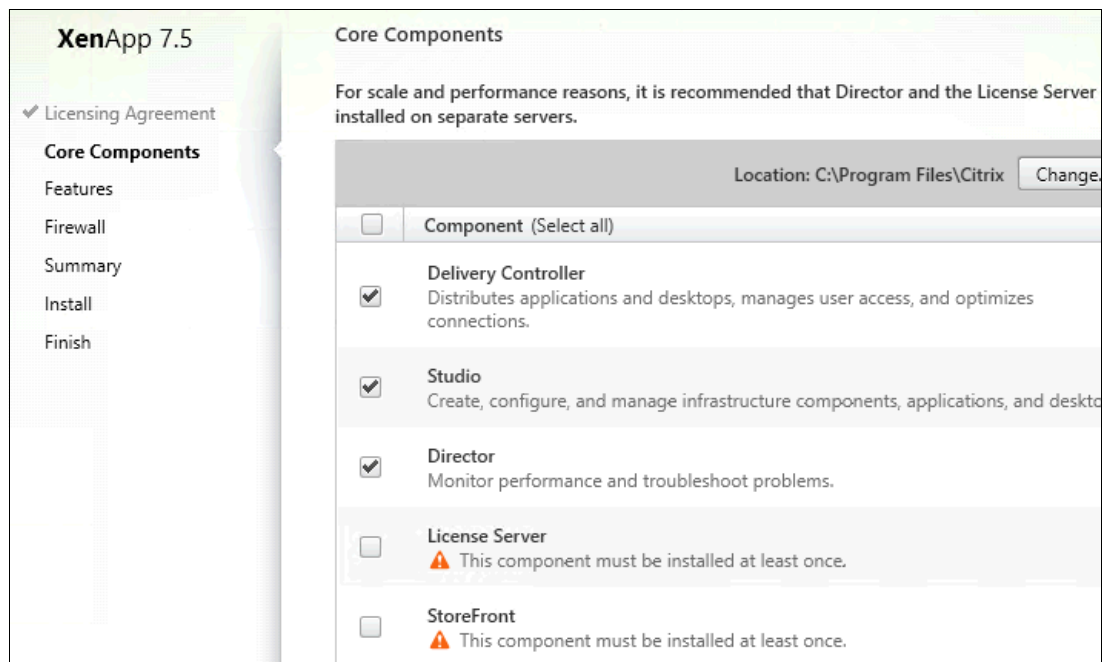


Figure 8-27 XenApp Add components

3. Clear the **Install Microsoft SQL Server Express 2012** option and leave **Install Remote Assistance** selected, as shown in Figure 8-28.

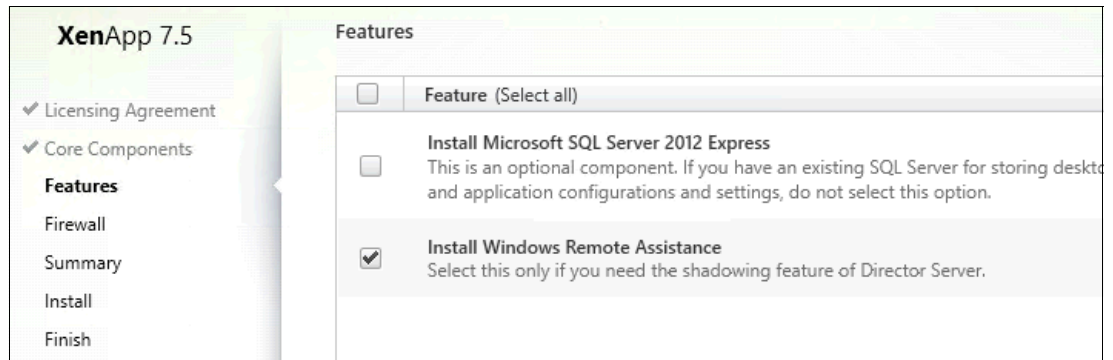


Figure 8-28 XenApp features

4. Leave **Automatically** selected for the firewall rules configuration, as shown in Figure 8-29. Click **Next**.

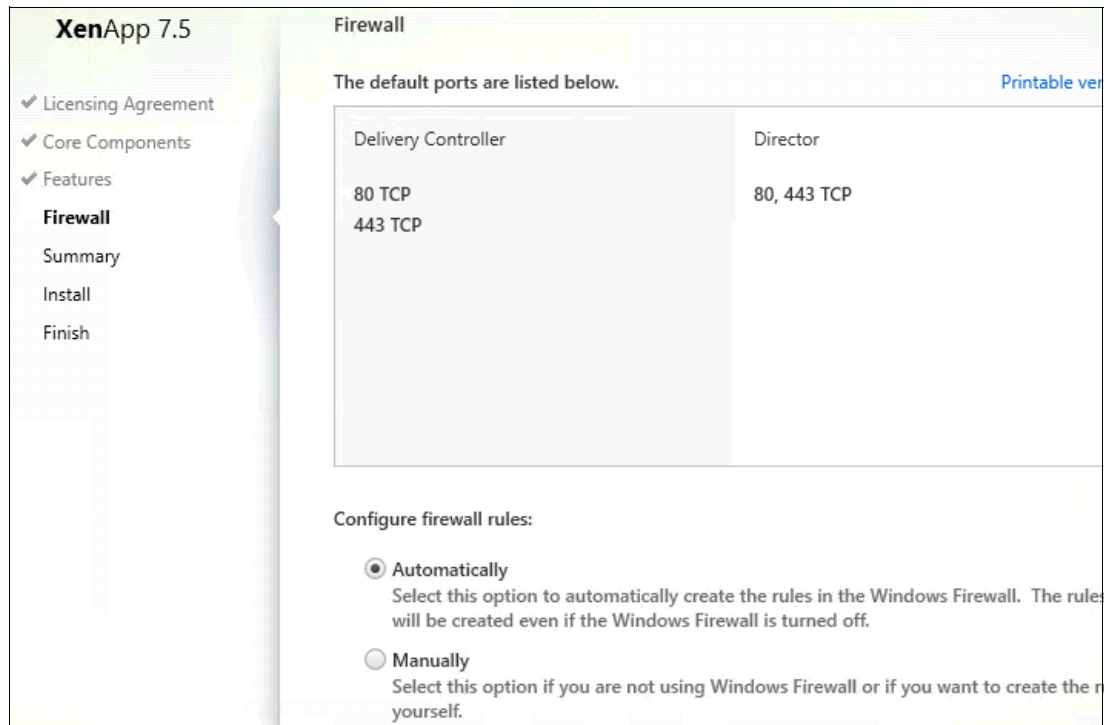


Figure 8-29 XenApp Firewall configuration

5. Review the Summary, as shown in Figure 8-30. Click **Install**.

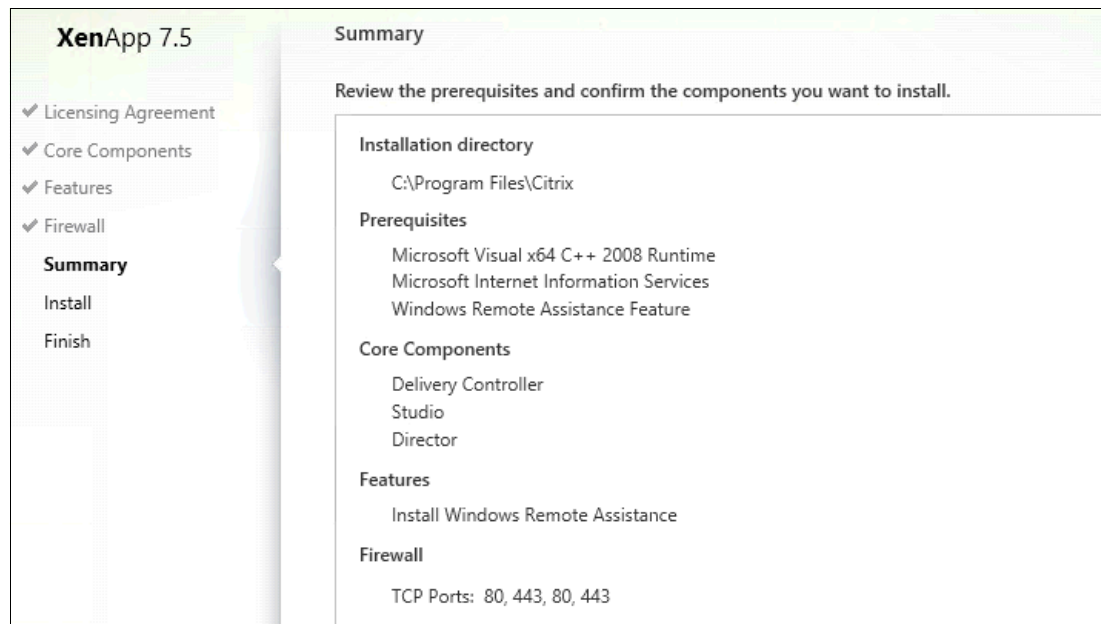


Figure 8-30 XenApp Summary

6. After the installation is complete, leave **Launch Studio** selected and then, click **Finish**, as shown in Figure 8-31.

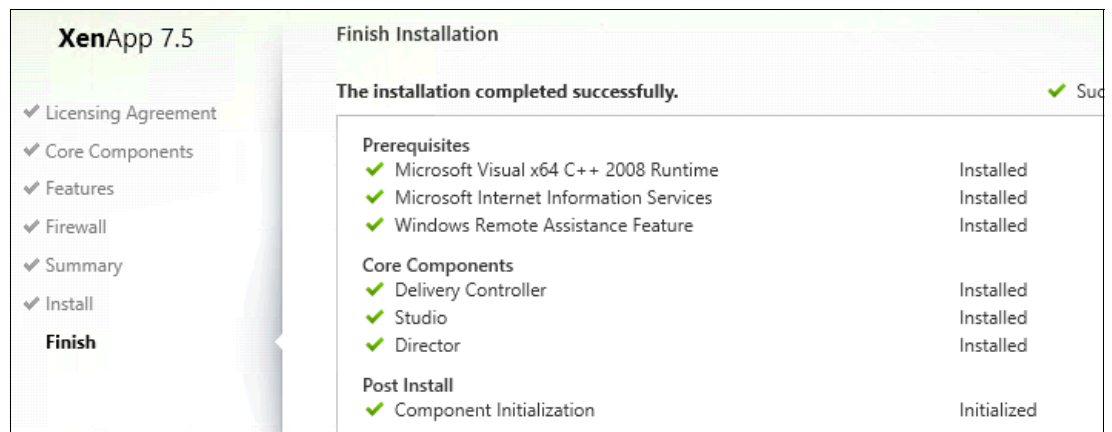


Figure 8-31 XenApp Finish Installation window

7. This controller now must be joined to the previously created site. Click **Connect this Delivery Controller to an existing Site**, as shown in Figure 8-32.

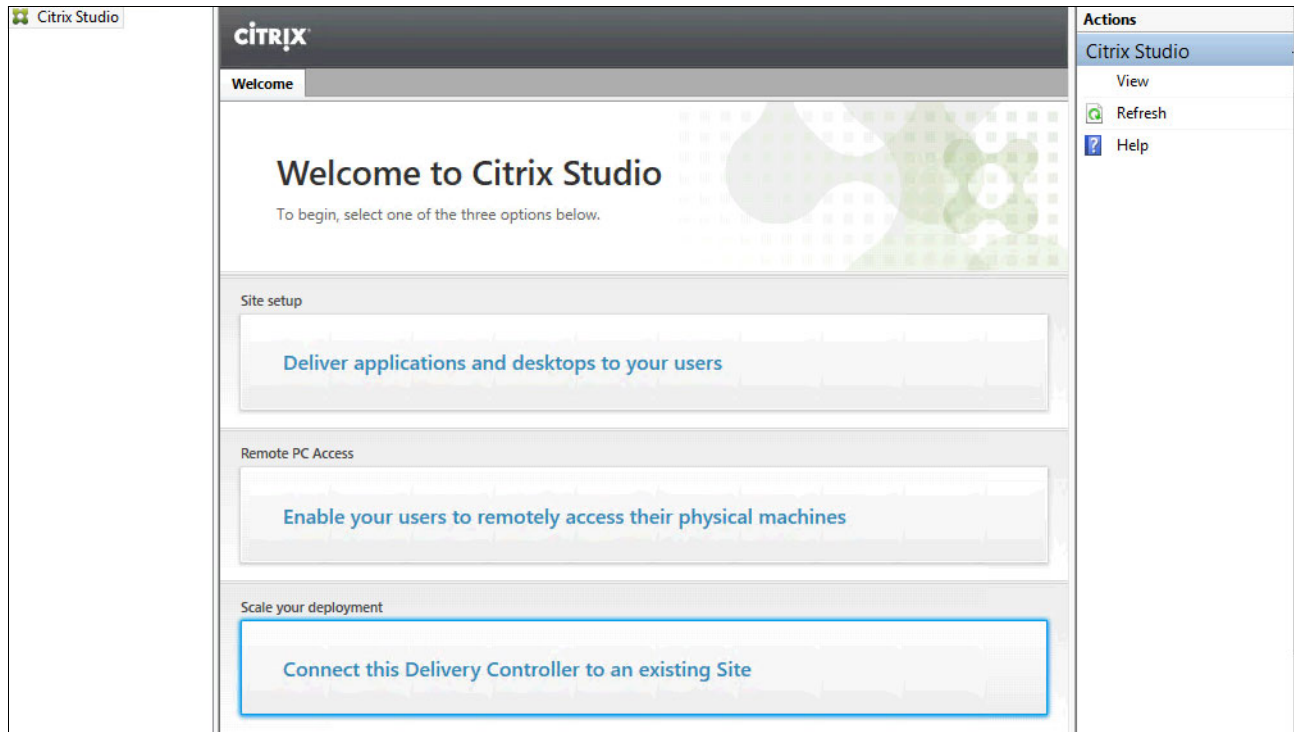


Figure 8-32 XenApp Connect Delivery Controller to a site

8. In the Select Site window, enter the name of delivery controller (in our example, `xdc.xenlab.local`), as shown in Figure 8-33.

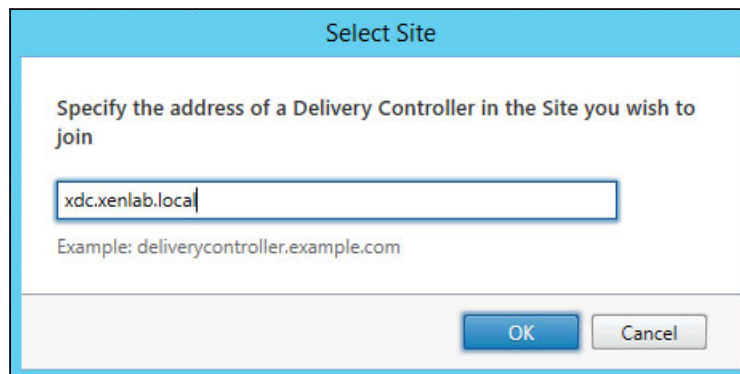


Figure 8-33 XenApp select site to join

9. Click **Yes** to the question “Would you like Studio to update the database automatically?”, as shown in Figure 8-34.

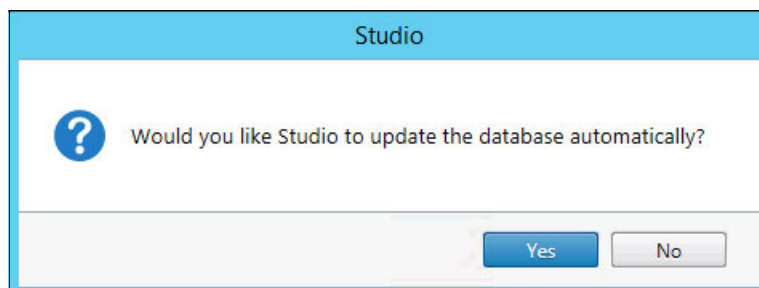


Figure 8-34 Database updated automatically by Studio

When the process is complete, the window that is shown in Figure 8-35 opens.

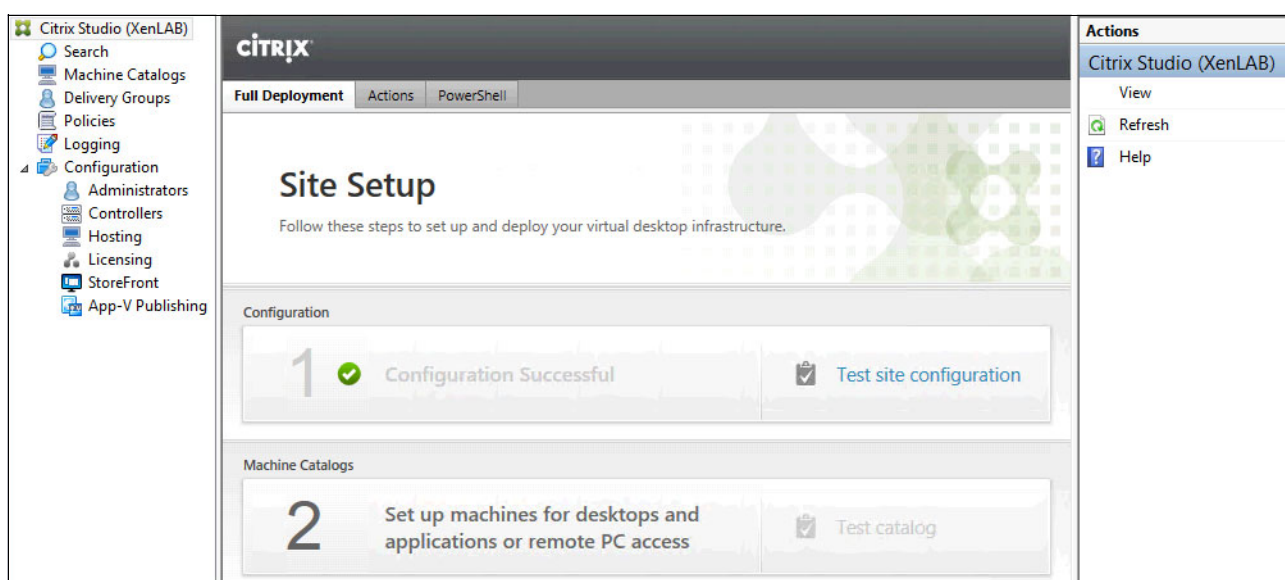


Figure 8-35 XenApp installed successfully

8.5 Installing Citrix StoreFront

The Citrix StoreFront is the preferred method to provide access to XenApp and XenDesktop resources. The Web Interface is still supported (version 5.4), but some features are available with StoreFront only.

Important: For more information about installing StoreFront version 2.5, see this website:
<http://bit.ly/1lW70XQ>

Complete the following steps to install Citrix Web Interface on a separate VM:

1. Mount the remote media on the XSF VM. After inserting the Citrix XenDesktop 7.5 installation media, the AutoRun window opens. Accept the license agreement.
2. Select the **StoreFront** component and clear the other options, as shown in Figure 8-36 on page 168. Click **Next**.

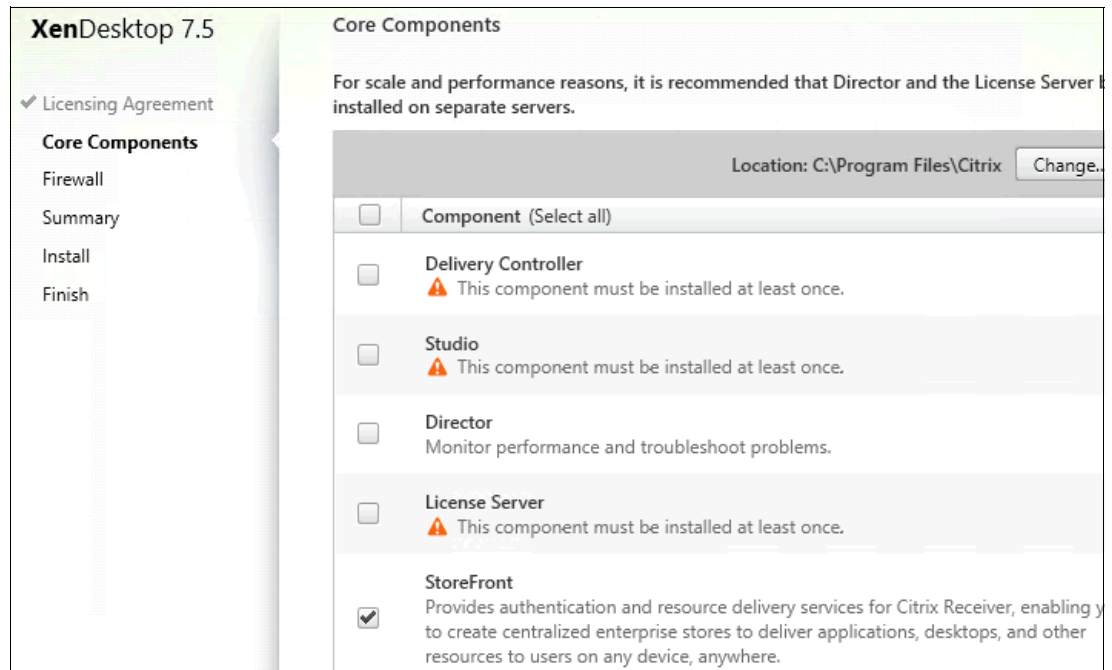


Figure 8-36 XenDesktop components to install

3. After accepting default settings for firewall, review the summary information and click **Install**, as shown in Figure 8-37.

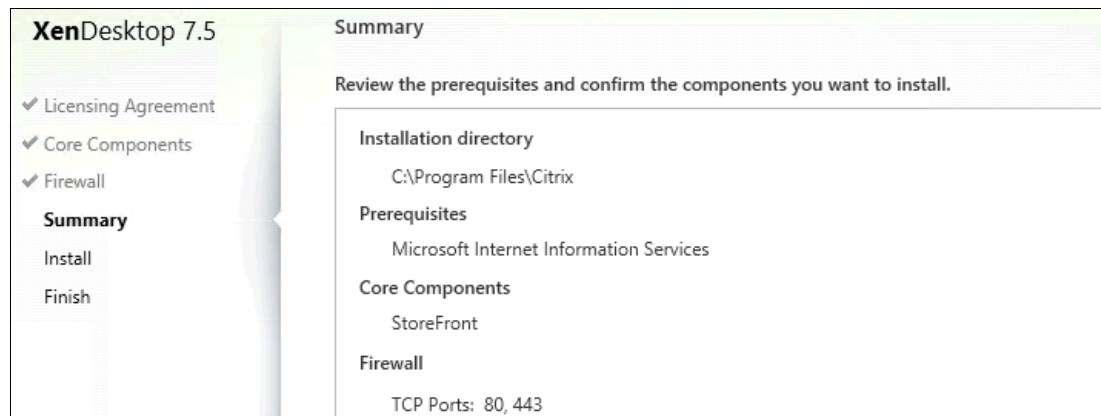


Figure 8-37 StoreFront review summary

4. After the installation process completes, click **Finish**, as shown in Figure 8-38.

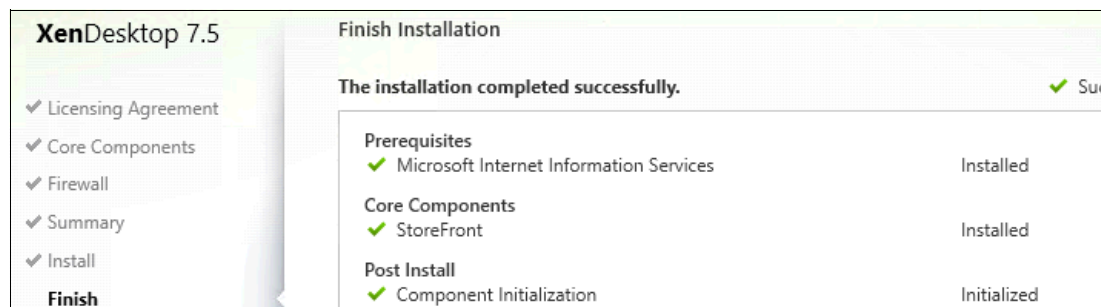


Figure 8-38 StoreFront Installation Successful

8.5.1 Configuring the StoreFront

Complete the following steps to configure the component:

1. Select **Start** → **Citrix StoreFront** to open the Citrix StoreFront console.
2. Click **Create a new deployment** to start the wizard for the initial configuration, as shown in Figure 8-39.

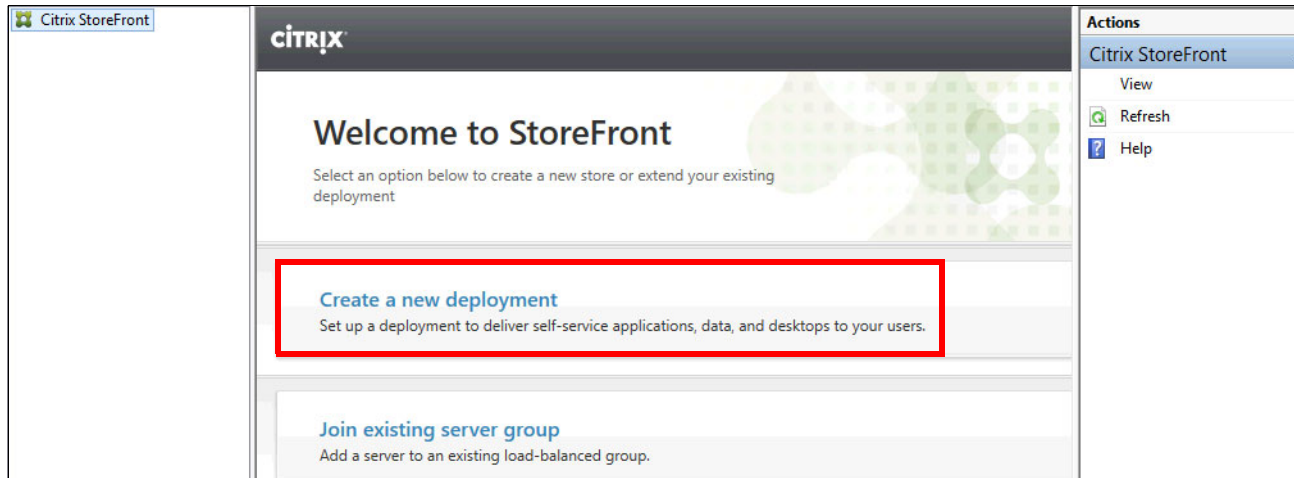


Figure 8-39 StoreFront console

3. In Base URL window, confirm the base URL (which can be changed later), if needed, as shown in Figure 8-40. Click **Next**.

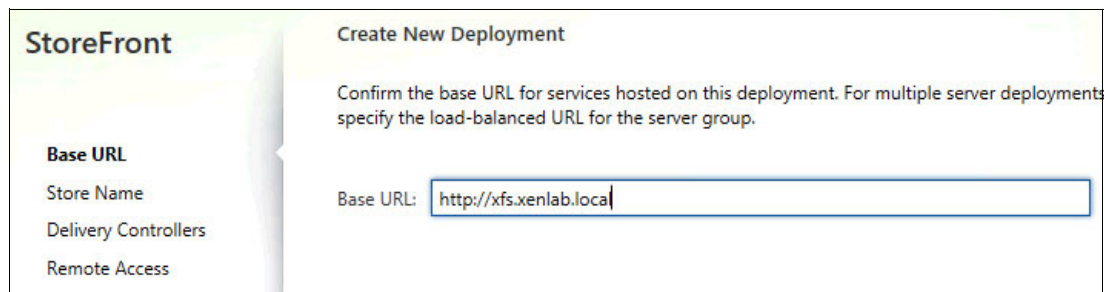


Figure 8-40 StoreFront Configure base URL

4. Enter the Store Name, as shown in Figure 8-41. Click **Next**.

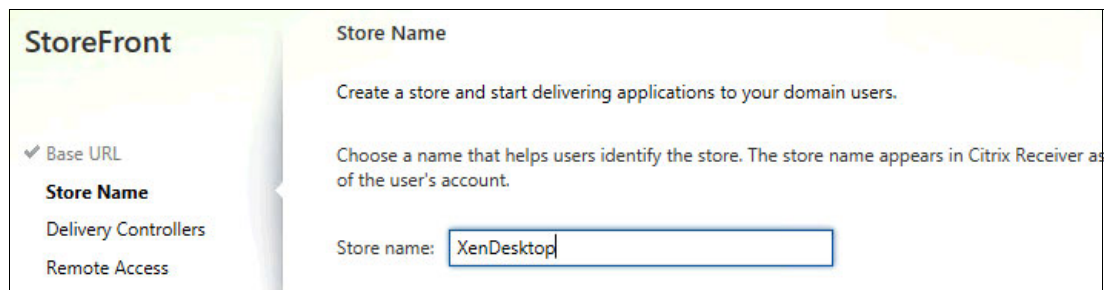
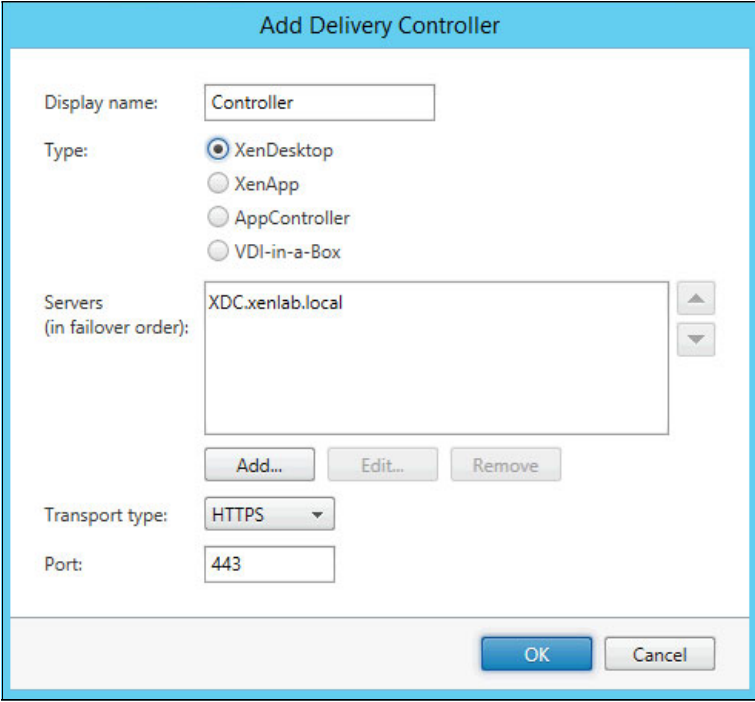


Figure 8-41 Store Name window

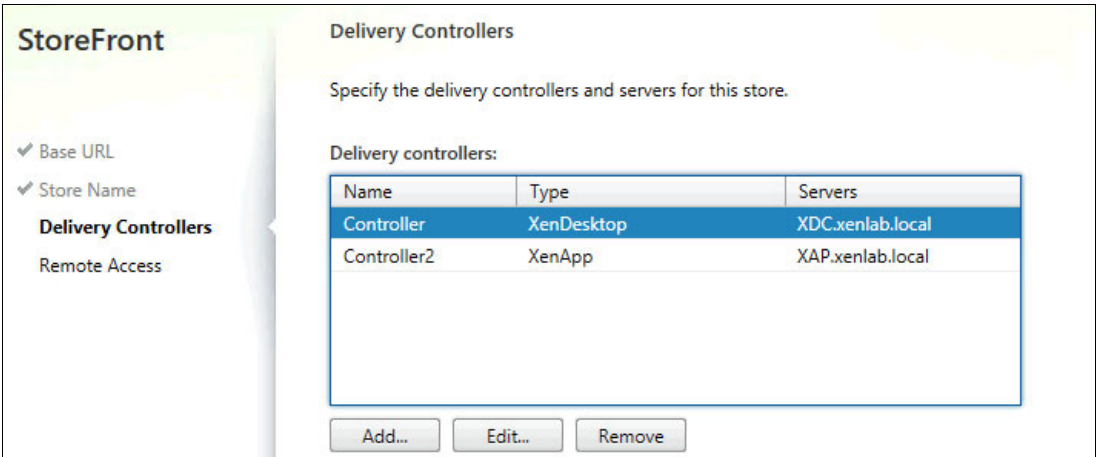
5. In the Add Delivery Controller window, add the servers, as shown in Figure 8-42. Click **OK**.



The 'Add Delivery Controller' dialog box is shown. It has a title bar 'Add Delivery Controller'. Inside, there are several fields: 'Display name' with the value 'Controller'; 'Type' with radio buttons for 'XenDesktop' (selected), 'XenApp', 'AppController', and 'VDI-in-a-Box'; 'Servers (in failover order):' with a text box containing 'XDC.xenlab.local' and up/down arrow buttons; 'Add...', 'Edit...', and 'Remove' buttons; 'Transport type' with a dropdown menu set to 'HTTPS'; and 'Port' with a text box containing '443'. At the bottom right are 'OK' and 'Cancel' buttons.

Figure 8-42 StoreFront Add Delivery Controller

6. Add the second controller and specify the type XenApp, as shown in Figure 8-43. Click **Next**.



The 'StoreFront' configuration window is shown. On the left is a sidebar with 'Base URL', 'Store Name', 'Delivery Controllers', and 'Remote Access'. The 'Delivery Controllers' section is active. The main area is titled 'Delivery Controllers' and contains the text 'Specify the delivery controllers and servers for this store.' Below this is a table with the following data:

Name	Type	Servers
Controller	XenDesktop	XDC.xenlab.local
Controller2	XenApp	XAP.xenlab.local

Below the table are 'Add...', 'Edit...', and 'Remove' buttons.

Figure 8-43 StoreFront Delivery Controller

7. In our example, we do not use Remote Access (see Figure 8-44). Click **Create**.

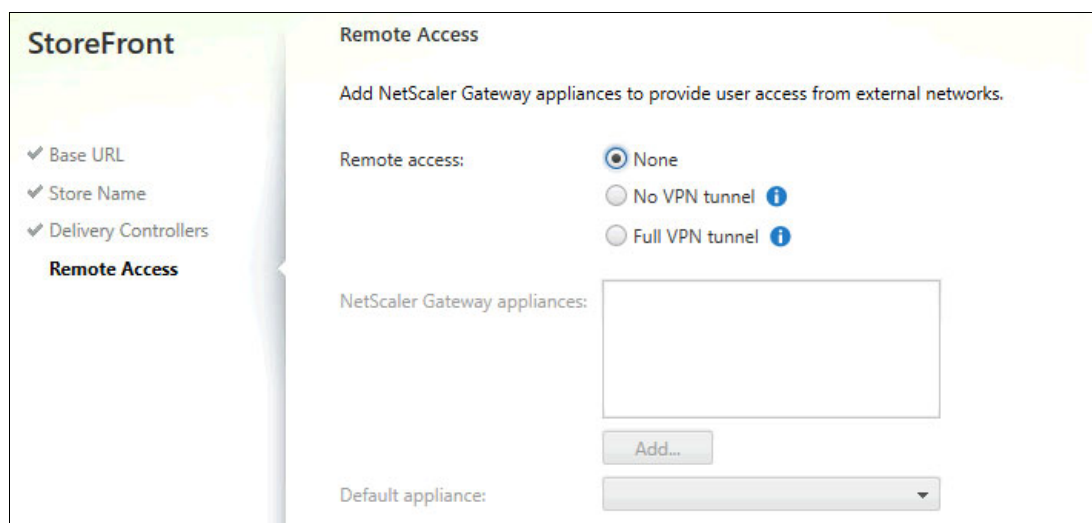


Figure 8-44 StoreFront Remote Access configuration

8. Wait until StoreFront is created successfully (see Figure 8-45). Click **Finish**.

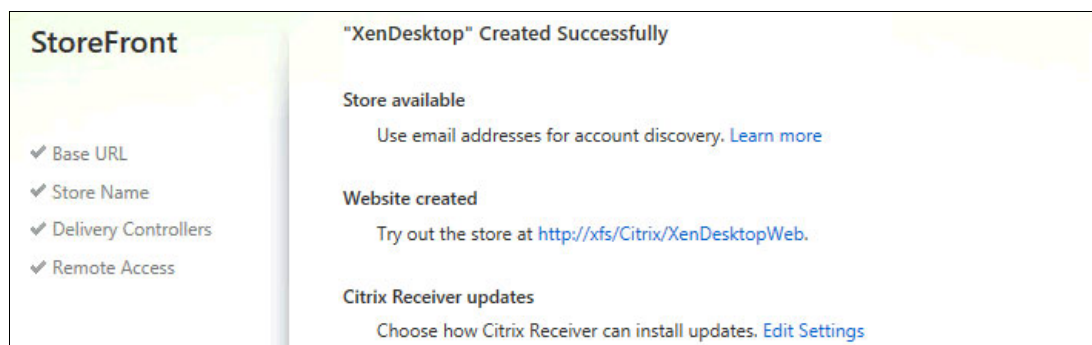


Figure 8-45 StoreFront successfully created

8.6 Installing Citrix Provisioning Services

As part of the XenDesktop implementation, Provisioning Services streaming technology is the main delivery method for non-persistent desktops in scalable implementations.

Note: For more information about installing and configuring Provisioning Services 7.x, see this website:

<http://bit.ly/lgTh1ro>

Complete the following steps to install Provisioning Services on a dedicated VM. The Dynamic Host Configuration Protocol (DHCP) service role is configured on this VM:

1. Mount the remote media on the PVS VM. After mounting the Provisioning Services installation media, the Provisioning Services AutoRun window opens. Select **Server Installation**.
2. The window that is shown in Figure 8-46 shows the prerequisite items that must be installed. Click **Install** to install all of the prerequisites.

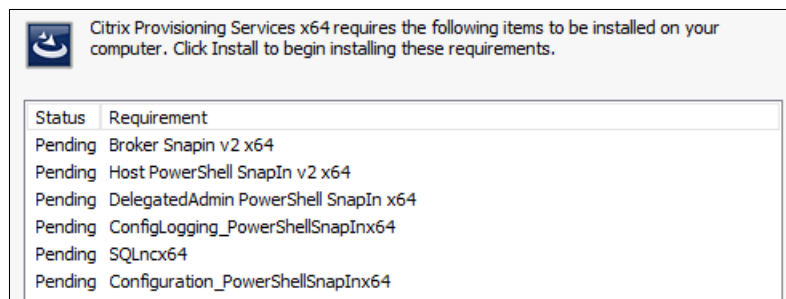


Figure 8-46 Provisioning Services prerequisites

3. If you do not have SQL Native Client installed, click **Yes** to install SQLnxc64, then click **Next**.
4. Accept the License Agreement and click **Next**.
5. Specify the company information and click **Next**.
6. Accept the default installation folder and the complete setup type. Review the summary and then click **Install**.
7. When the installation completes, the Provisioning Services Configuration Wizard opens. Click **Next**. The following steps assume that the Provisioning Console is installed (for more information, see 8.6.1, “Installing the Citrix Provisioning Console” on page 177).
8. Specify the DHCP services, as shown in Figure 8-47. Click **Next**.

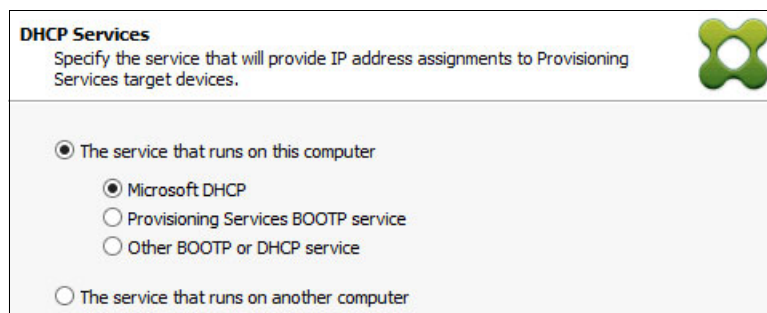


Figure 8-47 Provisioning Services DHCP Services window

9. Specify the Preboot Execution Environment (PXE) services, as shown in Figure 8-48. Click **Next**.

PXE Services
Specify which service will deliver this information to target devices.

During the PXE boot process the bootstrap file name and FQDN/IP address of the TFTP server hosting the bootstrap are delivered via a PXE service or DHCP options 66/67.

☒ The service that runs on this computer

☒ Microsoft DHCP
☐ Provisioning Services PXE service

☐ The service that runs on another computer

Figure 8-48 Provisioning Services PXE Services window

10. Select **Create farm**, as shown in Figure 8-49. Click **Next**.

Farm Configuration
Create a new Farm or join an existing Farm. Can be skipped if already configured.

☒ Create farm
☐ Join existing farm

Figure 8-49 Provisioning Services Farm Configuration window

11. Specify the database server name and the instance name, as shown in Figure 8-50. Click **Next**.

Database Server
Enter the Server and Instance names.

Server name: Browse...

Instance name:

Optional TCP port:

☐ Specify database mirror failover partner

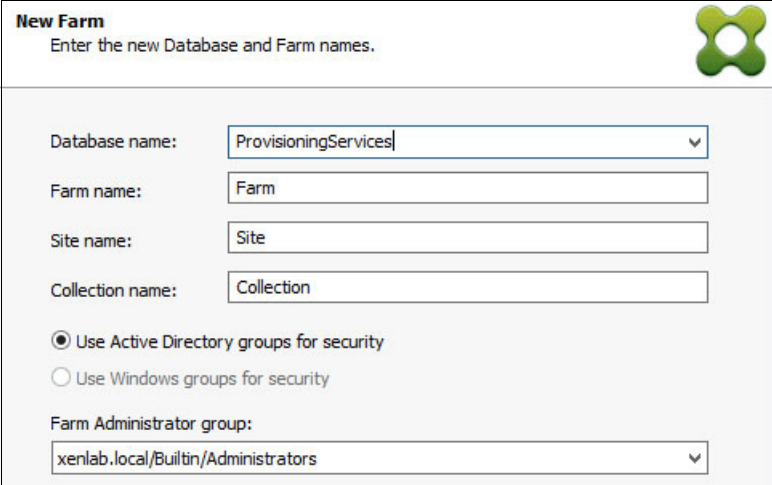
Server name: Browse...

Instance name:

Optional TCP port:

Figure 8-50 Provisioning Services Database Server window

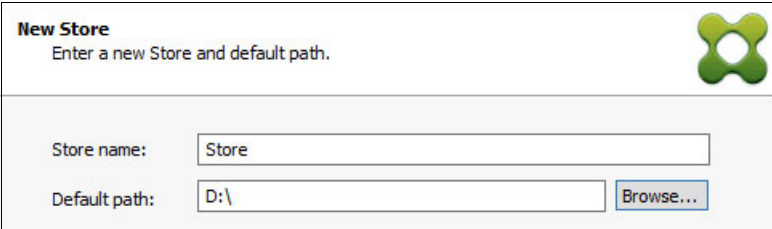
12. After entering the credentials that are needed for the database connection, specify the name for the database, farm, site, and collection, and the farm administrator group, as shown in Figure 8-51. Click **Next**.



The 'New Farm' window has a title bar with the text 'New Farm' and a subtitle 'Enter the new Database and Farm names.' in the top left corner. A green Citrix logo is in the top right corner. The main area contains several input fields: 'Database name:' with a dropdown menu showing 'ProvisioningServices'; 'Farm name:' with a text box containing 'Farm'; 'Site name:' with a text box containing 'Site'; 'Collection name:' with a text box containing 'Collection'; two radio buttons for security, with 'Use Active Directory groups for security' selected; and 'Farm Administrator group:' with a dropdown menu showing 'xenlab.local/Builtin/Administrators'.

Figure 8-51 Provisioning Services farm details

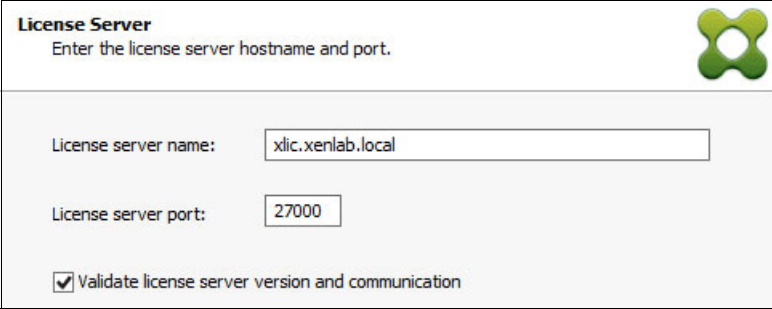
13. Define a new store name and the default path, as shown in Figure 8-52. Click **Next**.



The 'New Store' window has a title bar with the text 'New Store' and a subtitle 'Enter a new Store and default path.' in the top left corner. A green Citrix logo is in the top right corner. The main area contains two input fields: 'Store name:' with a text box containing 'Store'; and 'Default path:' with a text box containing 'D:\' and a 'Browse...' button to its right.

Figure 8-52 New store configuration window

14. Specify the License Server name and select **Validate license server version and communication**, as shown in Figure 8-53. Click **Next**.



The 'License Server' window has a title bar with the text 'License Server' and a subtitle 'Enter the license server hostname and port.' in the top left corner. A green Citrix logo is in the top right corner. The main area contains two input fields: 'License server name:' with a text box containing 'xlic.xenlab.local'; and 'License server port:' with a text box containing '27000'. At the bottom, there is a checked checkbox labeled 'Validate license server version and communication'.

Figure 8-53 License Server window

15. Set a user account to run the Stream and Soap Services. Select **Configure the database for the account**, as shown in Figure 8-54. Click **Next**.

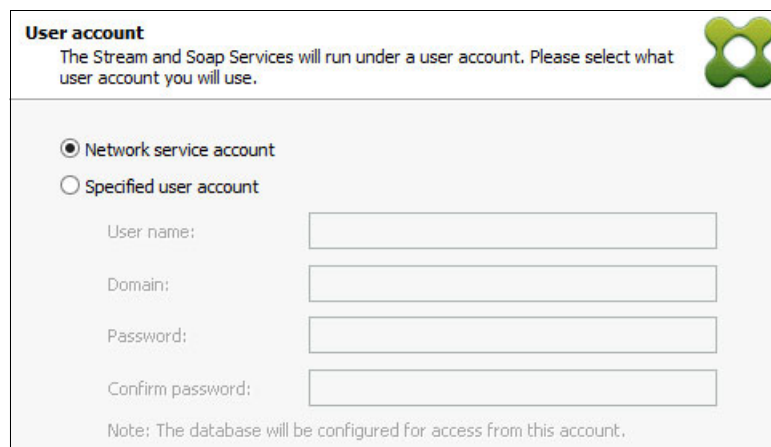


Figure 8-54 User account for Stream and Soap Services

16. Select **Automate computer account password updates** in the AD, as shown in Figure 8-55. Click **Next**.

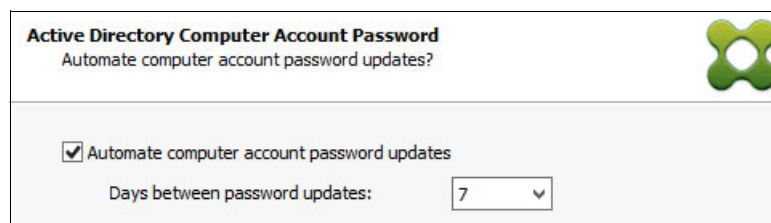


Figure 8-55 AD computer account password update

17. Specify the network cards for streaming services, as shown in Figure 8-56. Choose the NIC that is connected to the dedicated network for streaming (in our example, VM VLAN 30). Click **Next**.

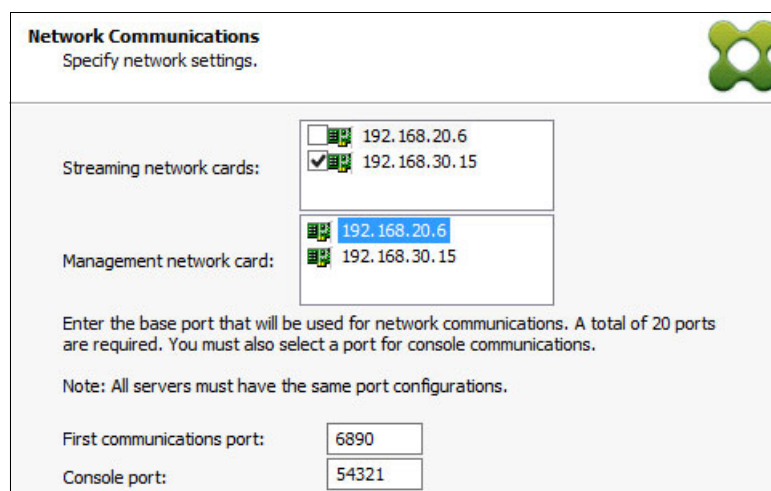


Figure 8-56 Network Communications window

18. Specify the Trivial File Transfer Protocol (TFTP) option and Bootstrap location, as shown in Figure 8-57. Click **Next**.

Figure 8-57 TFTP Option and Bootstrap Location window

19. Specify the servers to which the target devices can contact to complete the boot process, as shown in Figure 8-58. Click **Next**.

Server IP Address	Server Port	Device Subnet Mask	Device Gateway
192.168.30.15	6910	255.255.255.0	0.0.0.0

Figure 8-58 Stream Servers Boot List window

20. Review the summary information of the settings and select **Automatically Start Services**, as shown in Figure 8-59. Click **Finish**.

Figure 8-59 Finish

When the installation completes, the Finish window of the Configuration Wizard's opens, as shown in Figure 8-60. Click **Done**.

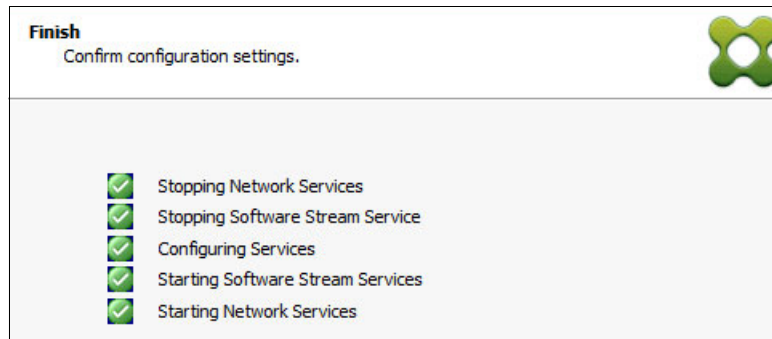


Figure 8-60 Provisioning Services configuration completed

8.6.1 Installing the Citrix Provisioning Console

Complete the following steps to install the Citrix Provisioning Console:

1. After inserting the product installation media, the Provisioning Services AutoRun window opens. Select **Console Installation**.
2. Accept the license agreement, and then select the installation path and setup type: full or custom. After the setup wizard is finished, click **Finish**.
3. Select **Start** → **Provisioning Services Console** to access the console, as shown in Figure 8-61.

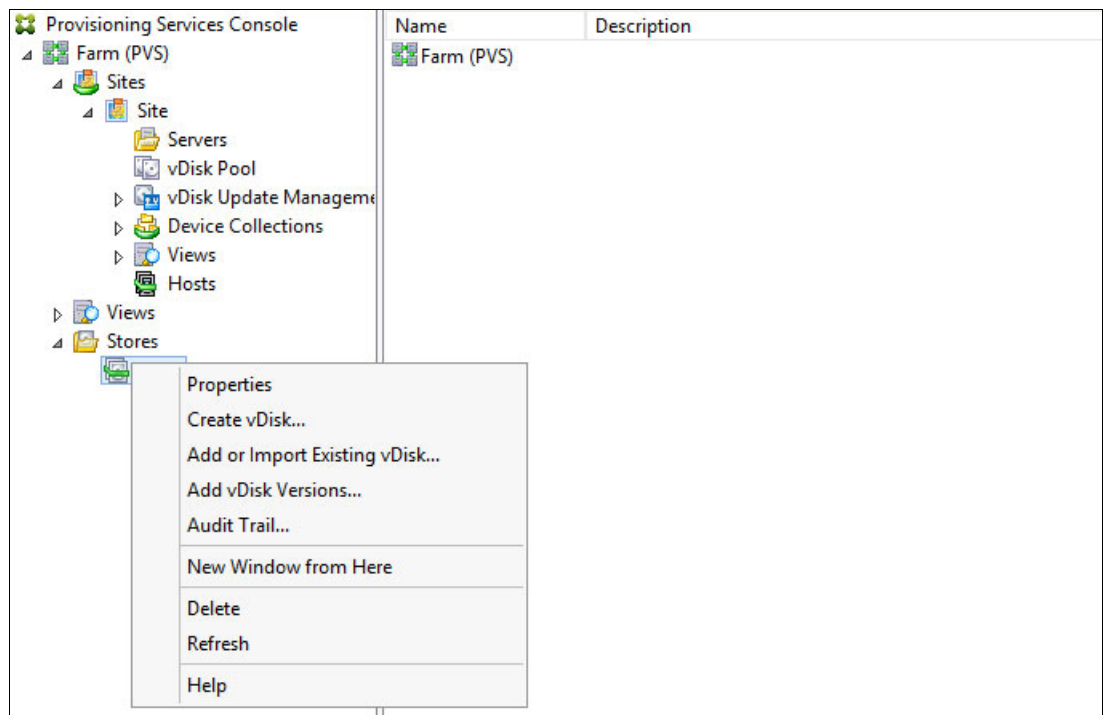


Figure 8-61 Provisioning Services Console

Operating Citrix XenDesktop

This chapter describes the steps to prepare your virtual desktop infrastructure (VDI) environment to deliver desktops to your users.

This chapter includes the following topics:

- ▶ 9.1, “Introduction” on page 180
- ▶ 9.2, “Configuring the gold image” on page 180
- ▶ 9.3, “Configuring desktop distribution” on page 198
- ▶ 9.4, “Roaming profiles and folder redirection” on page 225

9.1 Introduction

The VDI operations consist of the initial desktop installation and configuration in accordance with your company's business requirements and security policies.

This chapter describes how to prepare the desktop image (*gold image*) to be integrated with Provisioning Services and the XenDesktop Controller to be published to users.

This chapter also covers the user data that describes the profile and folder redirection and the integration with XenDesktop/XenApp infrastructure to provide the business applications to desktops.

9.2 Configuring the gold image

The concept of the gold image means the initial installation of a virtual desktop that contains all of the customizations to meet your company's directives and requirements.

One of the benefits of using a gold image installation is the reduction of administrative effort. As the VDI administrator, you perform the security and business applications' update at the gold image and then you schedule when this new version is available to your users.

The desktops that you deliver by using the Citrix XenDesktop infrastructure is a derivative of this gold image.

The following sections describe how to prepare the gold image and integrate it with your XenDesktop infrastructure.

9.2.1 Preparing the gold image for streaming services

The gold image preparation consists of installing your client operating system by using a virtual machine (VM) and all of the other required components before integrating with the Citrix infrastructure.

Integrating the gold image with the Citrix Infrastructure consists of the following main steps:

1. Include an extra hard disk to your VM to store the write cache file (which is used for streaming desktops). Make the following configuration changes to this disk:
 - Create a partition and format the new disk by using the New Technology File System (NTFS).
 - Assign the letter D: for this new disk.
 - Move the page file location from C: to the new disk (D).
2. Install the Citrix Profile Manager.

The Citrix Profile Manager is responsible for managing the user profile by loading the files when the user logs on to the desktop and saving the files when the user logs off from the desktop.

The product offers the following methods to set the parameters:

- By using the .INI file that is stored on the installation folder
- Through Group Policies that are created at the Active Directory (AD) level

In this scenario, we use the Group Policy to create the configuration and to apply it to the desktops.

For more information about the configuration, see 9.4, “Roaming profiles and folder redirection” on page 225.

3. Install the Virtual Desktop Agent.

The Virtual Desktop Agent is responsible for registering the provisioned desktop in XenDesktop Controller after the start. After this registration process, the XenDesktop Controller acts as a broker to deliver the desktops to the users. In this scenario, we chose to install the Virtual Desktop by running `AutoRun.exe`, which is on the Citrix XenApp/XenDesktop media.

Installing Provisioning Services target device

The Provisioning Services target device is the gold image to allow streaming desktops. Complete the following steps to install the Provisioning Services target device, configure the agent to communicate with the Provisioning Server (PVS), and to convert the gold image as a VDisk on PVS to be delivered to the desktops:

1. To start the installation window, run `Autorun` from the Provisioning Services media. When the PVS installer starts, click **Target Device Installation**.
2. Select **Target Device Installation** to proceed with installation, as shown in Figure 9-1.




Figure 9-1 Target device installation

3. In the Installation wizard, click **Next** to continue. In the next window, accept the license agreement to continue.
4. Enter the customer information and click **Next**.
5. In the Destination Folder window, click **Next** to continue.
6. Click **Install** to continue with installation.
7. Select **Launch Imaging Wizard** and click **Finish**, as shown in Figure 9-2.

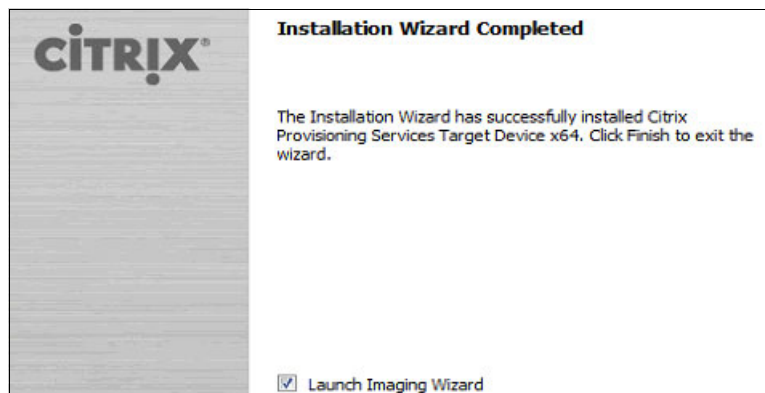


Figure 9-2 Wizard configuration

8. The Provisioning Services Imaging Wizard opens. Click **Next**.
9. Enter the name or IP address of the PVS Server, select the option for Credentials, and click **Next**, as shown in Figure 9-3.

Connect to Farm
Enter the name or address of a server in the farm to connect to.

Server information

Server:

Port:

Credentials

☒ Use my Windows credentials

☐ Use these credentials

User name:

Password:

Domain:

Figure 9-3 Connection configuration

10. Select **Create new VDisk** and click **Next**, as shown in Figure 9-4.

Select New or Existing vDisk
Choose whether you want to create a new vDisk or use an existing one.

☒ Create new vDisk

☐ Use existing vDisk

vDisk name:

Figure 9-4 Provisioning VDisk configuration

11. Specify the VDisk name, Store, VDisk type, and VDisk block size, as shown in Figure 9-5. Click **Next**.

New vDisk
Enter the details for the new vDisk.

vDisk name:

Store:
Accessible by server: PVS

vDisk type:

Figure 9-5 Provisioning new VDisk configuration

12. On the Microsoft Volume Licensing page, select the volume license option to use for target devices or select **None** if volume licensing is not being used. We selected **None** for our installation, as shown in Figure 9-6.

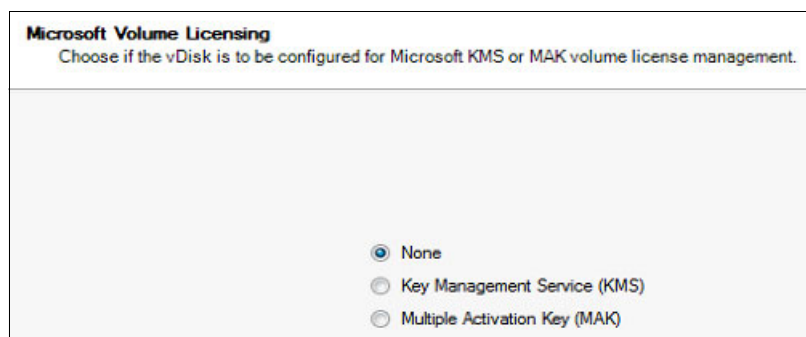


Figure 9-6 Microsoft VDisk volume licensing

Note: KMS activates computers on a local network, which eliminates the need for individual computers to connect to Microsoft.

Note: An MAK is used for one-time activation with Microsoft's hosted activation services. Each MAK has a predetermined number of allowed activations; this number is based on Volume Licensing agreements.

13. As shown in Figure 9-7, under Source Volume, the second disk (D:) and the CD-ROM drive must be changed to **None** so that they are not converted. Optionally, you can adjust the size for the C: partition. Click **Next**.

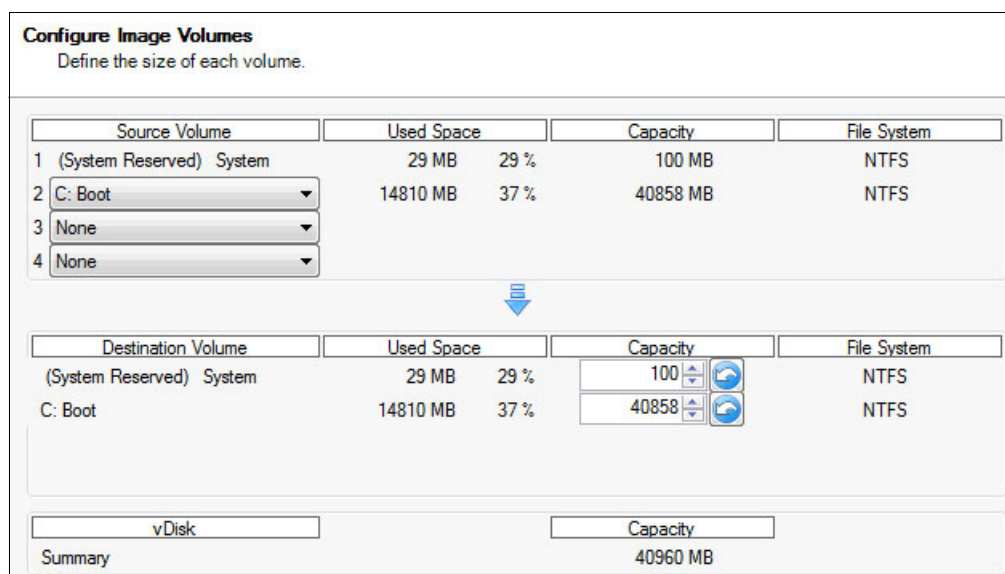


Figure 9-7 Configure image volumes

14. Create the target device at the Provisioning Server by providing the target device name, MAC, and collection. This target device configuration is an association between the Provisioning Server and the client that uses the Media Access Control (MAC) address from the desktop VM, as shown in Figure 9-8. Click **Next**.

Add Target Device
Add this device to the farm.

Target device name: WIN7BaseImage
Note: The target device name cannot be the same Active Directory name of this machine.

MAC: VLAN30 00-15-5D-AB-18-11

Collection: GoldImage

In the Site site of server: PVS

Figure 9-8 Add target device

15. Figure 9-9 shows the summary of farm changes before the configuration. Click **Optimize for Provisioning Services**.

Summary of Farm Changes
This page summarizes the changes to the farm.

The Wizard has enough information to create a new vDisk and add it to the farm.
Please review the information below and click Finish to create the vDisk.

- Name: Windows 7 BaseImage
- Store: Store
- Type: Dynamic
- Size: 40960
- VHD Block Size: 2 MB
- Microsoft Volume Licensing: None
- Volume: (System Reserved) , 29 MB used, 100 MB capacity, NTFS system
- Volume: C:, 12755 MB used, 40858 MB capacity, NTFS system

☒ Add this machine to the farm

- Device name: WIN7BaseImage
- MAC: 00-15-5D-AB-18-11
- Collection: GoldImage

Optimize for Provisioning Services

Figure 9-9 Provisioning Services target device installation summary

16. Review the options that are selected (see Figure 9-10) and click **OK** to return to the summary of changes.

<input checked="" type="checkbox"/> Disable Offline Files	<input checked="" type="checkbox"/> Disable Windows Autoupdate
<input checked="" type="checkbox"/> Disable Defrag BootOptimizeFunction	<input checked="" type="checkbox"/> Disable Background Layout Service
<input checked="" type="checkbox"/> Disable Last Access Timestamp	<input checked="" type="checkbox"/> Disable Hibernation
<input checked="" type="checkbox"/> Reduce DedicatedDumpFile DumpFileSize to 2MB	<input checked="" type="checkbox"/> Disable Indexing Service
<input checked="" type="checkbox"/> Disable Move to Recycle Bin	<input checked="" type="checkbox"/> Reduce Event Log Size to 64k
<input checked="" type="checkbox"/> Reduce IE Temp File	<input checked="" type="checkbox"/> Disable Clear Page File at Shutdown
<input checked="" type="checkbox"/> Disable Machine Account Password Changes	<input checked="" type="checkbox"/> Disable Windows SuperFetch
<input checked="" type="checkbox"/> Disable Windows Defender	<input checked="" type="checkbox"/> Disable Windows Search
<input checked="" type="checkbox"/> Disable ScheduledDefrag	<input checked="" type="checkbox"/> Disable System Restore
<input checked="" type="checkbox"/> Disable ProgramDataUpdater	<input checked="" type="checkbox"/> Run NGen ExecuteQueuedItems (new window)

Figure 9-10 Provisioning services device optimization tool

17. Depending on the .NET Framework versions that are installed on the VM, the optimization process can take from less than a second to over an hour. After the process completes, click **Finish** (see Figure 9-11).

Summary of Farm Changes
This page summarizes the changes to the farm.

The Wizard has enough information to create a new vDisk and add it to the farm.
Please review the information below and click Finish to create the vDisk.

- Name: Windows 7 BaselineImage
- Store: Store
- Type: Dynamic
- Size: 40960
- VHD Block Size: 2 MB
- Microsoft Volume Licensing: None
- Volume: (System Reserved) , 29 MB used, 100 MB capacity, NTFS system
- Volume: C:, 12755 MB used, 40858 MB capacity, NTFS system
- ☒ Add this machine to the farm
 - Device name: WIN7BaselineImage
 - MAC: 00-15-5D-AB-18-11
 - Collection: GoldImage

[Optimize for Provisioning Services](#)

Figure 9-11 VDisk configuration summary

18. You must adjust the boot order of your VM before it restarts. Your network adapter must be at the top of the list to boot by using the network. With this configuration, connect in the PVS and upload your image.

Figure 9-12 shows the configuration order for the boot options.

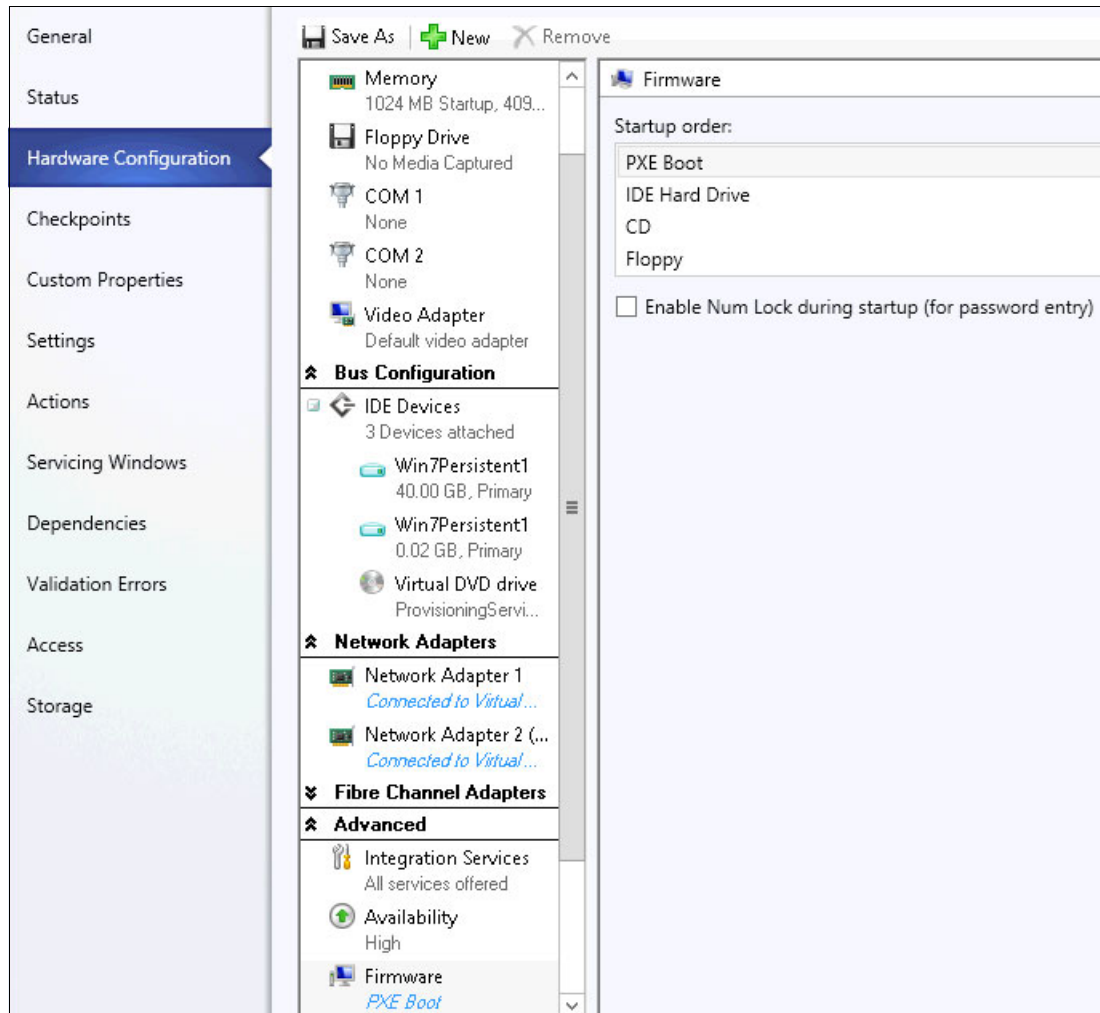


Figure 9-12 Boot order adjustment

19. Before we continue, first log on to the PVS server to verify that the VDisk is created, as shown in Figure 9-13.




Figure 9-13 Verify VDisk creation

20.A Target Device also was created with the MAC address of the VM, which is linked to the VDisk that was created and the Target Device is configured to boot from its hard disk because the VDisk is empty now (see Figure 9-14).

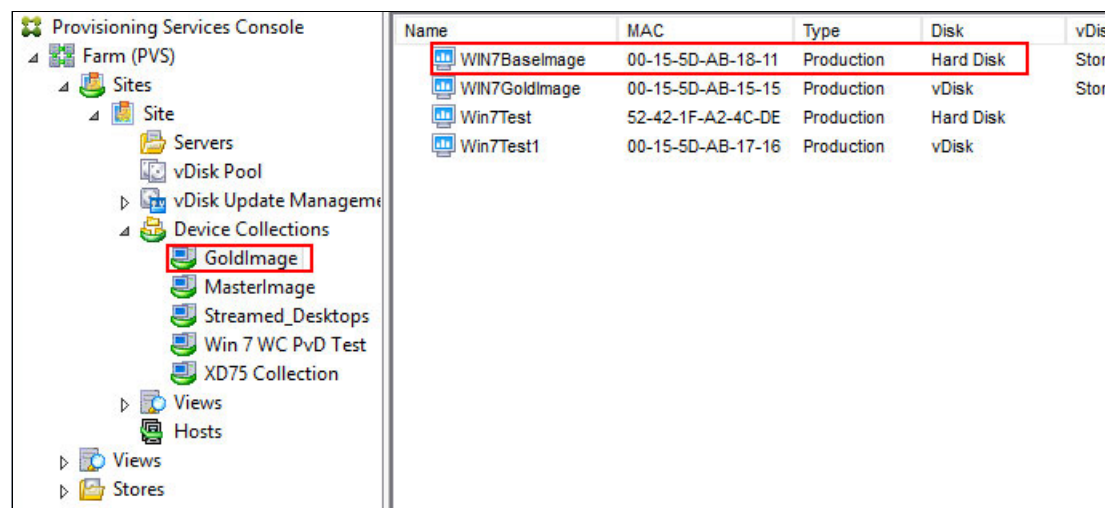


Figure 9-14 Target device verification

21.After the VM is configured to boot from the network first and the hard disk drive second, power on the VM. When the VM is at the logon window, log on with the same domain account and the Imaging Wizard process continues.

22.When the Imaging Wizard process is complete, click **Finish** and shutdown the VM, as shown in Figure 9-15.




Figure 9-15 Successful conversion

23.By using the Provisioning Services Console, you must change the target device to boot the VM from VDisk. Select the target device, right-click, and select **Properties**, as shown in Figure 9-16.

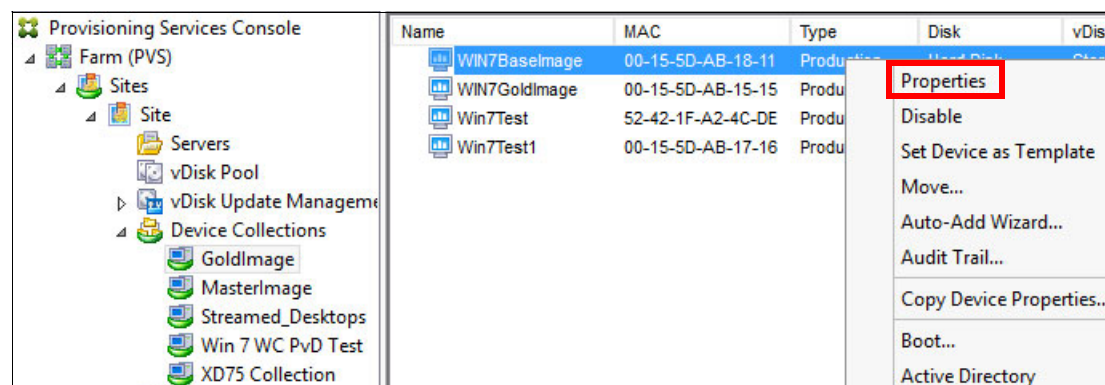
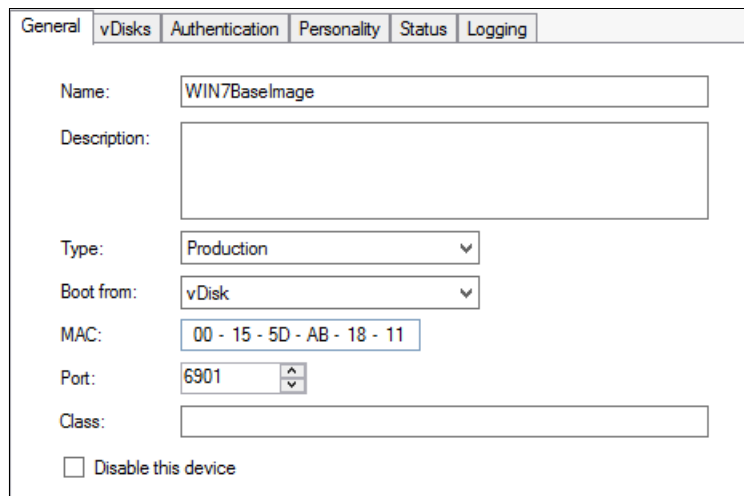


Figure 9-16 VDisk properties

24. Change the Boot from to **VDisk** and click **OK**, as shown in Figure 9-17.

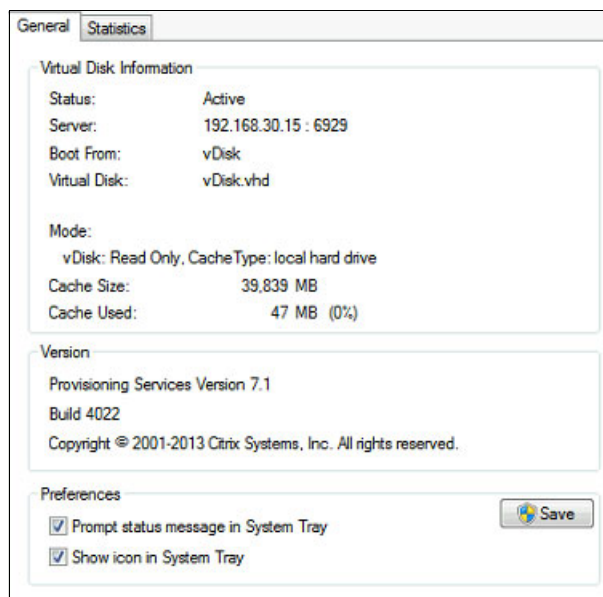


The screenshot shows the 'vDisks' tab in the Provisioning Services console. The configuration for the device 'WIN7BaseImage' is displayed. The 'Boot from' dropdown menu is set to 'vDisk'. Other fields include 'Name' (WIN7BaseImage), 'Description' (empty), 'Type' (Production), 'MAC' (00 - 15 - 5D - AB - 18 - 11), 'Port' (6901), and 'Class' (empty). There is a checkbox for 'Disable this device' which is currently unchecked.

Figure 9-17 Target Device boot adjustment

25. After adjusting the target device in the Provisioning Services Console, power on your VM normally. Currently, your gold image boots by using the streaming services. Your VDisk is running in private mode, which means that the changes that you perform in your VM are stored on the VDisk.

Figure 9-18 shows the VDisk status after the boot.



The screenshot shows the 'Statistics' tab in the Provisioning Services console. The 'Virtual Disk Information' section displays the following details: Status: Active, Server: 192.168.30.15 : 6929, Boot From: vDisk, Virtual Disk: vDisk.vhd. The 'Mode' section shows: vDisk: Read Only, Cache Type: local hard drive, Cache Size: 39,839 MB, and Cache Used: 47 MB (0%). The 'Version' section shows: Provisioning Services Version 7.1, Build 4022, and Copyright © 2001-2013 Citrix Systems, Inc. All rights reserved. The 'Preferences' section has two checked options: 'Prompt status message in System Tray' and 'Show icon in System Tray'. A 'Save' button is located to the right of the preferences.

Figure 9-18 VDisk status

Now, the XenDesktop 7.5 Virtual Delivery Agent (VDA) must be installed. This process is described next.

Installing the Virtual Desktop Agent

Use the following procedure to install the VDA:

1. After you insert the Citrix XenApp/XenDesktop 7.5 installation media, the AutoRun XenApp/XenDesktop window opens. Select **XenDesktop Deliver applications and desktops**.
2. Select **Virtual Delivery Agent for Windows Desktop OS**.
3. In the Environment window, select **Create a Master Image** and click **Next**, as shown in Figure 9-19.

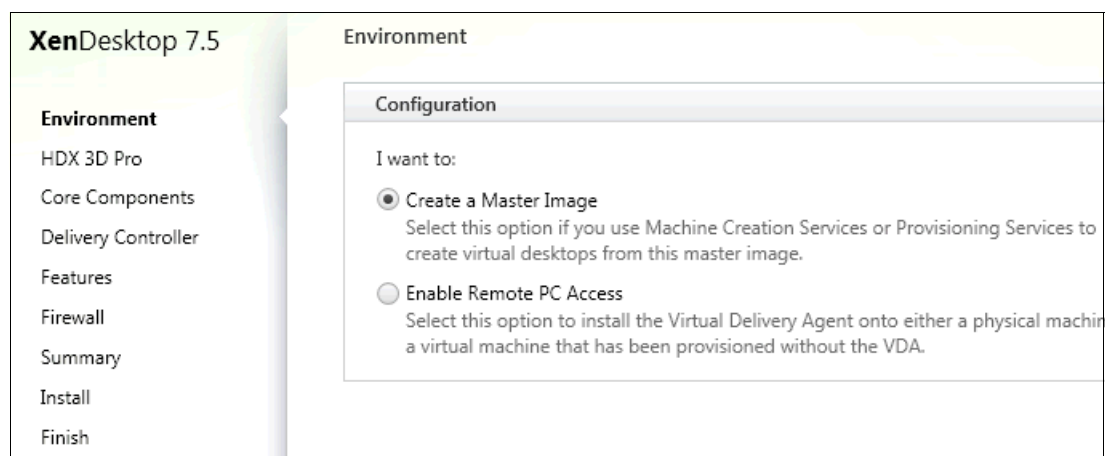


Figure 9-19 Create a Master Image

4. Because we do not access a graphics processor, we select **No, install standard VDA** and click **Next** (see Figure 9-20).

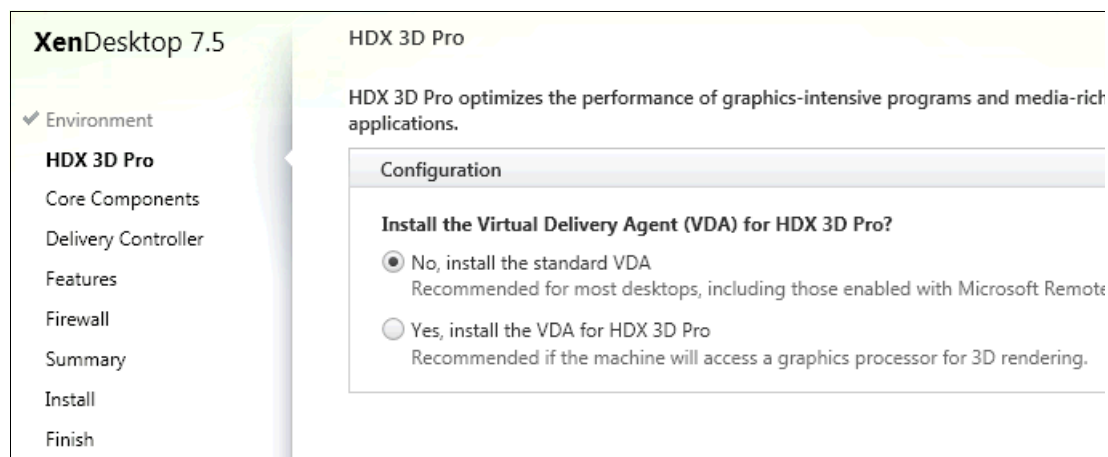


Figure 9-20 Install Standard VDA

5. Select **Citrix Receiver** and click **Next** (see Figure 9-21).

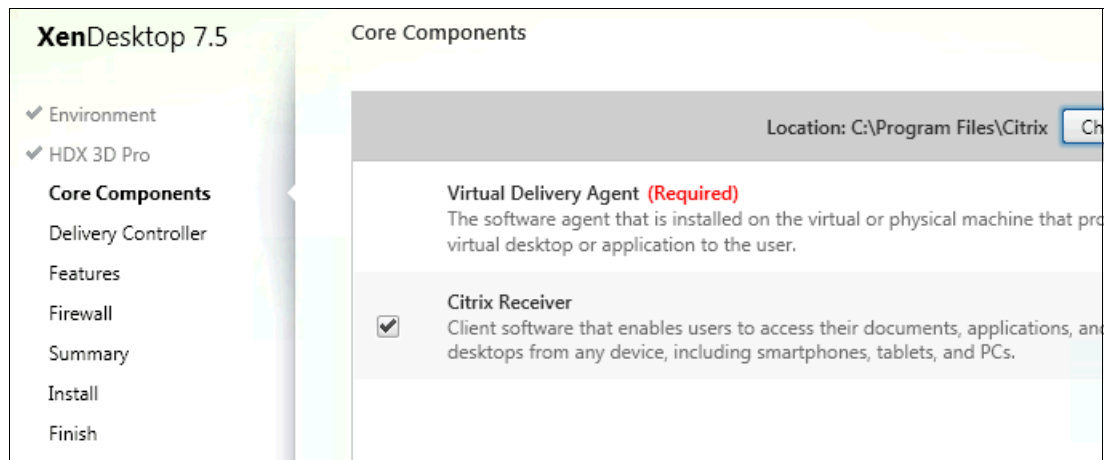


Figure 9-21 Components selection

6. Specify how XenDesktop locates the delivery controller (or controllers). For our installation, we selected the **Do it Manually** option, as shown in Figure 9-22.

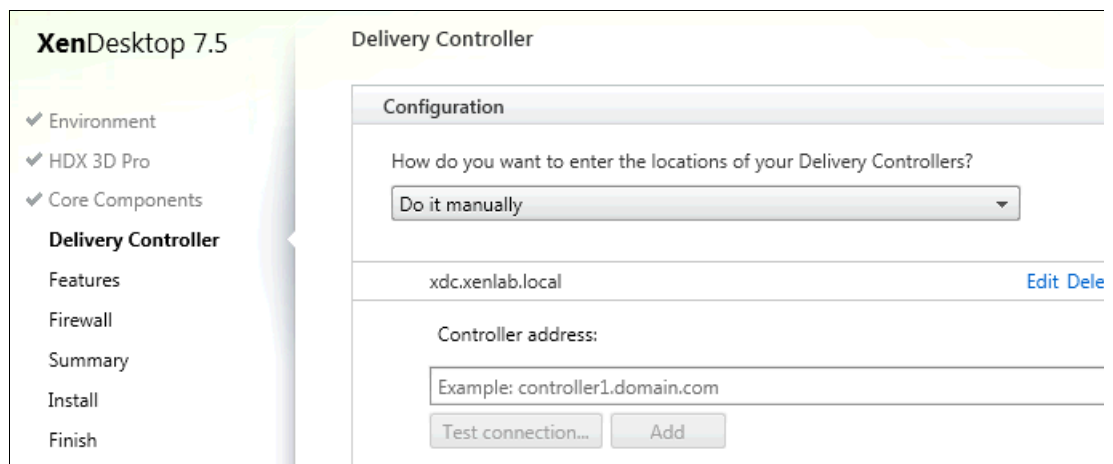


Figure 9-22 Configuring the XenDesktop Controller

We entered the name of our XenDesktop 7.5 Controller, then clicked **Test connection**. After the test is completed successfully, click **Add**. Repeat this step for other XenDesktop Controllers in your environment, then click **Next**.

- In the next window, **Optimize XenDesktop Performance**, **User Desktop Shadowing**, and **Real Time Monitoring** are preselected. Select **personal VDisk** to install this option on the gold image, as shown in Figure 9-23. Click **Next**.

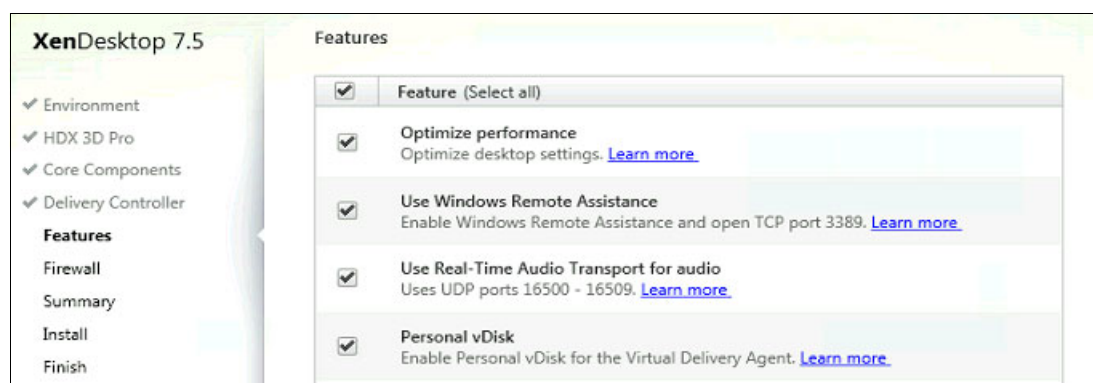


Figure 9-23 Virtual Desktop Configuration

This feature allows the users to customize their desktop and permits centralized administration from Provisioning Services. For more information about how to configure a personal VDisk, see “Configuring streaming desktops with personal VDisk” on page 208.

- The VDA installer offers to open the required ports in the Windows Firewall for you. Select the appropriate firewall rules option and click **Next** (see Figure 9-24).

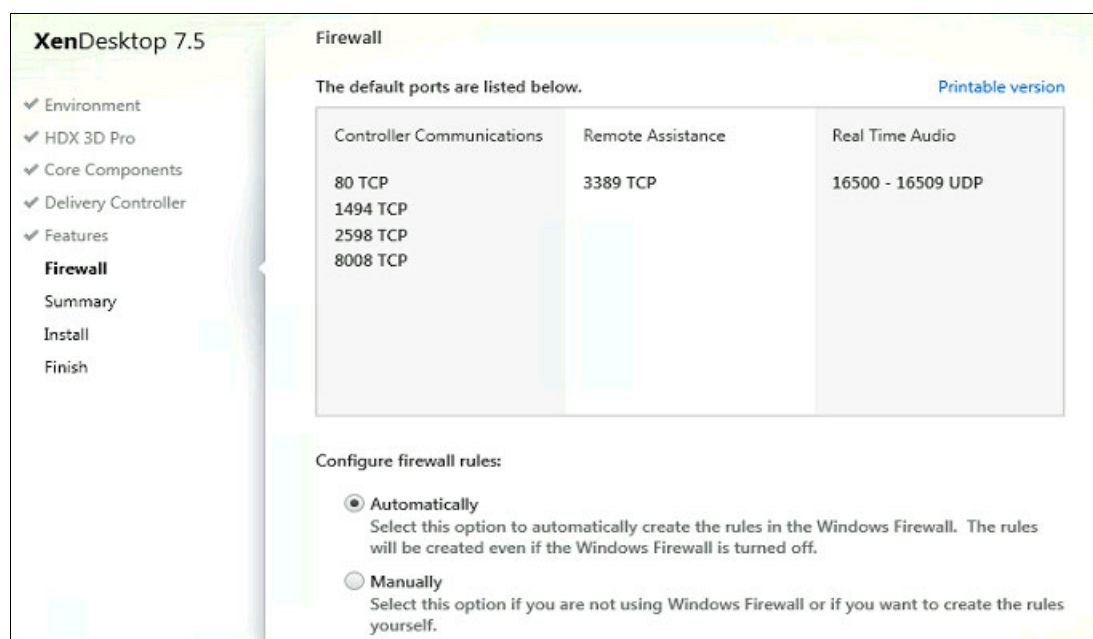


Figure 9-24 Configure firewall rules

9. Figure 9-25 shows the installation summary. Click **Install**.

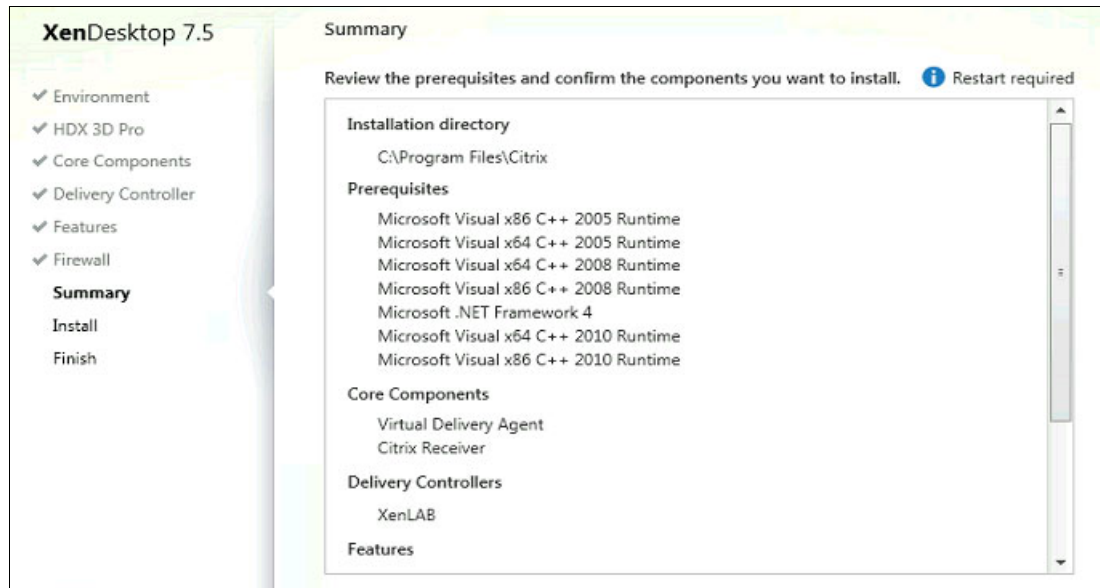


Figure 9-25 Installation Summary

10. Click **Finish** to complete the installation and restart your machine (see Figure 9-26).

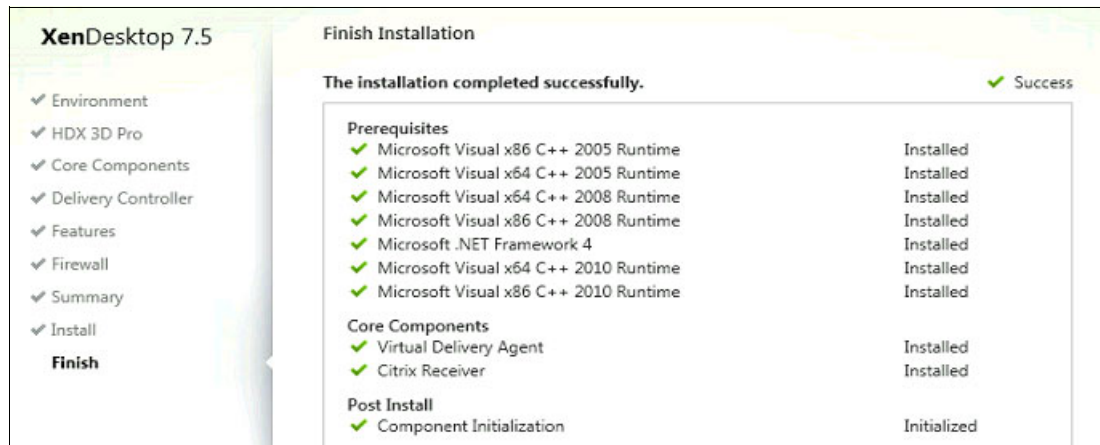


Figure 9-26 Finish installation

11. Manually run the Personal VDisk Inventory. Click **Start** → **All Programs** → **Citrix** → **Update personal VDisk**. The inventory progress is displayed, as shown in Figure 9-27.




Figure 9-27 Personal VDisk inventory

After the inventory completes, the VM is shut down.

Changing the image mode to Standard

The next step is to modify the VDisk to Standard Image mode. Complete the following steps:

1. Switch to PVS and start the Provisioning Services Console. Click **Sites** → **Site** → **VDisk Pool** and right-click the VDisk and select **Properties**, as shown in Figure 9-28.

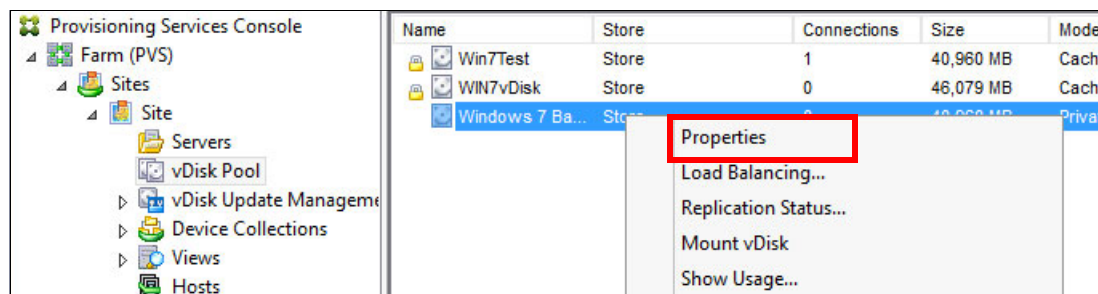


Figure 9-28 VDisk properties

2. Change Access mode from **Private Image** to **Standard Image** and Cache type to **Cache on device hard drive**, as shown in Figure 9-29. Click **OK**.

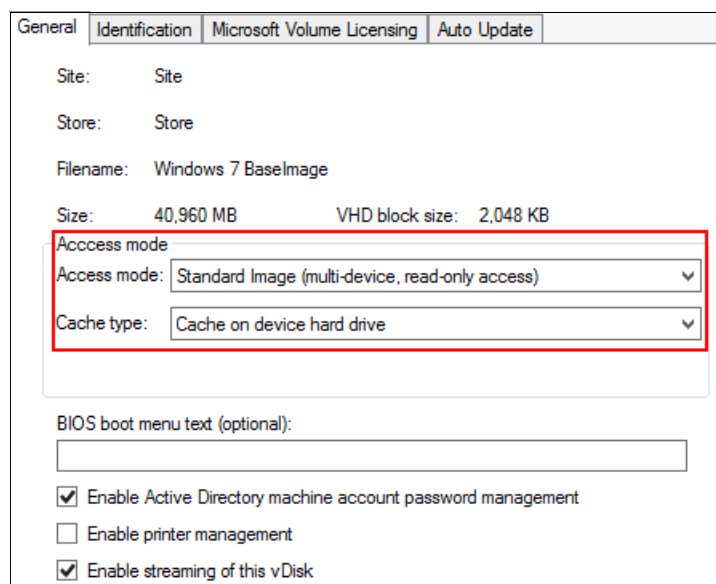


Figure 9-29 VDisk Access mode

3. After making these adjustments, power on your VM. After you log on, go to the D: drive to see the vdiskcache file that is created, as shown in Figure 9-30.

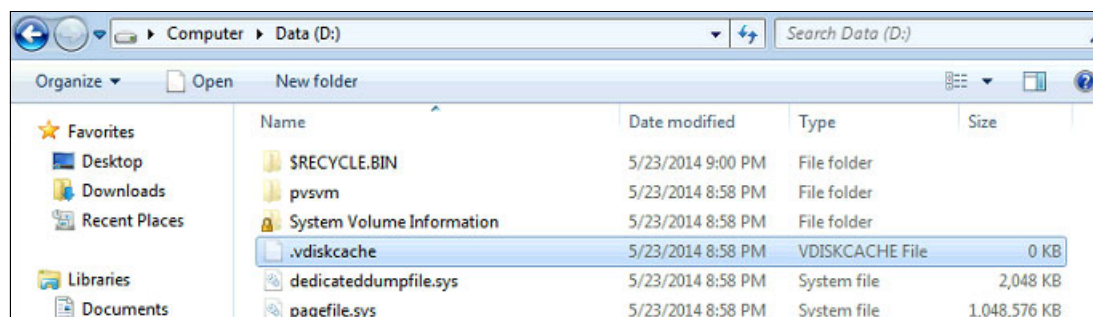


Figure 9-30 Write cache file verification

Note: From this point, any modification that you perform is lost after the VM restarts.

Before you create your VMs to deliver to the users, you must clone your gold image in a Hyper-v template. For more information about creating a template, see this website:

<http://technet.microsoft.com/en-us/library/hh427282.aspx>

You are now ready to create your desktops by integrating the XenDesktop and Provisioning Services. For more information about how to create the desktop catalogs and associate them for your domain users, see 9.3, “Configuring desktop distribution” on page 198.

9.2.2 Preparing the gold image for persistent desktops

The persistent desktop model uses Machine Creation Services instead of Provisioning Services to provision the desktop image. Therefore, you must create a separate gold image for dedicated desktops.

Complete the following steps to create a gold image for dedicated desktops:

1. After you insert the Citrix XenApp/XenDesktop 7.5 installation media, the AutoRun XenApp/XenDesktop window opens. Select **XenDesktop Deliver applications and desktops**.
2. Select **Virtual Delivery Agent for Windows Desktop OS**.
3. In the Environment window, select **Create a Master Image** (as shown in Figure 9-31) and click **Next**.

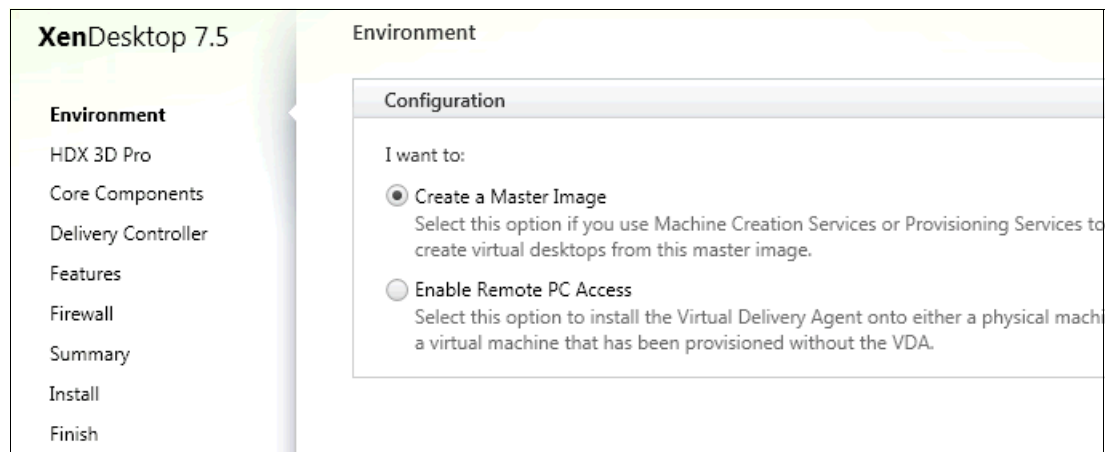


Figure 9-31 Creating a Master Image

4. Because we do not access a graphics processor, we select **No, install standard VDA** (as shown in Figure 9-32) and click **Next**.

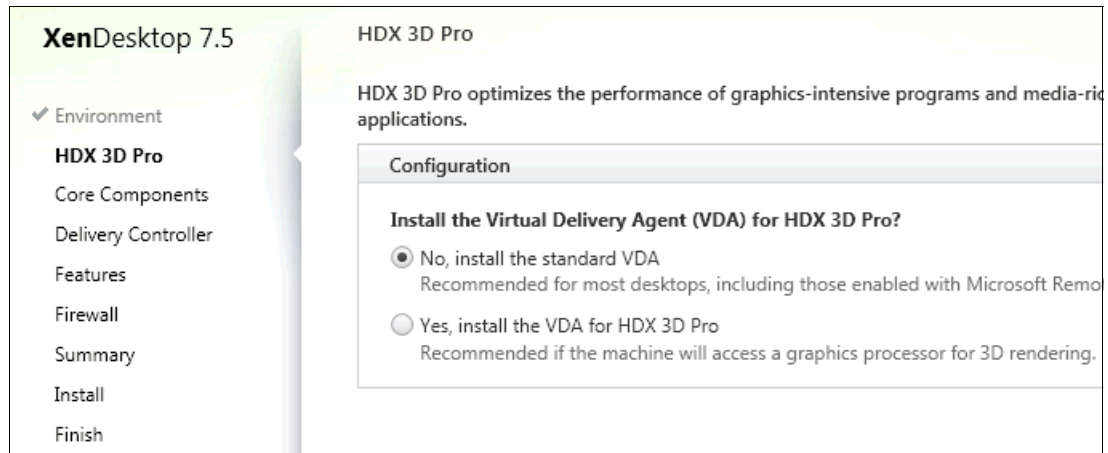


Figure 9-32 Install Standard VDA

5. Select **Citrix Receiver** (as shown in Figure 9-33) and click **Next**. The Citrix Receiver configuration is applied by using Group Policy Object (GPO), as described in 9.3, “Configuring desktop distribution” on page 198.

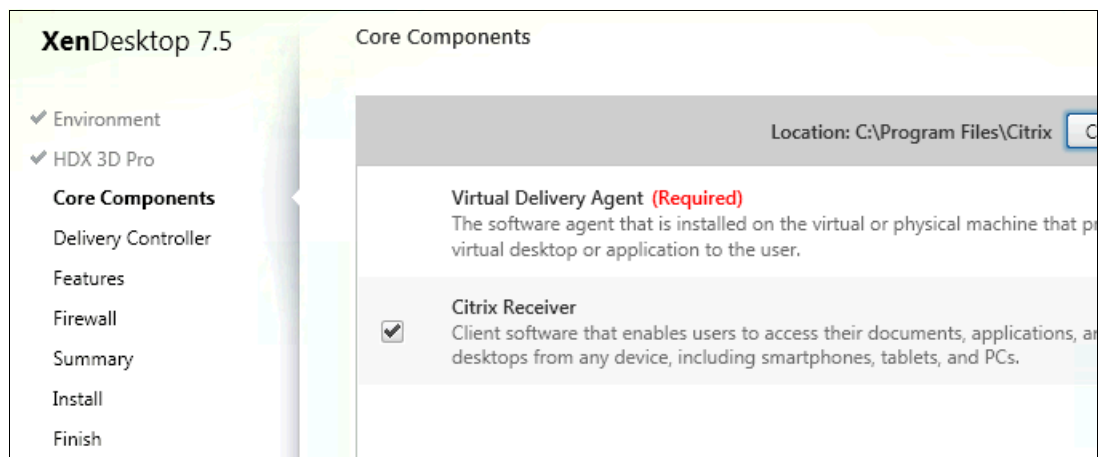


Figure 9-33 Component selection Citrix Receiver

6. Specify how XenDesktop locates the delivery controller (or controllers). For our installation, we selected **Do it Manually**, as shown in Figure 9-34.

The screenshot shows the 'XenDesktop 7.5' interface with a sidebar on the left containing links: Environment, HDX 3D Pro, Core Components, Delivery Controller (selected), Features, Firewall, Summary, Install, and Finish. The main panel is titled 'Delivery Controller' and contains a 'Configuration' section. It asks 'How do you want to enter the locations of your Delivery Controllers?' with a dropdown menu set to 'Do it manually'. Below this, there is a table with one entry: 'xdc.xenlab.local' with 'Edit' and 'Delete' links. Underneath the table is a 'Controller address:' label, an input field with the placeholder 'Example: controller1.domain.com', and two buttons: 'Test connection...' and 'Add'.

Figure 9-34 Configuring the XenDesktop Controller

We entered the name of our XenDesktop 7.5 Controller, then clicked **Test connection**. After the test is completed successfully, click **Add**. Repeat this step for other XenDesktop Controllers in your environment, then click **Next**.

7. For dedicated desktops, the personal VDisk is not used. Do not select the Personal VDisk option now, as shown in Figure 9-35. Click **Next**.

The screenshot shows the 'XenDesktop 7.5' interface with the sidebar on the left. The 'Features' section is selected in the sidebar and the main panel. The main panel is titled 'Features' and has a 'Feature (Select all)' checkbox at the top. Below it are four feature options, each with a checkbox and a 'Learn more' link: 'Optimize performance' (checked), 'Use Windows Remote Assistance' (checked), 'Use Real-Time Audio Transport for audio' (checked), and 'Personal vDisk' (unchecked). The 'Personal vDisk' option is described as 'Enable Personal vDisk for the Virtual Delivery Agent'.

Figure 9-35 Virtual Desktop Configuration

Optimize XenDesktop Performance, User Desktop Shadowing, and Real Time Monitoring are preselected.

8. The VDA installer offers to open the required ports in the Windows Firewall for you. Set configure firewall rules to **Automatically** to create rules in Windows Firewall, as shown in Figure 9-36. Click **Next**.

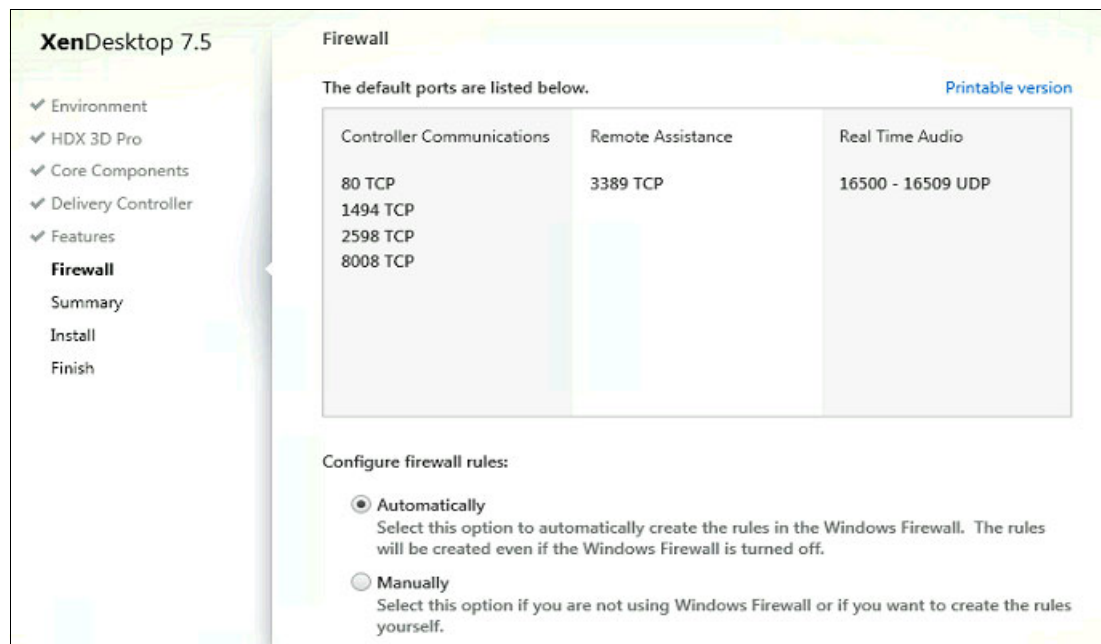


Figure 9-36 Configure Firewall rules

9. Figure 9-37 shows the installation summary. Click **Install** to complete the installation and restart your machine.

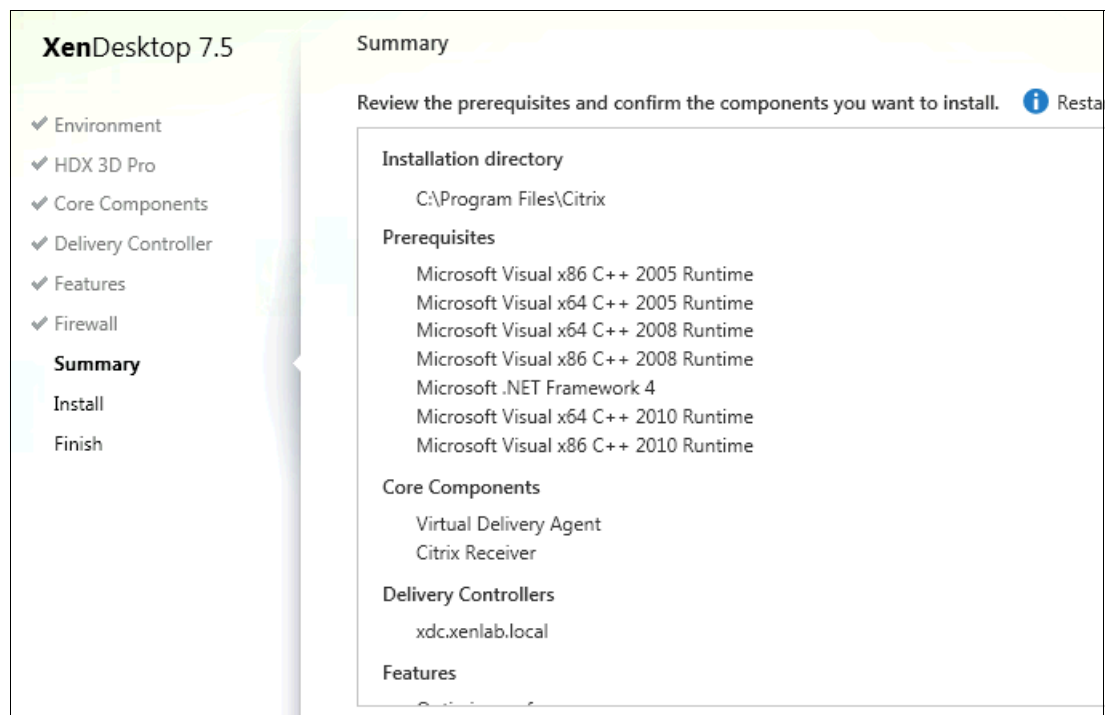


Figure 9-37 Installation Summary

Your gold image for persistent desktops is now ready.

For more information about the configuration process at the Desktop Studio to create and publish dedicated desktops, see 9.3.3, “Configuring persistent desktops” on page 216.

9.3 Configuring desktop distribution

The process to create desktop catalogs consists of creating a group of VMs that are based on the gold image that you created and making these VMs accessible to the users.

In this environment, we create the following types of desktop catalogs:

- Non-persistent streamed desktops

The catalog is *desktop streamed* at Citrix Desktop Studio. You create a catalog with a predetermined number of desktops that are integrated with Provisioning Services and associated to a group of users.

These desktops are available for use, but they are not fixed for these users. When the users log off and log on again, they can log on to any available desktop in the catalog.

Note: Because these desktops are non-persistent, any customizations that are made by the user are lost after the machine is restarted.

From a management perspective, if you modify your VDisk that is stored on PVS and release it for production, this new version is available for use the next time that your desktops restart.

- Non-persistent streamed with the personal VDisk desktops

The non-persistent streamed with personal VDisk (pvDisk) desktops are similar to the first catalog (they are integrated with Provisioning Services). However, in this catalog, a disk is created and associated with each desktop. On this disk, all customizations that are made by the users are stored to be available after the machine is restarted.

Another difference is that the desktop is associated to the user that logs on for the first time and is always associated with this user.

- Persistent desktops

Persistent (or dedicated) desktops consist of virtual desktops that are created by Machine Creation Services that are based on a template that is stored on your hypervisor.

This procedure creates a predetermined number of desktops that are available to a specific group of users.

When the user logs on to the desktop for the first time, the user is associated with this desktop and always uses this desktop.

From an administrative perspective, these desktops are not integrated with PVS, and new update requirements for security patches or business applications must be performed with the other tools. For more information, see *Endpoint Security and Compliance Management Design Guide Using IBM Tivoli Endpoint Manager*, SG24-7980.

9.3.1 Configuring streamed desktops

The process to configure the streamed desktop catalog starts at the Provisioning Services Console where you create the catalog and target devices. Then, the process finishes at the Desktop Studio where you grant the permission for a domain group to access these desktops.

Complete the following steps to configure this catalog:

1. In the Provisioning Services Console, create a device collection by selecting **Farms** → **Sites** → **your site name** → **Device Collections**. Enter the collection name and a description, as shown in Figure 9-38. Click **OK**.

General Security Auto-Add

Name:
Streamed_Desktops

Description:
Collection for Streamed Desktops

Figure 9-38 Device collection creation

2. In the Provisioning Services Console, run the XenDesktop wizard by right-clicking the site name that you created and selecting **XenDesktop Setup Wizard**, as shown in Figure 9-39.

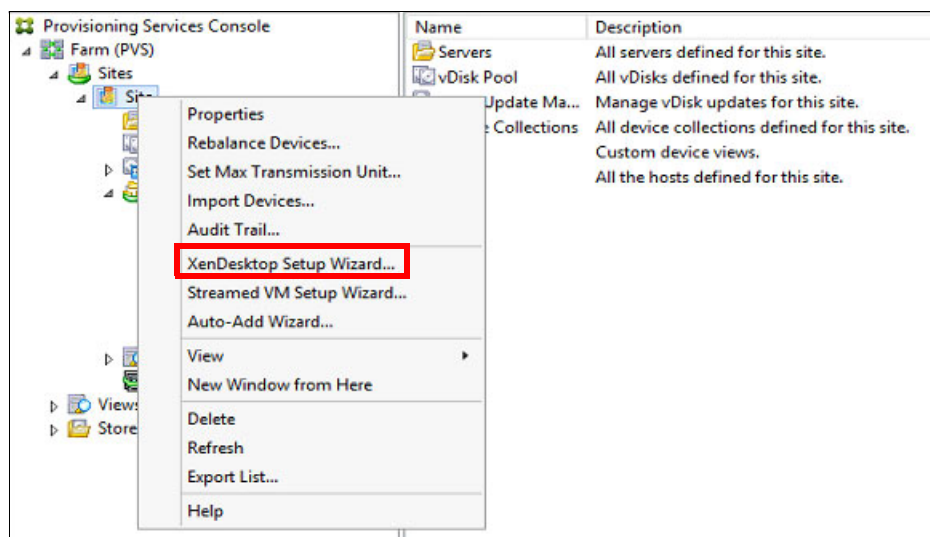


Figure 9-39 XenDesktop Setup Wizard

3. In the initial XenDesktop Setup window, click **Next**.
4. Specify the address of your XenDesktop Controller, as shown in Figure 9-40. Click **Next**.

XenDesktop Controller

Enter the address of the XenDesktop Controller you want to configure.

XenDesktop Controller address:
xdc.xenlab.local

Figure 9-40 XenDesktop Controller configuration

5. The wizard connects to your XenDesktop Host Resource, as shown in Figure 9-41. Click **Next**.




Figure 9-41 Host resource selection

6. Enter your user name and password in the Username and Password fields, as shown in Figure 9-42. Click **OK**.

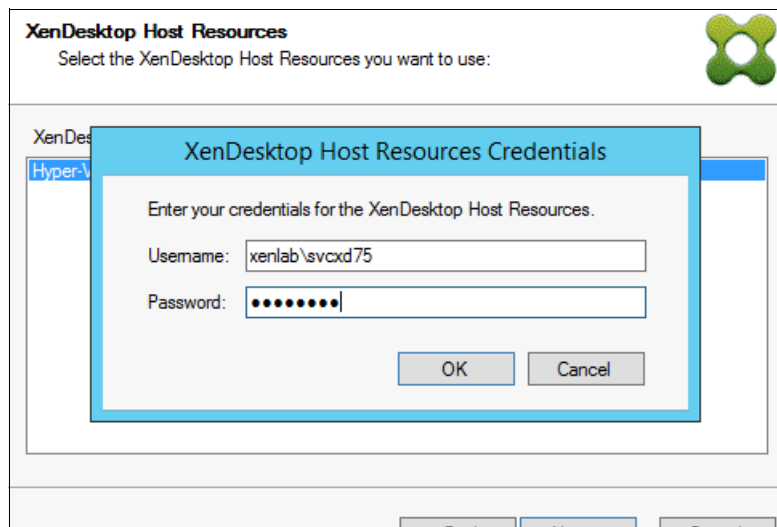


Figure 9-42 Host resources logon

7. Select a VM template, as shown in Figure 9-43. Click **Next**.

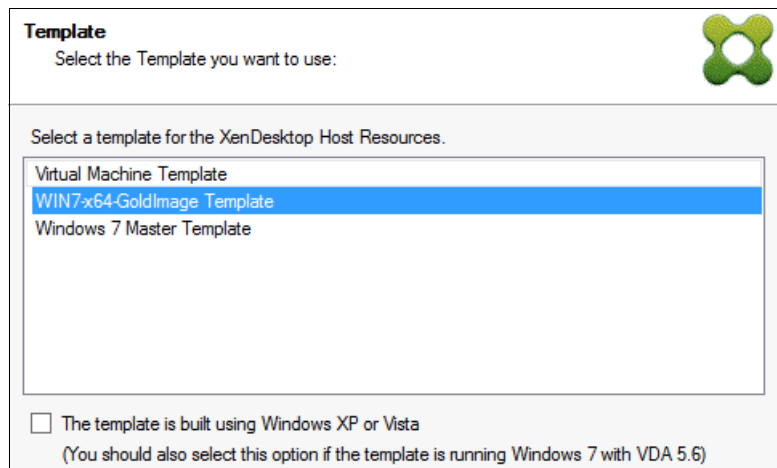


Figure 9-43 VM template selection

8. Select the VDisk, as shown in Figure 9-44. Click **Next**.




Figure 9-44 Device Collection selection

9. Select whether to Create a new catalog or Use an existing catalog, as shown in Figure 9-45. Complete the following fields:

- Catalog name: Specify the catalog name to be displayed in Desktop Studio.
- Description: Specify a description for the catalog that is created.

Click **Next**.

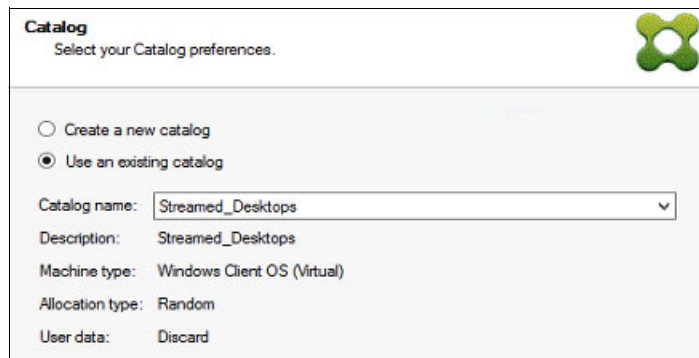


Figure 9-45 XenDesktop catalog creation

10. Select Windows Desktop Operating System, as shown in Figure 9-46. Click **Next**.

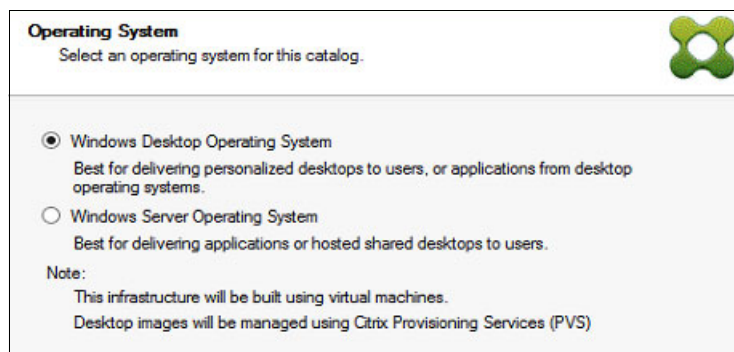
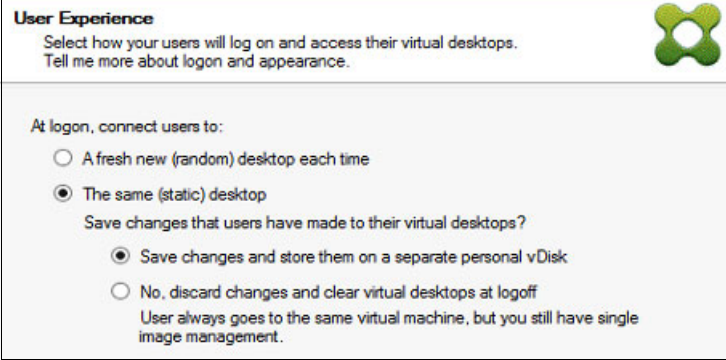


Figure 9-46 Catalog operating system selection

11. Because we are using PVD, select **The same (static) desktop**, also select **Save changes and store them on a separate personal VDisk**, as shown in Figure 9-47. Click **Next**.



The screenshot shows the 'User Experience' settings window. It has a title bar with the text 'User Experience' and a green Citrix logo. Below the title bar, there is a subtitle 'Select how your users will log on and access their virtual desktops. Tell me more about logon and appearance.' The main content area contains the following options:

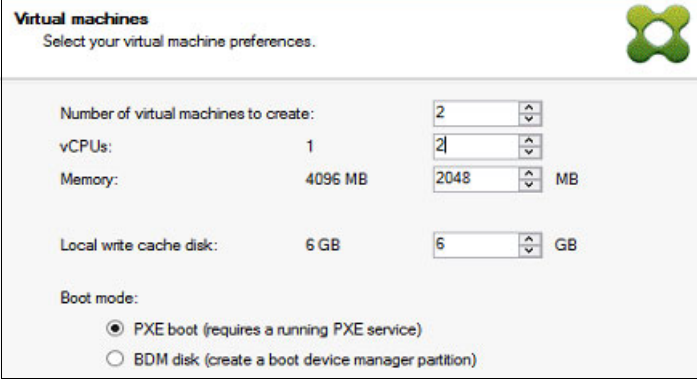
- At logon, connect users to:
 - ☐ A fresh new (random) desktop each time
 - ☒ The same (static) desktop
- Save changes that users have made to their virtual desktops?
 - ☒ Save changes and store them on a separate personal vDisk
 - ☐ No, discard changes and clear virtual desktops at logoff
User always goes to the same virtual machine, but you still have single image management.

Figure 9-47 Select random or static desktop

12. Define the following settings, as shown in Figure 9-48:

- Number of virtual machines to create: Select the number of desktops to create.
- VM characteristics:
 - vCPUs: Select the number of vCPUs.
 - Memory: Select the amount of memory.
- Active Directory computer accounts: Select whether computer accounts are created or accounts are reused (imported).

Click **Next**.



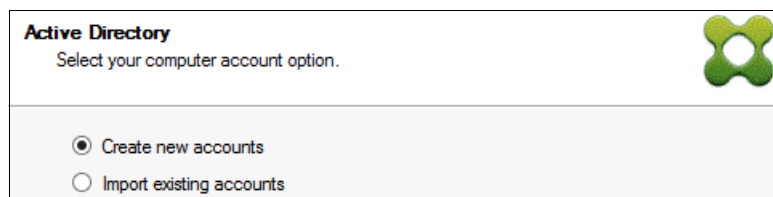
The screenshot shows the 'Virtual machines' settings window. It has a title bar with the text 'Virtual machines' and a green Citrix logo. Below the title bar, there is a subtitle 'Select your virtual machine preferences.' The main content area contains the following settings:

- Number of virtual machines to create: 2
- vCPUs: 1
- Memory: 4096 MB
- Local write cache disk: 6 GB
- Boot mode:
 - ☒ PXE boot (requires a running PXE service)
 - ☐ BDM disk (create a boot device manager partition)

Figure 9-48 Virtual machine preferences

Note: If you do not see the option Local write cache disk, you left the VDisk at the default of Cache on server. Exit this wizard, correct the VDisk properties, and rerun the wizard.

13. Select **Create new accounts** to have new Active Directory computer accounts created, as shown in Figure 9-49. Click **Next**



Active Directory

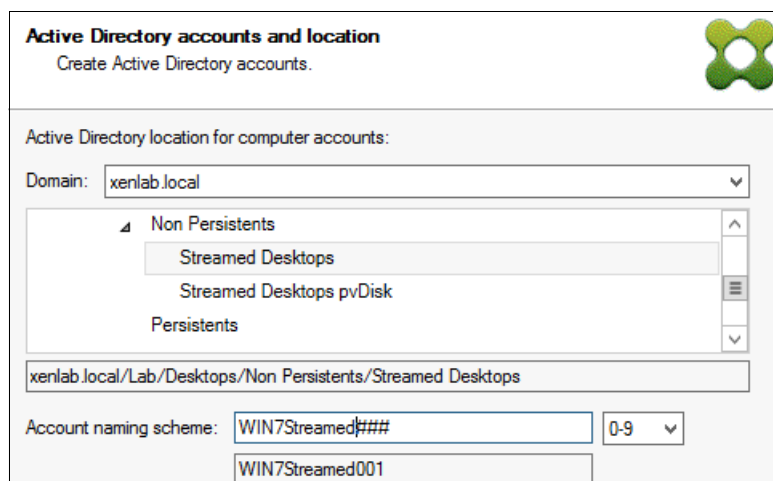
Select your computer account option.

☒ Create new accounts

☐ Import existing accounts

Figure 9-49 Select account option

14. To create Active Directory computer accounts, select the Domain, OU, Account naming scheme, as shown in Figure 9-50. Click **Next**.



Active Directory accounts and location

Create Active Directory accounts.

Active Directory location for computer accounts:

Domain: xenlab.local

Non Persistents

- Streamed Desktops
- Streamed Desktops pvDisk
- Persistents

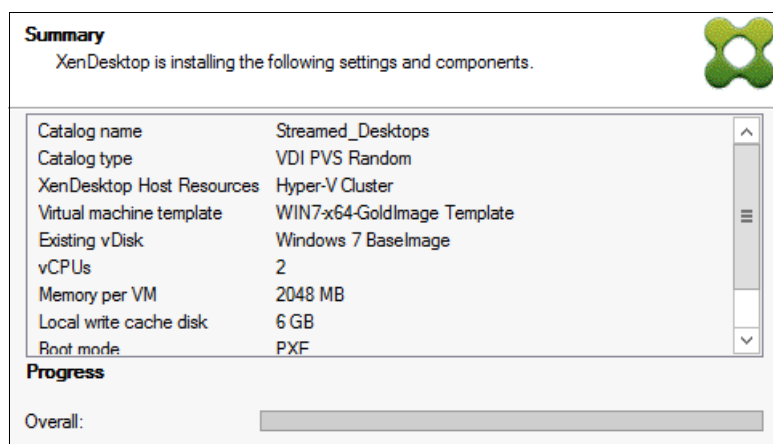
xenlab.local/Lab/Desktops/Non Persistents/Streamed Desktops

Account naming scheme: WIN7Streamed### 0-9

WIN7Streamed001

Figure 9-50 Active Directory accounts and location

15. At the Summary window, click **Finish** and the wizard creates the VMs, desktops, and target devices, as shown in Figure 9-51.



Summary

XenDesktop is installing the following settings and components.

Catalog name	Streamed_Desktops
Catalog type	VDI PVS Random
XenDesktop Host Resources	Hyper-V Cluster
Virtual machine template	WIN7-x64-GoldImage Template
Existing vDisk	Windows 7 BaseImage
vCPUs	2
Memory per VM	2048 MB
Local write cache disk	6 GB
Boot mode	PXF

Progress

Overall:

Figure 9-51 Summary window

16. When the wizard is complete, click **Done**. The setup is complete and a device is created (see Figure 9-52).




Figure 9-52 Execution process

17. Confirm the operation by refreshing the device collection that you created. Select **Provisioning Services Console** → **your site name** → **Device Collections** → **Win 7 WC PvD Test**. Right-click to select **Refresh**. Figure 9-53 shows the collection.

Name	MAC	Type	Disk	vDisk	IP Address
WIN7Streamed...	00-1D-D8-B7-1C-04	Production	vDisk	Store\Windows 7 Ba...	Down
WIN7Streamed...	00-1D-D8-B7-1C-03	Production	vDisk	Store\Windows 7 Ba...	Down

Figure 9-53 Target device list

18. Looking in Active Directory Users and Computers shows the new computer account, as shown in Figure 9-54.

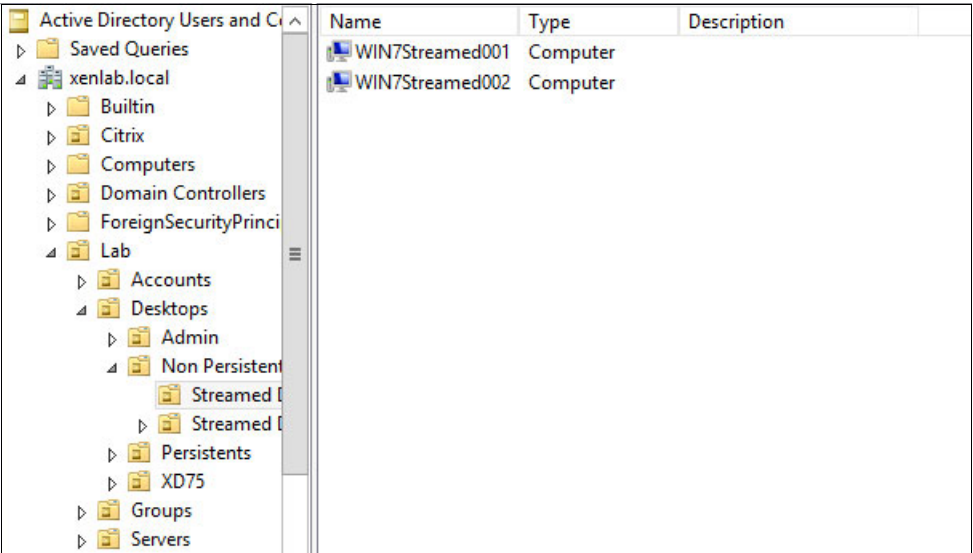


Figure 9-54 New computer account created

19. In the Desktop Studio, confirm the creation by selecting **Desktop Studio** → **Machines**. Right-click to select **Refresh**. Figure 9-55 shows the result.

Desktop OS Machines (2)		Server OS Machines (0)		Sessions (0)		
Name	Machine Catalog	Delivery Group	User	Maintenance M...	Persist User Cha...	Pow
WIN7Streamed...	Streamed_Desktops	-	-	Off	Discard	On
WIN7Streamed...	Streamed_Desktops	-	-	Off	Discard	On

Figure 9-55 Machine catalog

20. The next step is to associate the machine catalog that was created with a domain users group. Currently, there is no Delivery Group to deliver the desktops. Right-click the **Delivery Groups** in Citrix Studio and click **Create Delivery Group**, as shown in Figure 9-56.

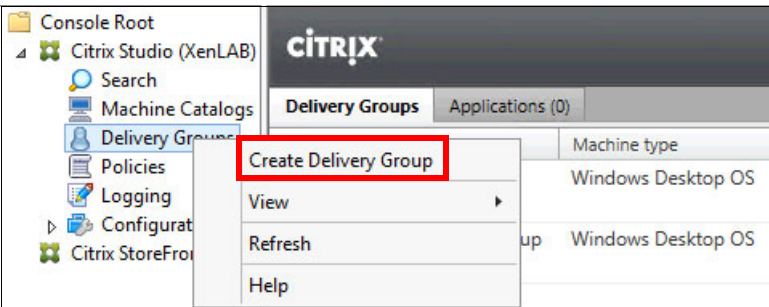


Figure 9-56 Create Delivery Group

21. In the Getting started window, click **Next** to continue.
22. Select the Machine Catalog and the number of machines to be added from the catalog to this delivery group, as shown in Figure 9-57. Click **Next**.

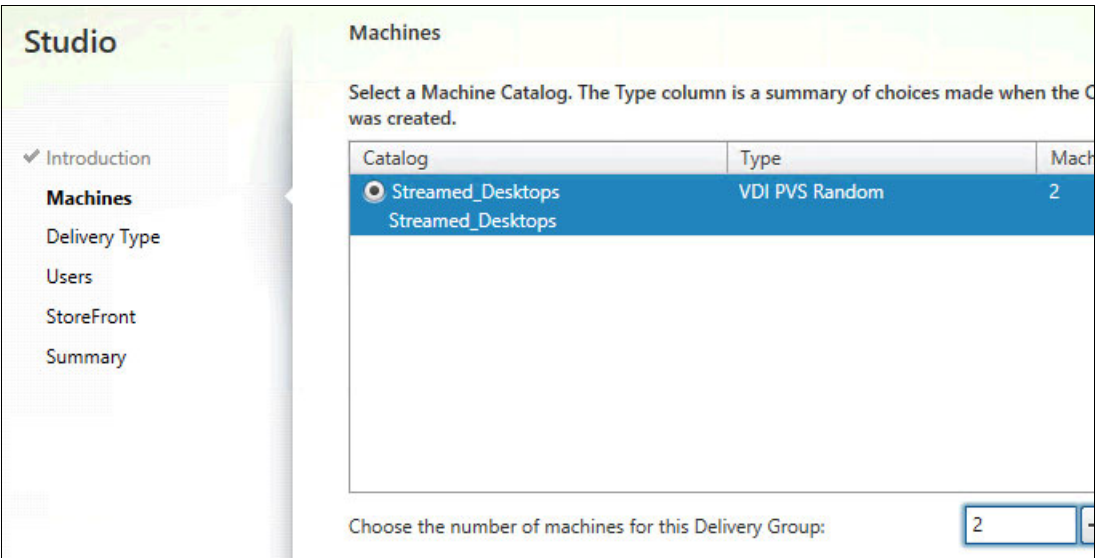


Figure 9-57 Catalog configuration

23. Select Desktops, as shown in Figure 9-58. Click **Next**.

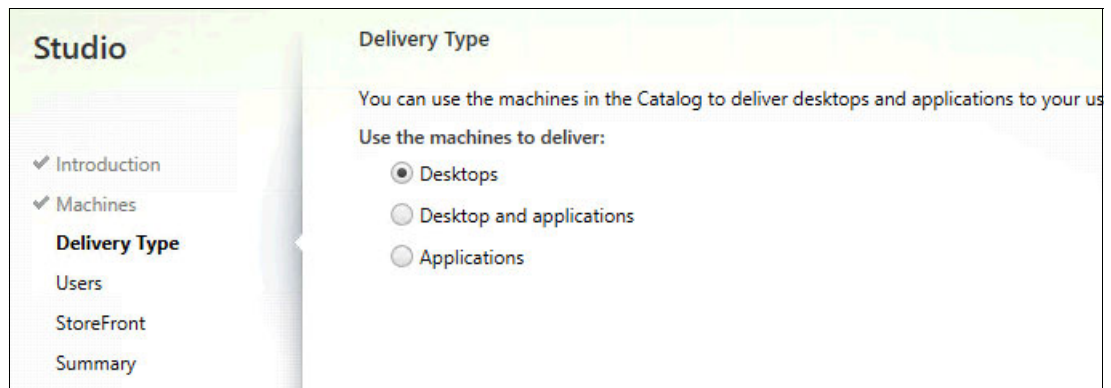


Figure 9-58 Delivery type

24. Click **Add users...**, as shown in Figure 9-59.

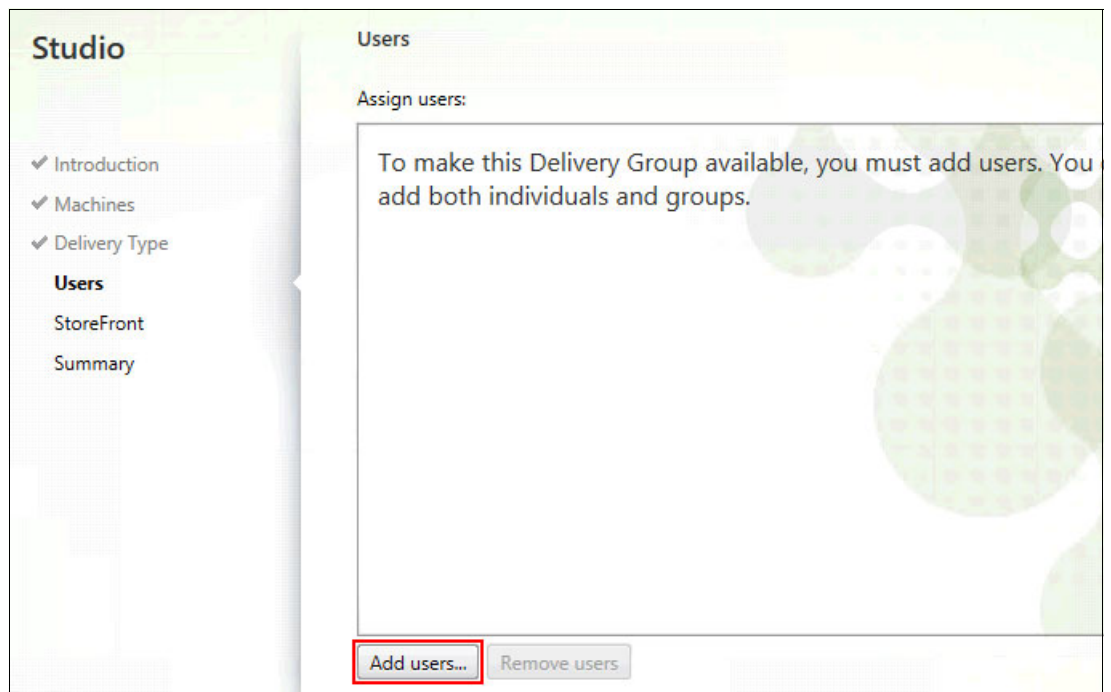


Figure 9-59 Assign users

25. Use the Select Users or Groups dialog to add users that can have access to the desktops, as shown in Figure 9-60. Click **Next**.

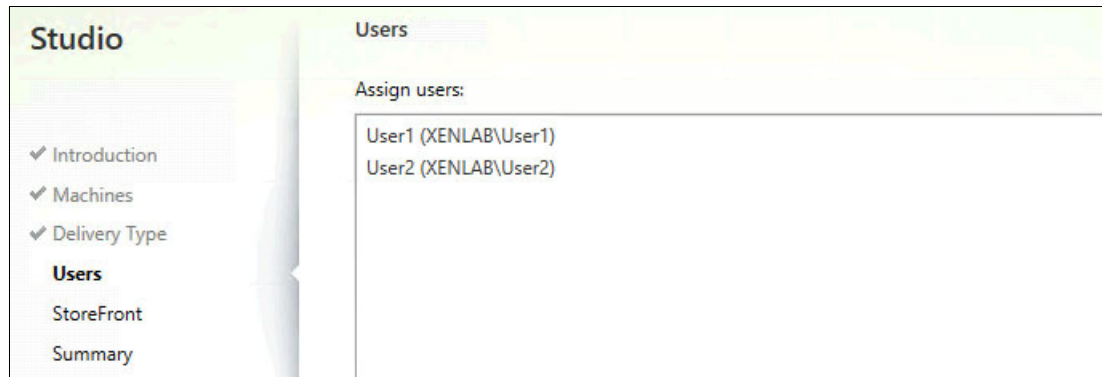


Figure 9-60 Domain users group selection

26. Select the appropriate **StoreFront** option, as shown in Figure 9-61. Click **Next**.

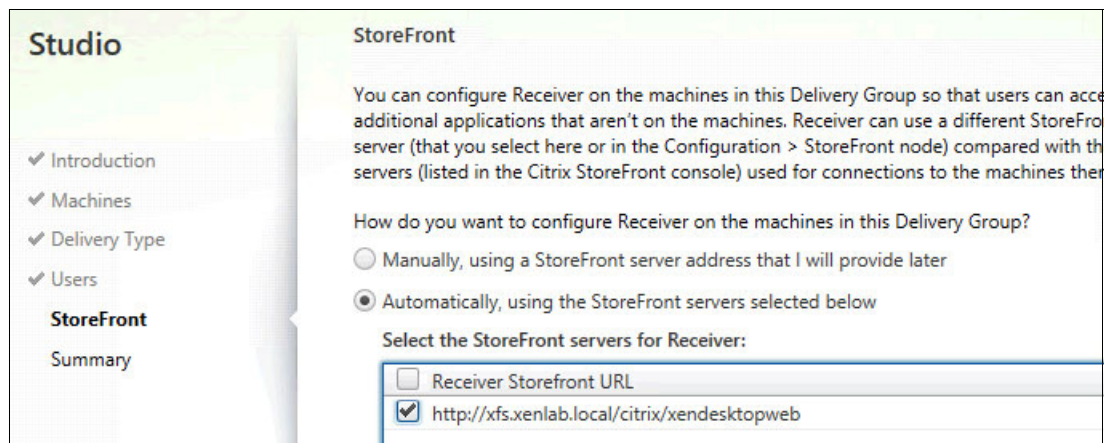


Figure 9-61 Select StoreFront option

27. Enter a Delivery Group name, Display name, an optional Delivery Group description for users, as shown in Figure 9-62. Click **Finish**.

Summary	
Machine Catalog:	Streamed_Desktops
Machine type:	Windows Desktop OS
Allocation type:	Random
Machines added:	2 unassigned
Delivery type:	Desktops
Users:	User1 (XENLAB\User1); User2 (XENLAB\User2)
Storefronts:	1
Scopes:	-
Delivery Group name:	
Streamed_Desktops	
Display name:	
WIN7_Streamed	
Delivery Group description, used as label in Receiver (optional):	
Streamed_Desktops	

Figure 9-62 Delivery group summary

28. The desktop that is created on Hyper-v is powered on and registers its Virtual Desktop Agent to the XenDesktop Controller to be available for users. Confirm this process by right-clicking **Delivery Group** and then selecting **Refresh**. Figure 9-63 shows the result.

CITRIX			
Delivery Groups		Applications (0)	
Delivery Group	Machine type	No. of machines	Sessions in use
Test delivery nonpersistent	Windows Desktop OS	0	
State: Enabled		Unregistered: 0	Disconnected
Win 7 PVD Test Deliverz Group	Windows Desktop OS	1	
State: Enabled		Unregistered: 0	Disconnected
Streamed_Desktops	Windows Desktop OS	2	
State: Enabled		Unregistered: 2	Disconnected

Figure 9-63 Delivery group status

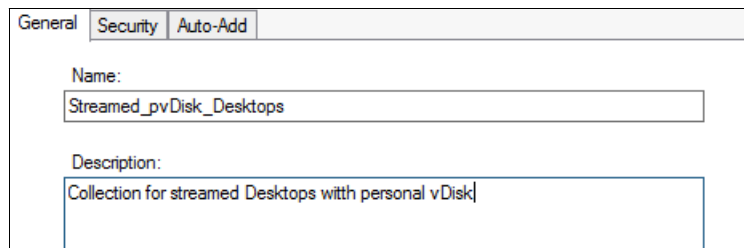
9.3.2 Configuring streaming desktops with personal VDisk

The process to configure the streamed desktop with personal VDisk is similar to configuring streaming desktops. The main difference is that to create streaming desktops with personal VDisk, the wizard creates another disk to store the user's customization.

Complete the following steps to configure streaming desktops with personal VDisk:

1. At the Provisioning Services Console, create a device collection by selecting **Farms/Sites** → **your site name** → **Device Collections**.

2. Enter a name and description for the device collection, as shown in Figure 9-64. Click **OK**.



General	Security	Auto-Add
Name: Streamed_pvDisk_Desktops		
Description: Collection for streamed Desktops with personal vDisk		

Figure 9-64 Device collection creation

3. After you create the device collection, run the XenDesktop wizard by right-clicking your site name and selecting **XenDesktop Setup Wizard**, as shown in Figure 9-65.

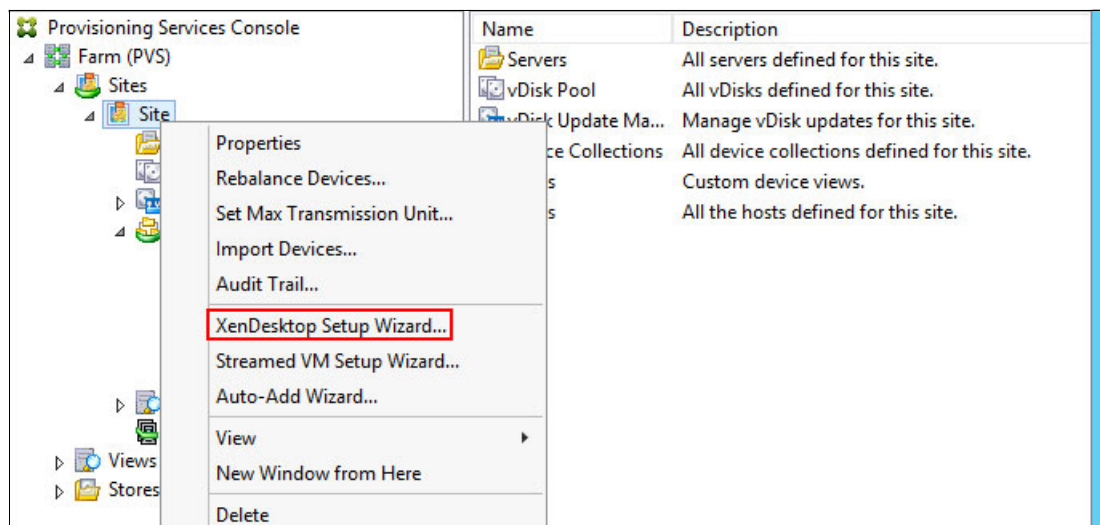
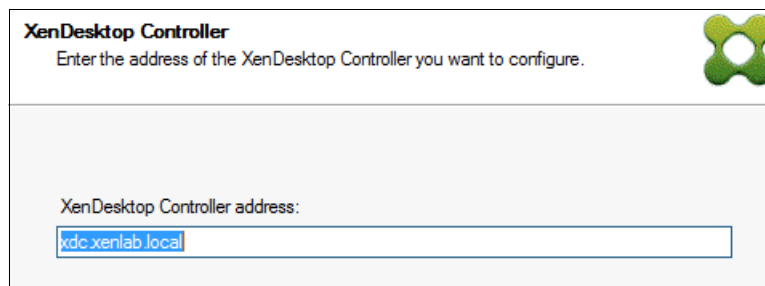


Figure 9-65 XenDesktop Setup Wizard

4. In the initial XenDesktop Setup window, click **Next**.
5. Enter the name of your XenDesktop Controller, as shown in Figure 9-66. Click **Next**.



XenDesktop Controller
Enter the address of the XenDesktop Controller you want to configure.

XenDesktop Controller address:
xdc.xenlab.local

Figure 9-66 XenDesktop Controller configuration

6. The wizard connects to your SCVMM to load the defined templates. Select the host resource (Hyper-V Cluster in our example), as shown in Figure 9-67. Click **Next**.

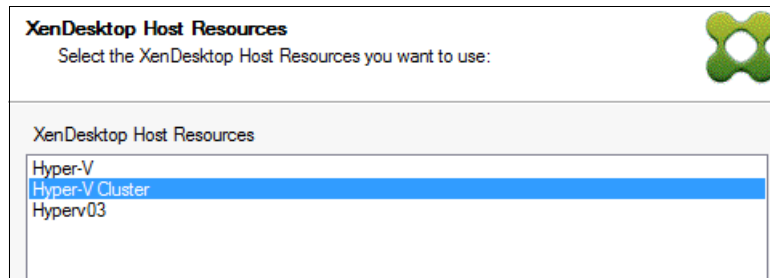


Figure 9-67 Host resource selection

7. Enter the logon credentials for the host resource and click **OK**, as shown in Figure 9-68.




Figure 9-68 Host resource credentials

8. Select the appropriate template, as shown in Figure 9-69. Click **Next**.

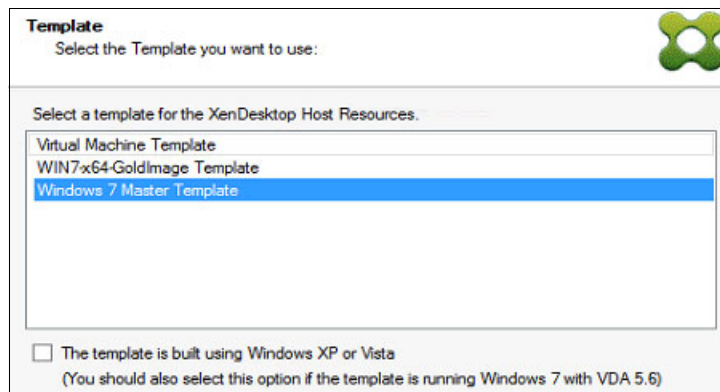


Figure 9-69 VM template selection

9. Select the VDisk, as shown in Figure 9-70. Click **Next**.




Figure 9-70 VDisk selection

10. Define your preferences to create the catalog in XenDesktop Controller. Select whether to **Create a new catalog** or **Use an existing catalog**, as shown in Figure 9-71. Click **Next**.

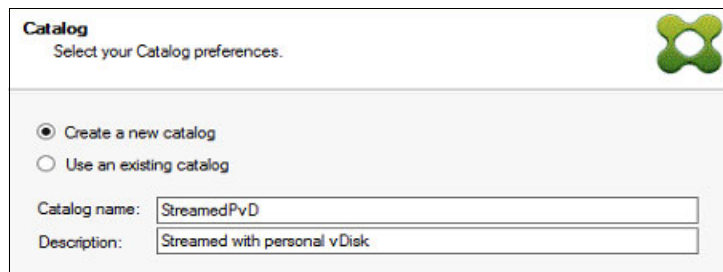


Figure 9-71 Catalog preferences

11. Make the appropriate choices. For this lab, we are creating two VMs with 1 vCPUs, 2 GB RAM, a 10 GB write cache disk, a 20 GB PvD disk, and changing the PvD drive to Y: as shown in Figure 9-72. Click **Next**.

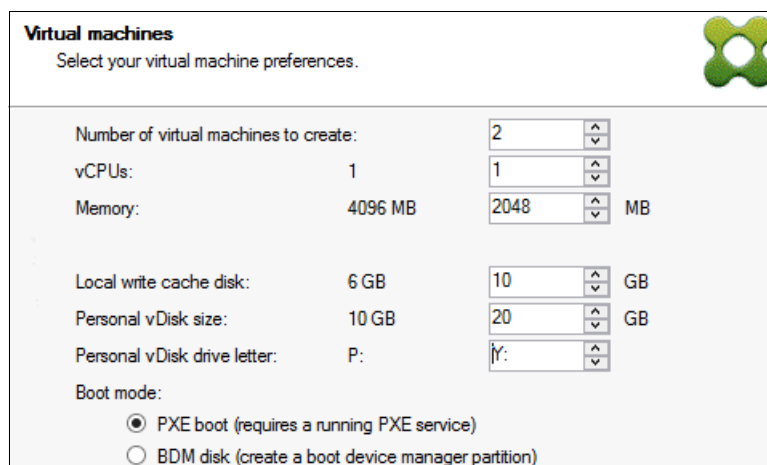


Figure 9-72 VM characteristics

12. Select **Create new accounts** to have the AD computer accounts created, as shown in Figure 9-73. Click **Next**.

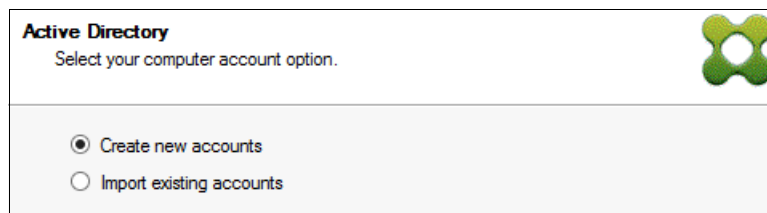


Figure 9-73 Account option

13. Select the Domain, OU, and Account naming scheme, as shown in Figure 9-74. Click **Next**.

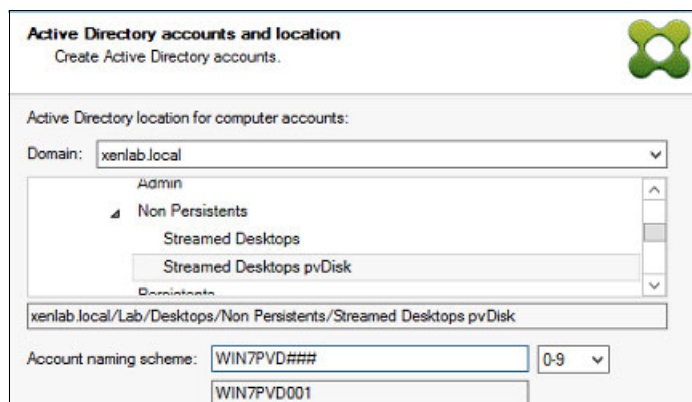


Figure 9-74 Active Directory location and computer naming scheme

14. Click **Finish** and the wizard creates the VMs, desktops, and target devices, as shown in Figure 9-75.




Figure 9-75 Wizard Summary

15. The wizard creates VMs at Hyper-v, the target devices at the Provisioning Server, and the computer accounts at the Active Directory (see Figure 9-76). When the wizard is complete, click **Done**.




Figure 9-76 VM Setup completed

16. Reviewing the Device Collection in the PVS console shows the two target devices, as shown in Figure 9-77.




Figure 9-77 Created devices

17. Looking in Active Directory Users and Computers shows the new computer accounts, as shown in Figure 9-78.

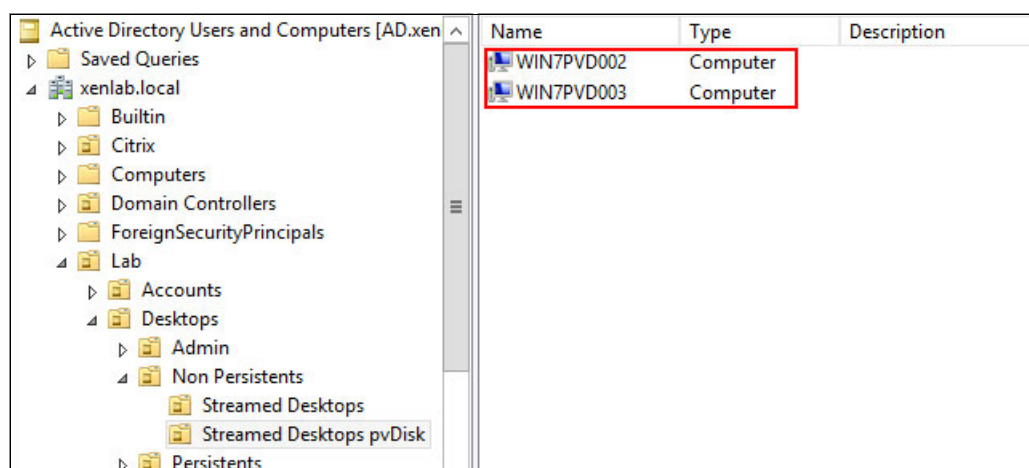


Figure 9-78 Created streamed desktops

18. The next step is to associate the machine catalog that was created with a domain users group. Currently, there is no Delivery Group to deliver the desktops. Right-click the **Delivery Groups** in Citrix Studio and select **Create Delivery Group**, as shown in Figure 9-79.

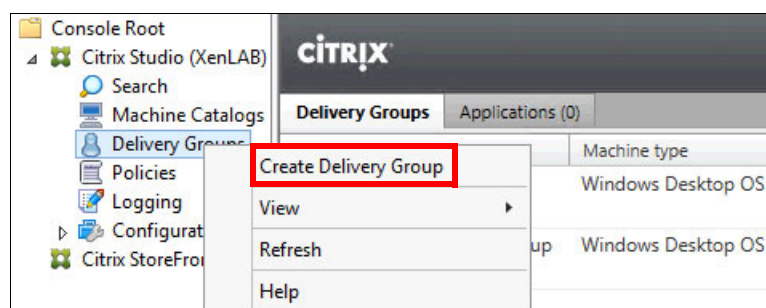


Figure 9-79 Create delivery group

19. At the Getting started window, click **Next** to continue.

20. Select the desktop catalog that you created. Then, define how many desktops are available for the users, as shown in Figure 9-80. Click **Next**.

Catalog	Type	Mach
Streamed PvD	VDI PVS Static Personal vDisk	2

Choose the number of machines for this Delivery Group: 2

Figure 9-80 Catalog configuration

21. Select Desktops, as shown in Figure 9-81. Click **Next**.

You can use the machines in the Catalog to deliver desktops and applications to your user.

Use the machines to deliver:

☒ Desktops

☐ Applications

Figure 9-81 Delivery type

22.To make this Delivery Group available, assign the users by clicking **Add Users.....**, as shown in Figure 9-82. Click **Next**.

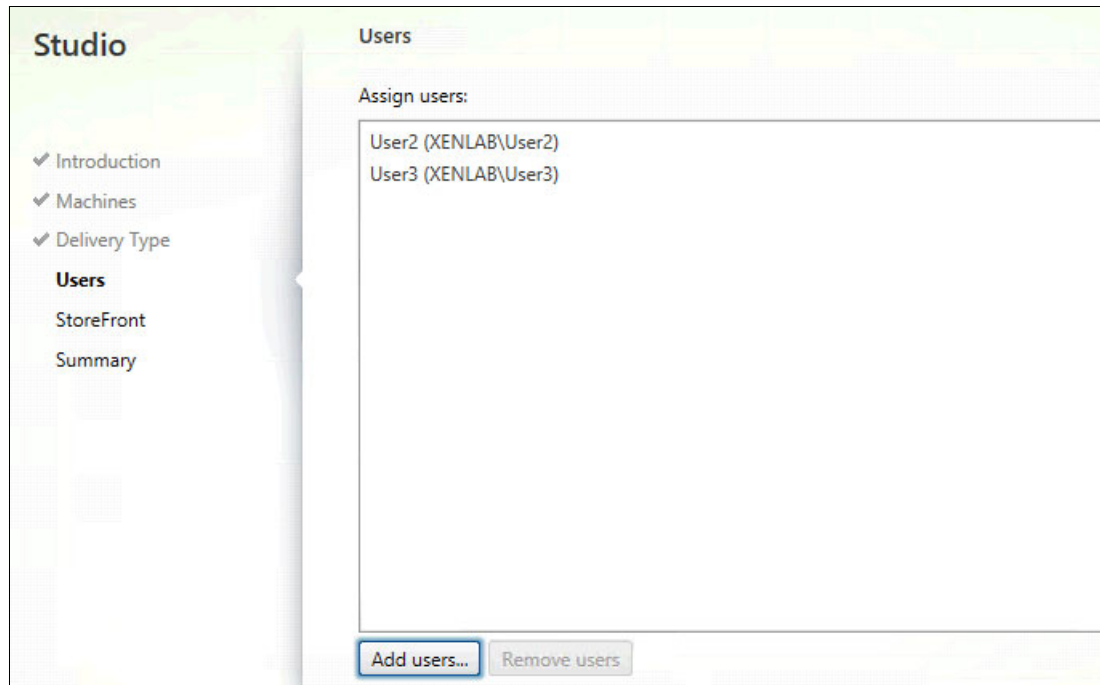


Figure 9-82 Assign users

23.Select the appropriate StoreFront option, as shown in Figure 9-83. Click **Next**.

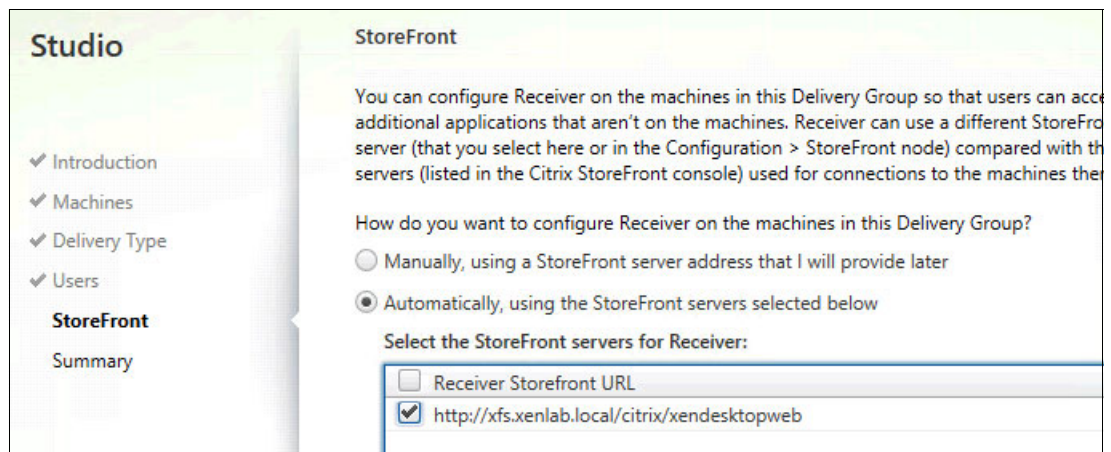


Figure 9-83 StoreFront option

24. Enter a Delivery Group name, Display name, an optional Delivery Group description for users, as shown in Figure 9-84. Click **Finish**.

Studio

- ✓ Introduction
- ✓ Machines
- ✓ Delivery Type
- ✓ Users
- ✓ StoreFront
- Summary**

Summary

Machine Catalog: Streamed PvD
Machine type: Windows Desktop OS
Allocation type: Static
Machines added: 2 unassigned
Delivery type: Desktops
Users: User2 (XENLAB\User2); User3 (XENLAB\User3)
Storefronts: 1
Scopes: -

Delivery Group name:
Streamed PvD

Display name:
PvD Desktops

Delivery Group description, used as label in Receiver (optional):
Streamed PvD

Figure 9-84 Desktop group summary

25. The desktop that is created on Hyper-v is powered on and registers its Virtual Desktop Agent to the XenDesktop Controller to be available for users. Confirm this process by right-clicking **Delivery Group** and then selecting **Refresh** (see Figure 9-85).

CITRIX

Console Root

- Citrix Studio (XenLAB)
 - Search
 - Machine Catalogs
 - Delivery Groups**
 - Policies
 - Logging
 - Configuration
 - Citrix StoreFront

Delivery Group	Machine type	No. of machines	Sessions in use
Streamed_Desktops	Windows Desktop OS	2	
State: Enabled		Unregistered: 2	Disconn
Streamed PvD	Windows Desktop OS	2	
State: Enabled		Unregistered: 0	Disconn
Win 7 PvD Test Deliverz Group	Windows Desktop OS	1	
State: Enabled		Unregistered: 0	Disconn

Figure 9-85 Desktop delivery group status

9.3.3 Configuring persistent desktops

In this task, we perform the steps necessary to create a catalog to be used with Machine Creation Services. Machine Creation Services uses a master VM within your XenDesktop environment to manage VMs, which enables you to easily administer and update target devices through one master image.

Use Desktop Studio to configure the collection for persistent desktops by using the gold image that was created in 9.2.2, "Preparing the gold image for persistent desktops" on page 194.

Complete the following steps to create the desktop collection and publish the collection for your users:

1. In Desktop Studio, right-click **Machines Catalog**, and select **Create Machine Catalog**, as shown in Figure 9-86.




Figure 9-86 Create Machine Catalog

2. Select **Windows Desktop OS** as the Operating System, as shown in Figure 9-87. Click **Next**.

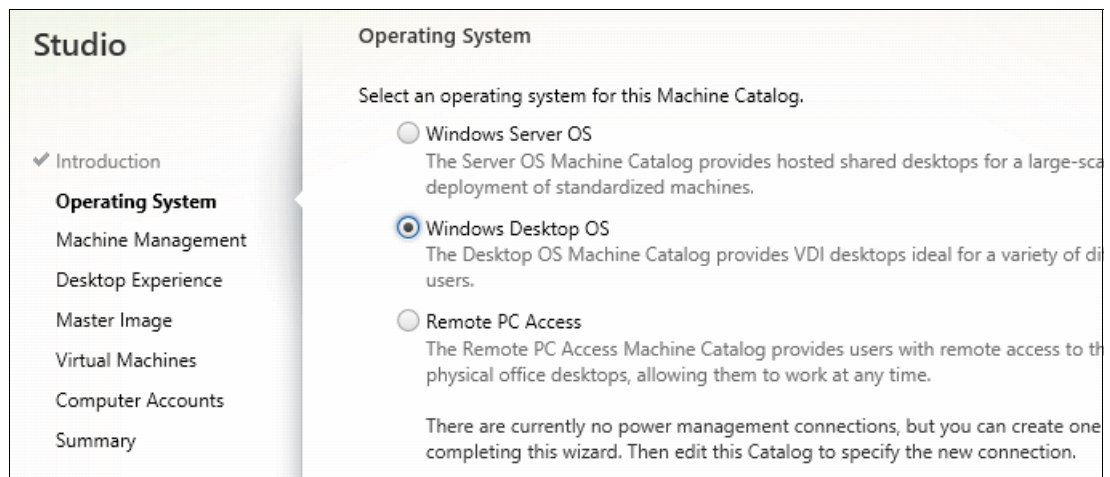


Figure 9-87 Collection type

3. In the Machine Management window, in the Deploy machines using section, select **Citrix Machine Creation Services (MCS)** and **Hyper-V** from the Resources pull-down menu, as shown in Figure 9-88. Click **Next**.

Figure 9-88 Select Machine Management

4. Select **I want users to connect to the same (static) desktop each time they log on** and **Yes, create a dedicated virtual machine and save changes on local disk**, as shown in Figure 9-89. Click **Next**.

Figure 9-89 Desktop experience

5. Select the Master Image, as shown in Figure 9-90. Click **Next**.

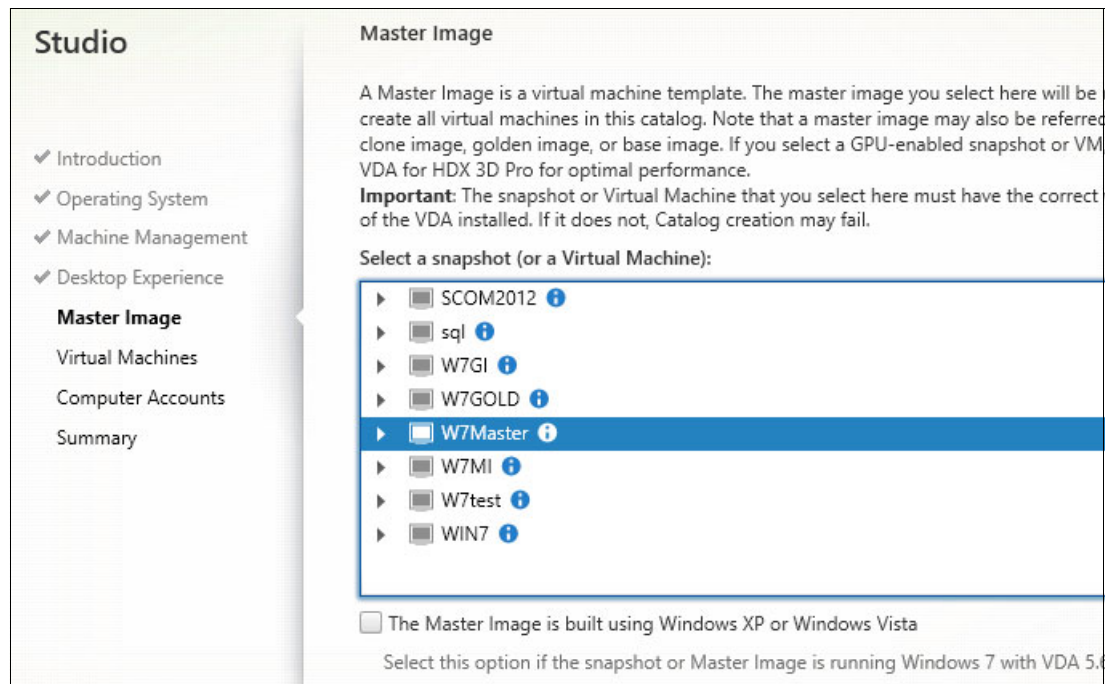


Figure 9-90 Select the Master Image

6. Be sure that the selected network is dedicated to client traffic (VLAN20), as shown in Figure 9-91. Click **Next**.

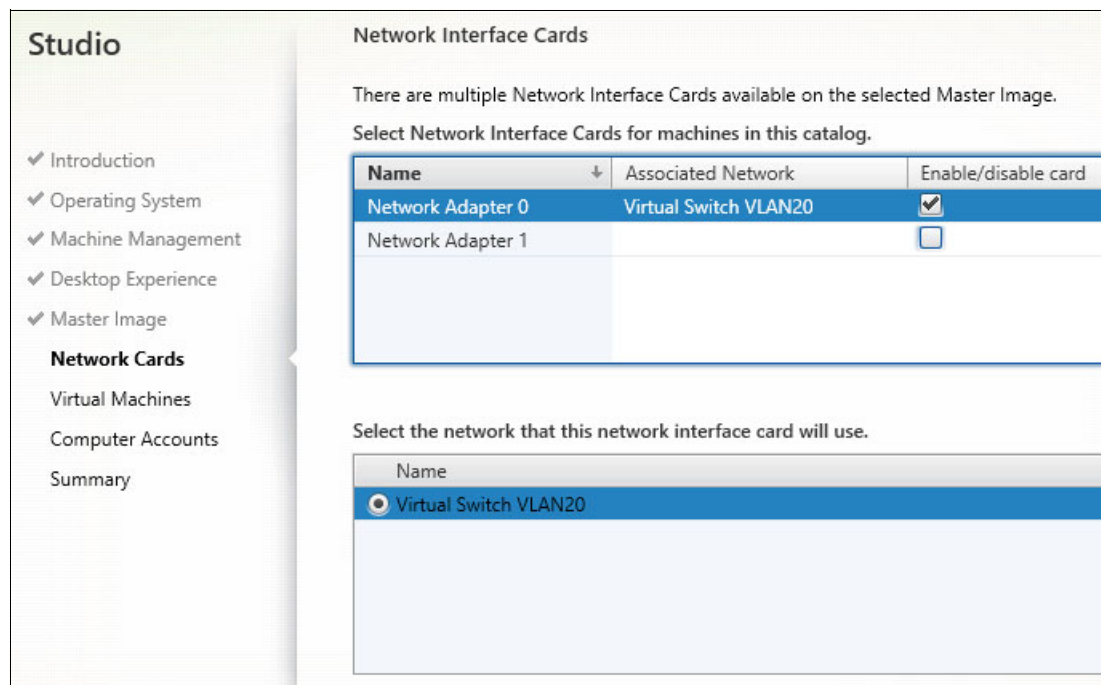


Figure 9-91 Select Network cards

7. Select the number of virtual desktops to create. Select the number of vCPUs and the amount of memory to allocate for these desktops, as shown in Figure 9-92. Click **Next**.

Studio

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- ✓ Network Cards
- Virtual Machines**
- Computer Accounts
- Summary

Virtual Machines

Number of virtual machines needed:

5

Configure your machines:

Name:	W7Master		
Virtual CPUs:	2		
Memory (MB):	1024		
Hard disk (GB):	40		

Figure 9-92 VM characteristics

8. Select the domain and then select the organizational unit (OU) and the account naming scheme for the new desktops, as shown in Figure 9-93. Click **Next**.

Studio

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- ✓ Network Cards
- ✓ Virtual Machines
- Computer Accounts**
- Summary

Active Directory Computer Accounts

Each machine in a Machine Catalog needs a corresponding Active Directory computer account.

Select an Active Directory account option:

☒ Create new Active Directory accounts

☐ Use existing Active Directory accounts

Active Directory location for computer accounts:

Domain: xenlab.local

- Default OU
- Citrix**
 - Computers
 - Domain Controllers
 - ForeignSecurityPrincipals

Selected location: OU=Citrix,DC=xenlab,DC=local

Account naming scheme:

Win7Persistent# 0-9

Win7Persistent0

Figure 9-93 Active Directory specifications

- In the Summary window, enter the Machine catalog name and an optional description, as shown in Figure 9-94. Click **Finish**.

Studio

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- ✓ Network Cards
- ✓ Virtual Machines
- ✓ Computer Accounts
- Summary**

Summary

Machine type:Windows Desktop OS

Machine management:Virtual

Provisioning method:Machine creation services (MCS)

Desktop experience:Users connect to the same desktop each time they log on
Save changes on the local disk

Resources:Hyper-V

Master Image name:W7Master
A snapshot of the Master Image VM will be created

Network interface cards:Network Adapter 0 - Using Virtual Switch VLAN20

Machine Catalog name:

Win7 Persistent

Machine Catalog description for administrators: (Optional)

Persistent Win 7

To complete the deployment, assign this Machine Catalog to a Delivery Group by selecting a Delivery Group and then Create or Edit a Delivery Group.

Figure 9-94 Machine Catalog Setup Summary window

- After the process completes, you should see the new Machine Catalogs, as shown in Figure 9-95.

Common Tasks

Use this screen to perform common maintenance tasks.

Site configuration		
Task	Administrator	Time
Create Machine Catalog 'WIN7NP'	XENLAB\svcXD75	5/22/2014 : 6:27:59 PM
Delete Machine Catalog 'WIN7NP'	XENLAB\svcXD75	5/22/2014 : 6:17:38 PM
Remove Catalog metadata	XENLAB\svcXD75	5/22/2014 : 6:14:13 PM

Machine catalogs

WIN7NP	5	Machines
--------	---	----------

Figure 9-95 Created Machine Catalogs listed

9.3.4 Assigning a catalog to a group

In this task, we create a delivery group to be used with the Machine Creation Services catalog of the desktop machines that were created.

Complete the following steps to assign a catalog to a group:

1. From Citrix Studio, right-click the **Delivery Groups** and click **Create Delivery Group**, as shown in Figure 9-96. In the Introduction window, click **Next** to continue.

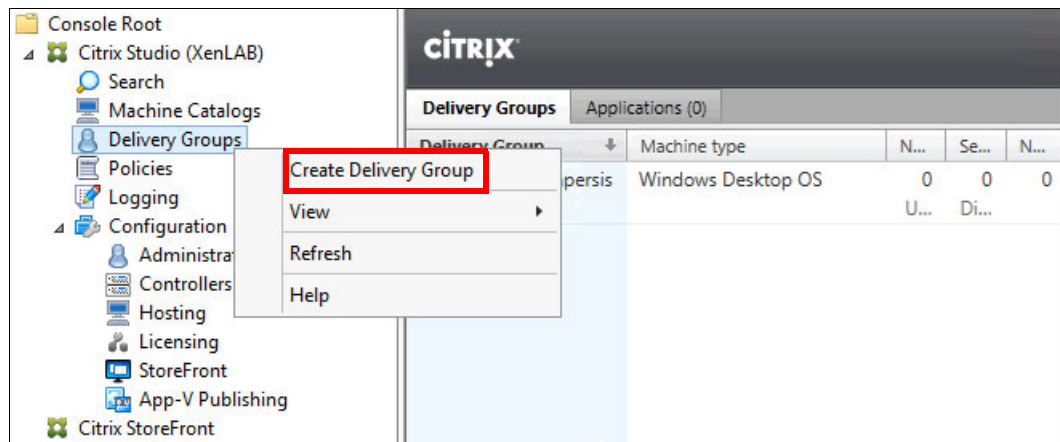


Figure 9-96 Create delivery group

2. Select the catalog that you created in the previous procedure and specify how many virtual desktops are available to the users, as shown in Figure 9-97. Click **Next**.

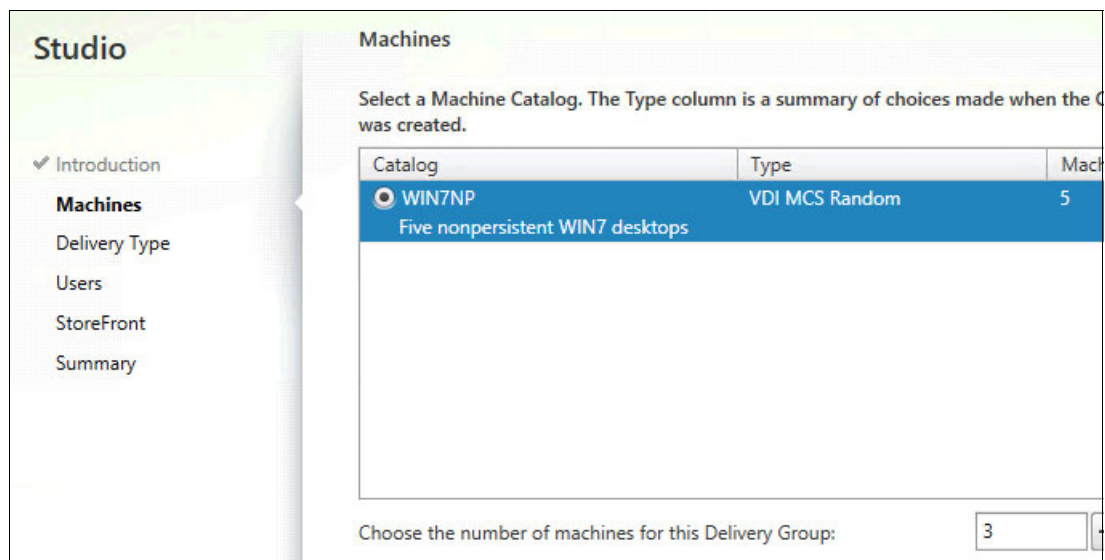


Figure 9-97 Select Machine Catalog

3. For Delivery Type, select **Desktops** and click **Next**, as shown in Figure 9-98.

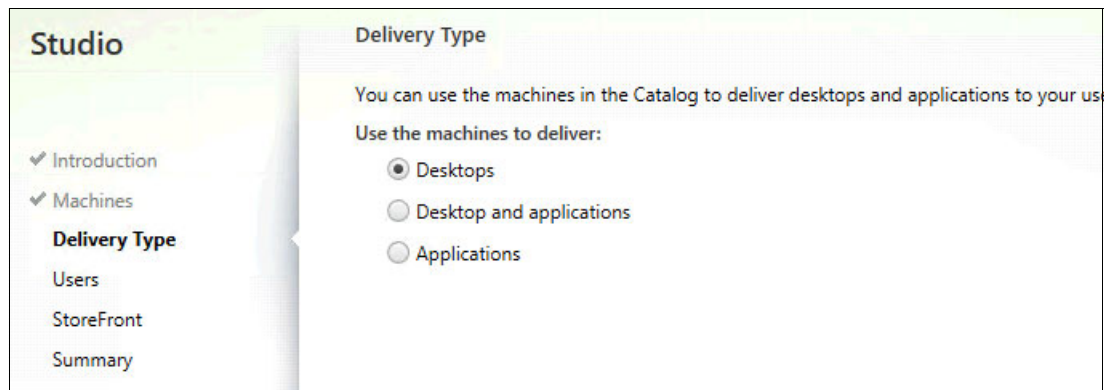


Figure 9-98 Delivery type

4. Click **Add users...** Add Domain Users, as shown in Figure 9-99. Click **OK** and then click **Next**.

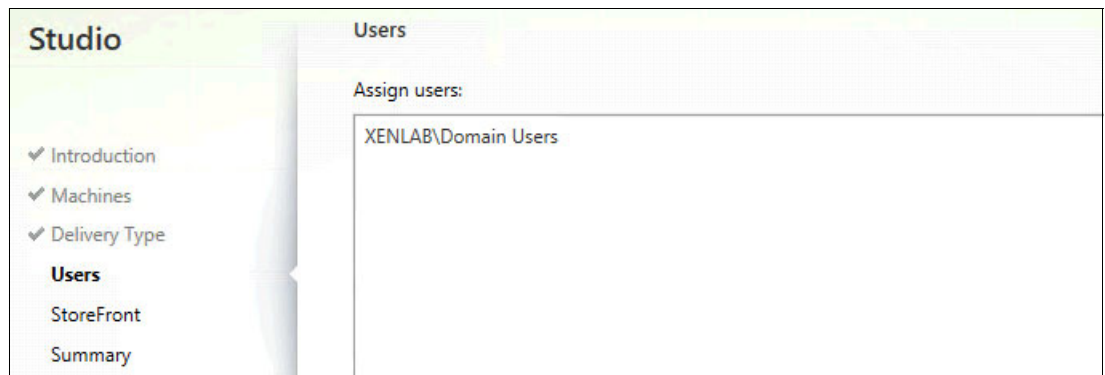


Figure 9-99 Add users

5. Select the appropriate StoreFront option, as shown in Figure 9-100. Click **Next**.

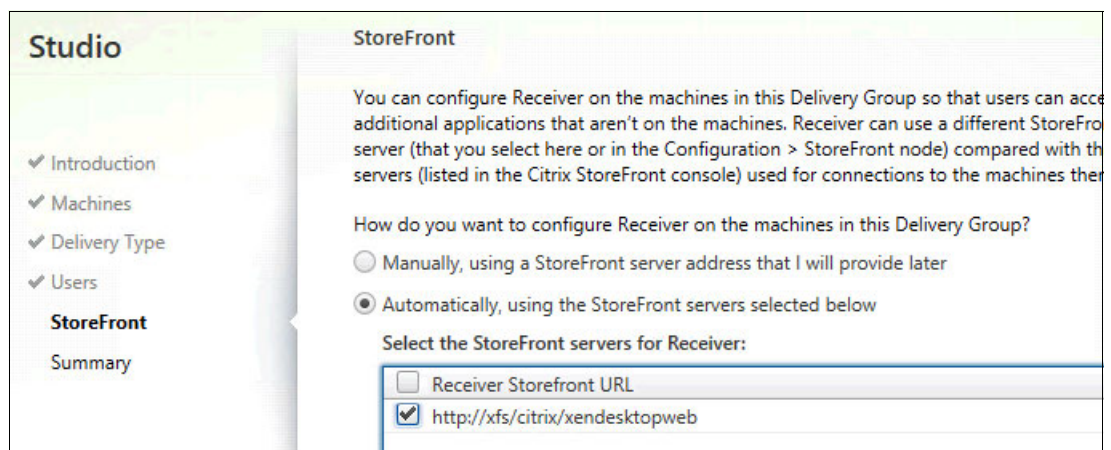


Figure 9-100 StoreFront option

6. Specify **Delivery Group name** and **Display name**, as shown in Figure 9-101. Click **Finish**.

The screenshot shows the 'Studio' configuration window with the 'Summary' tab selected. On the left, a navigation pane lists various configuration steps, with 'Summary' highlighted. The main area displays a summary of the configuration settings for a new Machine Catalog.

Summary	
Machine type:	Windows Desktop OS
Machine management:	Virtual
Provisioning method:	Machine creation services (MCS)
Desktop experience:	Users connect to the same desktop each time they log on Save changes on the local disk
Resources:	Hyper-V
Master Image name:	W7Master A snapshot of the Master Image VM will be created
Network interface cards:	Network Adapter 0 - Using Virtual Switch VLAN20

Machine Catalog name:

Win7 Persistent

Machine Catalog description for administrators: (Optional)

Persistent Win 7

To complete the deployment, assign this Machine Catalog to a Delivery Group by selecting Delivery Groups and then Create or Edit a Delivery Group.

Figure 9-101 Group name

After the process completes, confirm that the new virtual desktops are available to the users, as shown in Figure 9-102.

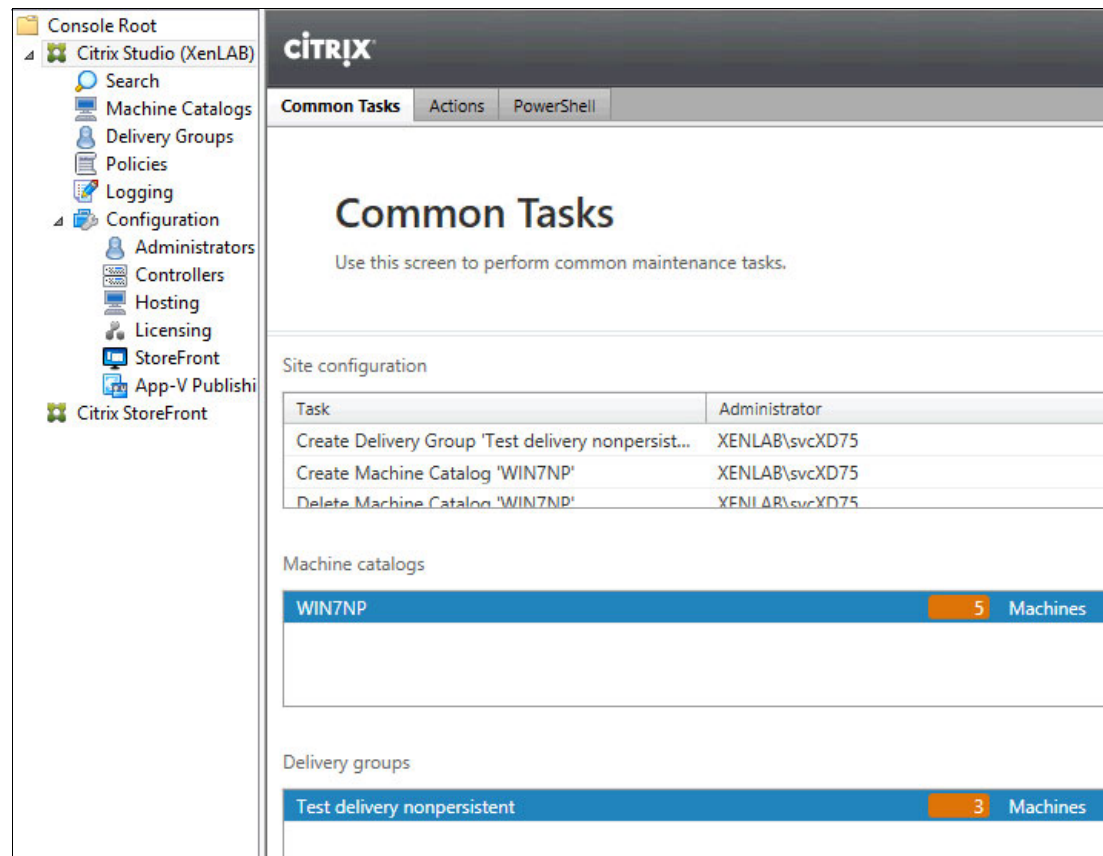


Figure 9-102 Desktop group status

9.4 Roaming profiles and folder redirection

Defining roaming user profiles is an important part of the VDI configuration because it allows profiles to be stored in a centralized server. When a user logs on from a different desktop, that user's profile is loaded.

The concept is similar for folder redirection. Redirection occurs to the same centralized server for accessing personal folders, such as My Documents, Favorites, and Desktop, is important for the following main reasons:

- ▶ If the user logs on from a different desktop.
- ▶ If the administrator modifies the desktop image, the user's files are not lost because the files are stored on a centralized server.
- ▶ For business applications, you can redirect the application settings folder to the same centralized server, and the configuration is loaded without reconfiguring.

To configure the roaming profile and folder redirection, we used the Group Policy on the Active Directory to centralize the configuration and to apply it at the organizational unit (OU) level to standardize the configuration for the environment.

The following sections describe the procedures for implementing roaming profiles and folder redirection.

9.4.1 Configuring the roaming profile

To start roaming profile functionality, we used the Citrix Profile Manager UPM that is installed on the desktop gold image and on the Citrix XenApp servers. Citrix provides an administrative template (ADM) that must be imported on the Group Policy Object to configure how the Profile Manager works.

Complete the following steps to configure the GPO:

1. Log in to My Citrix and browse to **Downloads** → **XenApp** → **Components** to download the Citrix Profile Management software by using the following website:
<https://www.citrix.com/downloads/xenapp/components.html>
2. After the download completes, extract the file and run `profilemgt_x64.msi` on your XenApp server.
3. In the Welcome window, click **Next**.
4. Accept the license agreement and choose the destination folder.
5. Click **Install**.
6. After the setup process is complete, click **Finish** and restart the server, as shown in Figure 9-103.

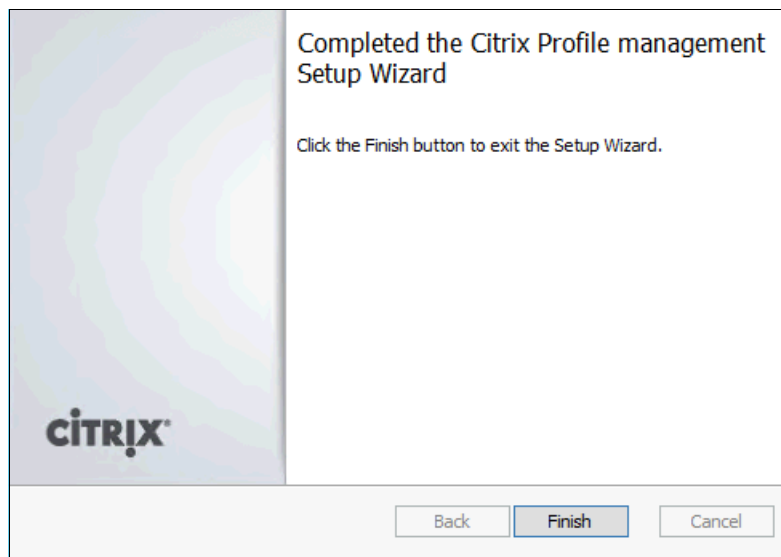


Figure 9-103 Setup completed

7. In the Group Policy Management, right-click your OU and select **Create a GPO in this domain, and Link it here...**, as shown in Figure 9-104.

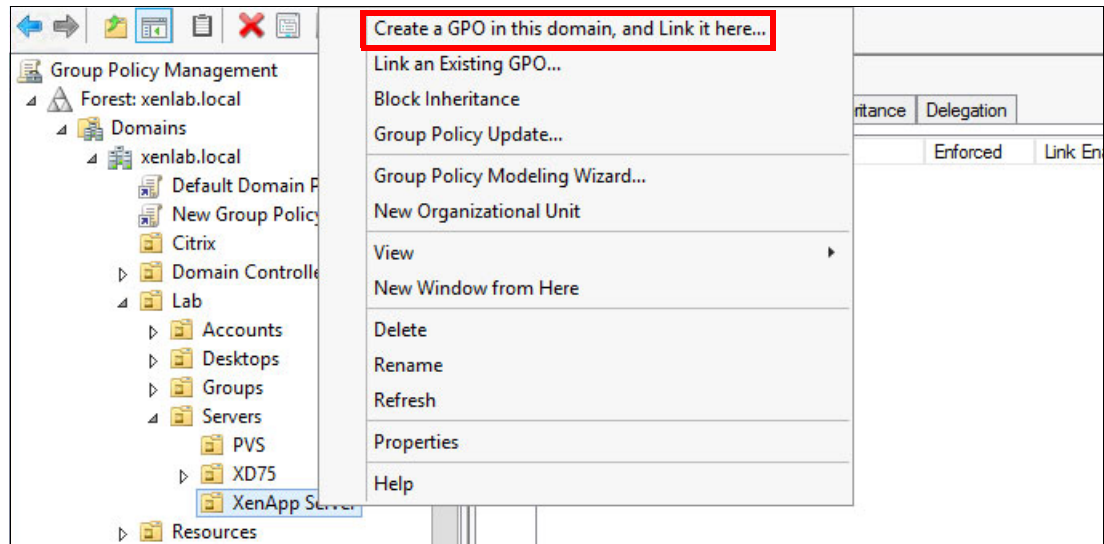


Figure 9-104 Create a GPO

8. Enter a name for the policy; for example, ProfilesManagement, as shown in Figure 9-105.

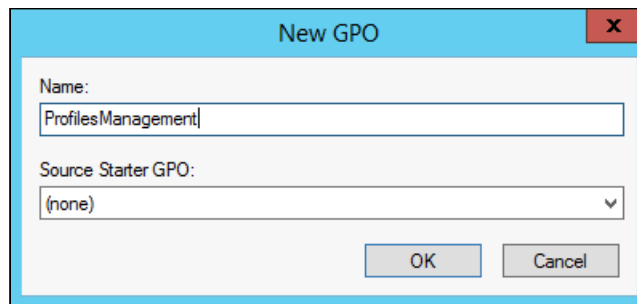


Figure 9-105 New GPO

9. Right-click the new policy, then select **Edit**, as shown in Figure 9-106.

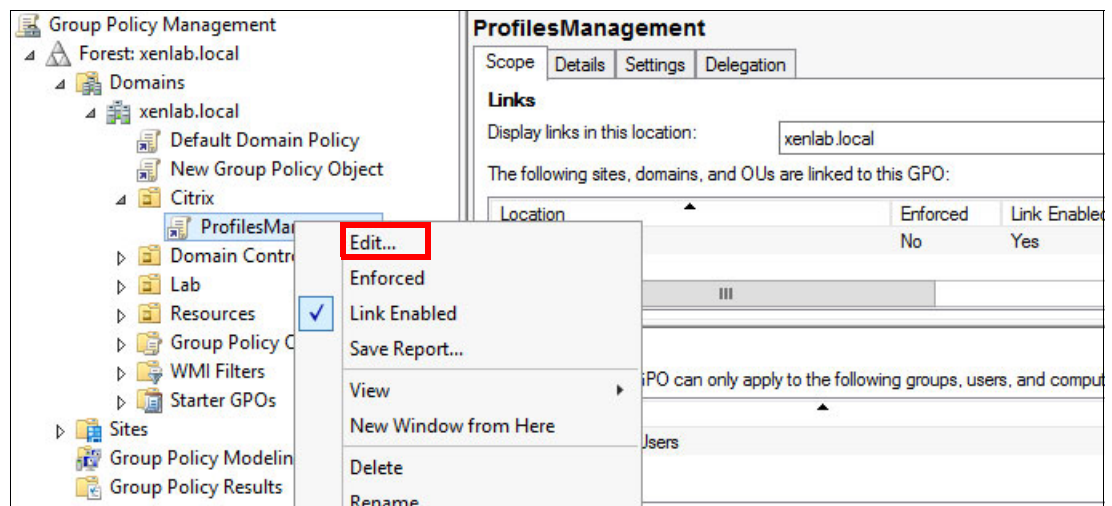


Figure 9-106 Edit new policy

10. Click **Computer Configuration**, then right-click **Administrative Templates** and select **Add/Remove Templates...** to import the Citrix Profile Manager ADM template, as shown in Figure 9-107.

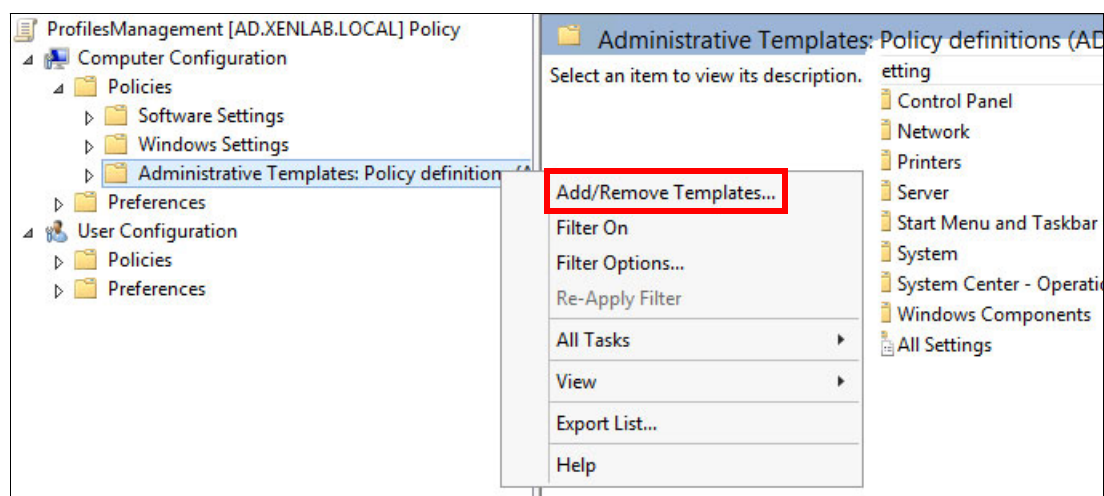


Figure 9-107 GPO ADM template import

11. In the opened window, click **Add** to import the templates.
12. Open the `ctxprofile5.1.1.adm` file that is in the subfolder `\GPO_Templates\en`, as shown in Figure 9-108.

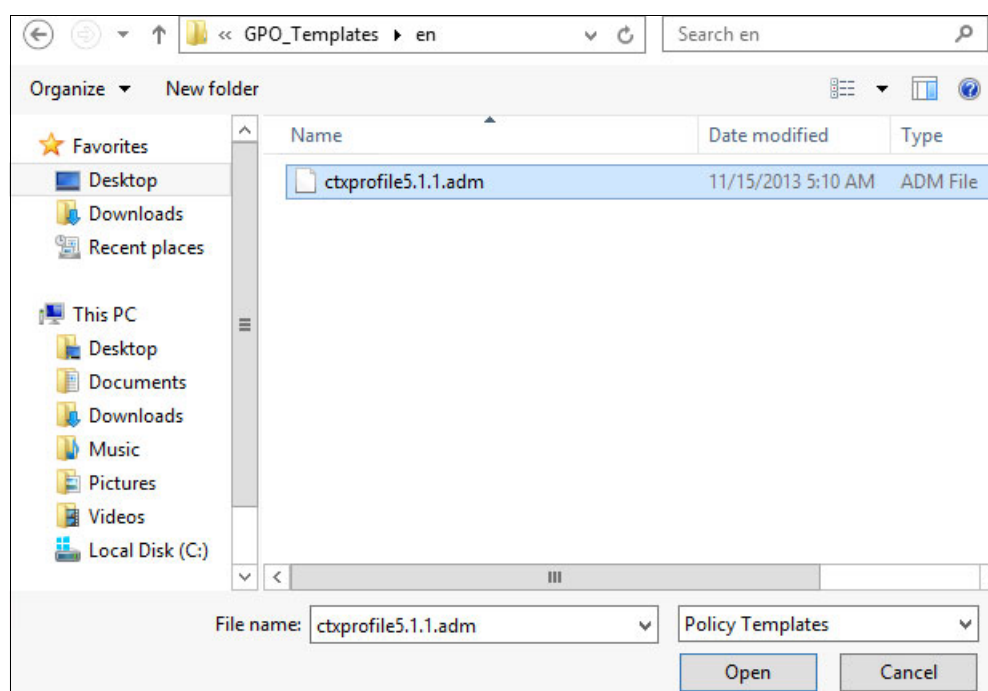


Figure 9-108 GPO template selection

13. The template appears in the Add/Remove Templates window, as shown in Figure 9-109. Click **Close**.




Figure 9-109 Template added

After the import process is complete, you must customize the template according to your requirements. For this implementation, the Citrix Profile Manager is configured to process all logons from a group that is called UPMUsers and to store the users' profiles at a centralized file server.

14. As shown in Figure 9-110, we selected **UPM_FolderRedirectionGPO Policy** → **Computer Configuration** → **Policies** → **Administrative Templates: Policy definitions (ADMX files)** → **Classic Administrative Templates (ADM)** → **Citrix** → **Profile Management** to find the new Profile Management section.

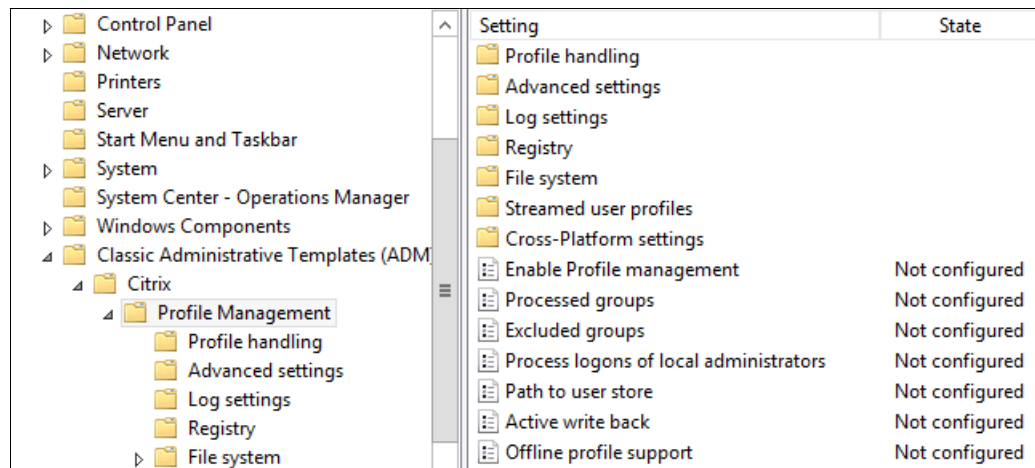


Figure 9-110 Citrix Profile Manager: Policies

15. To enable parameters, select the **Profile Management** folder, then right-click the parameters that you want to enable and select **Edit**, as shown in Figure 9-111.

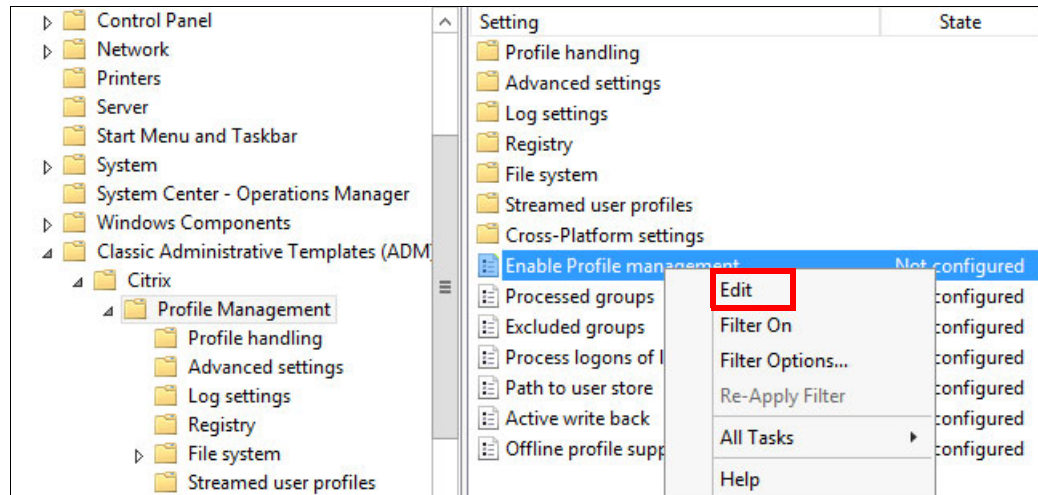


Figure 9-111 Profile Management folder

16. The first parameter to enable is the User Profile Manager (UPM) process (see Figure 9-112). Select **Enabled** and click **Apply**. Click **Next Setting**.

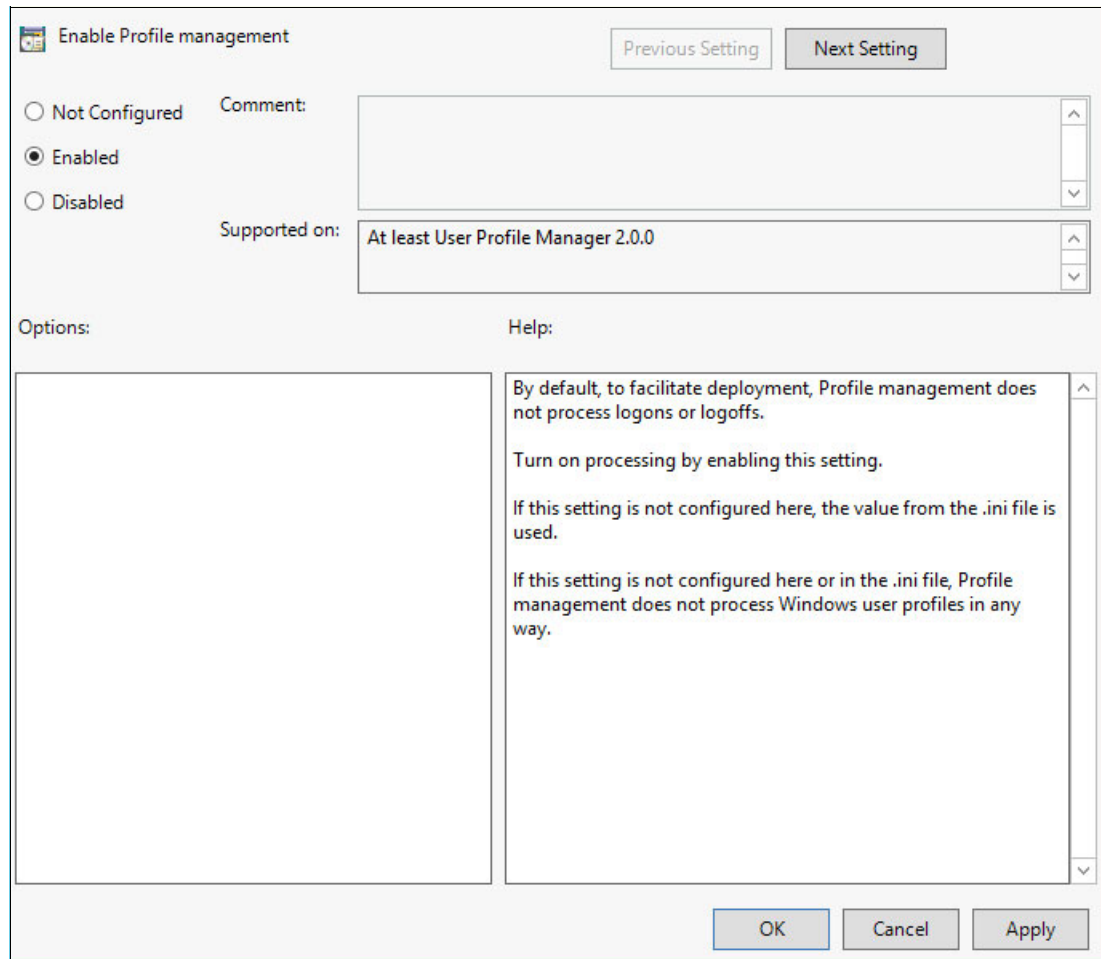


Figure 9-112 Enable Citrix Profile Management

17. Determine the group that the UPM processes. In our example, a group that is named UPMUsers was created to filter the users that must process the UPM, as shown in Figure 9-113.

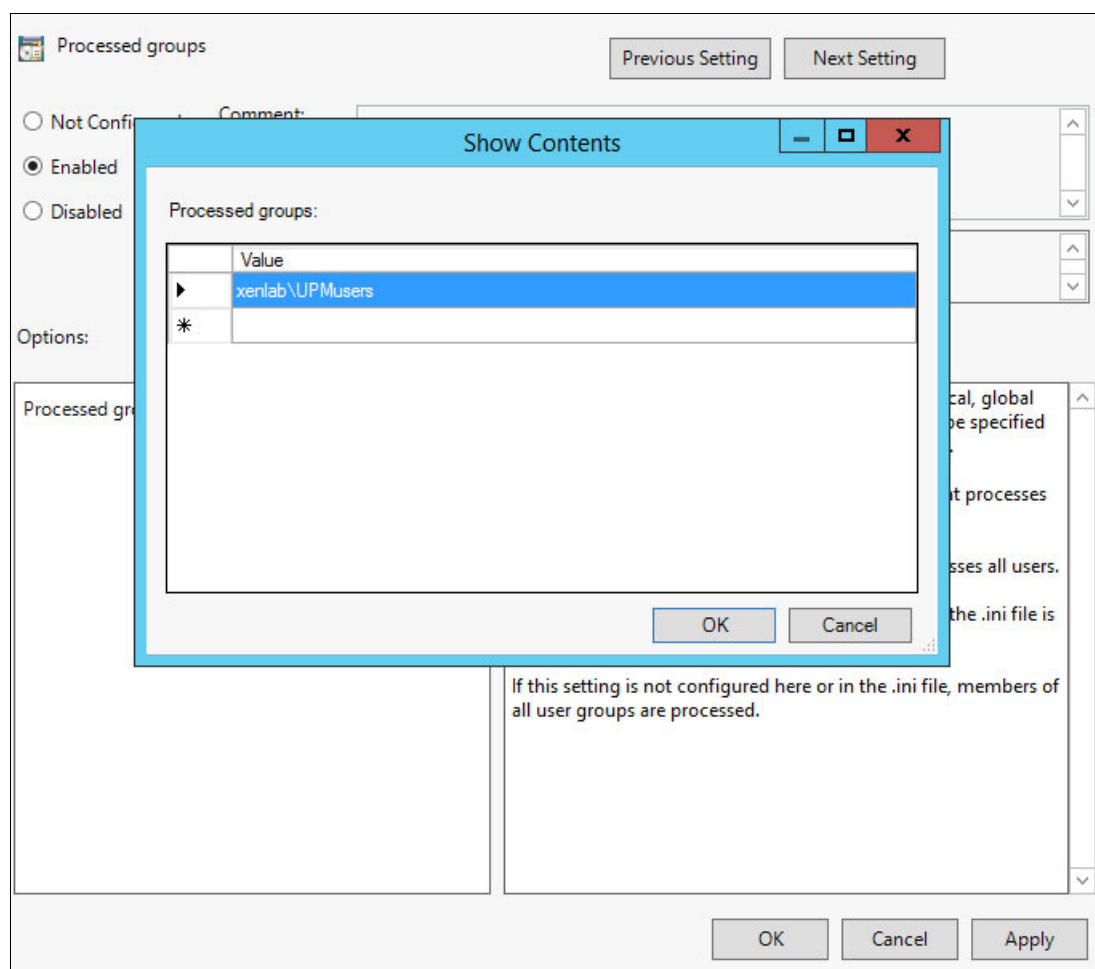


Figure 9-113 Domain group filter

18. The next setting is Process logons of local administrators. Figure 9-114 shows the exception that was created to not process the local administrator's logon, which is helpful when you are troubleshooting. Select **Disabled**, then click **Apply**. Click **Next Setting**.

Process logons of local administrators [Previous Setting] [Next Setting]

☐ Not Configured Comment:
☐ Enabled
☒ Disabled

Supported on:

Options:

Help:

Specifies whether logons of members of the local group "Administrators" are processed by Profile management.

If this setting is disabled, logons by local administrators are not processed by Profile management.

If this setting is not configured here, the value from the .ini file is used.

If this setting is not configured here or in the .ini file, administrators will not be processed.

Figure 9-114 Disable the administrator's profile process

Before the next step, you should create an area to store the user profiles. Follow the guidelines that are available at this website:

<http://support.citrix.com/proddocs/topic/user-profile-manager-sou/upm-create-user-store-c-den.html>

19. The next configuration defines where the profiles are created. In our example, we created a file share on a centralized server to store the profiles. We used the variable #SAMAccountName# to create the profile folder according to the user name, as shown in Figure 9-115.

Path to user store [Previous Setting] [Next Setting]

☐ Not Configured Comment:
☒ Enabled
☐ Disabled

Supported on:

Options:

Help:

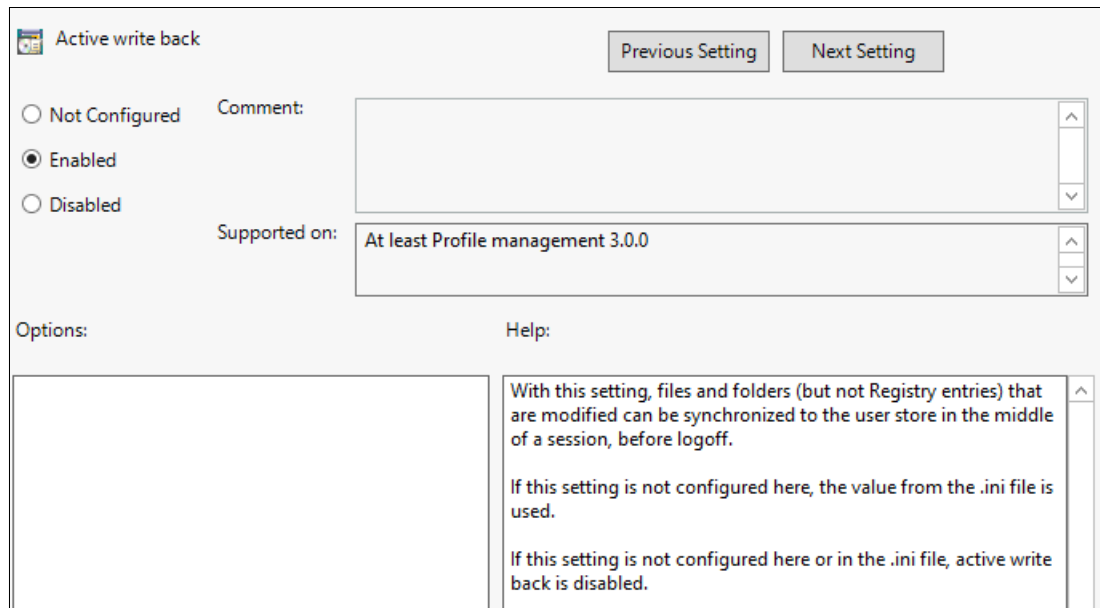
Sets the path to the directory in which the user settings (registry changes and synchronized files) are saved (user store).

The path can be an absolute UNC path or a path relative to the home directory.

Figure 9-115 Define centralized profile store

Select **Enabled**, specify the options, click **Apply**, then click **Next Setting**.

20. The active write back is enabled to reduce the time of synchronization during the logoff (see Figure 9-116). Select **Enabled**, click **Apply**, then click **OK**.



The 'Active write back' configuration window shows the 'Enabled' radio button selected. The 'Supported on:' field is set to 'At least Profile management 3.0.0'. The 'Help' section contains the following text: 'With this setting, files and folders (but not Registry entries) that are modified can be synchronized to the user store in the middle of a session, before logoff. If this setting is not configured here, the value from the .ini file is used. If this setting is not configured here or in the .ini file, active write back is disabled.'

Figure 9-116 Active write back configuration

21. To enable Local profile conflict handling, select the profile handling folder, right-click the setting and click **Edit**, as shown in Figure 9-117.

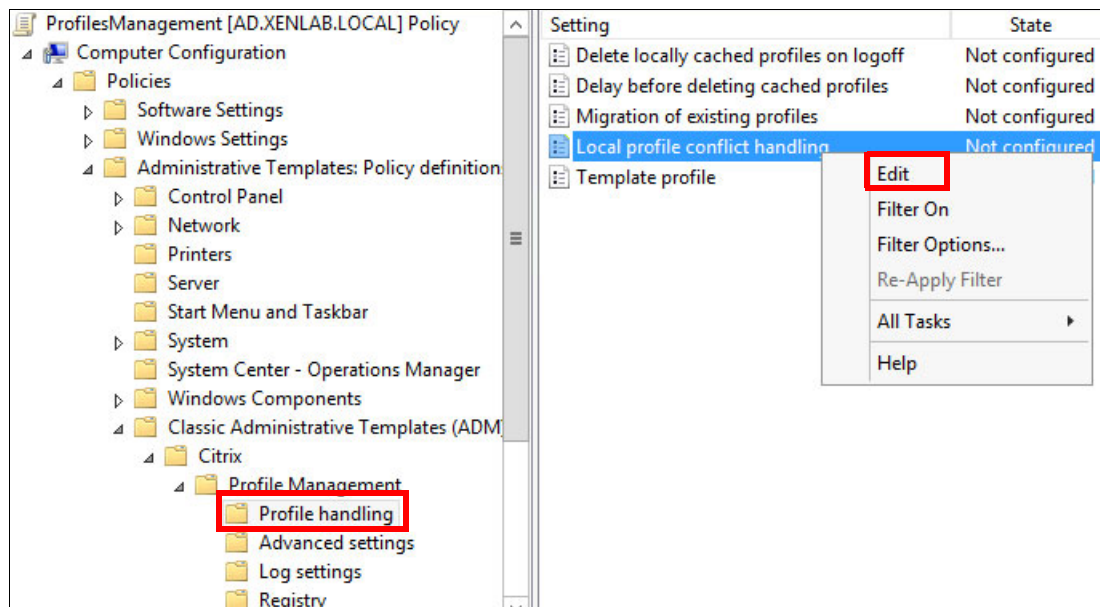


Figure 9-117 Profile handling configuration

22. The next configuration sets how the UPM processes the user if the user has a local profile at the desktop. In our example, we delete the local profile and load the profile from the central file server (see Figure 9-118). Select **Enabled**, select **Delete local profile** option, then click **Apply**. Click **Next Setting**.

Local profile conflict handling Previous Setting Next Setting

☐ Not Configured Comment:
☒ **Enabled**
☐ Disabled

Supported on: At least User Profile Manager 2.0.0

Options: Help:

If both a local Windows user profile and a Citrix user profile in the user store both exist:

Delete local profile

This setting configures what Profile management does if both a profile in the user store and a local Windows user profile (not a Citrix user profile) exist.

If this setting is disabled or set to the default value of "Use local profile", Profile management uses the local profile, but does not change it in any way.

Figure 9-118 Local profile conflict handling configuration

23. Configure the location of the template profile. UPM uses this template folder to create profiles. The second part of configuration specifies for UPM that the template overwrites the local and roaming profiles, as shown in Figure 9-119.

Template profile Previous Setting Next Setting

☐ Not Configured Comment:
☒ **Enabled**
☐ Disabled

Supported on: At least User Profile Manager 2.0.0

Options: Help:

Path to the template profile:

\\FS\\homefolders\\ProfileTemplate

☒ Template profile overrides local profile
☒ Template profile overrides roaming profile
☐ Template profile used as a Citrix mandatory profile for all logons

By default, new user profiles are created from the default user profile on the computer where a user first logs on. Profile management can alternatively use a centrally stored template when creating new user profiles. Template profiles are identical to normal profiles in that they reside in any file share on the network. Use UNC notation to specifying paths to templates. Users need read access to a template profile.

If this setting is disabled, templates are not used.

Figure 9-119 Template profile configuration

24. Select **Enabled**, specify options for the setting, and click **Apply**. Click **OK**.

25. Enable profile streaming. The streaming profile feature loads the user's profile when it is needed. This setting is in the Streaming user profiles folder that is under the Profile management subdirectory (as shown in Figure 9-121).

This feature can reduce the logon time for users (see Figure 9-120). Enable Profile streaming, so that files are synced only when they are needed.

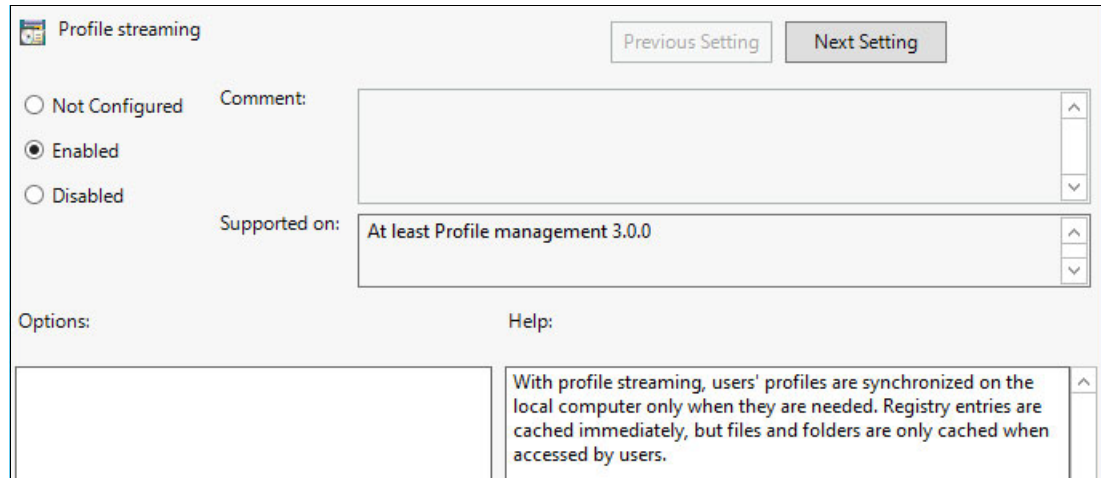


Figure 9-120 Streaming profile configuration

26. To enable Streamed user profile groups, select **Streamed user profile groups folder**, right-click the setting, and click **Edit**, as shown in Figure 9-121.

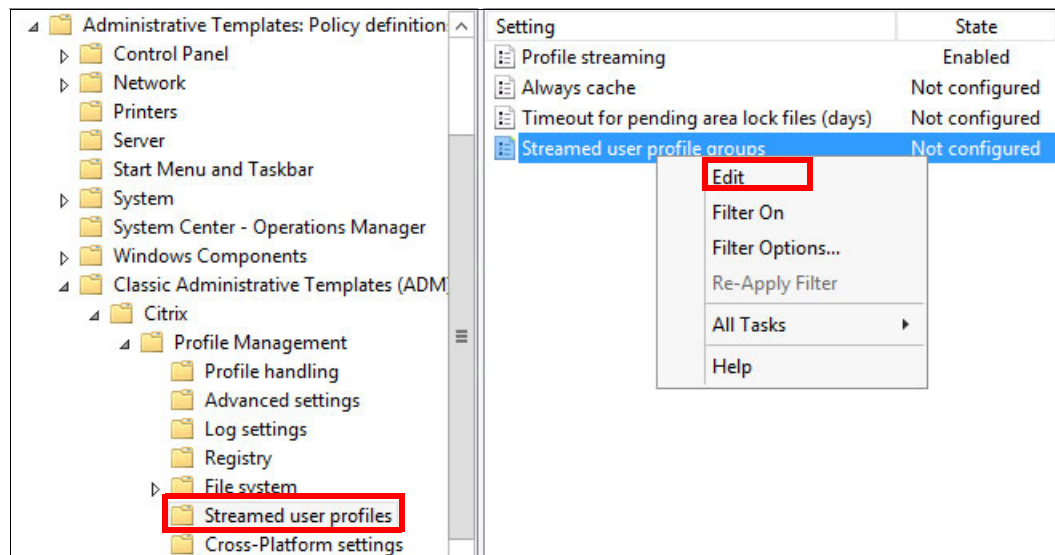


Figure 9-121 Streamed user profiles

27. Enable the UPMUsers profile group so that their profiles are processed by using the streaming profile feature (see Figure 9-122).

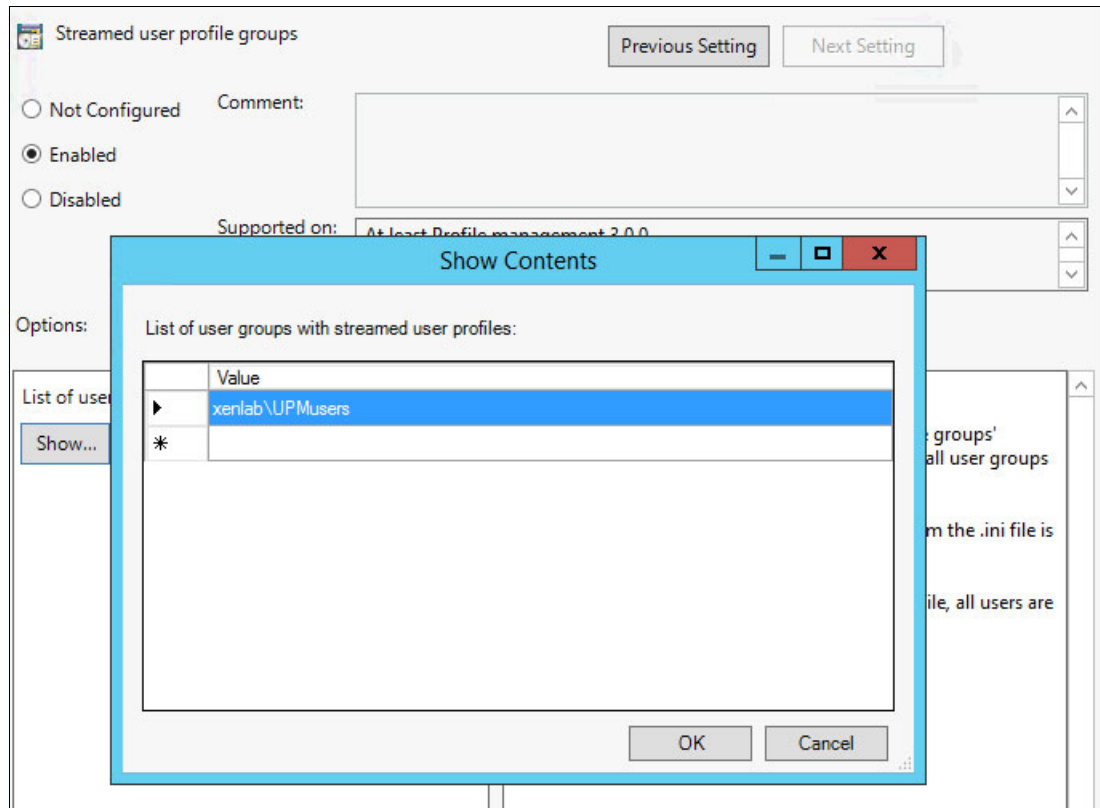


Figure 9-122 Streaming profile filter

9.4.2 Configuring folder redirection

The folder redirection policy moves the default location of the important user's folders from the local disk to a centralized store. Examples of these folders are My Documents, Favorites, and Desktop.

By configuring the folder redirection policy, a link that is created by the user is available if the user logs on to a different desktop.

In our example, we use the same GPO to configure UPM and folder redirection.

To access the folder redirection policies, click **User Configuration** → **Windows Settings** → **Folder Redirection**, as shown in Figure 9-123.



Figure 9-123 Folder Redirection policy location

Figure 9-124 shows the parameters that are used to redirect the Desktop folder. We specified the centralized file server by using the same location and domain group that was used to configure the UPM.

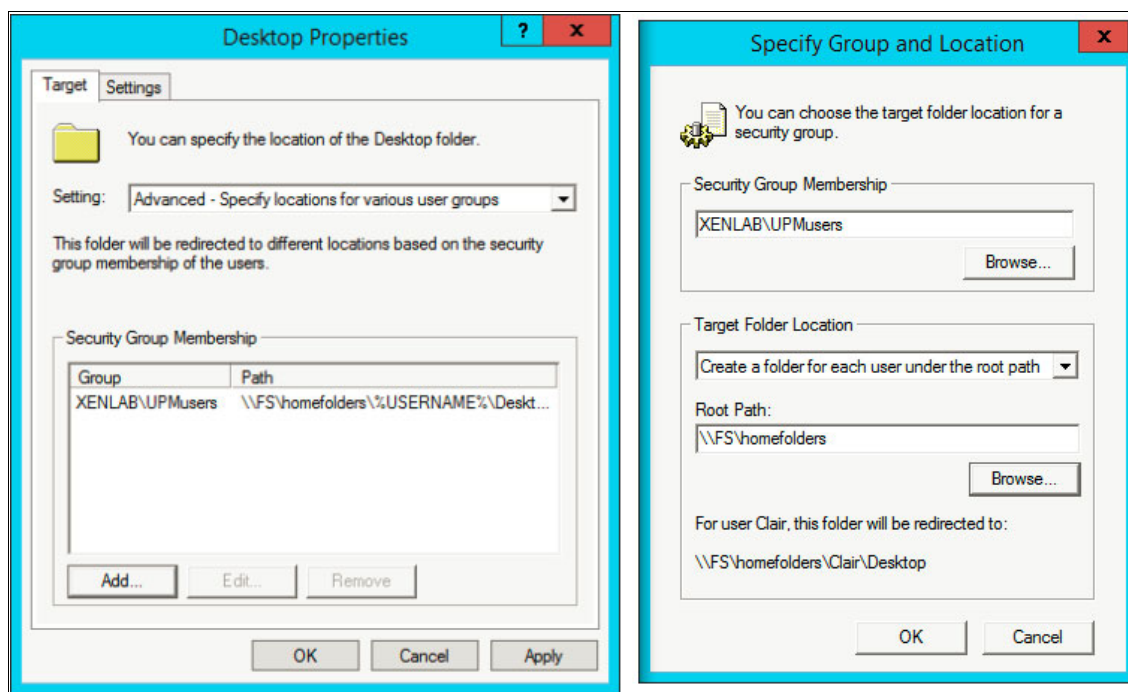


Figure 9-124 Desktop folder redirection configuration

We used a similar configuration to redirect the Documents, Favorites, and Contacts folders, as shown in Figure 9-125, Figure 9-126 on page 238, and Figure 9-127 on page 238.

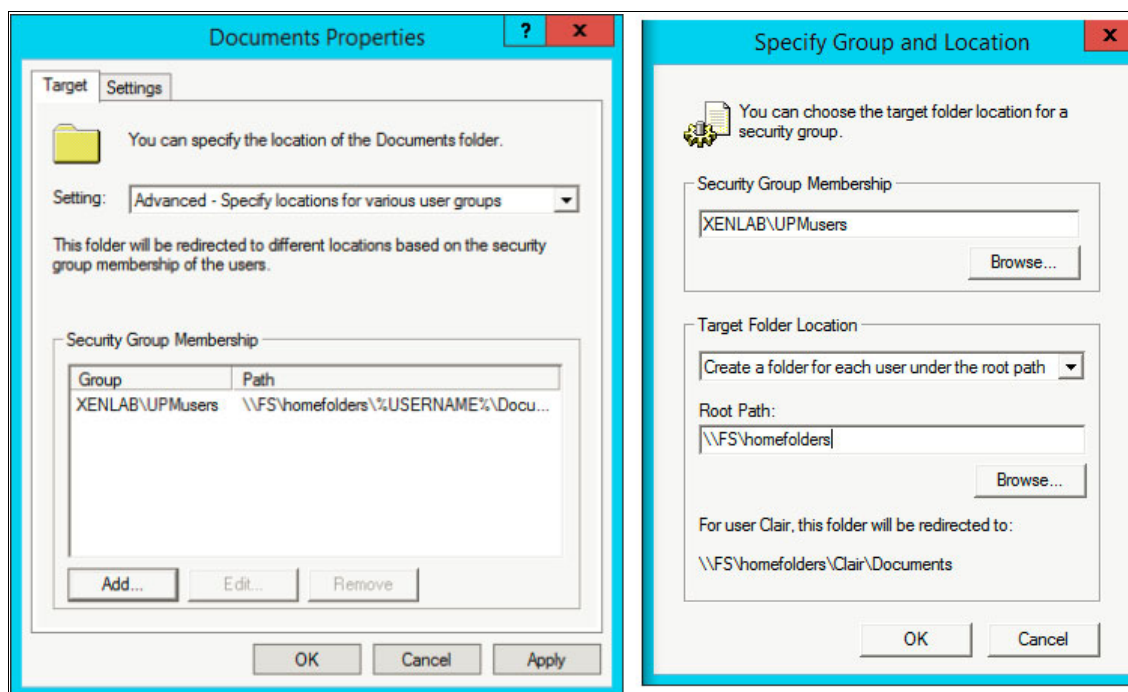


Figure 9-125 Documents folder redirection configuration

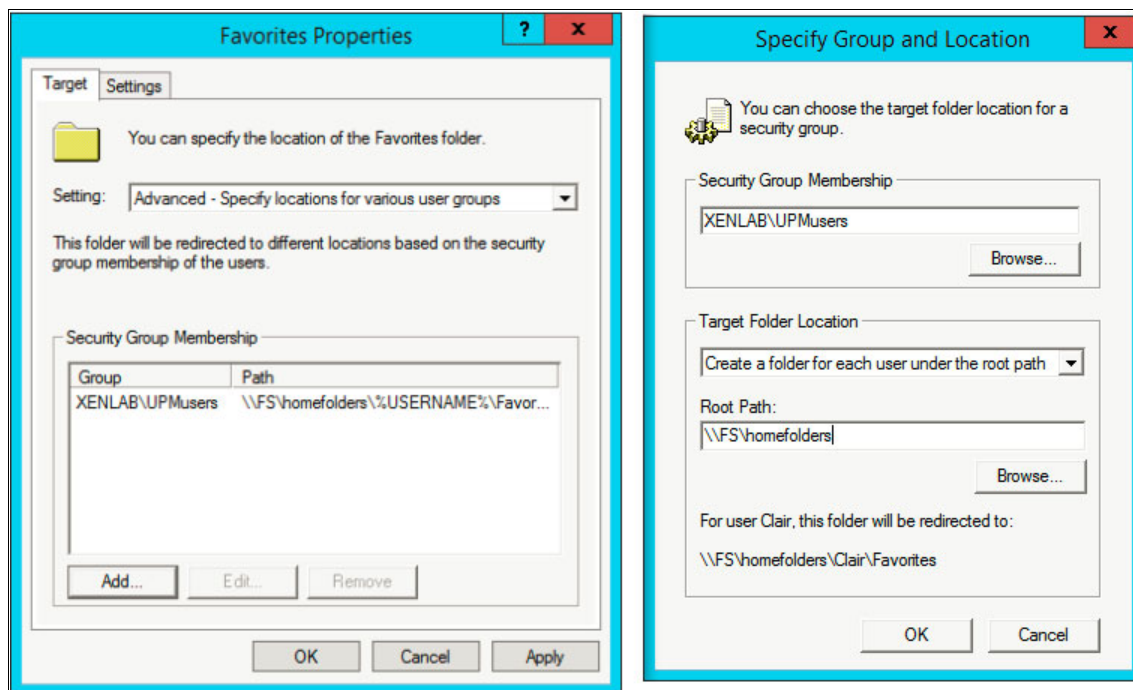


Figure 9-126 Favorites folder redirection configuration

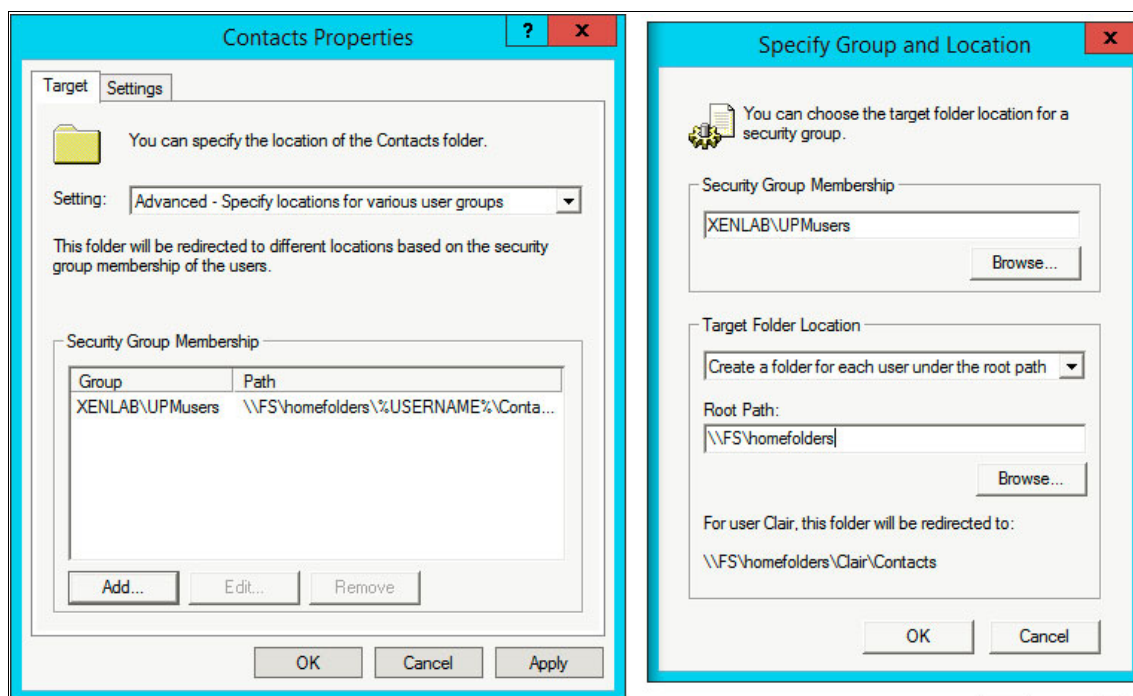


Figure 9-127 Contacts folder redirection configuration

9.4.3 Configuring the Citrix Receiver

To automatically configure Citrix Receiver or Online Plug-in with the Web Server Address on client computers, use Group Policy Preferences to create the necessary registry key.

After the key is created, you must link this GPO to all OUs in which you have your virtual desktop accounts.

Complete the following steps:

1. Create or modify a GPO and click **Computer Configuration** → **Preferences** → **Windows Settings** → **Registry** → **New** → **Registry Item**, as shown in Figure 9-128.

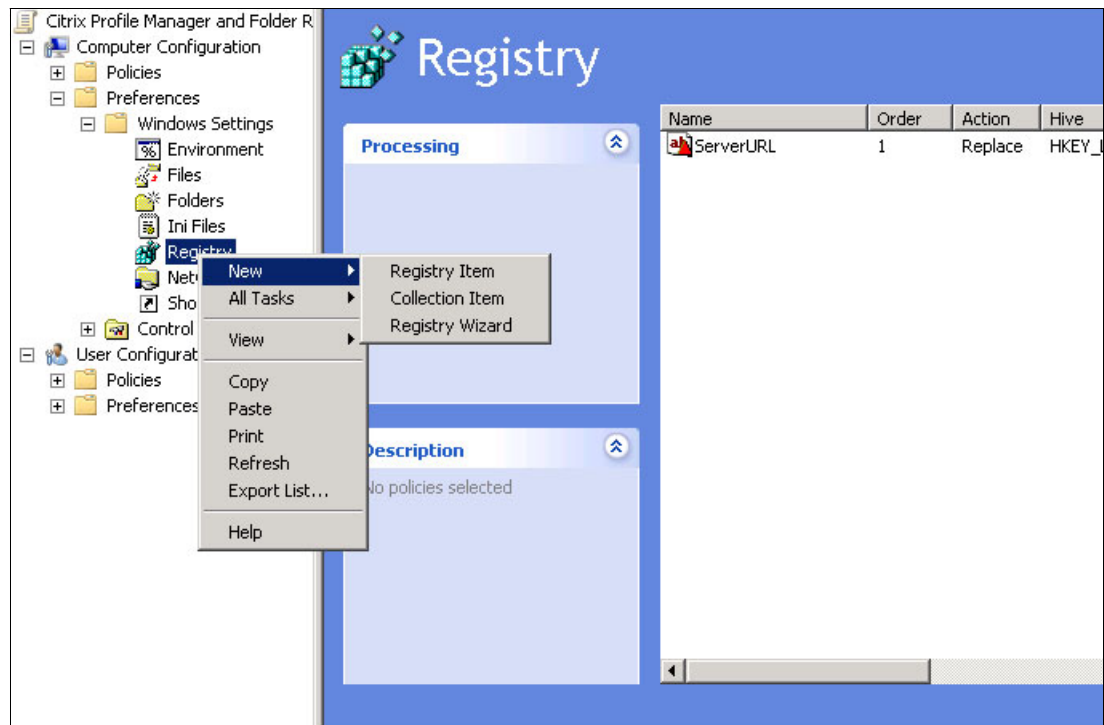


Figure 9-128 Group policy editor

2. Specify the following parameters for the new registry key, as shown in Figure 9-129 on page 240:

- Action: Replace
- Hive: HKEY_LOCAL_MACHINE
- Key Path:
 - 32-bit client computers: SOFTWARE\Citrix\PNAgent
 - 64-bit client computers: SOFTWARE\Wow6432Node\Citrix\PNAgent
- Value name: ServerURL
- Value type: REG_SZ
- Value data: `http://your web interface server hostname/Citrix/PNAgent/config.xml`

Click **OK**.

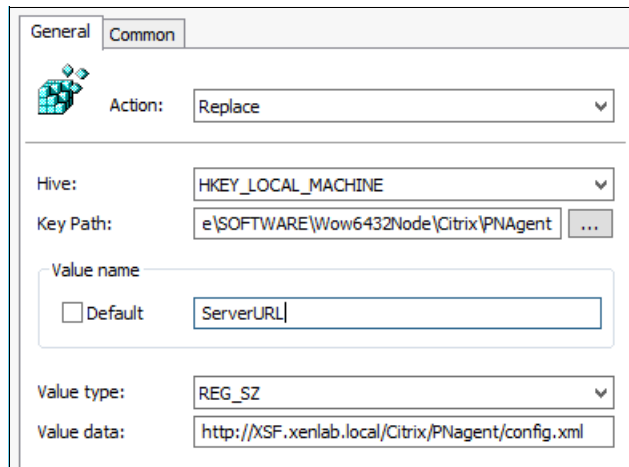


Figure 9-129 New registry key properties

- After the GPO is refreshed, check the Citrix Receiver configuration by accessing your virtual desktop by clicking **Start** → **Apps published** → **your applications**, as shown in Figure 9-130.




Figure 9-130 Accessing published applications at the virtual desktop

9.4.4 Group Policy Object link

The GPO was linked on a specific OU to store all computer accounts that were created for VDI. Figure 9-131 shows the structure that was created to store these desktop computer accounts.




Figure 9-131 Organization units for XenDesktop accounts

The GPO was linked on the following OUs:

- **XenApp Computers**

This OU stores the XenApp servers. The GPO was linked to this OU to ensure that when users open an application that is published on XenApp, the same profiles are loaded and the users' folders are available to open files.

- **XenDesktop NonPersistents**

This OU stores non-persistent computer accounts. The GPO was linked to ensure that all virtual desktops receive the same settings for UPM and folder redirection.

- **XenDesktop Persistents**

This OU stores persistent computer accounts. The GPO was linked to ensure that all virtual desktops receive the same settings for UPM and folder redirection.

Figure 9-132 shows the GPO that is linked on these three OUs.

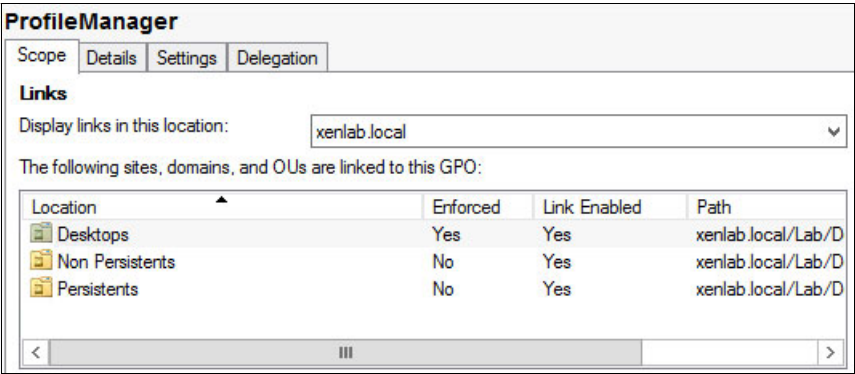


Figure 9-132 Group policy link

9.4.5 Configuring application distribution

In this scenario, the business applications are not installed on the desktop base.

All applications are centralized at Citrix XenApp and delivered to desktops by using the Citrix Receiver that was installed and configured in this chapter.

Managing System x and Flex System hardware in a VDI environment

This chapter describes the integration of Lenovo hardware management into the VDI environments that are based on VMware vSphere and Microsoft Hyper-V. Specifically, we describe the use of UIM for VMware vCenter to manage vSphere-based BladeCenter and Flex System environment, and UIM for Microsoft System Center to manage Microsoft Windows Server based physical and virtual environments.

This chapter includes the following topics:

- ▶ 10.1, “Managing a vSphere environment with UIM” on page 244
- ▶ 10.2, “Managing a Windows Server environment with UIM” on page 261

10.1 Managing a vSphere environment with UIM

By using UIMs for VMware vSphere, administrators can integrate the management features of the System x, BladeCenter, and Flex System with VMware vCenter. It also expands the virtualization management capabilities of VMware vCenter with System x hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration.

UIMs also provide the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

For more information about UIMs for VMware vSphere, see this website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?ln docid=migr-vmware>

In our scenario, it is assumed that you installed UIMs on your VMware vSphere environment.

Consideration: Consider the use of UIM for VMware vSphere for the unified hardware and software management of the VDI environment that is running VMware vSphere.

10.1.1 Enabling UIMs for a newly added ESXi host

Complete the following steps to enable UIM on your newly added ESXi host:

1. Log in to VMware vSphere web client.
2. Enter Hosts and Clusters view.
3. Click the cluster to which your ESXi host belongs.
4. Select the **Manage** tab and click **Upward Integration**.
5. In Overview tab of Upward Integration, select **Cluster Overview**.
6. From the list of ESXi hosts, select the host that you want to enable.
7. From the drop-down list of ESXi hosts, select **Request Host Access**.
8. Enter the credentials when prompted.

These steps are shown in Figure 10-1.

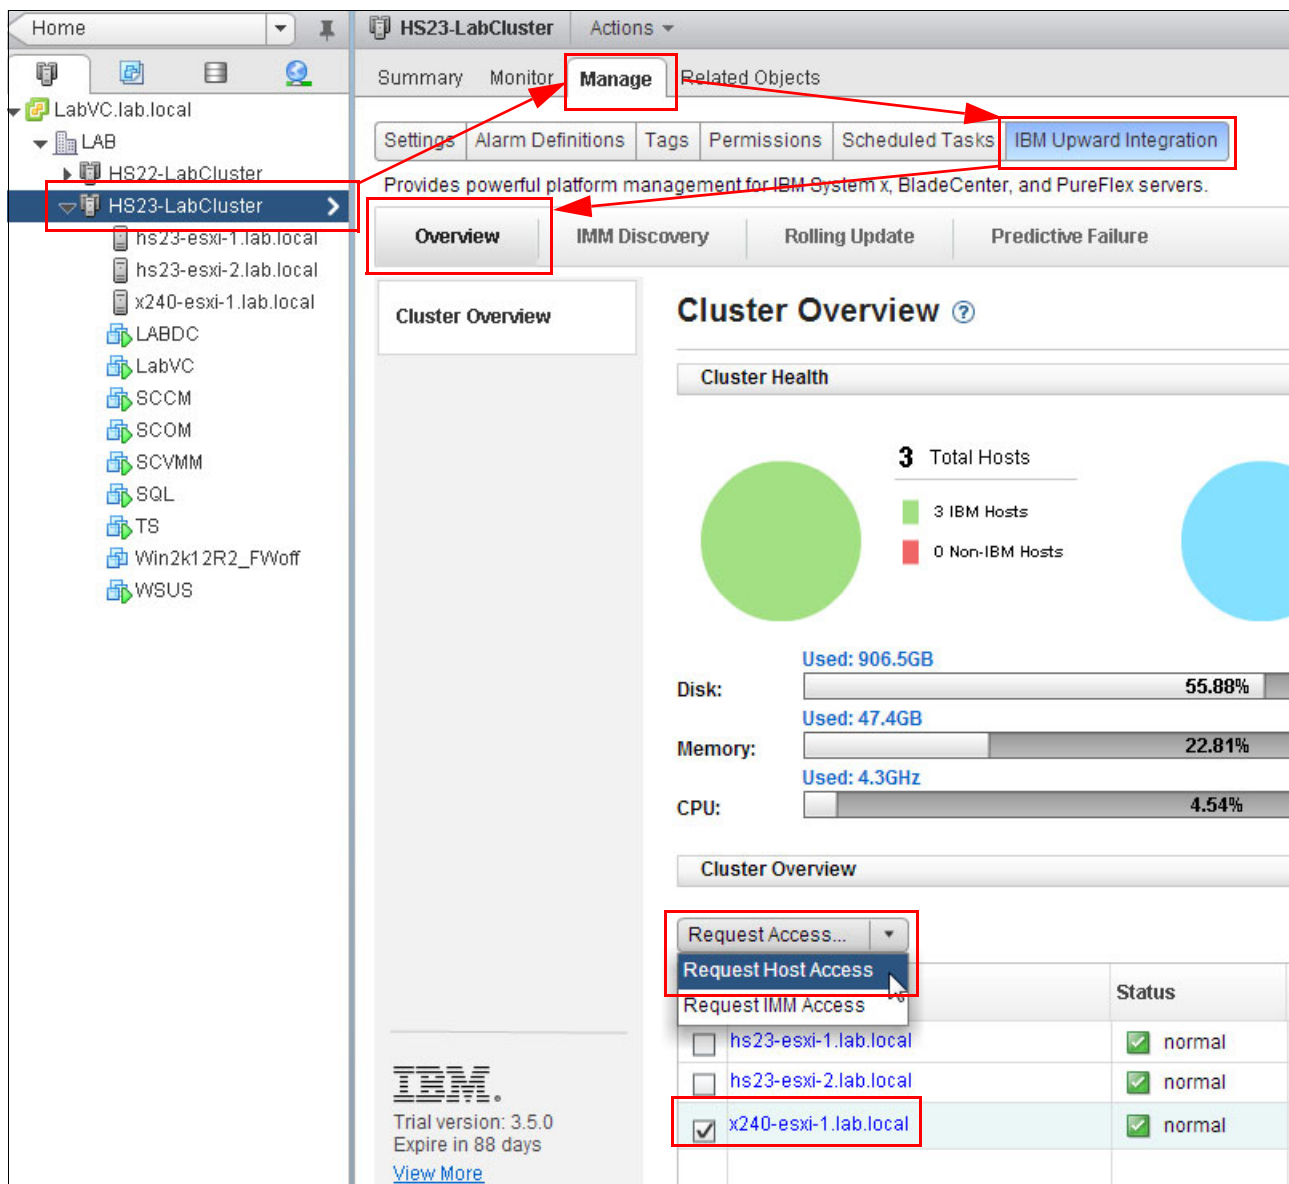


Figure 10-1 Request Host Access

10.1.2 Collecting system inventory with UIM

Complete the following steps to see the available information that is related to your VMware ESXi host:

1. Log in to VMware vSphere web client.
2. Enter the Hosts and Clusters view.
3. Click the ESXi host for which you want to gather the information.
4. Select the **Manage** tab and click **Upward Integration**.
5. In the System tab of Upward Integration, click **Collect** to collect hardware and software details. (The collection process can take several minutes.)

These steps are shown in Figure 10-2.

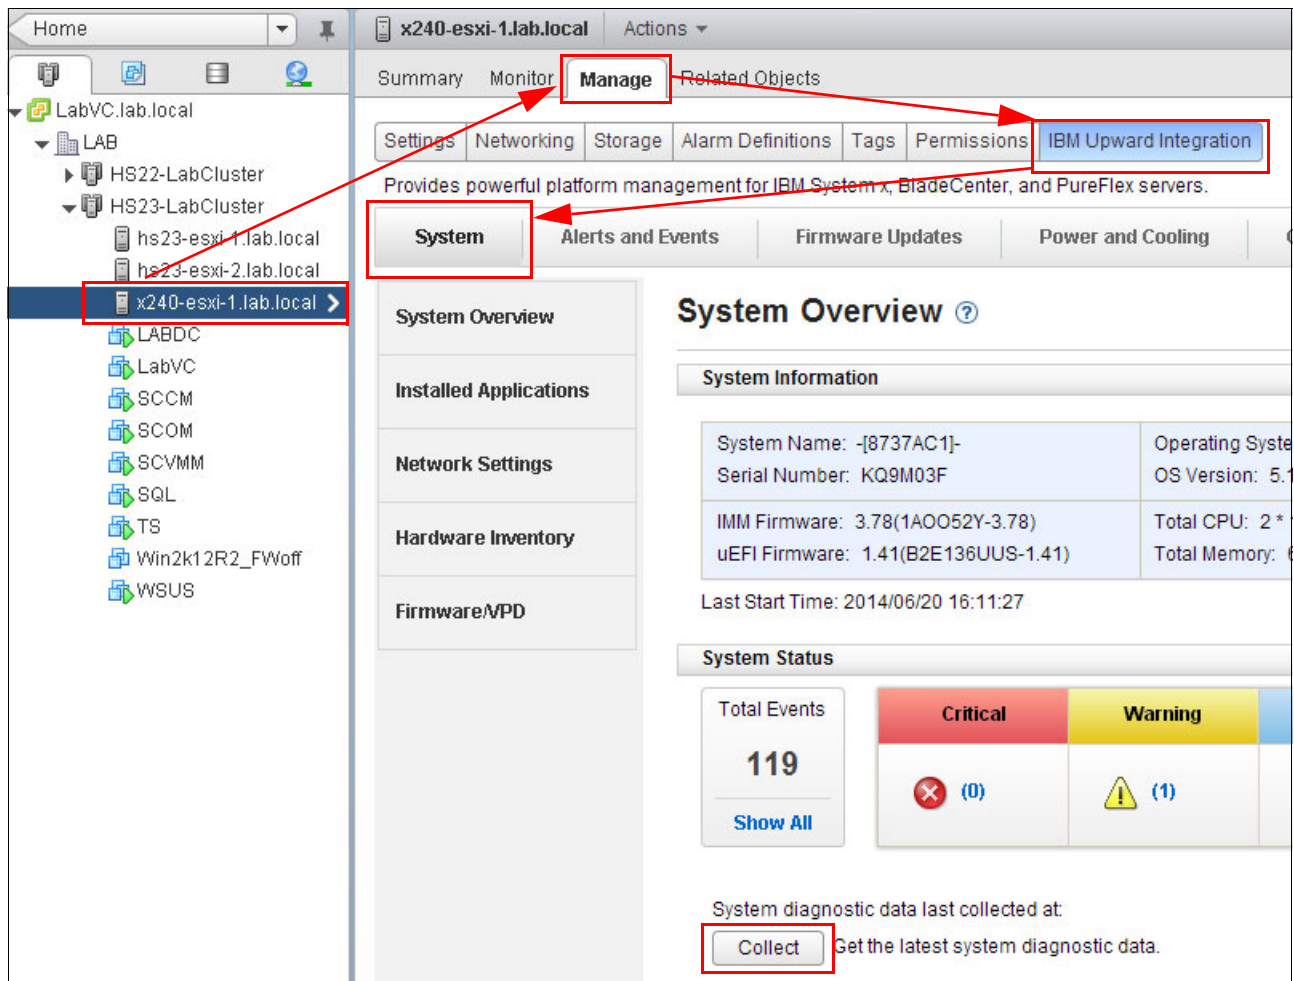


Figure 10-2 Collecting details about hardware and software of ESXi host

After the collection process is finished, you can access hardware and software details by clicking menu items on the left side of the Upward Integration page.

The following views are available:

- The Installed Applications view is shown in Figure 10-3.

System Overview

Installed Applications

Network Settings

Hardware Inventory

Firmware/VPD

Installed Applications ⓘ

Name	Version	Caption	Install Date
brcm	500.2.0.3-000000	brcm	2014061314332
misc-cnrc-register	1.72.1.v50.2-10EM.500.0.0.472	misc-cnrc-register	2014061314332
net-bnx2	2.2.3e.v50.1-10EM.500.0.0.472	net-bnx2	2014061314332
net-bnx2x	1.74.22.v50.1-10EM.500.0.0.472	net-bnx2x	2014061314332
net-cnrc	1.74.04.v50.3-10EM.500.0.0.472	net-cnrc	2014061314332
net-tg3	3.135b.v50.1-10EM.500.0.0.472	net-tg3	2014061314332
scsi-bnx2fc	1.74.02.v50.2-10EM.500.0.0.472	scsi-bnx2fc	2014061314332
scsi-bnx2i	2.74.07.v50.1-10EM.500.0.0.472	scsi-bnx2i	2014061314332
brcdprovider	3.2.0.0-0	brcdprovider	2014061314332
net-bna	3.2.0.0-10EM.500.0.0.472560	net-bna	2014061314332
scsi-bfa	3.2.0.0-10EM.500.0.0.472560	scsi-bfa	2014061314332
emulex-cim-provider	3.8.21.1-01	emulex-cim-provider	2014062016054
ima-be2iscsi	4.6.142.2-10EM.500.0.0.47262	ima-be2iscsi	2014061314332
net-be2net	4.6.142.10-10EM.510.0.0.8022	net-be2net	2014061314332
scsi-be2iscsi	4.6.142.2-10EM.500.0.0.47262	scsi-be2iscsi	2014061314332
scsi-lpfc820	8.2.4.151.65-10EM.500.0.0.472	scsi-lpfc820	2014061314332
concreteish	500.24CE22C1S	concreteish	2014062016054

IBM
Trial version: 3.5.0
Expire in 88 days
[View More](#)
©2013, All Rights Reserved

Figure 10-3 Installed Applications view

- The Network Settings view is shown in Figure 10-4.

x240-esxi-1.lab.local

Actions ▾

Summary

Monitor

Manage

Related Objects

Settings

Networking

Storage

Alarm Definitions

Tags

Permissions

IBM Upward Integration

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System

Alerts and Events

Firmware Updates

Power and Cooling

Configuration

Help

System Overview

Installed Applications

Network Settings

Hardware Inventory

Firmware/VPD

Network Settings ?

Physical Network Ports

Name	vmnic0	vmnic1
DeviceID	vmnic0	vmnic1
OtherIdentifyingInfo	vmklinux,0x12,0x0,0x0,0x19a2,0x 710	vmklinux,0x12,0x0,0x1,0
LinkTechnology	Ethernet	Ethernet
PermanentAddress	3440B5BE7D00	3440B5BE7D04
NetworkAddresses	3440B5BE7D00	3440B5BE7D04
ActiveMaximumTransmissionUnit	1.5 Kilobytes	1.5 Kilobytes
EnabledState	Enabled	Enabled
FullDuplex	true	true

IPv4 Endpoint

Name	vmk0	vmk1	vmk2
TransitioningToState	Not Applicable	Not Applicable	Not Applicable
SubnetMask	255.255.254.0	255.255.255.0	255.255.255.0
RequestedState	No Change	No Change	No Change
ProtocolIFType	IPv4	IPv4	IPv4
IPv4Address	9.42.171.26	169.254.95.120	10.30.30.26

IBM

Trial version: 3.5.0

Expire in 88 days

[View More](#)

©2013, All Rights Reserved

Figure 10-4 Network Settings view

- ▶ The Hardware Inventory view is shown in Figure 10-5.

x240-esxi-1.lab.local Actions ▾

Summary Monitor **Manage** Related Objects

Settings Networking Storage Alarm Definitions Tags Permissions **IBM Upward Integration**

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System Alerts and Events Firmware Updates Power and Cooling Configuration ? Help

System Overview

Installed Applications

Network Settings

Hardware Inventory

Firmware/VPD

Hardware Inventory ?

Memory				
Manufacturer	Samsung	Samsung	Samsung	Samsu
Capacity	8589934592	8589934592	8589934592	858993
BankLabel	Bank 1	Bank 4	Bank 9	Bank 12
SerialNumber	33F8CC8A	33F8CCD5	33F8CC89	33F8CE
Model	DDR3	DDR3	DDR3	DDR3
Speed	1600	1600	1600	1600
PartNumber	M393B1K70DH0-CK0	M393B1K70DH0-CK0	M393B1K70DH0-CK0	M393B1
Description	DIMM 1	DIMM 4	DIMM 9	DIMM 12

Processor		
Name	Processor 1	Processor 2
Family	179	179
CPUStatus	1	1
NumberOfEnabledCores	8	8
CurrentClockSpeed	2000	2000
OtherFamilyDescription	Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz	Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz

Trial version: 3.5.0
Expire in 88 days

[View More](#)

©2013, All Rights Reserved

Figure 10-5 Hardware Inventory view

- The Firmware/VPD view is shown in Figure 10-6.

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System | Alerts and Events | Firmware Updates | Power and Cooling | Configuration | ? Help

Firmware/VPD ?

Software Identity

Description	ElementName	IdentityInfoType	IdentityInfoValue
IMM2 Firmware	IMM2	SoftwareID;SoftwareStatus	1A00;2,6
IMM2 Backup Firmware	IMM2-Backup	SoftwareID;SoftwareStatus	1A00;6
UEFI Firmware/BIOS	UEFI	SoftwareID;SoftwareStatus	B2E1;2,6
UEFI Backup Firmware/B	UEFI-Backup	SoftwareID;SoftwareStatus	B2E1;6
DSA Diagnostic Software	DSA	SoftwareID;SoftwareStatus	DSYT;2,3,6

Figure 10-6 Firmware/VPD view

10.1.3 Monitoring hardware status

By using UIMs, the vSphere administrator can get a detailed view of the hardware's health. You can view your Hardware event logs directly from your vSphere web client, and there is no need to log in to IMM.

The System Health view is in the Alerts and Events tab of the UIM, as shown in Figure 10-7.

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System | **Alerts and Events** | Firmware Updates | Power and Cooling | Configuration | ? Help

System Health ?

Critical(0)
 Warning(1)
 Information(499)

Filter by: All

Message ID	Severity	Time Stamp	Message Detail
PLAT0188	Information	2014-06-13 09:10:55	The System IBM Flex System x240 with 10Gb...
IMM0001	Information	2014-06-13 09:13:24	Management Controller SN# Network Initial...
IMM0025	Information	2014-06-13 09:13:30	LAN: Ethernet[IMM:ep1] interface is now activ...

Figure 10-7 System Health view

Information and statistics about ESXi host power usage can be found in the Power and Cooling tab of the UIM, as shown in Figure 10-8.

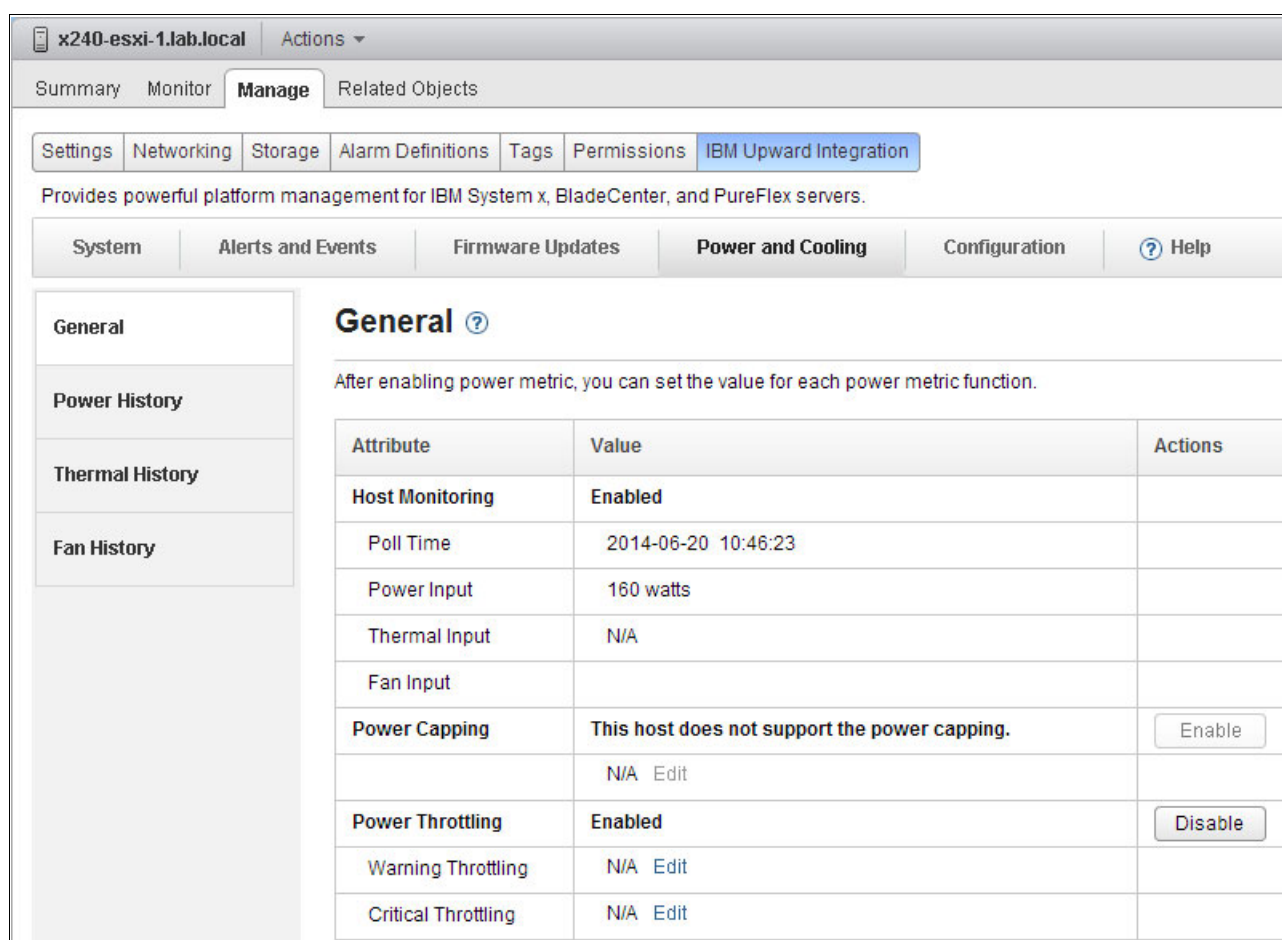


Figure 10-8 General View of Power and Cooling tab

10.1.4 Using PFA alert to move VMs to another ESXi host

In this section, we describe how to use predictive failure management on the vSphere web client to protect your running workload. By using the Policy and Rules page, you can set management policies for a server that is based on a hardware Predictive Failure Alert (PFA).

Based on a defined policy, the Upward Integration for VMware vSphere evacuates virtual machines (VMs) from the server to other hosts in the cluster in response to a PFA. You can view PFAs from the server and the triggered policy history on the Predictive Failures page.

Before you begin

Before predictive failure management is used, verify that the following prerequisites are met:

- ▶ The predictive failure management policy can be set until you discover the IMMs and request the IMMs access.
- ▶ Predictive failure management relies on the hardware PFA capability. The IMM of the server must send out Predictive Failure Alerts when a failure is detected.

- Proper configuration of the network management policy on the vCenter server is required to enable TCP on the https port that you selected when IVP was installed (the default port is 9500). Upward Integration for VMware vSphere listens on this port for incoming indications.
- The host must be put in a properly configured cluster. There must be a host available with vMotion enabled in this cluster. Upward Integration evacuates VMs to other hosts in the cluster, and then puts the host in maintenance mode.

Setting a new policy

You can set an RAS policy on each supported server in the cluster. A policy defines the hardware event categories that you want to monitor and the corresponding action when the event occurs.

To implement this task, click your cluster object in the Hosts and Clusters view, select the **Manage** tab, and click **Upward Integration**. Then, select the **Predictive Failure** tab and you the Policy and rules page opens.

Complete the following steps to set up a policy:

1. Select one or more nodes.
2. Click **Set policy**. The Manage RAS Policy page is displayed, as shown in Figure 10-9.

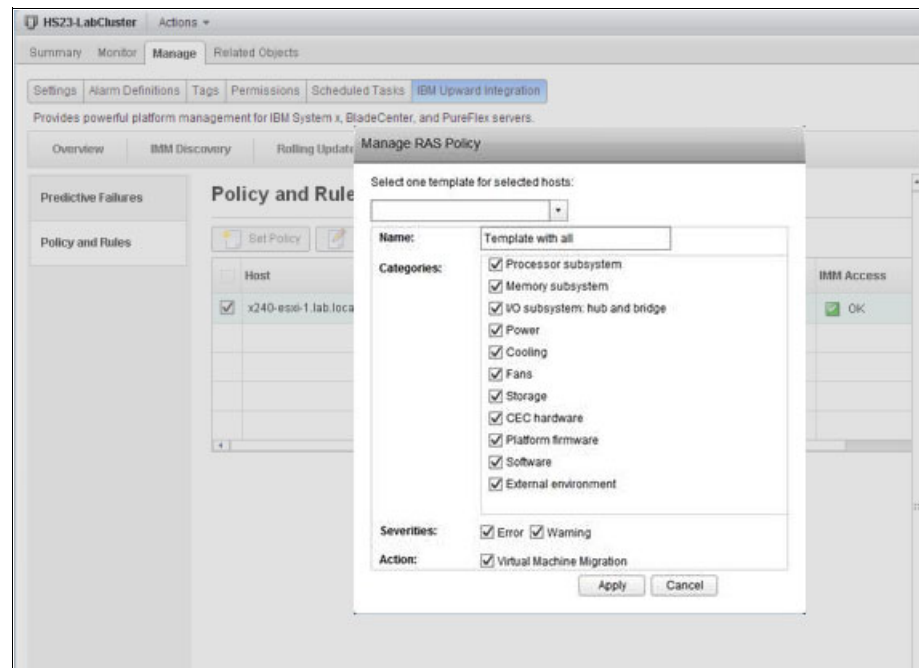


Figure 10-9 Manage RAS policy

3. Select the following event categories, severities, and action:

– Event categories

The Table 10-1 lists the Predictive Failure Alert Event categories that are used on the Manage RAS Policy page.

Table 10-1 Predictive Failure Alert Event categories

PFA Event	Description
Processor subsystem	Processor subsystem includes the CPU and its internal circuits, such as cache, the bus controller, and external interface.
Memory subsystem	Memory subsystem includes the memory controller, memory buffer, memory bus interface, memory card, and DIMM.
I/O subsystem	I/O subsystem includes: IO Hub, IO bridge, IO bus, IO processor, IO adapters for various IO protocols, such as PCI and InfiniBand.
Power	Power includes the power supply and power control hardware.
Cooling	All thermal-related events.
Fans	Includes the fan and blower.
Storage	Includes the storage enclosure, storage controller, raid controller, and media (disk, flash).
Platform firmware	Platform firmware includes IMM and uEFI.
Software	Operating system software and application software.
External environment	All events of an external-related environment including: AC power source, Room ambient temperature, and user error.

– Event severity

Table 10-2 lists the PFA Event severity levels.

Table 10-2 Predictive Failure Alert severity levels.

Severity	Description
Warning	An indication of a failure, which can have no effect on performance. Service action is necessary.
Error	A failure that causes a loss of performance and can cause machines to be inoperable. Immediate service action is necessary.

– Action

The Virtual Machine Migration action evacuates all of the VMs from the server and puts the server in maintenance mode.

After setting the event categories and corresponding action, click **Apply** to apply the policy to the host.

Note: The new policy is saved as a template automatically so that for any other hosts, you can choose a template from the top template drop-down list to apply the same policy.

Editing a policy

You can modify a policy that is defined on a host by using the Edit policy function. Complete the following steps:

1. Select a host.
2. Click **Edit policy**.

Note: When the policy is modified and the policy also is used by other hosts, a warning message is displayed with which you can apply the changes to other hosts or save the changed policy with a different policy name.

Disabling a policy

You can remove a policy from one or more hosts by using the Disable policy function. Complete the following steps:

1. Select one or more hosts.
2. Click **Disable policy**.
3. Click **Disable** to confirm the deletion of the policy from the hosts.

Viewing predictive failure alert events and the Action History table

Upward Integration for VMware vSphere with vSphere Client monitors PFAs from the IMM. All predictive failure events are listed in the Event Log table. When the conditions of a rule are met, the defined action of the rule is started on the managed endpoint. All of the triggered rules and action results are listed in the Action History table, as shown in Figure 10-10.

Getting Started Summary Monitor **Manage** Related Objects

Settings Alarm Definitions Tags Permissions Scheduled Tasks **IBM Upward Integration**

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

Overview IMM Discovery Rolling Update **Predictive Failure**

Predictive Failures

Policy and Rules

Predictive Failures ?

View Predictive Failure event log and action history.

Event Log

Host	Message ID	Severity	Time Stamp
2002:97b:c2bb:830:20a:f7ff:fe26:9a32	PLAT0138	Error	04:46:43 05/0

Action History

Host	Message ID	Status	Start Time
2002:97b:c2bb:830:20a:f7ff:fe26:9a32	PLAT0138	✓ Success Detail...	12:56:27 0

IBM
Version information: 3.5.0

Figure 10-10 Viewing Predictive Failures

10.1.5 Rolling firmware upgrades

You can upgrade your firmware by using Update manager by using one of two methods: you can manually upgrade each ESXi host individually, or you can schedule a rolling update so that update is pushed to the servers at a scheduled time. UIM manages evacuating the ESXi host before the firmware is updated.

Complete the following steps to create a rolling update:

1. In your vSphere web client, browse to the Hosts and Clusters view, click your cluster, select the **Manage** tab, click **Upward Integration**, and then select the **Rolling Update** tab, as shown in Figure 10-11. Click **Create**.

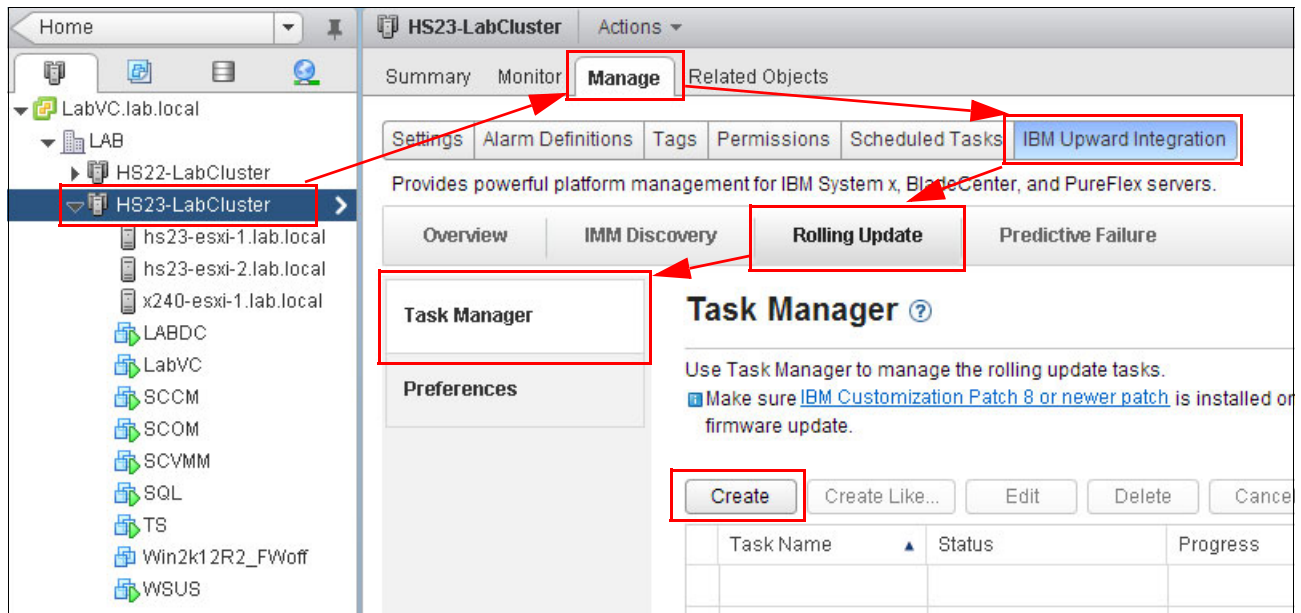


Figure 10-11 Rolling Update in UIM

2. A wizard opens. Complete the following steps:
 - a. Enter a Task Name for the rolling update job. Select the Task Type and click **Next**, as shown in Figure 10-12.

Rolling System Update

1. Name and Type

2. Select hosts and firmware

3. Update options and schedules

Task Name: Rolling_Update1

Task Type: ☒ Update and Reboot ☐ Update Only ☐ Reboot Only

Figure 10-12 Select Name and Type of Rolling update.

- b. Select the ESXi host the updates that you want to apply, as shown in Figure 10-13. Click **Next**.

Rolling System Update

1. Name and Type **2. Select hosts and firmware** 3. Update options and schedules

▼ ☐ -[7875AC1]-

☐ hs23-esxi-2.lab.local

☐ hs23-esxi-1.lab.local

▼ ☐ -[8737AC1]-

☒ x240-esxi-1.lab.local (3 selected items)

Available firmware for x240-esxi-1.lab.local

Firmware Name	New Versions	1 ▲ Install
<input checked="" type="checkbox"/> ▼ UXSP		
<input checked="" type="checkbox"/> IBM Dynamic System Analysis (DE	DSYTE0R-9.60	DSYTE
<input checked="" type="checkbox"/> IBM Flex System x240 UEFI Flash	B2E142A-1.50	B2E14
<input checked="" type="checkbox"/> Integrated Management Module 2	1A0058R-4.20	1A005
<input type="checkbox"/> ▼ Individual		
<input type="checkbox"/> IBM Dynamic System Analysis (DE	DSYTC4P-... ▼	DSYTE
<input type="checkbox"/> IBM Flex System x240 UEFI Flash	B2E136U-1... ▼	B2E14
<input type="checkbox"/> Integrated Management Module 2	1A0056G-... ▼	1A005

Figure 10-13 Select host and firmware

- c. You can update several hosts at the same time if your cluster resources can manage the workload. To do so, select the **Update Parallelization** option and enter the number of hosts that you want to update at the same time. You can force the downgrade of the firmware by selecting the **Force Downgrade** option. If you want to schedule this update instead of running it immediately, select the **Schedule** option and enter the date and time that you want to run the update, as shown in Figure 10-14. Click **Next**.

Rolling System Update

1. Name and Type 2. Select hosts and firmware **3. Update options and schedules**

☐ **Update Parallelization**

Scale: Make sure the value is set according to the current available system resources of the cluster.

☐ **Force downgrade**

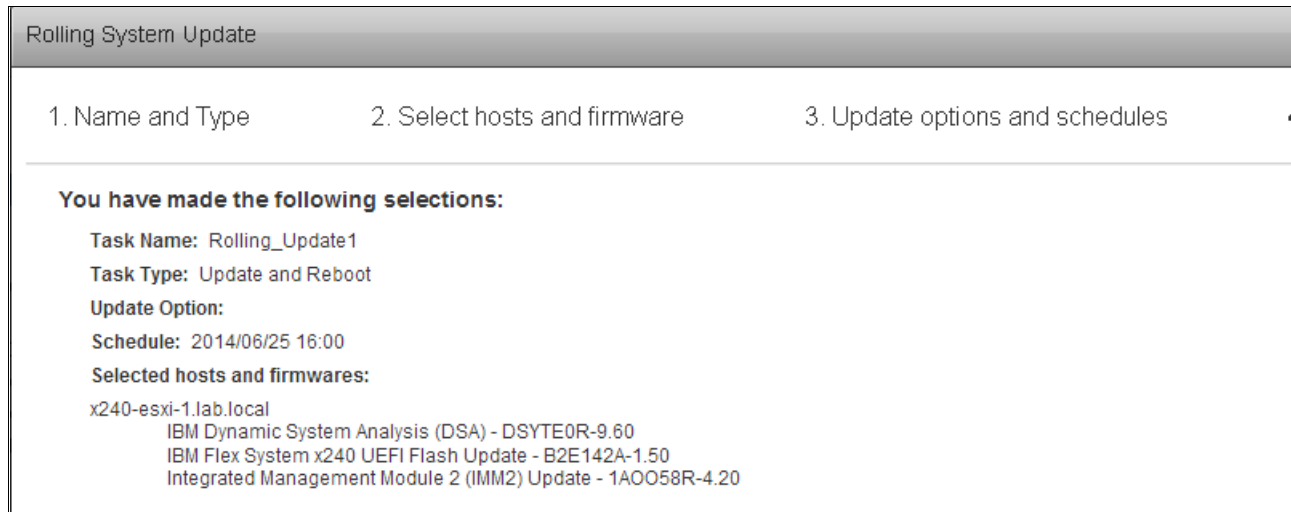
☒ **Schedule**

☐ Now

☒ Schedule

Figure 10-14 Update options and schedules

- d. On the last page, review the summary of the created job and click **Finish**, as shown in Figure 10-15.



Rolling System Update

1. Name and Type 2. Select hosts and firmware 3. Update options and schedules

You have made the following selections:

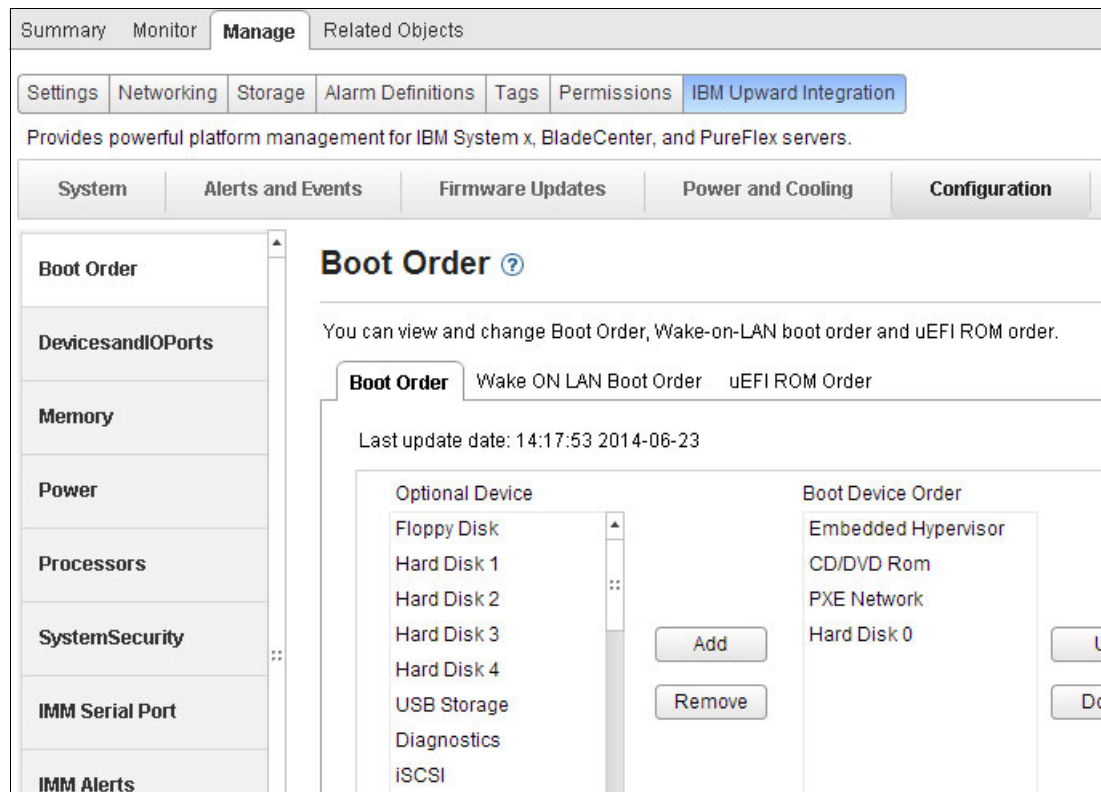
Task Name: Rolling_Update1
 Task Type: Update and Reboot
 Update Option:
 Schedule: 2014/06/25 16:00
 Selected hosts and firmwares:
 x240-esxi-1.lab.local
 IBM Dynamic System Analysis (DSA) - DSYTE0R-9.60
 IBM Flex System x240 UEFI Flash Update - B2E142A-1.50
 Integrated Management Module 2 (IMM2) Update - 1AO058R-4.20

Figure 10-15 Rolling Update job creation

10.1.6 Changing IMM and UEFI configuration

By using UIM, you can change some of the IMM and UEFI parameters. To do so, browse to the Hosts and Clusters view, click your ESXi host, select the **Manage** tab, click **Upward Integration**, and then select **Configuration** tab, as shown in the following examples:

- Edit host boot order is shown in Figure 10-16.



Summary Monitor **Manage** Related Objects

Settings Networking Storage Alarm Definitions Tags Permissions **IBM Upward Integration**

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System Alerts and Events Firmware Updates Power and Cooling **Configuration**

Boot Order ?

You can view and change Boot Order, Wake-on-LAN boot order and uEFI ROM order.

Boot Order Wake ON LAN Boot Order uEFI ROM Order

Last update date: 14:17:53 2014-06-23

Optional Device	Boot Device Order
Floppy Disk	Embedded Hypervisor
Hard Disk 1	CD/DVD Rom
Hard Disk 2	PXE Network
Hard Disk 3	Hard Disk 0
Hard Disk 4	
USB Storage	
Diagnostics	
iSCSI	

Add Remove

Figure 10-16 Edit Boot Order window

- Manage Devices and IO ports is shown in Figure 10-17.

Item	Value
ActiveVideo	Add-in Device
COMPort1	Enable
COMPort2	Enable
Com1ActiveAfterBoot	Disable
Com1BaudRate	115200
Com1DataBits	8
Com1FlowControl	Disable

Figure 10-17 Devices and IO Ports window

- The Manage your Memory modules configuration window is shown in Figure 10-18.

Setting	Value
CKEThrottling	[Dropdown]
CKSelfRefresh	[Dropdown]
DIMM10onProcessor1	Enable
DIMM11onProcessor1	Enable
DIMM12onProcessor1	Enable
DIMM13onProcessor2	Enable
DIMM14onProcessor2	Enable
DIMM15onProcessor2	Enable

Figure 10-18 Memory settings

- Manage Power management settings is shown in Figure 10-19.

Summary	Monitor	Manage	Related Objects
---------	---------	---------------	-----------------

Settings	Networking	Storage	Alarm Definitions	Tags	Permissions	IBM Upward Integration
----------	------------	---------	-------------------	------	-------------	------------------------

Provides powerful platform management for IBM System x, BladeCenter, and PureFlex servers.

System	Alerts and Events	Firmware Updates	Power and Cooling	Configuration
--------	-------------------	------------------	-------------------	----------------------

<div> <div>Boot Order</div> <div>Devices and I/O Ports</div> <div>Memory</div> <div>Power</div> <div>Processors</div> <div>System Security</div> <div>IMM Serial Port</div> <div>IMM Alerts</div> <div>IMM Port Assignments</div> </div>	<h2>Power ?</h2> <p>Please save the changes when you finish the setting to make them effective.</p> <p> <input type="button" value="Save"/> <input type="button" value="Refresh"/> Last update date: 14:22:28 2014-06-23 </p> <table border="1"> <tr> <td>ActiveEnergyManager</td> <td>Capping Enabled ▼</td> </tr> <tr> <td>PlatformControlledType</td> <td>Efficiency - Favor Perfor... ▼</td> </tr> <tr> <td>PowerPerformanceBias</td> <td>Platform Controlled ▼</td> </tr> <tr> <td>S3Enable</td> <td>▼</td> </tr> <tr> <td>WorkloadConfiguration</td> <td>Balanced ▼</td> </tr> </table>	ActiveEnergyManager	Capping Enabled ▼	PlatformControlledType	Efficiency - Favor Perfor... ▼	PowerPerformanceBias	Platform Controlled ▼	S3Enable	▼	WorkloadConfiguration	Balanced ▼
ActiveEnergyManager	Capping Enabled ▼										
PlatformControlledType	Efficiency - Favor Perfor... ▼										
PowerPerformanceBias	Platform Controlled ▼										
S3Enable	▼										
WorkloadConfiguration	Balanced ▼										

Figure 10-19 Power management

- Manage SNMP configuration of the IMM is shown in Figure 10-20.

IMM SNMP	
Please save the changes when you finish the setting to make them effective.	
<input type="button" value="Save"/> <input type="button" value="Refresh"/> Last update date: 14:32:43 2014-06-23	
SNMP Agent Port	161
SNMP Trap Port	162
SNMP Traps	Disabled
SNMPv3 Access Type	Set
SNMPv3 Authentication Protocol	HMAC-SHA
SNMPv3 Privacy Protocol	AES
SNMPv3 Trap Hostname	9.42.171.38

Figure 10-20 Configure SNMP on IMM

10.2 Managing a Windows Server environment with UIM

For managing Microsoft Windows server environment that is hosted on System x and Flex System servers, you can use the System x UIM for Microsoft System Center.

Important: Consider the use of UIM for Microsoft System Center for the unified hardware and software management of the VDI environment that is based on Hyper-V infrastructure.

Lenovo expands Microsoft System Center server management capabilities by integrating System x hardware management functionality, which provides affordable, basic management of physical and virtual environments to reduce the time and effort that is required for routine system administration. It also provides the discovery, configuration, monitoring, event management, and power monitoring that is needed to reduce cost and complexity through server consolidation and simplified management.

For more information about UIM for Microsoft System Center, see this website:
<http://www-947.ibm.com/support/entry/portal/docdisplay?lnocid=SYST-MANAGE>

10.2.1 Enabling Hardware Monitoring on the Flex System

In this section, we describe how to discover a Flex System in Microsoft System Center Operations Manager 2012 (SCOM).

Setting up Flex System Chassis Management Module for discovery

Before you can monitor the hardware status of the Flex chassis components in SCOM, you must configure SNMP in the Chassis Management Module (CMM). Complete the following steps:

1. Log in to the CMM console as Administrator.
2. To change the SNMP settings, click **Mgt Module Management** → **Network** → **SNMP**. Select **Enabled for SNMPv3 Agent**. (You also can enter the Contact and Location information). Click **Apply**, as shown in Figure 10-21.

The screenshot shows the 'Network Protocol Properties' window. At the top, there's an 'Apply' button. Below it are tabs for 'Ethernet', 'SNMP', 'DNS', 'SMTP', 'LDAP Client', and 'TCP Command Mode'. The 'SNMP' tab is active. Underneath are 'Port Assignments' and 'CIM' sub-tabs. The main section is titled 'Simple Network Management Protocol (SNMP)'. It contains two checkboxes: 'Enable SNMPv1 Agent' (unchecked) and 'Enable SNMPv3 Agent' (checked). Below these are 'Contact' and 'Traps' sub-tabs. The 'Contact' sub-tab is active, showing a 'Contact and Location' section with a note: 'Contact and location information are required in order to successfully enable both SNMPv1 and SNMPv3'. It includes two text boxes: 'Contact person:' with the value 'No Contact Configured', and 'Chassis location (site, geographical coordinates, etc):' with the value 'No Location Configured'.

Figure 10-21 Enable SNMPv3 Agent

Note: There are two SNMP agent versions that can be selected for the SCOM to manage the Flex System chassis: SNMPv1 and SNMPv3. In our example, we show SNMPv3, which provides more security than SNMPv1.

To receive events from the management modules, a network connection must exist between the management module and the Microsoft SCOM. You also must configure the management module to send events.

3. To define the SNMP recipient, click **Event** → **Event Recipients**.

4. Click **Create** → **Create SNMP Recipient**.
5. In the Create SNMP Recipient dialog box, enter the IP address of the SCOM server in Descriptive name field.
6. Select **Enable this recipient**.
7. Select **Use the global settings** or **Only receive critical alerts**, as shown in Figure 10-22. Click **OK** to return to the Event Recipients page.

Figure 10-22 Create SNMP Recipient

If you selected **Use the global settings**, the Event Recipient Global Settings dialog box opens, as shown in Figure 10-23. Click **OK**.

Monitored Event Table	Critical Events	Warning Events	Informational Events
Chassis/System Management	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Cooling Devices	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Power Modules	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Compute Nodes	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
I/O Modules	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Event Log		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Power On/Off			<input checked="" type="checkbox"/>
Inventory change			<input checked="" type="checkbox"/>
Network change			<input checked="" type="checkbox"/>
User activity			<input checked="" type="checkbox"/>

Figure 10-23 Event: Recipient Global Settings window

8. To define the SNMPv3 user, click **Mgt Module Management** → **User Accounts**.

9. Click the existing user or **Create** to create a user.
10. In the General tab, enter the user name and password and click the **SNMPv3** tab.
11. Specify the security settings that are based on your company security policy. Set the Access type to Set and enter the IP address of the SCOM server for traps, as shown in Figure 10-24.

The screenshot shows the 'User Properties' dialog box with the 'SNMPv3' tab selected. The 'Context name' field contains 'context2'. The 'Authentication Protocol' dropdown is set to 'Hash-based Message Authentication Code (HMAC) - Secure Hash Algorithm (SHA)'. The 'Use a privacy protocol' checkbox is checked, and the 'Encryption Method' dropdown is set to 'Advanced Encryption Standard (AES)'. The 'Privacy password' and 'Confirm privacy password' fields are masked with dots. The 'Access type' dropdown is set to 'Set'. The 'IP address or host name for traps' field contains '9.42.171.38'. At the bottom are 'OK' and 'Cancel' buttons.

Figure 10-24 SNMPv3 User Properties window

Setting up System Center Operations Manager 2012 for Discovery

There is only one discovery rule for network devices per SCOM management server allowed. Because we are integrating Flex System chassis into the existing environment, we describe modifying the existing rule in this section.

Note: If you are using dynamic discovery, your Flex System chassis might be discovered automatically if the CMM is in the previously defined discovery range with same the SNMP credentials that were assigned to the discovery rule.

Hardware monitoring by using SCOM requires the Lenovo Hardware Management Pack for Microsoft SCOM to be imported in SCOM. Complete the following steps:

1. Log in to the Microsoft SCOM operations console as Administrator.

Note: This feature supports a CMM IP address only. Do not use an IMM IP address.

2. Click **Administration** → **Network Management** → **Discovery Rules** to see the list of discovery rules.
3. Double-click a rule that you want to modify. In our example, we use the rule for BladeCenter AMM discovery, as shown in Figure 10-25 on page 265.

Network Devices Discovery Wizard

General Properties

General Properties

Discovery Method

Default Accounts

Devices

Schedule Discovery

Summary

Completion

Specify general properties

Name:

AMM+CMM

Description (optional):

Select a management or gateway server

Select an Operations Manager management server or gateway server to run the discovery. A server can run only one network discovery. Servers that already run a network discovery do not appear in the list.

Available servers:

SCOM.lab.local

Select a resource pool

Create Resource Pool

Select an Operations Manager resource pool for monitoring of discovered network devices.

Available pools:

All Management Servers Resource Pool

Figure 10-25 Edit Discovery Rule

4. Edit the name (if wanted) and click **Next** twice to open the Devices page.
5. Click **Add**. The Add a Device window opens, as shown in Figure 10-26.

Add a Device

Specify the settings for the network device you want to discover.

Name or IP address:

9.42.170.215

Access mode:

ICMP and SNMP

SNMP version:

v3

Port number:

161

SNMP V3 Run As account:

Select account

Add SNMP V3 Run As Account

Figure 10-26 Add a Device window

Specify the IP address of the CMM. Set the Access mode to ICMP and SNMP or SNMP and then select SNMP version **v3**. Select **Run As account** or **Add new** if you have different credentials for each device.

Complete the following steps:

- a. To define a new Run As account, click **Add SNMP V3 Run As Account**. Then, click **Next** in the Introduction page and enter the name and description of the new account, as shown in Figure 10-27. Click **Next**.

The screenshot shows a web interface with a left sidebar containing three tabs: 'Introduction', 'General Properties', and 'Credentials'. The 'General Properties' tab is selected and highlighted in blue. The main content area is titled 'Specify general properties for the Run As account'. Below the title, there is a text instruction: 'Select the type of Run As account that you want to create, and then provide a display name and description.' The form contains three fields: 'Run As account type:' with a dropdown menu showing 'SnmpV3Account'; 'Display name:' with a text input field containing 'snmp3user'; and 'Description (optional):' with a large text area.

Figure 10-27 Define display name

- b. Specify the credentials that were configured in CMM for SNMPv3 and click **Create**, as shown in Figure 10-28.

The screenshot shows the same web interface as Figure 10-27, but now the 'Credentials' tab in the sidebar is selected and highlighted in blue. The main content area is titled 'Provide account credentials'. Below the title, there is a text instruction: 'Provide credentials for this Run As account for SNMPv3 devices.' The form contains several fields: 'User name:' with a text input field containing 'snmp3user'; 'Context (optional):' with a text input field containing 'context2'; 'Authentication protocol:' with a dropdown menu showing 'SHA'; 'Privacy protocol:' with a dropdown menu showing 'AES'; 'Authentication key:' with a password input field (masked with dots); 'Privacy key:' with a password input field (masked with dots); 'Confirm authentication key:' with a password input field (masked with dots); and 'Confirm privacy key:' with a password input field (masked with dots).

Figure 10-28 Credentials for SNMPv3 CMM account

6. You can add more devices or you can continue by clicking **Next** (see Figure 10-29).

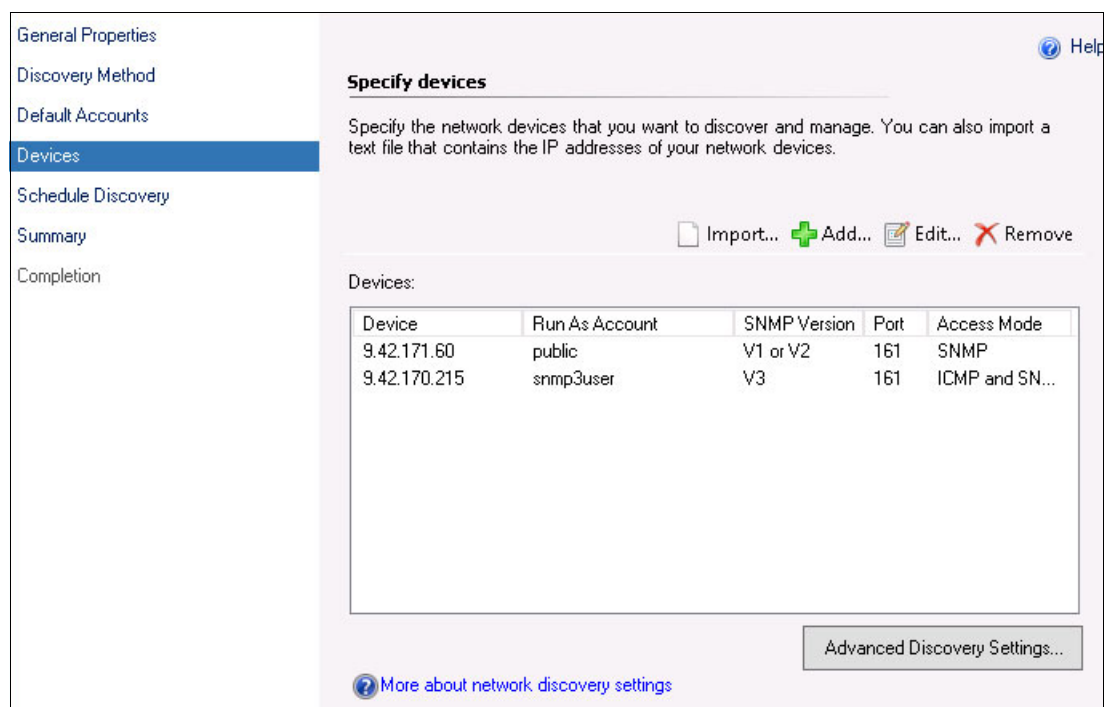


Figure 10-29 Specify devices window

7. Review the Schedule Discovery and Summary sections, or continue by clicking **Next**.
8. On the Completion page, select **Run the network discovery rule after the wizard is closed**. Click **Close**.

Note: It can take several hours for a new device to be discovered with all monitors enabled in SCOM. You can check whether the device discovery was successful in the Operations Manager logs that are in Windows Event Viewer.

After the discovery is completed, you see your discovered Flex System chassis in the Network Devices view of the Administration panel in the SCOM, as shown in Figure 10-30.

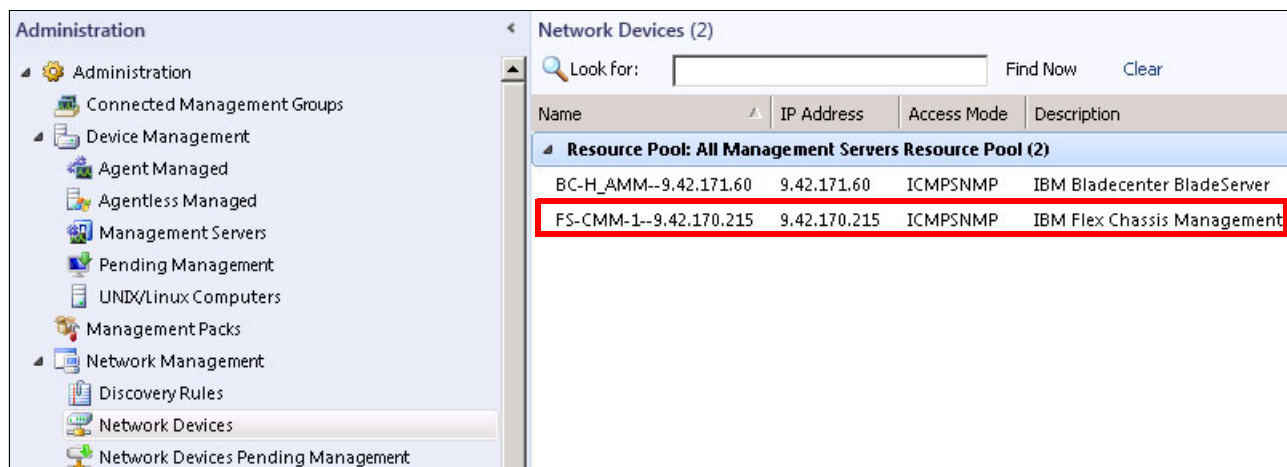


Figure 10-30 Network Devices view after the discovery completed

10.2.2 Deploying System Center agents for hardware monitoring

To enable operating system monitoring, enable the Lenovo Hardware Performance and Resource Optimization Pack for Microsoft System Center Virtual Machine Manager (SCVMM) or Lenovo Inventory Tool for Microsoft System Center Configuration Manager (SCCM). More management agents must be deployed to the Windows Operating system that is installed on a Flex System compute node.

Deploying Microsoft SCOM agent

SCOM agent is required to enable operating system monitoring with enabling Performance and Resource Optimization (PRO) tips in SCVMM.

Complete the following steps to install SCOM agent:

1. Log in to the Microsoft SCOM operations console as Administrator.
2. Click **Administration**. Right-click **Device Management** then click **Discovery Wizard**, as shown in Figure 10-31.

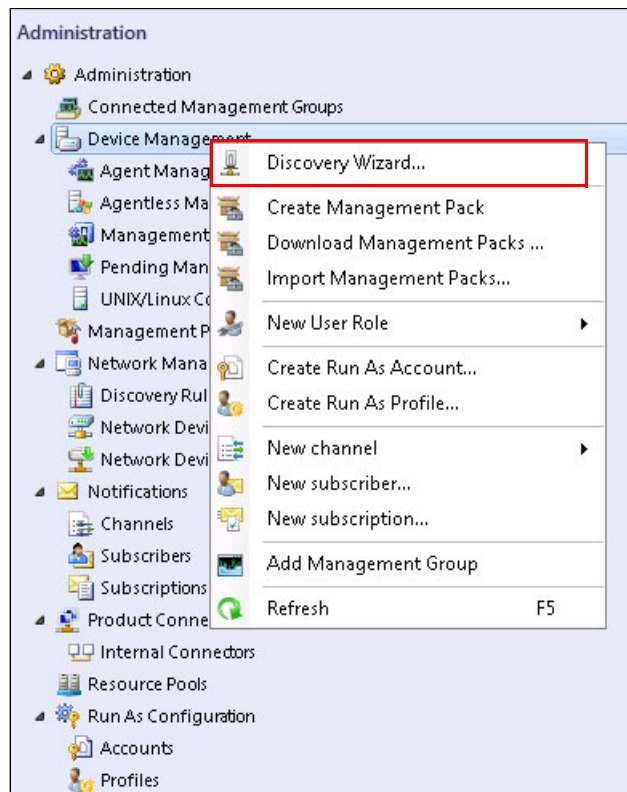


Figure 10-31 Selecting Discovery Wizard option

3. Select **Windows Computers** and click **Next**, as shown in Figure 10-32.

Figure 10-32 Select Windows computers

4. Specify the discovery method. In larger environments, it might be faster to select **Advanced discovery**, as shown in Figure 10-33. Click **Next**.

Figure 10-33 Discovery Method

5. Select **Scan Active Directory** and click **Configure**. Enter the computer name or prefix, as shown in Figure 10-34. Click **OK**, then click **Next**.

The screenshot shows the 'Find Computers' dialog box with the 'Advanced' tab active. The 'Computer name' field contains 'x240'. The 'Owner' field is empty, and the 'Role' dropdown is set to 'Any'. In the background, the 'Discovery Method' section of the wizard has 'Scan Active Directory' selected, and the 'Configure...' button is visible.

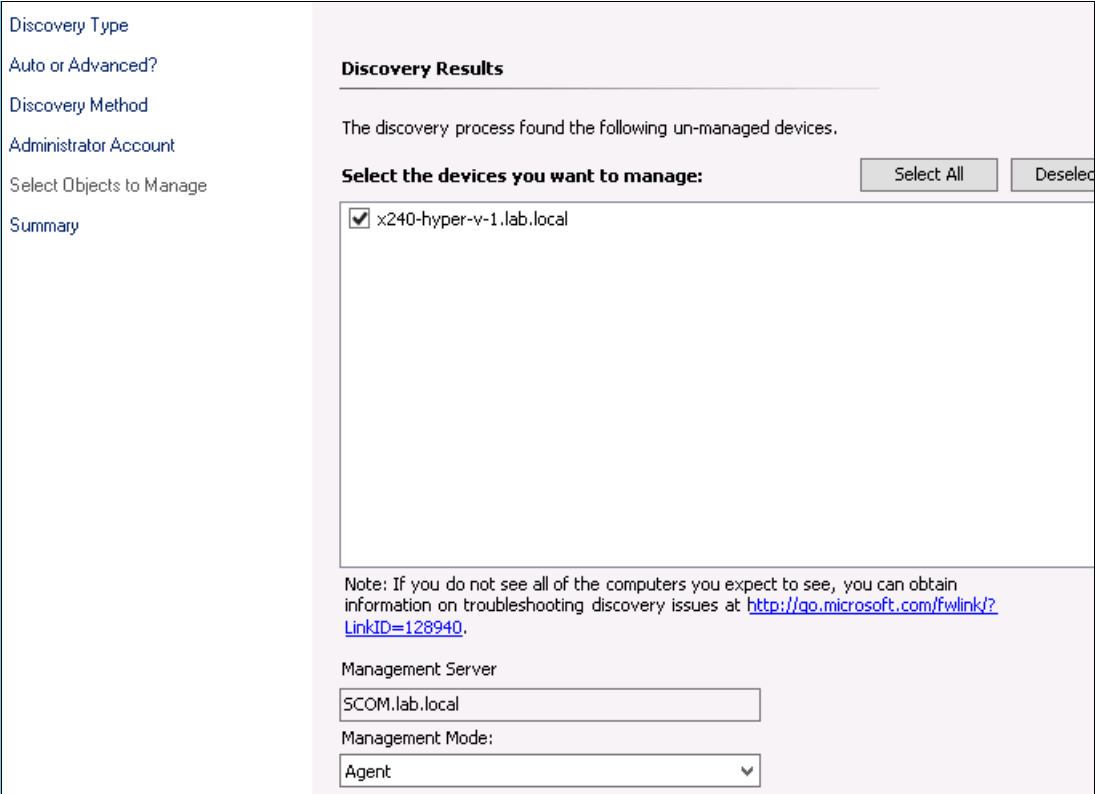
Figure 10-34 Specify computer name or prefix

6. Click **Discover** or specify another user account for discovery and agent installation, as shown in Figure 10-35. The user must have administrator privileges on the target server.

The screenshot shows the 'Administrator Account' section of the wizard. The 'Use selected Management Server Action Account' radio button is selected. The 'Other user account' section has fields for 'User name', 'Password', and 'Domain' (set to 'LAB').

Figure 10-35 Specify Administrator Account

7. Select discovered servers for agent installation and click **Next**, as shown in Figure 10-36.



Discovery Type

Auto or Advanced?

Discovery Method

Administrator Account

Select Objects to Manage

Summary

Discovery Results

The discovery process found the following un-managed devices.

Select the devices you want to manage: Select All Deselect All

☒ x240-hyper-v-1.lab.local

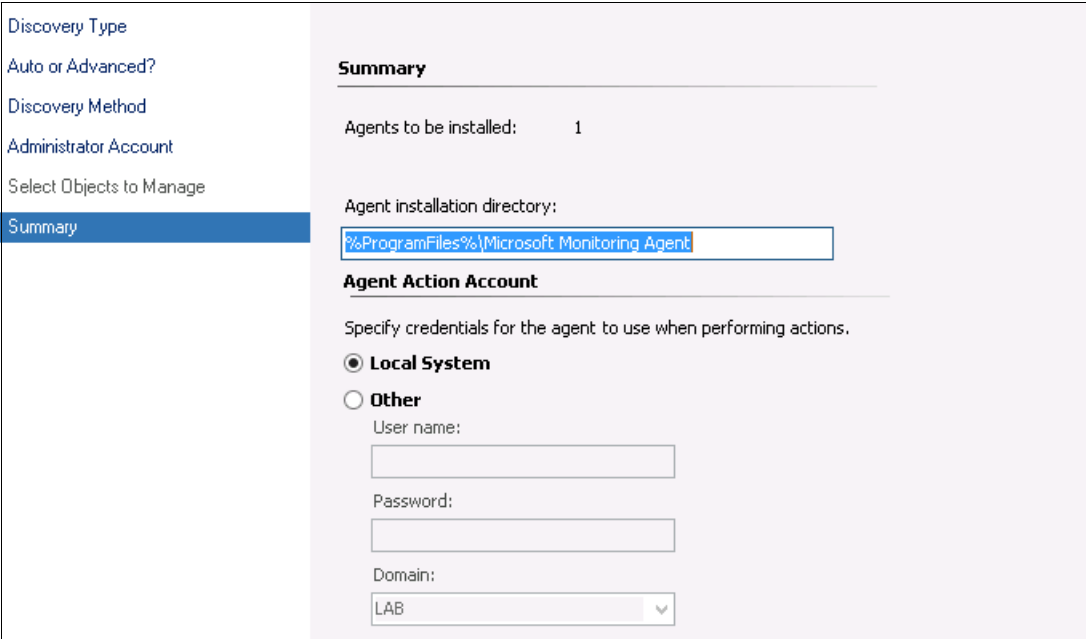
Note: If you do not see all of the computers you expect to see, you can obtain information on troubleshooting discovery issues at <http://go.microsoft.com/fwlink/?LinkID=128940>.

Management Server
SCOM.lab.local

Management Mode:
Agent

Figure 10-36 Select discovered servers for agent deployment

8. Specify the agent installation folder and run as account for the agent that is based on your preferences and internal policies, as shown in Figure 10-37. Click **Finish**.



Discovery Type

Auto or Advanced?

Discovery Method

Administrator Account

Select Objects to Manage

Summary

Summary

Agents to be installed: 1

Agent installation directory:
%ProgramFiles%\Microsoft Monitoring Agent

Agent Action Account

Specify credentials for the agent to use when performing actions.

☒ **Local System**

☐ **Other**

User name:

Password:

Domain:
LAB

Figure 10-37 Installation Path and Run As policies

9. Monitor the deployment status, as shown in Figure 10-38.

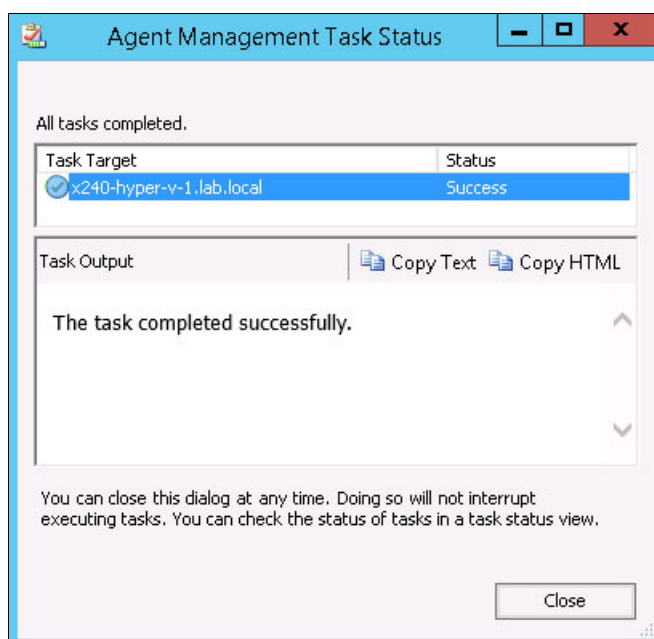


Figure 10-38 Agent deployed

If the target computer is a member of the Microsoft cluster, some management packs require management agent to be enabled in proxy mode. Complete the following steps:

1. Click **Administration** → **Device Management** → **Agent managed**. Right-click the wanted computer and click **Properties**.
2. Click **Security** and select **Allow this agent to act as a proxy and discover managed objects on other computers**, as shown in Figure 10-39.

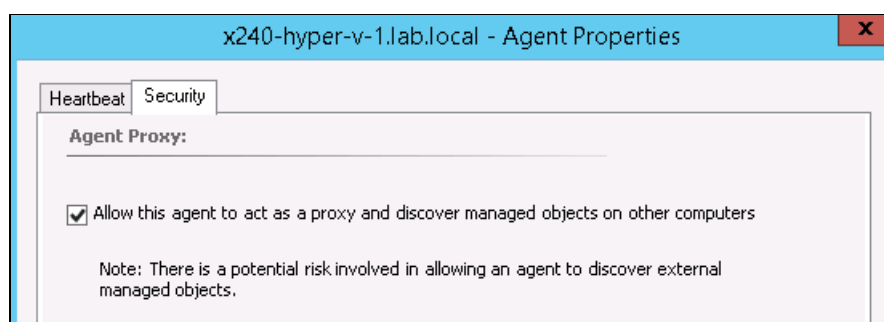


Figure 10-39 Allow agent to act as a proxy

Important: You must install Systems Director Platform agent for System x to enable monitoring of some Lenovo hardware components in SCOM.

For more information about Systems Director agents releases, see this website:

<http://www-03.ibm.com/systems/director/downloads/agents.html>

10.2.3 Monitoring hardware status in SCOM

After you deploy monitoring agents onto Flex System node's operating system, you can monitor the status of the systems hardware components for BladeCenter and Flex System in SCOM monitoring.

Note: We are showing a few views only for demonstration purposes. For more information, see the Lenovo Hardware Management Pack for Microsoft System Center Operations Manager User's Guide, which is available at this website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5082204>

In the Navigation pane, click **Monitoring** and expand **Lenovo Hardware**. Here, you can select from various monitors to see important information about your environment health.

Click **Lenovo Licensed System Group** under the Monitoring to see the list of your Flex System managed servers, as shown in Figure 10-40.

State	Name	Lenovo Platform ...	Lenovo M/T and S/N	Lenovo Product Family
Healthy	HS22-Hyper-V-1...	Blade	7870-06BT218	BladeCenter HS22
Healthy	HS22-Hyper-V-2...	Blade	7870-06RPN99	BladeCenter HS22
Healthy	x240-Hyper-V-1....	Compute Node	8737-KQ9M03F	IBM Flex System x240 ..
Healthy	x240-Hyper-V-2....	Compute Node	8737-KQ9M03G	IBM Flex System x240 ..

Figure 10-40 Lenovo Licensed System Group

Figure 10-40 also shows the following available Lenovo groups to monitor Flex System hardware components:

- ▶ Lenovo Flex System Chassis and Modules
- ▶ Lenovo SCVMM-Managed Licensed Hosts
- ▶ Lenovo System x and x86/x64 Blade Servers

For example, expand **Lenovo Flex System Chassis(s) and Modules** and click **Lenovo Flex System Chassis(s)** to check the status of the Flex System chassis, as shown in Figure 10-41.

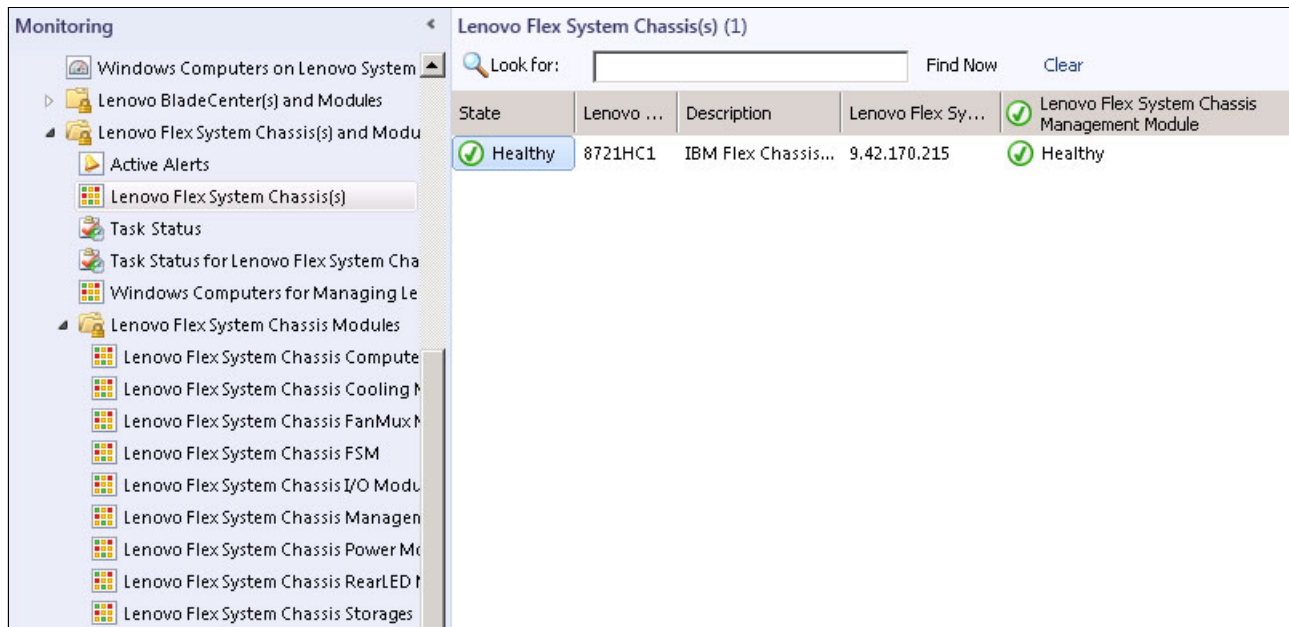


Figure 10-41 Flex System chassis status

You can check the status of other Flex System components by clicking the respective group. For example, the Flex System compute node status is shown under the **Lenovo Flex System Compute Nodes** group (as shown in Figure 10-42) and the I/O module status is shown under the **Lenovo Flex System I/O Modules** group, as shown in Figure 10-43 on page 275.

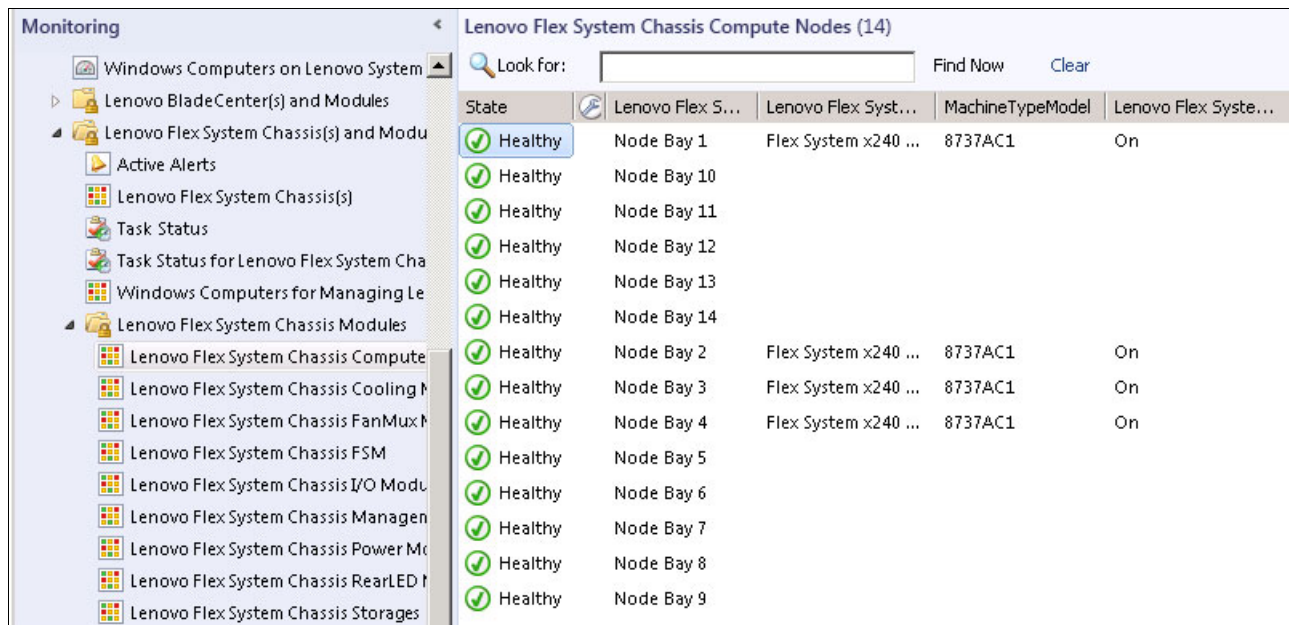


Figure 10-42 Flex System compute nodes status

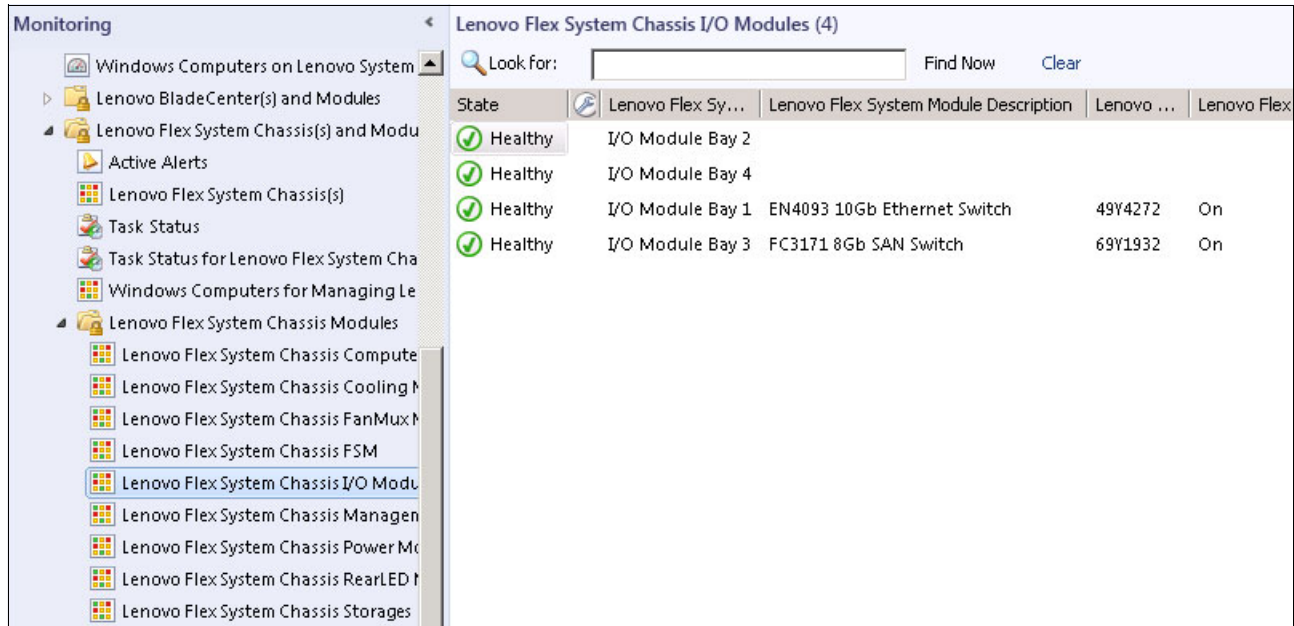


Figure 10-43 Flex System I/O modules status

For a single view of all Lenovo x86 systems, including BladeCenter servers and Flex System compute nodes and the status of their hardware components, expand **Lenovo System x and x86/x64 Blade Servers** and click **All Lenovo System x and x86/x64 Blade Servers**, as shown in Figure 10-44.

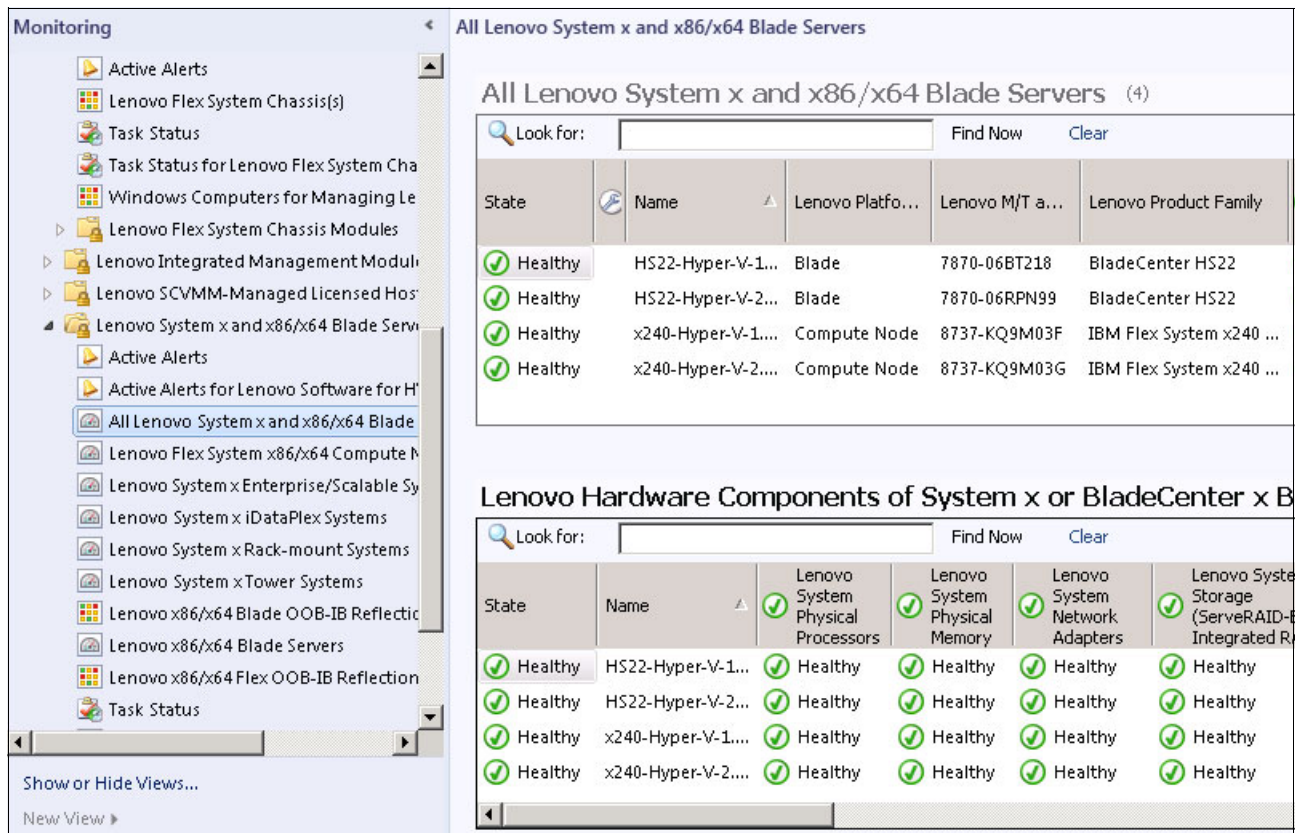


Figure 10-44 All System x and x86/x64 Blade Servers

10.2.4 Lenovo Hardware Performance and Resource Optimization Pack

By using the Lenovo Hardware PRO for Microsoft SCVMM, you can monitor and manage alerts for the physical host resources in a virtualized environment.

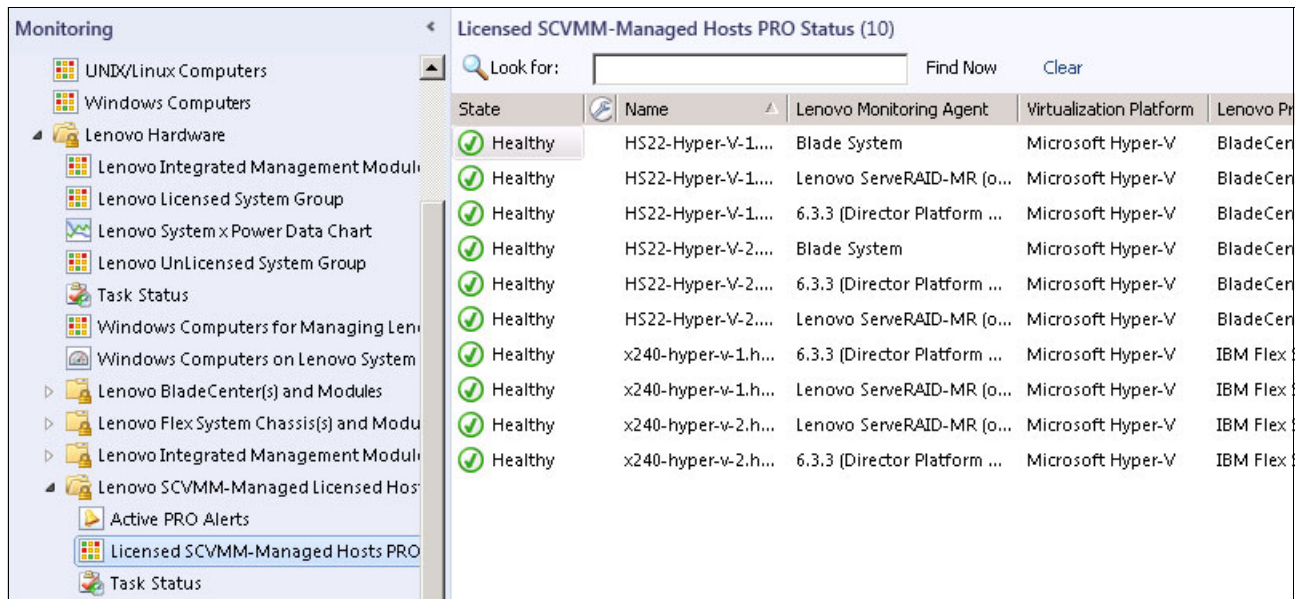
PRO includes the following key features:

- ▶ Automated VM Migration support. This support is based on hardware failure events or power consumption threshold exceptions for UEFI or IMM System x servers and blades that are running Windows 2012, Windows 2008 and 2008 R2, Hyper-V, or Virtual Server.
- ▶ Advisory PRO tips if existing or predictive hardware problems occur that warrant VMM administrative operations.

Note: For more information, see the following Lenovo Hardware Performance and Resource Optimization Pack for Microsoft System Center Virtual Machine Manager website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnidocid=MIGR-5082203>

If PRO Monitors were enabled on the cluster or Host Group in SCVMM, they are enabled automatically after SCOM agent is deployed. You can verify the PRO status in the SCOM console by clicking **Monitoring** → **Lenovo Hardware** → **Lenovo SCVMM-Managed Licensed Hosts (PRO Views)** → **Licensed SCVMM-Managed Hosts PRO Status**, as shown in Figure 10-45.



State	Name	Lenovo Monitoring Agent	Virtualization Platform	Lenovo Pr
Healthy	HS22-Hyper-V-1...	Blade System	Microsoft Hyper-V	BladeCen
Healthy	HS22-Hyper-V-1...	Lenovo ServeRAID-MR (o...	Microsoft Hyper-V	BladeCen
Healthy	HS22-Hyper-V-1...	6.3.3 (Director Platform ...	Microsoft Hyper-V	BladeCen
Healthy	HS22-Hyper-V-2...	Blade System	Microsoft Hyper-V	BladeCen
Healthy	HS22-Hyper-V-2...	6.3.3 (Director Platform ...	Microsoft Hyper-V	BladeCen
Healthy	HS22-Hyper-V-2...	Lenovo ServeRAID-MR (o...	Microsoft Hyper-V	BladeCen
Healthy	x240-hyper-v-1.h...	6.3.3 (Director Platform ...	Microsoft Hyper-V	IBM Flex
Healthy	x240-hyper-v-1.h...	Lenovo ServeRAID-MR (o...	Microsoft Hyper-V	IBM Flex
Healthy	x240-hyper-v-2.h...	Lenovo ServeRAID-MR (o...	Microsoft Hyper-V	IBM Flex
Healthy	x240-hyper-v-2.h...	6.3.3 (Director Platform ...	Microsoft Hyper-V	IBM Flex

Figure 10-45 Licensed SCVMM-managed hosts status

Based on the SCVMM console settings, new PRO tips can appear as pop-up windows, as shown in Figure 10-46.

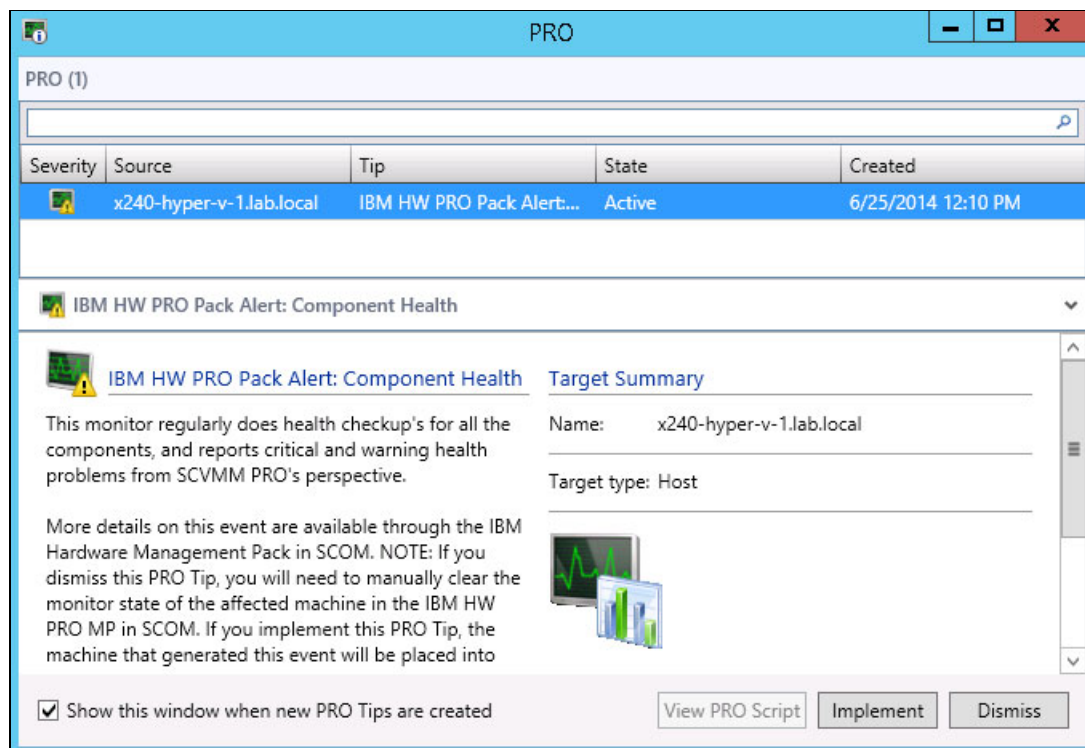


Figure 10-46 SCVMM PRO tip pop-up window

10.2.5 Rolling firmware upgrades by using UIM for System Center VMM

Upward Integration Modules Add-in for Microsoft System Center Virtual Machine Manager (VMM) provides nondisruptive system firmware updates in clustered environment.

Note: For more information, see the following Upward Integration Modules Add-in for Microsoft System Center Virtual Machine Manager website:

<http://www-947.ibm.com/support/entry/portal/docdisplay?lnodocid=MIGR-5095711>

After you join a new system to the cluster, you must set up the authentication information for the new hosts. Complete the following steps:

1. Start SCVMM console as OS administrator and login as SCVMM administrator.
2. In the Fabric view, select the wanted cluster and click the UIM icon.
3. Click the host that is marked red. Then, click **Set Auth Info**.

This process is shown in Figure 10-47.

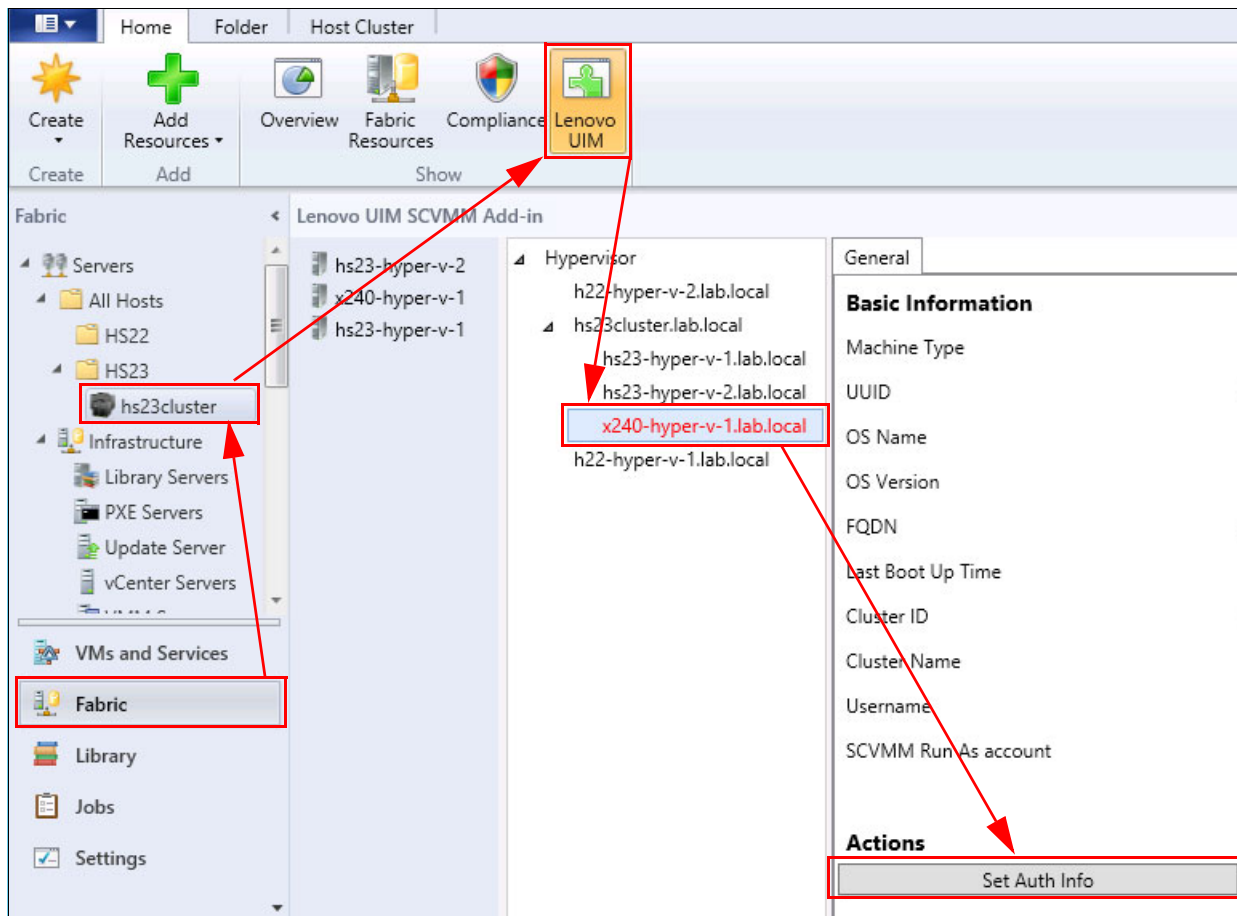


Figure 10-47 New system in UIM for SCVMM

4. Specify Run As Account for SCVMM job and administrator account for new server. You can apply these credentials for All Hosts, Hosts in Cluster, or selected Host only, as shown in Figure 10-48. Click **OK**.

Set Authentication Information

Run As Account

savmm2manage

Username

lab\savmm2manage

Password

••••••••

Confirm

••••••••

☐ Apply to the selected host only.
 ☒ Apply to the hosts of the cluster.
 ☐ Apply to all hosts.

Figure 10-48 Authentication Information window

5. If not already done, you must specify preferences for the local repository folder, including access credentials and firmware download schedule. Click the cluster name, then click **Preferences**, as shown in Figure 10-49 on page 279. Click **Save**.

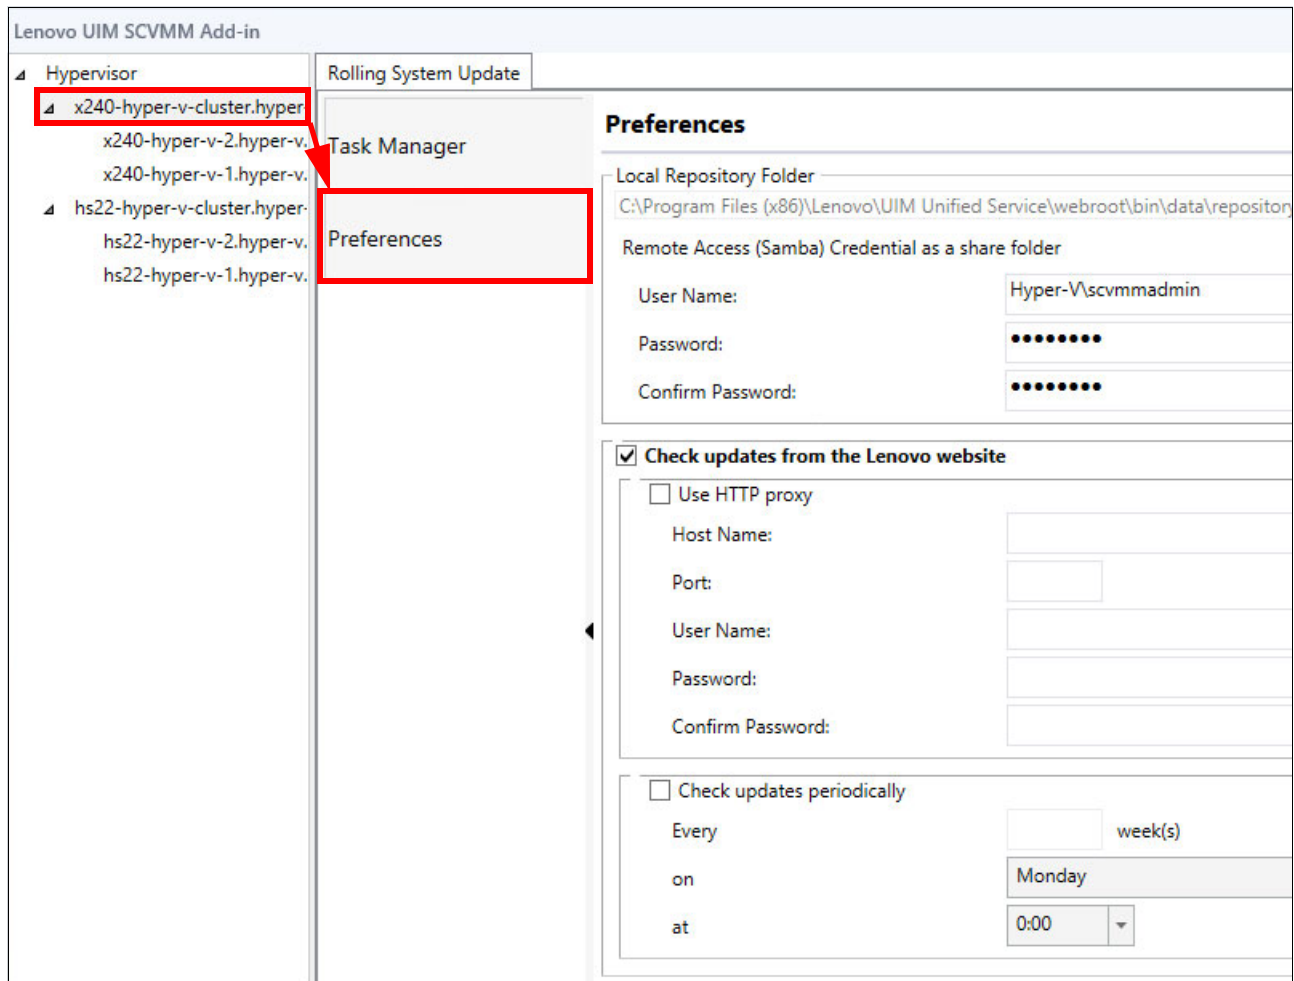


Figure 10-49 Rolling System Update preferences

To update firmware on the Flex System compute nodes in the cluster, complete the following steps:

1. In UIM, click the cluster name. Then, click **Task Manager** and click **Create**, as shown in Figure 10-50.

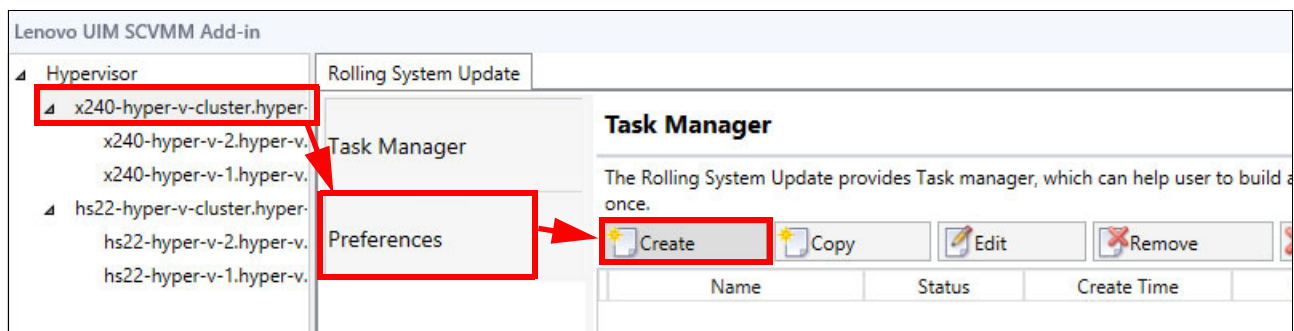


Figure 10-50 Create new task

2. Enter a task name. Select **Task type**, as shown in Figure 10-51. Click **Next**.

1. Name and Type

Task Name:

Task type: ☒ Update and Reboot ☐ Update Only ☐ Reboot Only

Figure 10-51 Task name and type

3. Select the wanted hosts to update and select the firmware and versions to update. You can specify per host or per host model, as shown in Figure 10-52.

2. Select hosts and firmwares

Lenovo System x Server

- ☒ 8737
 - ☒ x240-hyper-v-2.hyp
 - ☒ x240-hyper-v-1.hyp

Available firmware for 8737

UXSP Package

<input type="checkbox"/>	Firmware Name	Installed Version	New Version
<input checked="" type="checkbox"/>	Emulex HBA (LPe1600x) Firmware Update for		ibm14a-10.2.261.3
<input checked="" type="checkbox"/>	Integrated Management Module 2 (IMM2) U		1a0064l-4.50
<input checked="" type="checkbox"/>	Online Broadcom NetXtreme and NetXtreme		2.4.1d4
<input checked="" type="checkbox"/>	IBM Flex System x240 UEFI Flash Update		b2e142a-1.50
<input checked="" type="checkbox"/>	IBM Dynamic System Analysis (DSA)		dsyte2f-9.61
<input checked="" type="checkbox"/>	LSI 2004 SAS Controller BIOS and Firmware U		x240-1.18.01
<input checked="" type="checkbox"/>	IBM Online SAS/SATA Hard Disk Drive Updat		sas-1.14.04-1

Figure 10-52 UIM for SCVMM: Select hosts and firmware

4. Define the wanted update options and schedule, as shown in Figure 10-53. Click **Next**.

3. Update options and schedule

☐ Update Parallelization
Scale: Make sure the value is set according to the current available system resources of the cluster

☐ Force Downgrade

☒ Schedule

☒ Now

☐ Schedule Time
 : :

Figure 10-53 Update options and schedule

5. Review the Summary page, as shown in Figure 10-54. Click **Save**.

4. Summary

You have made following selections:

Task name: Cluster Firmware Update

Task type: Update and reboot both

Update option:

Schedule: Now

Selected hosts and firmwares:

x240-hyper-v-2.hyper-v.lenovopresslab.local:

- IBM Online SAS/SATA Hard Disk Drive Update Program
- IBM Flex System x240 UEFI Flash Update
- Online Broadcom NetXtreme and NetXtreme II Firmware Utility for Windows 2.4.1d4
- Mellanox WinOF update for Windows 2012 R2 Server x86_64
- Brocade BootCode Update for 16G FC HBA
- Emulex OCe11xxx UCNA Firmware Update for Windows
- Emulex HBA (LPe1600x) Firmware Update for Windows
- IBM Flex System FC3172 2
- IBM Flex System FC5172 2
- Emulex HBA (LPe1205/LPe1200x) Firmware Update for Windows
- IBM Dynamic System Analysis (DSA)

Figure 10-54 Review summary

You can monitor the progress and check the details of your task by clicking **Task** in Task Scheduler View window, as shown in Figure 10-55.

Task name: Cluster Firmware Update

Status: Running

Update Details:

Step 1: Download firmware

Status	Progress	Message	Start Time	End Time
Finished	100%	Download Completed	2/18/2015 11:03:02 AM	2/18/2015 11:04:20 AM

Step 2: Update progress

▸ x240-hyper-v-1.hyper-v.lenovopresslab.local Not Started

▾ x240-hyper-v-2.hyper-v.lenovopresslab.local Running Updating

Firmware Name	Installed Version	New Version	State	Message
IBM Online SAS/SATA Hard Disk Drive Updat	Undetected	sas-1.14.04-1	Not Start	The device is not
IBM Flex System x240 UEFI Flash Update	B2E142AUS-1.50	B2E142A-1.50	Not Start	The package ver
Emulex HBA (LPe1600x) Firmware Update for	Undetected	ibm14a-10.2.261.36-1	Not Start	The device is not
Firmware Update for ServeRAID M5115 PSoC	Undetected	m5115-68-1	Not Start	The device is not
Online Broadcom NetXtreme and NetXtreme	Undetected	2.4.1d4	Not Start	The device is not
IBM Flex System FC3172 2	Undetected	3.11af.d-8g-flex	Not Start	The device is not
Integrated Management Module 2 (IMM2) U	1A0058R-4.20	1A0064L-4.50	Running	Start Calling iFlas
Mellanox WinOF update for Windows 2012 R	Undetected	4.61.50000p4	Not Start	The device is not
IBM Flex System FC5172 2	Undetected	3.80.09-16g-flex	Not Start	The device is not
LSI 2004 SAS Controller BIOS and Firmware U		x240-1.18.01	Running	Package installat

Figure 10-55 UIM for SCVMM: Task Status details

When the cluster update task is completed successfully, it is reflected in the task status, as shown in Figure 10-56.

Task name: Cluster Firmware Update

Status: Finished

Update Details:

Step 1: Download firmware

Status	Progress	Message	Start Time	End Time
Finished	100%	Download Completed	2/18/2015 11:03:02 AM	2/18/2015 11:04:20 AM

Step 2: Update progress

▸ x240-hyper-v-2.hyper-v.lenovopresslab.local

Finished Success

▸ x240-hyper-v-1.hyper-v.lenovopresslab.local

Finished Success

Figure 10-56 Task completed successfully

You can perform the same firmware update actions for another clusters, if required.

Abbreviations and acronyms

AD	Active Directory
ATMs	Automated teller machines
BE3	BladeEngine 3
BYOD	Bring-your-own-device
CAD	Computer-aided design
CIFS	Common Internet File System
CIM	Common Information Model
CMM	Chassis Management Module
COM	Component Object Model
DCOM	Distributed component object model
DDC	Desktop Delivery Controller
DPM	Distributed Power Management
DRS	Distributed Resource Scheduler
FT	Fault tolerance
FoD	Features On Demand
GPO	Group Policy Object
GPU	Graphics processing unit
GUI	Graphical user interface
HA	High availability
HDD	Hard disk drive
HVD	Hosted virtual desktop
IBM	International Business Machines Corporation
ICA	Independent Channel Architecture
IMM2	Integrated Management Module II
IOP	Input/output operation
IPC	Interprocess communication
ITSO	International Technical Support Organization
LAN	Local area network
LOM	LAN-on-motherboard
LRO	Large receive offload
LUN	Logical unit number
MCS	Machine Creation Services
MDisk	Managed disk
MSDE	Microsoft Data Engine
MSRP	Microsoft Roaming Profile

NAS	Network-attached storage
NFS	Network File System
NIC	Network interface card
NPIV	N_Port ID virtualization
OU	Organizational unit
PVS	Provisioning Services
RA	Reference Architecture
SAN	Storage area network
SAS	Serial-attached SCSI
SEN	Storage Expansion Node
SLC	Single level cell
SLP	Service Location Protocol
SNIA	Storage Networking Industry Association
SNMP	Simple Network Management Protocol
SSD	Solid-state drive
SSH	Secure Shell
TCO	Total cost of ownership
TOE	TCP offload engine
TSO	TCP segmentation offload
UIM	Upward Integration Module
VDA	Virtual Desktop Agent
VLAN	Virtual LAN
VM	Virtual machine
eMLC	Enterprise multi-level cell
pNIC	Physical NIC
pvDisk	Personal vDisk
vDisk	Virtual disk
vNIC	Virtual NIC

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this book.

Lenovo Press publications

The following Lenovo Press publications provide more information about the topics in this document:

- ▶ *Flex System Products and Technology*, SG24-8255
- ▶ *NIC Virtualization in Flex System Fabric Solutions*, SG24-8223
- ▶ *Flex System Networking in an Enterprise Data Center*, REDP-4834

You can search for, view, or download these documents and other books, papers, and product guides at the following website:

<http://lenovopress.com>

Online resources

The following websites also are relevant as further information sources:

- ▶ Lenovo Reference Architecture for Lenovo Client Virtualization
<http://lenovopress.com/tips1275>
- ▶ Lenovo Reference Architecture for Citrix XenDesktop
<http://lenovopress.com/tips1278>



Implementing Lenovo Client Virtualization with Citrix XenDesktop



(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages



Implementing Lenovo Client Virtualization with Citrix XenDesktop

Introduces Lenovo x86 servers and Citrix XenDesktop offerings

Reviews design, planning, and deployment considerations

Provides step-by-step configuration guidance

Describes VMware vSphere and Microsoft Hyper-V implementation scenarios

The Lenovo Client Virtualization offers robust, cost-effective, and manageable virtual desktop solutions for a wide range of clients, user types, and industry segments. These solutions help to increase business flexibility and staff productivity, reduce IT complexity, and simplify security and compliance. Based on a reference architecture approach, this infrastructure supports various hardware, software, and hypervisor platforms.

The Lenovo Client Virtualization solution with Citrix XenDesktop that is running on System x rack and blade servers offers tailored solutions for every business, from the affordable all-in-one Citrix VDI-in-a-Box for simple IT organizations to the enterprise-wide Citrix XenDesktop. XenDesktop is a comprehensive desktop virtualization solution with multiple delivery models that is optimized for flexibility and cost-efficiency.

This Lenovo Press publication provides an overview of the Lenovo Client Virtualization solution, which is based on Citrix XenDesktop that is running on System x rack and blade servers. It highlights key components, architecture, and benefits of this solution. It also provides planning and deployment considerations, and step-by-step instructions about how to perform specific tasks.

This book is intended for IT professionals who are involved in the planning, design, deployment, and management of the Lenovo Client Virtualization that is built on System x family of servers that are running Citrix XenDesktop.



**BUILDING
TECHNICAL
INFORMATION
BASED ON
PRACTICAL
EXPERIENCE**

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.