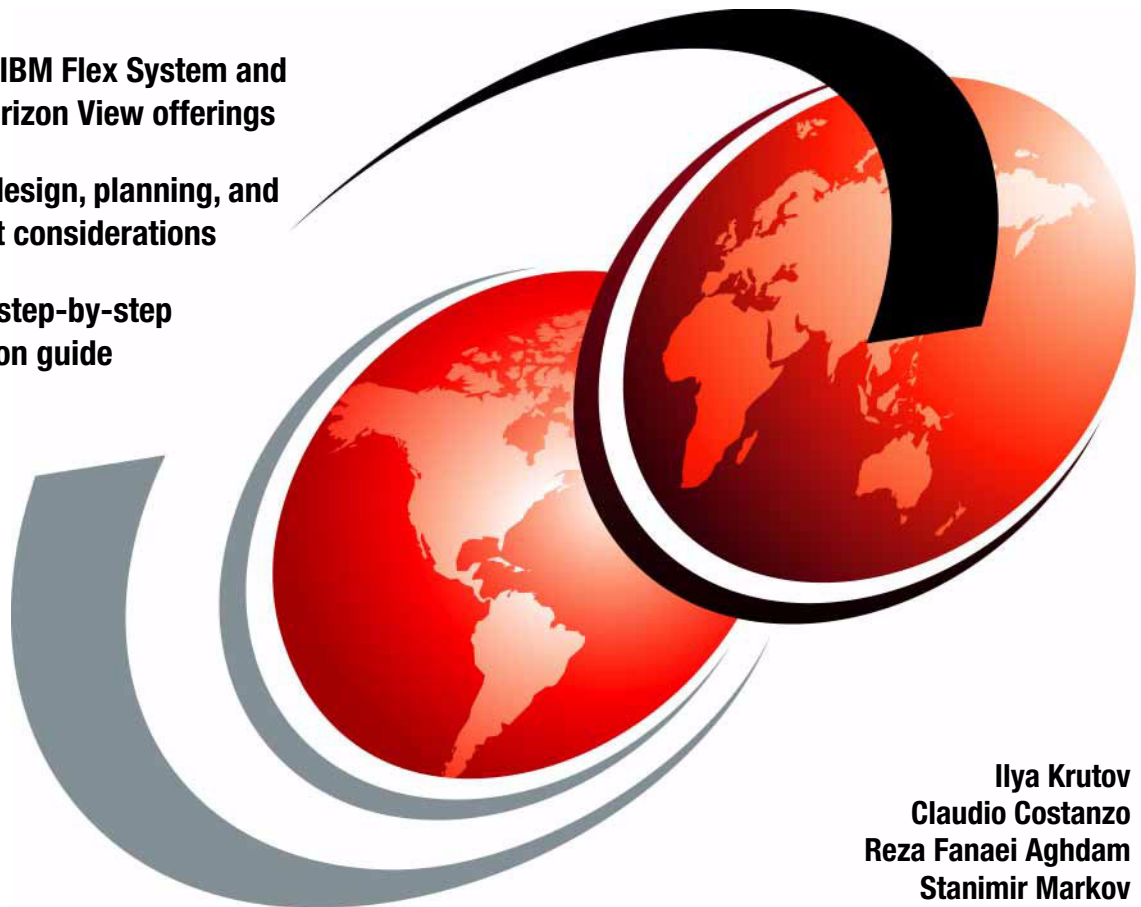


# Implementing VMware Horizon View on IBM Flex System

**Introduces IBM Flex System and VMware Horizon View offerings**

**Describes design, planning, and deployment considerations**

**Provides a step-by-step configuration guide**



**Ilya Krutov  
Claudio Costanzo  
Reza Fanaei Aghdam  
Stanimir Markov**





International Technical Support Organization

## **Implementing VMware Horizon View on IBM Flex System**

April 2014

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

**First Edition (April 2014)**

This edition applies to IBM Flex System and VMware Horizon View 5.2.

**© Copyright International Business Machines Corporation 2014. All rights reserved.**

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



# Contents

- Notices** ..... vii
- Trademarks** ..... viii
- Preface** ..... ix
  - Authors ..... x
  - Now you can become a published author, too! ..... xii
  - Comments welcome ..... xiii
  - Stay connected to IBM Redbooks publications ..... xiii
- Chapter 1. IBM SmartCloud Desktop Infrastructure overview** ..... 1
  - 1.1 Virtual desktop infrastructure overview ..... 2
  - 1.2 IBM SmartCloud Desktop Infrastructure ..... 3
  - 1.3 IBM Flex System ..... 7
  - 1.4 VMware Horizon View ..... 10
  - 1.5 Integration with other IBM software products ..... 13
- Chapter 2. IBM Flex System components for VDI** ..... 15
  - 2.1 Introduction to IBM Flex System ..... 16
  - 2.2 Planning for IBM Flex System components ..... 17
  - 2.3 IBM Flex System Enterprise Chassis ..... 18
  - 2.4 IBM Flex System Compute Nodes ..... 20
    - 2.4.1 IBM Flex System x222 Compute Node ..... 21
    - 2.4.2 IBM Flex System x240 Compute Node ..... 22
    - 2.4.3 IBM Flex System x440 Compute Node ..... 24
    - 2.4.4 IBM Flex System PCIe Expansion Node ..... 25
    - 2.4.5 VMware ESXi 5.1 embedded hypervisor ..... 26
  - 2.5 Storage considerations ..... 27
    - 2.5.1 IBM Flex System V7000 ..... 27
    - 2.5.2 IBM Storwize V7000 Unified System ..... 29
    - 2.5.3 IBM Flex System Storage Expansion Node ..... 31
    - 2.5.4 IBM FlashSystem 820 and IBM FlashSystem 720 ..... 32
    - 2.5.5 SSDs compared to HDDs ..... 33
    - 2.5.6 RAID considerations ..... 34
  - 2.6 Network considerations ..... 36
    - 2.6.1 IBM Flex System 10GbE network switches ..... 37
    - 2.6.2 Network adapters ..... 42
  - 2.7 Flex System Fibre Channel switches ..... 47
    - 2.7.1 FC5022 16Gb SAN Scalable Switch ..... 47
    - 2.7.2 Fibre Channel adapters ..... 48

2.8 IBM Flex System Manager functions and considerations . . . . .	50
2.8.1 Management network . . . . .	52
2.8.2 Chassis Management Module. . . . .	54
2.8.3 Integrated Management Module II . . . . .	54
2.8.4 Configuration Patterns . . . . .	55
2.8.5 Storage connectivity selection guidance . . . . .	56
<b>Chapter 3. VMware vSphere design considerations . . . . .</b>	<b>65</b>
3.1 Compute servers layer . . . . .	66
3.1.1 ESXi hypervisor. . . . .	66
3.1.2 VMware vCenter Server . . . . .	67
3.1.3 vMotion and Storage vMotion . . . . .	68
3.1.4 Distributed Resource Scheduler . . . . .	69
3.1.5 High Availability considerations. . . . .	70
3.1.6 vSphere licensing considerations . . . . .	72
3.1.7 Flex System integration with VMware . . . . .	73
3.2 Networking considerations . . . . .	74
3.2.1 Virtual switches . . . . .	74
3.2.2 Ports and port groups . . . . .	75
3.2.3 Uplink ports . . . . .	75
3.3 Storage considerations . . . . .	76
3.3.1 Local or shared storage . . . . .	76
3.3.2 Tiered storage . . . . .	76
3.3.3 Load balancing . . . . .	77
3.3.4 Redundancy . . . . .	78
<b>Chapter 4. VMware Horizon View design considerations . . . . .</b>	<b>79</b>
4.1 VMware Horizon View components . . . . .	81
4.2 Choosing a desktop protocol . . . . .	87
4.3 VMware View provisioning . . . . .	89
4.3.1 <b>Dedicated and floating desktop pools . . . . .</b>	<b>90</b>
4.3.2 Provisioning by using full and linked clones . . . . .	93
4.4 Storage configuration . . . . .	98
4.5 Network configuration . . . . .	103
4.6 Choosing desktop and application delivery model . . . . .	105
4.7 Operational model and sizing guidelines. . . . .	106
4.7.1 Workload definition for the IBM RA test environment . . . . .	107
4.7.2 VDI compute node configuration. . . . .	108
4.7.3 Management services configuration . . . . .	111
4.7.4 Shared storage configuration . . . . .	113
<b>Chapter 5. IBM Flex System and VMware View lab environment . . . . .</b>	<b>117</b>
5.1 Lab environment . . . . .	118
5.2 VMware View solution overview . . . . .	119

5.3 IBM Flex System chassis overview . . . . .	121
5.4 Storage configuration overview . . . . .	124
5.5 Network configuration overview . . . . .	125
5.6 VDI solution planning . . . . .	128
5.6.1 Management Cluster component model . . . . .	131
5.6.2 VDI Cluster component model . . . . .	132
5.6.3 Desktop pool consideration . . . . .	134
<b>Chapter 6. Deploying IBM Flex System . . . . .</b>	<b>135</b>
6.1 Initial configuration of Chassis Management Module . . . . .	137
6.1.1 Connecting to the CMM . . . . .	137
6.1.2 Using the initial setup wizard . . . . .	139
6.1.3 Configuring IP addresses for the chassis components . . . . .	154
6.2 IBM Flex System Manager Setup wizard . . . . .	155
6.3 Selecting a chassis to manage . . . . .	175
6.4 Discovery and inventory collection . . . . .	179
6.4.1 Discovery . . . . .	180
6.4.2 I/O modules . . . . .	194
6.5 IBM Flex System Fabric EN4093 10Gb configuration . . . . .	204
6.6 IBM Flex System x240 compute node configuration . . . . .	222
6.7 IBM Flex System V7000 Storage Node configuration . . . . .	236
6.7.1 IBM Flex System V7000 Storage Node initial configuration . . . . .	236
6.7.2 IBM Flex System V7000 Storage Node Setup Wizard . . . . .	240
6.7.3 MDisk configuration . . . . .	249
6.7.4 Zoning configuration . . . . .	253
6.7.5 Configuring volumes . . . . .	265
6.7.6 Configuring hosts . . . . .	268
6.8 VMControl activation . . . . .	274
<b>Chapter 7. Deploying VMware Horizon View infrastructure . . . . .</b>	<b>277</b>
7.1 Installing vSphere components and infrastructure services . . . . .	278
7.1.1 Configuring ESXi . . . . .	278
7.1.2 Installing infrastructure services . . . . .	279
7.1.3 Creating vCenter data source name . . . . .	283
7.1.4 Installing vCenter Server . . . . .	291
7.1.5 Installing vSphere Web Client . . . . .	305
7.2 Configuring vSphere . . . . .	310
7.3 Installing View Composer . . . . .	315
7.4 Installing View Connection Server . . . . .	327
7.5 Configuring View Connection Server initially . . . . .	335
<b>Chapter 8. Operating IBM Flex System . . . . .</b>	<b>349</b>
8.1 Configuring VMControl for vSphere integration . . . . .	350
8.2 Navigating the vSphere environment by using FSM . . . . .	352

8.3 Automating tasks . . . . .	354
8.4 Introducing IBM FSM Explorer . . . . .	363
8.5 Monitoring and logging in Flex System . . . . .	365
<b>Chapter 9. Operating VMware Horizon View infrastructure . . . . .</b>	<b>373</b>
9.1 Preparing the base Microsoft Windows 7 operating system x64 image to deploy . . . . .	374
9.1.1 Creating a customization specification file . . . . .	374
9.1.2 Creating vCenter folders for a VDI . . . . .	386
9.1.3 VMware View Administrator check . . . . .	388
9.1.4 Provisioning a full VM image. . . . .	390
9.1.5 Provisioning a linked clone virtual desktop image. . . . .	393
9.2 Installing the VMware Horizon View Agent . . . . .	394
9.3 Configuring active directory policies . . . . .	398
9.3.1 Configuring View Persona Management active directory policies . . . . .	400
9.3.2 Allowing a connection to a remote desktop users group policy . . . . .	411
9.4 VMware View Manager and desktop pools. . . . .	413
9.4.1 Configuring Event DB . . . . .	413
9.4.2 Provisioning a linked clone virtual desktop . . . . .	416
9.4.3 Provisioning a full virtual machine virtual desktop. . . . .	440
9.5 Operating View Composer . . . . .	463
9.5.1 Performing a desktop refresh operation . . . . .	463
9.5.2 Performing a desktop recompose operation . . . . .	468
9.5.3 Performing a desktop rebalance operation . . . . .	473
9.5.4 Migrating virtual desktops to another data store . . . . .	476
<b>Abbreviations and acronyms . . . . .</b>	<b>483</b>
<b>Related publications . . . . .</b>	<b>485</b>
IBM Redbooks . . . . .	485
Online resources . . . . .	485
Help from IBM . . . . .	486

# Notices

This information was developed for products and services offered in the U.S.A.

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

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

*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION 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 states 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. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results 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.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


## **COPYRIGHT LICENSE:**

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM 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 IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®	IBM Flex System Manager™	Redbooks (logo)  ®
DB2®	IBM SmartCloud®	Storwize®
DS8000®	POWER®	System Storage®
Easy Tier®	Power Systems™	System x®
FlashCopy®	POWER7®	Tivoli®
FlashSystem™	PureFlex™	VMready®
Global Technology Services®	RackSwitch™	X-Architecture®
IBM®	Real-time Compression™	
IBM Flex System™	Redbooks®	

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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.

ITIL is a registered trademark, and a registered community trademark of The Minister for the Cabinet Office, and is registered in the U.S. Patent and Trademark Office.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

# Preface

The IBM® SmartCloud Desktop Infrastructure offers robust, cost-effective, and manageable virtual desktop solutions for various clients, user types, and industry segments. These solutions can help 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.

IBM SmartCloud® Desktop Infrastructure with VMware Horizon View simplifies desktop and application management and increases security and control. Horizon View delivers a personalized, high-fidelity experience for users across sessions and devices. It also enables higher availability and agility of desktop services that are unmatched by traditional PCs, reducing the total cost of desktop ownership is reduced. Users can enjoy new levels of productivity and the freedom to access desktops from more devices and locations with IT greater policy control.

This IBM Redbooks® publication provides an overview of the SmartCloud Desktop Infrastructure solution that is based on VMware Horizon View that is running on IBM Flex System™. 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 planning, design, deployment, and management of the IBM SmartCloud Desktop Infrastructure that is built on IBM Flex System that is running VMware Horizon View.

## Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.



**Ilya Krutov** is a Project Leader at the ITSO Center in Raleigh. Ilya has more than 15 years of experience in the IT industry, and he has been with IBM since 1998. Before he joined the ITSO, Ilya served in IBM as a Team Leader, Portfolio Manager, Brand Manager, IT Specialist, and Certified Instructor. Ilya has expertise in IBM System x®, BladeCenter®, and Flex System products; server operating systems, and networking solutions. He authored over 130 books, papers, Product Guides, and Solution Guides. He has a Specialist's degree in computer engineering from the Moscow Engineering and Physics Institute.



**Claudio Costanzo** is an IBM Certified IT Specialist with more than 13 years of experience with IBM and in the IT field. He joined IBM in 2000 and performed the following job roles: System Administrator, Client Technical Leader, Virtualization Specialist, and Solution Designer. He supports client's projects that are related to user devices, desktops management, Microsoft Network Solutions and Microsoft Operating systems, VDI infrastructures, VMware, and Citrix virtualization solutions. He has many Microsoft Certifications that are related to messaging and OS, and he is a Citrix Certified Administrator. He is passionate about IT in general, with a particular focus on virtualization technologies and IBM Cloud Solutions.





**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, IBM Certified Specialist), Midrange Storage Technical Support, and VMware Certified Professional.



**Stanimir Markov** is a Technical Leader of VMware Center of Excellence in IBM Global Technology Services® and a core member of the IBM virtualization development team. He leads the virtualization efforts in large-scale transformation projects for enterprise customers. He also has a key role in defining VMware best practices across IBM GTS Delivery. In addition to his work on design, implementation, and support of complex virtual infrastructures, Stanimir delivers authorized VMware classes as a VMware Certified Instructor. Stanimir has the highest VMware certification: VMware Certified Design Expert. He also holds a bachelor degree in computer science, and is certified in Microsoft and Citrix technologies and ITIL. Stanimir has been with IBM since 2006.

Special thanks to Matthew Darlington, VDI lead in the Advanced Technical Skills team, who made a significant contribution to the development of this book by extensively consulting and guiding the team on VDI topics.

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

- ▶ Kevin Barnes
- ▶ Tamikia Barrow
- ▶ David Bennin
- ▶ Ella Buslovich
- ▶ Mary Comianos
- ▶ Richard Conway
- ▶ Shari Deiana
- ▶ Cheryl Gera
- ▶ David Watts
- ▶ Debbie Willmschen

International Technical Support Organization, Raleigh Center

- ▶ Amy Freeman
- ▶ Britni Coble
- ▶ Cam-Thuy Do
- ▶ Andreas Groth
- ▶ Michael Perks

IBM

- ▶ Kae Chy Chen
- ▶ Matt Coppinger
- ▶ John Dodge
- ▶ Rasmus Jensen
- ▶ Andrew Johnson
- ▶ Paul Kohler

VMware

## **Now you can become a published author, too!**

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at this website:

<http://www.ibm.com/redbooks/residencies.html>

## Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:  
<http://www.ibm.com/redbooks>
- ▶ Send your comments in an email to:  
[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)
- ▶ Mail your comments to:  
IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks publications

- ▶ Find us on Facebook:  
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:  
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:  
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:  
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:  
<http://www.redbooks.ibm.com/rss.html>





# IBM SmartCloud Desktop Infrastructure overview

In this chapter, we introduce IBM SmartCloud Desktop Infrastructure and describe one of its solutions, VMware Horizon View on IBM Flex System.

This chapter includes the following topics:

- ▶ Virtual desktop infrastructure overview
- ▶ IBM SmartCloud Desktop Infrastructure
- ▶ IBM Flex System
- ▶ VMware Horizon View
- ▶ Integration with other IBM software products

## 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 limiting them to an office. Personal tablets, smartphones, and other mobile devices 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 the device in a secure manner. The resources are centralized and users can move between locations while accessing the applications and data. By using this access method, administrators have better control over the management of the desktop and tighter security.

The idea of having a centralized infrastructure has been around since the day of mainframe and terminal clients. In the early 1990s, this centralized infrastructure shifted to a client/server model to meet the need for more flexibility by the user. This shift led to the idea of having a centralized infrastructure for back-end processing and gave users the ability to save programs and files locally on hard disk drives.

As the workforce changed from office-oriented to more mobile and on demand, the need for flexibility grew and VDI provides a flexible solution for many businesses. The market for VDI changed how vendors are marketing their solutions. Traditional IT shops can build out their infrastructure piece by piece with the software or hypervisor. This type of solution tends to increase the amount of time that is needed to manage the storage, servers, and network environment.

A new market emerged with the introduction of complete solutions of all aspects that are needed to implement, deploy, and maintain a virtual desktop solution. IBM PureSystems leads the way with the only homogeneous vendor infrastructure that provides software, servers, storage, and networking in a single management system.

One of the most important aspects of deploying a virtual desktop solution is to control costs while providing a familiar user experience and functions. 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 out grow the initial deployment and find it hard 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.

IBM VDI solutions are consolidated under the SmartCloud Desktop Infrastructure umbrella.

## 1.2 IBM SmartCloud Desktop Infrastructure

The IBM SmartCloud Desktop Infrastructure offers robust, cost-effective, and manageable virtual desktop solutions for various 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 SmartCloud Desktop Infrastructure solution with VMware Horizon View running on IBM Flex System simplifies IT manageability and control. It delivers high fidelity user experiences across devices and networks. The features of Horizon View that are included in the SmartCloud Desktop Infrastructure solution provide enhanced security, high availability, centralized management and control, and scalability.

The hosted virtual desktop (HVD) approach, which is combined with the application streaming, 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.

Virtual desktops in today's business climate include the following key features:

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

IBM SmartCloud Desktop Infrastructure 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
- ▶ Better supports growth initiatives for mobility and flexible work locations

The IBM SmartCloud Desktop Infrastructure solution with VMware Horizon View running on IBM Flex System includes the following components:

- ▶ Virtual infrastructure software: VMware Horizon View
- ▶ Hardware platform:
  - IBM Flex System
  - IBM System Storage®
- ▶ Integration services:
  - Assess and plan
  - Design
  - Implement
  - Operate and manage



Figure 1-1 shows the functional components of the SmartCloud Desktop Infrastructure solution.

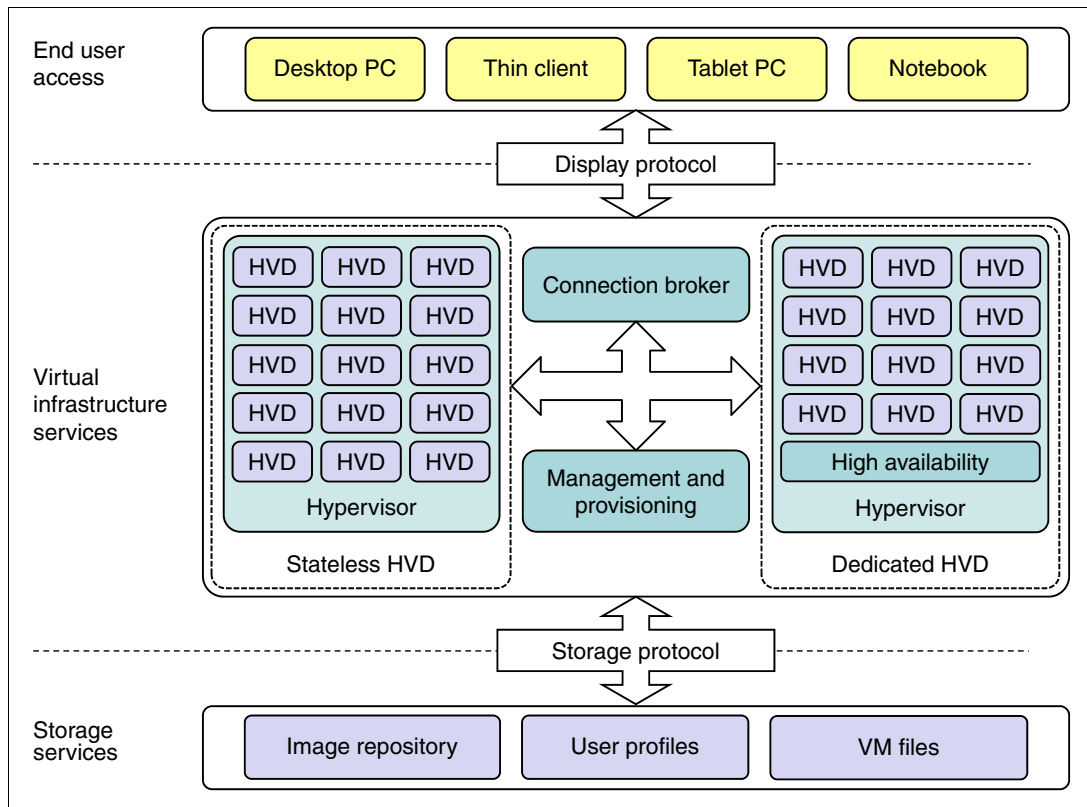


Figure 1-1 SmartCloud Desktop Infrastructure functional components

The SmartCloud Desktop Infrastructure 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 RDP and PCoIP display protocols are available in Horizon View solution.

- ▶ 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 that are supported by Horizon View include Network File System (NFS), Common Internet File System (CIFS), iSCSI, and Fibre Channel.

The virtual infrastructure services layer has 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* (HVD).

- ▶ Hosted virtual desktops

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

- ▶ Connection broker

The connection 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 are allowed 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 enable the centralized management of the virtual infrastructure, which provides a single console to manage multiple tasks. They also 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.

There are two types of the assignment models for the user HVDs: persistent and non-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 every time 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.

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 new 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 must provide 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

IBM Flex System can be used in the future and is an integrated platform that satisfies these requirements.

## 1.3 IBM Flex System

IBM Flex System is an integrated platform that delivers custom-tuned, client-specific configurations for optimum flexibility. IBM Flex System combines compute nodes, networking, storage, and management into a complete data center building block that is built for the future and heterogeneous data centers with flexibility and open choice of architectures, hypervisors, and environments.

Figure 1-2 shows IBM Flex System.



*Figure 1-2 IBM Flex System*

IBM Flex System offers the following unique capabilities that make this platform an exceptional choice for the deployment of the SmartCloud Desktop Infrastructure solution:

► Compute nodes

Compute nodes provide sufficient processing capacity for the most demanding SmartCloud Desktop Infrastructure deployments.

IBM Flex System x240 is a dual-socket Intel Xeon processor E5-2600 product family-based compute node. It supports the most powerful 135 W Intel Xeon processor E5-2690, up to 768 GB of memory, and up to 16 physical I/O connections to provide scalable, high-density HVD deployments.

The x240 compute node also supports local solid-state drives to address VDI IOPS performance questions, and it supports GPU adapters through the Flex System PCIe Expansion Node for true high-performance graphics user experience.

IBM Flex System x222 Compute Node is a high-density dual-server offering that has two independent dual-socket servers in one mechanical package. Each server has two 10 GbE Virtual Fabric ports, and it supports up to 384 GB of memory. The x222 can be used as a dense VDI compute node for virtual desktops that do not require large amounts of memory.

► Networking

SmartCloud Desktop Infrastructure requires sufficient network bandwidth and efficient traffic management to host as many VMs as possible to ensure that all computing resources are not underutilized. When integrated into a chassis, IBM Flex System networking with IBM Virtual Fabric capabilities can help to reduce communication latency and provide the required bandwidth with 10 Gb Ethernet LAN connectivity that has 40 Gb uplinks and 8 Gb or 16 Gb FC SAN connectivity.

Virtual Fabric Adapters offer virtual network interface card (NIC) capability to allow up to 32 logical ports on a single compute node, with controllable bandwidth allocation to manage traffic prioritization. vNIC capability helps to simplify deployment and bandwidth management for VDI hosts by providing flexible network configuration capabilities.

► Management

IBM Flex System Manager™ is a systems management appliance that drives efficiency and cost savings in the data center. Flex System Manager provides a pre-integrated and virtualized management environment across servers, storage, and networking that is easily managed from a single interface. A single focus point for seamless multichassis management provides an instant and resource-oriented view of chassis and chassis resources for IBM System x and IBM Power Systems™ compute nodes.

Flex System Manager allows centralized management of the ESXi hypervisors that are used in the IBM's architecture for Horizon View. It also supports configuration patterns to simplify deployment of VDI hosts.

► Storage

As virtualized storage systems, integrated IBM Flex System V7000 Storage Node or external IBM Storwize® V7000 complement virtual desktop environments. These system offer robust enterprise-class storage capabilities, which include thin provisioning, automated tiering, internal and external virtualization, clustering, replication, multiprotocol support, and a next-generation graphical user interface (GUI). These features can be applied in virtual desktop environments to optimize storage capacity and performance and to simplify desktop user profile management and backup. These systems are flexible enough to support entry virtual desktop environments, but can also be scaled to support enterprise virtual desktop environments.

In summary, IBM Flex System in a SmartCloud Desktop Infrastructure 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 utilization of VDI server resources with flexible local SSD support.
- ▶ Transparent support for high-performance remote graphics through PCIe Expansion Node with GPU adapters installed.
- ▶ Lower communication latency because of integrated switching capabilities for a better user experience.
- ▶ Simplified deployment and management of physical and virtual infrastructures because of integrated design and IBM Flex System Manager capabilities.

## 1.4 VMware Horizon View

IBM SmartCloud Desktop Infrastructure with VMware Horizon View simplifies desktop and application management and increases security and control. Horizon View delivers a personalized high fidelity experience for users across sessions and devices. It also enables higher availability and agility of desktop services that are unmatched by traditional PCs, while reducing the total cost of desktop ownership. Users can enjoy new levels of productivity and the freedom to access desktops from more devices and locations with IT greater policy control.

The following VMware View features provide a familiar experience for the user:

- ▶ Use multiple monitor support for RDP and PCoIP; with PCoIP, you can adjust the display resolution and rotation separately for each monitor.
- ▶ Print from a virtual desktop (in a Microsoft Windows environment) to any local or networked printer.
- ▶ Access USB devices and other peripheral devices that are connected to the local device that displays your virtual desktop.
- ▶ Manage profiles by using View Persona Management to preserve user profiles and data between sessions and to dynamically synchronize them to a remote CIFS share at configurable intervals. View Persona Management can work with or without Windows roaming profiles.

VMware View offers several levels of security features, including the following features:

- ▶ Two-factor authentication, such as RSA SecurID or RADIUS, or smart cards
- ▶ Pre-created Active Directory accounts to provision View desktops in environments that have read-only access policies for Active Directory
- ▶ SSL tunneling to ensure that all connections are encrypted

The following VMware View features provide centralized administration and management:

- ▶ Microsoft Active Directory
- ▶ Web-based administrative console
- ▶ Use of a template, or master image, to quickly create and provision pools of desktops virtual desktops updates and patches

The following scalability features depend on the VMware virtualization platform to manage both desktops and servers:

- ▶ You can integrate with VMware vCenter to achieve cost-effective densities, high levels of availability, and advanced resource allocation control for your virtual desktops.
- ▶ You can use View Composer to quickly create desktop images that share virtual disks with a master image. By using linked clones in this way, you conserve disk space and simplify the management of patches and updates to the operating system.

VMware Horizon View software components are shown in Figure 1-3.

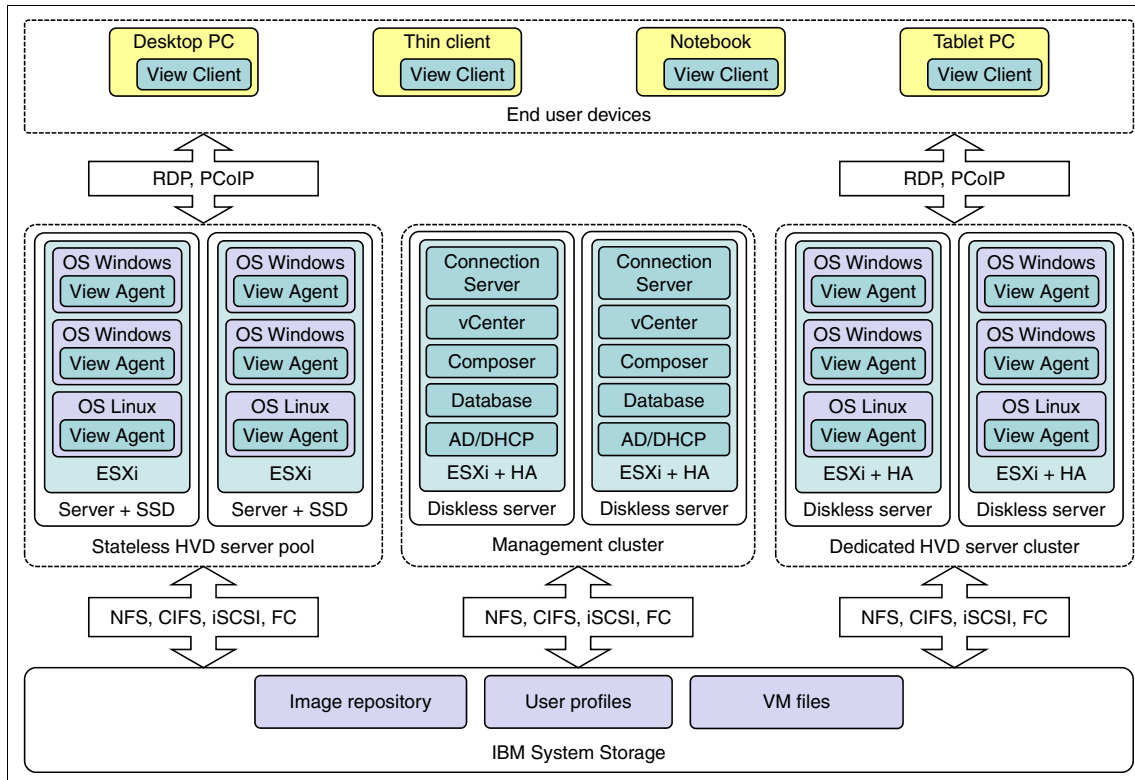


Figure 1-3 VMware Horizon View software components

The VMware View core services have the following software components:

- **View Client**  
View Client is client software to access View virtual desktops. View Client can run on a tablet; on a Windows, Linux, or Mac PC or notebook; on a thin client, and on other devices.
- **View Agent**  
View Agent communicates with View Client to provide features, such as connection monitoring, virtual printing, View Persona Management, access to locally connected USB devices, and single sign-on (SSO) capabilities.
- **View Connection Server**  
View Connection Server is a software service that acts as a broker for client connections. View Connection Server authenticates users through Windows Active Directory and directs the request to the appropriate VM.



- ▶ VMware vCenter

VMware vCenter service acts as a central administrator for VMware ESX/ESXi servers that are connected on a network. vCenter Server provides the central point for configuring, provisioning, and managing VMs in the data center.

- ▶ View Composer

View Composer can create a pool of linked clones from a specified parent VM. Each linked clone acts similar to an independent desktop, with a unique host name and IP address, yet the linked clone requires less storage because it shares a base image with the parent.

Users can access their personalized virtual desktop from a company notebook, their home PC, a thin client device, a Mac, or a tablet. From tablets and from Mac, Linux, and Windows notebooks and PCs, users open View Client to see their View desktop. Thin client devices use View Thin Client software. They can be configured so that the only application that users can start directly on the device is View Thin Client.

## 1.5 Integration with other IBM software products

IBM SmartCloud Desktop Infrastructure enables easy integration with optional security and endpoint management technologies, including the following technologies:

- ▶ IBM Security Access Manager for Enterprise Single Sign-On offers streamlined user access with automated sign-on and sign-off plus a single password for all applications. This technology can reduce help desk costs, improve productivity, and strengthen security for virtualized desktops.
- ▶ IBM Tivoli® Endpoint Manager combines endpoint and security management into a single solution. With this solution, your team can see and manage physical and virtual endpoints, such as servers, desktops, roaming notebooks, and specialized equipment such as point-of-sale devices, automated teller machines (ATMs), and self-service kiosks.





# IBM Flex System components for VDI

In this chapter, we describe the IBM Flex System components to consider when you are designing a virtual desktop infrastructure (VDI) solution that is based on VMware Horizon View.

This chapter includes the following topics:

- ▶ Introduction to IBM Flex System
- ▶ Planning for IBM Flex System components
- ▶ IBM Flex System Enterprise Chassis
- ▶ IBM Flex System Compute Nodes
- ▶ Storage considerations
- ▶ Network considerations
- ▶ Flex System Fibre Channel switches
- ▶ IBM Flex System Manager functions and considerations

## 2.1 Introduction to IBM Flex System

IBM Flex System is a custom-build infrastructure solution that integrates Intel x86 and IBM Power System compute nodes and storage systems, such as Flex System V7000 storage node, standard-based flexible enhanced networking, and management appliance, in a single chassis. It meets the increasing demand of computing capacity, integration, manageability, optimization, scalability, security, cost-efficiency, and flexibility. Flex System extends compute and networking choices to interoperate with existing environments.

Flex System is also designed to support emerging technologies. It supports up to four 40 Gbps Ethernet ports, which improves the Ethernet bandwidth to support higher speed devices in the future. With its design to support future technologies, it offers investment protection. With Flex System, you can increase bandwidth and storage capacity without compromise and without replacing existing Flex System components by applying features, such as pay as you grow scalability and capacity on demand. Flex System offers unmatched flexibility for you to customize your own chassis that is based on your own requirements of computing and storage capacity, network bandwidth, and so on, to meet rapidly changing IT demands.

Another advantage of Flex System is the management appliance, which provides the management of compute nodes, network, storage, and virtualization from a single management console. The management appliance is designed to manage multiple chassis in a single console. The ease of use, simplicity, and integration of the Flex System management appliance enables you to reduce costs for IT administration.

Flex System can reduce management costs up to 50% by integrating resource pools across compute, storage, and network. In addition, it reduces energy costs up to 40%. Further, it reduces software licensing costs by licensing fewer needed cores than previous generations. It can also reduce network latency up to 50% by enabling node-to-node traffic, which avoids the top-of-rack (TOR) switches.

## 2.2 Planning for IBM Flex System components

To design the VMware Horizon View infrastructure, you must determine what resources are needed by your infrastructure servers and your persistent and non-persistent desktops. Each category of user operates a specific software platform with a workload that involves different hardware resources. The assessment of how those resources are used is based on categories, such as CPU, memory, or I/O, along with the following types of characteristics:

- ▶ Size
- ▶ Percentage of read/write
- ▶ Type of access, such as random or sequential
- ▶ Size of user data and user profile
- ▶ Graphic utilization profile

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

When you are sizing compute nodes, keep in mind the following considerations:

- ▶ Do not overcommit memory because disk swapping can deteriorate 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. As a preferred practice, allow for an overhead of 20% in memory and processor to support these additional VMs without reaching the compute node bounds.
- ▶ A hypervisor uses 3 GB - 6 GB of compute node memory and 1 CPU core.

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

- ▶ The storage for infrastructure servers: Shared storage is the best solution.
- ▶ The storage for stateful desktops: Consider shared storage.
- ▶ The storage for stateless desktops: Consider the use of a local storage or shared storage with high I/O performance.

As a general purpose, consider for redundancy on I/O modules both Ethernet network switches and storage SAN switches.

## 2.3 IBM 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. Furthermore, it meets the needs of varying workloads with scalable IT resource pools for higher utilization 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 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 Intel x86 and IBM POWER® compute nodes. You can use both one-bay and two-bay compute nodes to meet your specific hardware needs. It can also support three 4-bay storage nodes or storage expansion enclosures. Additionally, the rear of the chassis has four high-speed networking switches bays. The compute nodes and storage nodes share common resources, such as power, cooling, management, and I/O resources in the chassis.

The chassis can support 40 Gb speed interconnecting compute, storage, and networking nodes that use a high-performance scalable mid-plane. 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. With support for IBM POWER7® and Intel processor-based nodes, you can choose the architecture that you need. IT can meet the needs of the business by using a single system for multiple architectures and 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, storage, 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 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.

Figure 2-1 shows the front of Enterprise Chassis.



*Figure 2-1 Front of the Flex System Enterprise Chassis*

Figure 2-2 shows the rear of IBM Flex System Enterprise Chassis.




Figure 2-2 Rear of the Flex System Enterprise Chassis

## 2.4 IBM Flex System Compute Nodes

The choice for computer nodes is wide, which is designed for multiple generations of technology. The following available Flex System compute nodes offer high performance for virtualization:

- ▶ Flex System x222
- ▶ Flex System x240
- ▶ Flex System x440

The choice of compute nodes depends on the requirement of the hosted VMs. Flex System x222 is designed for virtualization, dense cloud deployments, and hosted clients. It is a good choice if you want to virtualize workloads while maximizing the density of computing resources.

Flex System x240 is a good choice for VDI. It can host a good density of memory-demanding VMs.

For high resource-using VMs, Flex System x440 brings massive compute power and memory resources. A high VM density on a compute node might not be a goal to reach. Effects for users in a compute node failure is proportional.



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

Table 2-1 Comparison of x222, x240, and x440 compute node features

Feature	x222 (one half)	x240	x440
Processor	E5-2400	E5-2600	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

To help you to make a selection, the next sections describe the following compute nodes:

- ▶ IBM Flex System x222 Compute Node
- ▶ IBM Flex System x240 Compute Node
- ▶ IBM Flex System x440 Compute Node

### 2.4.1 IBM Flex System x222 Compute Node

The IBM Flex System x222 Compute Node is a high-density dual-server that is designed for virtualization, dense cloud deployments, and hosted clients. The x222 has two independent servers in one mechanical package. It 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 if you want to virtualize workloads while maximizing the density of your computing resources.

Figure 2-3 shows a Flex System x222 Compute Node.



Figure 2-3 IBM Flex System x222 Compute Node

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

- ▶ Processor

Includes the Intel Xeon Processor E5-2400 with up to eight cores per processor and up to 2.4 GHz core speeds, depending on the CPU's number of cores, with up to 20 MB of L3 cache and QPI interconnect links of up to 8 GTps. Up to four processors in a standard (half-width) Flex System form factor, 32 cores, and 64 threads maximize the concurrent running of multi-threaded applications.

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

- ▶ Memory

Includes up to 24 DIMM sockets in a standard (half-width) Flex System form factor. Each server up to 12 DIMM sockets DDR3 ECC memory RDIMMs provides 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

Includes up to 16 virtual I/O ports per compute node with integrated 10 Gb Ethernet ports (for more information, see “LAN-on-motherboard” on page 43), offering the choice of Ethernet, Fibre Channel, iSCSI, or FCoE connectivity.

- ▶ Disk

Each server, one 2.5-inch simple-swap SATA drive bay, supports SATA and solid-state drives (SSD). Includes an optional SSD mounting kit to convert a 2.5-inch simple-swap bay into two 1.8-inch hot-swap SSD bays.

- ▶ Operating system

Supports VMware ESXi 5.1 Embedded hypervisor.

## 2.4.2 IBM 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.

Figure 2-4 shows a Flex System x240 Compute Node.



Figure 2-4 Flex System x240 Compute Node

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

- ▶ **Processor**  
Includes the Intel Xeon Processor E5-2600 with up to six cores per processor and up to 3.3 GHz core speeds, depending on the CPU's number of cores, with up to 20 MB of L3 cache and QPI interconnect links of up to 8 GTps. Up to 2 processors, 16 cores, and 32 threads maximize the concurrent running of multi-threaded applications.
- ▶ **Memory**  
Includes up to 24 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**  
Includes up to 16 virtual I/O ports per compute node with integrated 10 Gb Ethernet ports (for more information, see "LAN-on-motherboard" on page 43), offering 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 SSD.
- ▶ **Operating system**  
Supports VMware ESXi 5.1 Embedded hypervisor.

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

### 2.4.3 IBM 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 high networking bandwidth.

Figure 2-5 shows a Flex System x440 Compute Node.



*Figure 2-5 Flex System x440 Compute Node*

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

- ▶ **Processor**  
Includes the Intel Xeon processor E5-4600 with up to eight cores per processor and up to 2.9 GHz core speeds, up to 20 MB of L3 cache, and up to two 8 GTps QPI interconnect links. Up to four processors, 32 cores, and 64 threads maximize the concurrent execution of multithreaded applications.
- ▶ **Memory**  
Includes 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**  
Includes up to 32 virtual I/O ports per compute node with integrated 10 Gb Ethernet ports, offering the choice of Ethernet, Fibre Channel, iSCSI, or FCoE connectivity. For models without integrated 10 Gb ports, you can have up to 64 virtual I/O ports by installing four CN4054 10 Gb Virtual Fabric Adapters.
- ▶ **Disk**  
Two 2.5-inch hot-swap SAS/SATA drive bays support SAS, SATA, and SSD.

- ▶ Operating system  
Supports VMware ESXi 5.1 Embedded hypervisor.

## 2.4.4 IBM Flex System PCIe Expansion Node

For VDI purposes, you can use the IBM Flex System PCIe Expansion Node to attach next-generation graphics processing units (GPU) to x240 compute nodes. The PCIe Expansion Node supports up to four PCIe adapters and two other Flex System I/O expansion adapters.

Figure 2-6 shows the PCIe Expansion Node that is attached to a compute node.



*Figure 2-6 IBM Flex System PCIe Expansion Node attached to a compute node*

The PCIe Expansion Node has 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 (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 attached compute node.

Table 2-2 lists the PCIe GPU adapters that can be used in the VDI solutions.

Table 2-2 Supported adapters

Part number	Description	Maximum supported
47C2120	NVIDIA GRID K1 for IBM Flex System PCIe Expansion Node	1 <sup>a</sup>
47C2121	NVIDIA GRID K2 for IBM Flex System PCIe Expansion Node	1 <sup>a</sup>

a. If installed, only this adapter is supported in the system. No other PCIe adapters can be installed.

NVIDIA GRID K1 and K2 are designed for VDI. NVIDIA GRID cards can be shared between multiple users, with up to 100 concurrent users in GPU sharing configuration for K1. K2 is intended to support heavy 3D applications, such as two power users in GPU pass through configuration.

### 2.4.5 VMware ESXi 5.1 embedded hypervisor

IBM offers versions of VMware vSphere Hypervisor (ESXi) that are customized for select IBM hardware to provide online platform management, including updating and configuring firmware, platform diagnostics, and enhanced hardware alerts. This option, which is delivered on a USB flash drive, is compatible with Flex System compute nodes and IBM System x. At the time of this writing, the last version that is provided by IBM is VMware vSphere Hypervisor (ESXi) 5.1.

**Download information:** You can download the most up-to-date VMware vSphere Hypervisor (ESXi) with IBM Customization from this website:

<http://www.ibm.com/systems/x/os/vmware/>

Choosing this option on the compute nodes that compose the VDI infrastructure produces the following results:

- ▶ Reduces server deployment time. Flex System Management integrates the management of the VMware vSphere Hypervisor (ESXi).
- ▶ Uses disk less compute node, which reduces cost and security exposure.
- ▶ Uses compute node local disks to host non-persistent virtual desktops.

## 2.5 Storage considerations

This section presents some of the storage options to consider for the VDI storage design.

### 2.5.1 IBM Flex System V7000

Flex System V7000 is integrated into IBM PureFlex™ Systems. It is a scalable internal storage system that supports the compute nodes of the Flex System environment. Flex System V7000 is a mid-range storage solution that combines simplicity and outstanding performance with a compact and modular design. It integrates the IBM SAN Volume Controller technology from the high-end IBM System Storage DS8000® family and provides the ability to virtualize internal storage and external SAN-attached storages.

Figure 2-7 shows a Flex System V7000 Storage Node.



*Figure 2-7 Flex System V7000 Storage Node*

One key feature is IBM System Storage Easy Tier®. The system automatically and non-disruptively moves frequently accessed data from hard disk drive (HDD) MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

The following sections provide a quick overview of the hardware and software of Flex System V7000.

## **Hardware overview**

Flex System V7000 consists of a set of drive enclosures. Control enclosures contain disk drives and two node canisters. A collection of up to four control enclosures that are managed as a single system is a Flex System V7000 clustered system.

Expansion enclosures contain drives and are attached to a control enclosure. You can connect a maximum of nine expansion enclosures to a control enclosure. The expansion enclosures can be the Flex System V7000 expansion enclosure or the IBM Storwize V7000 expansion enclosures, or both. Up to two Flex System V7000 expansion enclosures can be connected to a control enclosure. These expansion enclosures must be in the same Flex System chassis as the control enclosure. Up to nine IBM Storwize V7000 expansion enclosures can be connected to the control enclosure. These Storwize V7000 expansion enclosures should be mounted in the rack next to the Flex System chassis where the control enclosure is installed.

Expansion canisters include the serial-attached SCSI (SAS) interface hardware that enables the node canisters to use the drives of the expansion enclosures. An expansion enclosure cannot be connected to more than one control enclosures at the same time.

## **Software overview**

The Flex System V7000 Storage Node 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).

Advantages of the Flex System V7000 Storage Node include greater integration of server and storage management to automate and streamline provisioning.

The Flex System V7000 software performs the following functions for the Compute Nodes that attach to Flex System V7000:

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



The Flex System V7000 system also provides the following functions:

- ▶ Large scalable cache
- ▶ Copy Services:
  - IBM FlashCopy® (point-in-time copy) function, including thin-provisioned FlashCopy to make multiple targets affordable
  - Metro Mirror (synchronous copy)
  - Global Mirror (asynchronous copy)
  - Data migration
  - Volume mirroring
- ▶ Space management
  - IBM System Storage Easy Tier to migrate the most frequently used data to higher performing storage
  - Metering of service quality when combined with IBM Tivoli Storage Productivity Center
  - Thin-provisioned logical volumes
  - Compressed volumes to consolidate storage

## 2.5.2 IBM Storwize V7000 Unified System

The Storwize V7000 Unified system is a virtualizing RAID storage system that provides block and file storage volumes over iSCSI, Fibre Channel, and NFS to hosts. A Storwize V7000 Unified system is made up of a Storwize V7000 storage system and two Storwize V7000 file modules. The Storwize V7000 storage system enables you to improve application flexibility, responsiveness, and availability, while reducing storage usage and complexity through storage virtualization.

Figure 2-8 shows the Storwize V7000 Unified Storage system.



Figure 2-8 IBM Storwize V7000 Unified Storage

One important feature of the Storwize V7000 system is the ability to manage storage that is provided by internal and external storage systems. The Storwize V7000 system acts as the virtualization layer between the host and external storage system.

Figure 2-9 shows Storwize V7000 unified components.

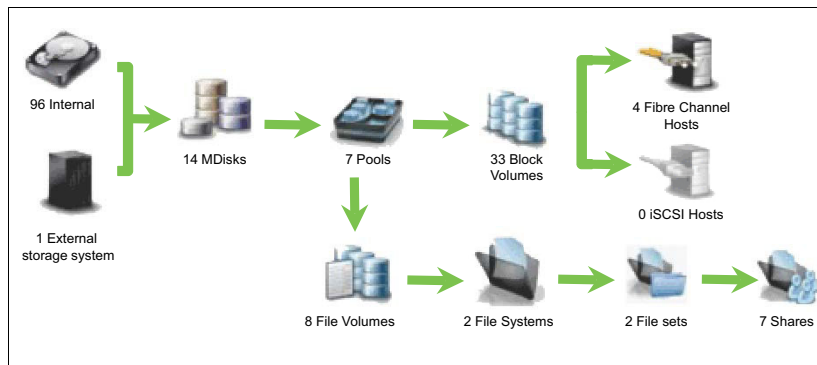


Figure 2-9 Storwize V7000 unified components


Storwize V7000 Unified NFS-based storage for the storage of VMware virtual disks and storage that is used by virtual machines includes the following features:

- ▶ NFS data stores do not have SCSI reservation performance issues. As a result, the use of large NFS data stores is much more practical. Preferred practices for block-based storage include minimizing data store size and the number of VMs per block data store where possible.
- ▶ Data sharing between multiple VMs or multiple operating systems is less complicated when shared through CIFS or NFS. This includes home directories for VDI, which are best shared through Active Directory authenticated CIFS share.
- ▶ Data that is shared through CIFS or NFS can scale gracefully without being bound by 2 TB VMware Virtual Machine Disk (VMDK) limitation. NAS shares are as scalable as your Storwize V7000 Unified (up to 720 TB per cluster).
- ▶ NAS share maximum capacity can be dynamically increased without requiring any client/vSphere side interaction or downtime.
- ▶ Storwize V7000 Unified uses 1 Gb or 10 Gb Ethernet, which is less expensive to implement and easier for most system administration professionals to use. Overall, 10 Gb Ethernet is faster than most today's 4 Gb and 8 Gb FCAL implementations.

### 2.5.3 IBM 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 2.5-inch drive bays and an LSI RAID controller. It connects to the compute node by using its PCIe expansion connector.

Figure 2-10 on page 32 shows a Storage Expansion Node that is attached to a x240 compute node.



*Figure 2-10 Storage Expansion Node that is attached to a x240 compute node*

The x240 Compute Node with the Storage Expansion Node can be used as an entry-level NAS-only or unified server storage in VDI deployments.

The following features are retained for VDI purposes:

- ▶ Support for 6 Gbps SAS and SATA drives (HDD and SSD)
- ▶ 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 Features on-Demand upgrades

## 2.5.4 IBM FlashSystem 820 and IBM FlashSystem 720

IBM FlashSystem™ storage systems deliver advanced performance, scalability, reliability, security, and energy-efficiency features. FlashSystem 720 and FlashSystem 820 storage systems are the appropriate choice for mission critical enterprise environments with the following characteristics:

- ▶ High storage performance requirements, such as low latency (microseconds as opposed to milliseconds)
- ▶ High bandwidth (gigabytes per second)
- ▶ High I/O operations per second (IOPS), hundreds of thousands

FlashSystem storage systems deliver over 500,000 read IOPS and up to 5 Gbps bandwidth with less than 100 microseconds latency, while they provide up to 24 TB of total usable capacity or up to 20 TB of 2D Flash RAID protected data storage in 1U of rack space.

Table 2-3 on page 33 lists the IOPS specifications.

Table 2-3 IOPS specification

	FlashSystem 720	FlashSystem 820
Write IOPS	400,000	280,000
Read IOPS	525,000	525,000

Based on enterprise multilevel cell (eMLC) flash, FlashSystem 820 is targeted to read-heavy workloads, where workload is distributed across multiple servers. Based on single-level cell flash, FlashSystem 720 is targeted to write-heavy enterprise workloads. It completes the Flex System infrastructure by providing the best performance solution for standard shared primary data storage devices, even compared to those that incorporate SSD or flash technology.

These storage options can be integrated with Flex System V7000 to be used as the top tier of storage alongside traditional arrays that are provided by the IBM Easy Tier functionality.

Figure 2-11 shows FlashSystem 720 and FlashSystem 820.



Figure 2-11 IBM FlashSystem 720 and FlashSystem 820

### 2.5.5 SSDs compared to HDDs

SSDs use non-volatile flash memory rather than spinning magnetic media to store data. The main advantage for VDI is the lower access times and latency rates that are 10 times faster than the spinning disks in an HDD. HDD is a proven technology with excellent reliability and performance, especially when the physical limitations of its spinning platters and moving arms are considered.

All of the Flex System compute nodes and the IBM Flex System V7000 support SSDs within the internal drive bay.

The use of SSDs works well in the following situations:

- ▶ When you want to provide the best performance for the non-persistent VDI hosts by installing two SSDs that are configured in as RAID-0
- ▶ When you want to implement Easy Tier function on the IBM Flex System V7000 to increase its IOPS performance on the most frequently accessed data

SSD technology includes the following types of cells:

- ▶ Single-level cell SSD  
Single-level cell flash memory stores data in arrays of floating-gate transistors, or *cells*, with 1 bit of data to each cell. This single bit per cell methodology results in faster transfer speeds, higher reliability, and lower power usage than that provided by HDDs. Single-level cell SSDs are two-to-three times more expensive to manufacture than multi-level cell devices.
- ▶ Multi-level cell SSD  
The basic difference between single-level cell flash memory and multi-level cell flash memory technologies is storage density. In comparison with single-level cell flash memory (which allows only two states to be stored in a cell, which stores only one bit of data per cell), multi-level cell flash memory can store up to four states per cell, which yields two bits of data that is stored per cell.

A comparison of IBM high-performance SSDs with traditional enterprise-level HDDs demonstrates a dramatic increase in overall I/O operations per second (IOPS), as shown in Table 2-4.

Table 2-4 IOPS comparison

IOPS	HDD (3.5-inch 15k)	HDD (2.5-inch 15k)	MLC SSD
Write IOPS	300	250	40,000
Read IOPS	390	300	60,000

### 2.5.6 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 require 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:

- ▶  $\text{Raw IOPS} = \text{Disk Speed IOPS} * \text{Number of disks}$
- ▶  $\text{Functional IOPS} = (\text{Raw IOPS} * \text{Write\%} / \text{RAID Penalty}) + (\text{RAW IOPS} * \text{Read\%})$

Table 2-5 provides the write penalty for RAID configuration.

*Table 2-5 RAID penalty*


RAID	Write Penalty
0	1
1	2
5	4
6	6
DP	2
10	2

On the storage part of the VDI infrastructure, RAID design depends on the following requirements that are gathered during assessments of users:

- ▶ For read-intensive workload, prefer RAID 0, 1, 5, 10 levels that spread read operations across multiple disk simultaneously. If the volume of data is important, you can also privilege a RAID level that optimize disk usability.
- ▶ For write-intensive workload, prefer a RAID level that offers a low write penalty, such as RAID 0 and 10.

## 2.6 Network considerations

The compute nodes are connected to I/O nodes through the I/O expansion adapters. Half-wide servers have two I/O expansion adapters, full-wide nodes have four adapters. Figure 2-12 shows the location of the adapters on a Flex System x240 compute node.



*Figure 2-12 I/O adapters slots in the IBM Flex System x240 compute node*



Each I/O expansion adapter is connected to switch bay by four links. Figure 2-13 shows the connections between the adapter in the compute nodes to the switch bays in the chassis.

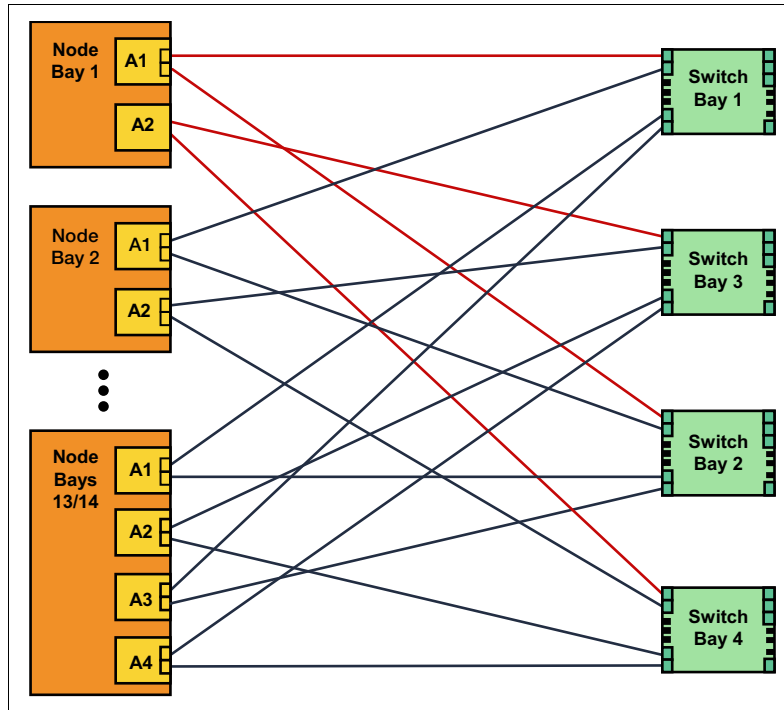


Figure 2-13 Logical layout of the interconnects between I/O adapters and I/O modules

## 2.6.1 IBM Flex System 10GbE network switches

The following switches can be used in a VDI environment:

- ▶ EN4093 and EN4093R 10Gb Scalable Switches
- ▶ CN4093 10Gb Converged Scalable Switch
- ▶ SI4093 System Interconnect Module
- ▶ EN4091 10Gb Ethernet Pass-thru Module

### EN4093 and EN4093R 10Gb Scalable Switches

The Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches provide unmatched scalability and performance, which address various networking concerns today and provide capabilities that help you prepare for the future. These switches can support up to 64 10 Gb Ethernet connections while offering Layer 2/3 switching.

These switches can help clients migrate to a 10 Gb or 40 Gb Ethernet infrastructure and offer virtualization features, such as Virtual Fabric and IBM VMready® and the ability to work with IBM Distributed Virtual Switch 5000V.

Figure 2-14 shows the IBM Flex System Fabric EN4093.



Figure 2-14 Flex System Fabric EN4093

The EN4093 and EN4093R switches initially are licensed for 14 10 Gb internal ports enabled and 10 10 Gb external uplink ports enabled. You can enable further ports when needed by purchasing more licenses, including 14 internal ports and two 40 Gb external uplink ports with Upgrade 1 and 14 other internal ports and four more SFP+ 10 Gb external ports.

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

- ▶ Integrated network management  
EN4093 and EN4093R 10 Gb Scalable Switches are tightly integrated and managed through the IBM Flex System Manager.
- ▶ Optimized network virtualization with virtual NICs  
IBM Virtual Fabric provides a way for companies to carve up 10 Gb ports into virtual NICs.
- ▶ Increased performance  
The EN4093 and EN4093R are the embedded 10 GbE switches for a server chassis to support aggregated throughput of 1.28 Tbps, while also delivering full line rate performance. These switches are ideal for managing dynamic workloads across the network. They also provide a rich Layer 2 and Layer 3 feature set and offer industry-leading uplink bandwidth by being the first integrated switches to support 40 Gb uplinks.

**Note:** Internal layer 2 switches provide a more effective approach for communication between co-resident server by using an east-west approach. Communication between nodes use an internal, active layer 2 switch to pass traffic to one other. By containing network traffic within the Flex System chassis, latency is improved by 50%, compared to a north-south approach. In a north-south approach, all of the traffic is routed to the top-of-rack switch; the flow goes up to the top-of-rack switch and down to co-located server.

► VM-aware networking

IBM System Networking's Distributed Virtual Switch 5000V (which is sold separately) enables network administrators to simplify management by having a consistent virtual and physical networking environment. The 5000V virtual and physical switches use the same configurations, policies, and management tools. Network policies migrate automatically with VMs to ensure that security, performance, and access remain intact as VMs move from server to server.

### CN4093 10Gb Converged Scalable Switch

The Flex System Fabric CN4093 10Gb Converged Scalable Switch provides 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 converged integrated solution. It is designed to install within the I/O module bays of the Flex System Enterprise Chassis. The switch can help you migrate to a 10 Gb or 40 Gb converged Ethernet infrastructure. It offers virtualization features, such as Virtual Fabric and VMready, plus the ability to work with IBM Distributed Virtual Switch 5000V.

Figure 2-15 shows the Flex System Fabric EN4093.



Figure 2-15 Flex System CN4093 Converged Switch

The CN4093's 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. You then have 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 4 x 10 GbE SFP+ DAC uplinks on each QSFP+ port) when you need them by using IBM Features on Demand licensing capabilities that provide “pay as you grow” scalability without the need for more hardware.

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

- ▶ Integrated network management  
The CN4093R 10Gb Scalable Switch is tightly integrated and managed through the Flex System Manager.
- ▶ Optimized network virtualization with virtual NICs  
IBM 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 10 Gb Virtual Fabric Adapters in each compute node.
- ▶ 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. Furthermore, it 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 IBM 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.

### **SI4093 System Interconnect Module**

The IBM Flex System Fabric SI4093 System Interconnect Module enables simplified integration of Flex System into your existing networking infrastructure. This module requires no management for most data center environments. It eliminates the need to configure each networking device or individual ports, which reduces the number of management points. It provides a low latency, loop-free interface that does not rely upon spanning tree protocols and 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.

Figure 2-16 shows the SI4093 System Interconnect Module.



*Figure 2-16 IBM Flex System Fabric SI4093 System Interconnect Module*

The SI4093 System Interconnect Module is initially licensed for 14 10 Gb internal ports enabled and 10 10 Gb external uplink ports enabled. You can enable further ports, including 14 more internal ports and two 40 Gb external uplink ports by using IBM Features on-Demand 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, and it enables server to server communication for optimum performance (for example, vMotion).
- ▶ **Optimized network virtualization with virtual NICs**  
IBM 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**  
Flex System SI4093 simplifies management and automates VM mobility by making the network VM aware with IBM VMready, which works with all the major hypervisors. Network policies migrate automatically along with virtual machines (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. Furthermore, it offers industry-leading uplink bandwidth by being the integrated switch to support 40 Gb uplinks.

### 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 top-of-rack data center switch. This module can function at both 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.

Figure 2-17 shows the Flex System EN4091 10Gb Ethernet Pass-thru Module.



Figure 2-17 Flex System EN4091 10Gb Ethernet Pass-thru Module

The Flex System EN4091 offers the following key features:

- Offers intelligent workload deployment and management for maximum business agility.
- Delivers high-speed performance complete with integrated servers, storage, and networking.
- The flexible design meets the needs of varying workloads with independently scalable IT resource pools for higher usage and lower cost per workload.

## 2.6.2 Network adapters

The following network adapters are described:

- LAN-on-motherboard
- IBM Flex system CN4054 10Gb Virtual Fabric Adapter

## LAN-on-motherboard

Some models of the Flex System x240 compute node have an Ethernet LAN-on-motherboard controller that is integrated on the system board. The LAN-on-motherboard is installed on the I/O expansion adapter 1 (A1) of the compute node.

The I/O expansion adapter A1 routes to two switch bays for redundancy and performance. The first port is linked to the I/O module 1 within the chassis. The second port is connected to I/O module 2.

**Installation note:** With LAN-on-motherboard enabled, the Ethernet I/O module can be installed only on bays 1 and 2 on the chassis. Integrated NICs ports are routed to these bays with a specialized periscope connector.

LAN-on-motherboard offers the following operational mode choices:

- ▶ One-port physical NIC mode (pNIC), multichannel disabled, which is the default

In this mode, the adapter operates as a standard dual-port 10 Gbps Ethernet adapter, and it functions with any 10 GbE switch.

- ▶ Virtual NIC mode (vNIC), multichannel enabled

This mode enables up to four virtual NIC interfaces per 10 Gb physical port (eight total for the LAN-on-motherboard). The adapter works with any 10 Gb Ethernet switch.

You can also use the following vNIC linking options:

- IBM Virtual Fabric mode works with IBM Flex System EN4093, EN4093R, and CN4093 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 is performed by both 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 offers the same capabilities as IBM Virtual Fabric Mode in terms of the number of vNICs and the bandwidth each can be configured to have. Switch Independent Mode extends the existing VLANs to the virtual NIC interfaces.

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

The IEEE 802.1Q VLAN tag is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch. The VLAN tags are added to the packet by the applications or drivers at each endstation rather than by the switch.

Consider configuring the LAN-on-motherboard to the vNIC mode with Switch Independent Mode to distribute the 10 GbE network bandwidth differently to the VLANs that are used within the VDI infrastructure.

### **IBM Flex system CN4054 10Gb Virtual Fabric Adapter**

The Flex System CN4054 10Gb Virtual Fabric Adapter is a 4-port, 10 Gb converged network adapter that can scale to up to 16 virtual ports and that supports Ethernet, iSCSI, and FCoE. Because this adapter supports up to 16 virtual NICs, where each physical 10 Gb port can be divided into four virtual ports, you can see benefits in bandwidth flexibility, virtualization specially for HA and VMotion operations and cost.



Figure 2-18 shows CN4054 connectivity to the switches.




Figure 2-18 CN4054 connectivity to the switches.

**Upgrade note:** You can upgrade the Flex System CN4054 10Gb Virtual Fabric Adapter to run storage protocols iSCSI and FCoE by applying the upgrade license.

The Flex System CN4054 10Gb Virtual Fabric Adapter offers the following modes of operation:

- One-port pNIC mode, multichannel disabled, which is the default  
In this mode, the adapter operates as a standard quad-port 10 Gbps or 1 Gbps 4-port Ethernet adapter, and it functions with any 10 GbE switch.
- vNIC mode, multichannel enabled  
vNIC mode enables up to four virtual NIC interfaces per 10 Gb physical port (16 total for the CN4054). The adapter works with any 10 Gb Ethernet switch.

Consider the following points:

- IBM Virtual Fabric mode works with IBM Flex System EN4093, EN4093R, and CN4093 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 is 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 offers the same capabilities as IBM Virtual Fabric Mode in terms of the number of vNICs and the bandwidth each can be configured to have. Switch Independent Mode extends the existing VLANs to the virtual NIC interfaces.

vNIC bandwidth allocation and metering is only performed by the adapter. In such a case, a unidirectional virtual channel is established where 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 available physical port bandwidth, as there is no metering that is performed on a switch side.

The IEEE 802.1Q VLAN tag is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch. The VLAN tags are added to the packet by the applications or drivers at each endstation rather than by the switch.

Consider configuring the CN4054 Adapter in vNIC mode with Switch Independent Mode to distribute the 10 GbE network bandwidth differently to the VLANs that are used within the VDI infrastructure.

## 2.7 Flex System Fibre Channel switches

The following Fibre Channel switches are described next:

- ▶ FC5022 16Gb SAN Scalable Switch
- ▶ FC3171 8Gb SAN Switch and Pass-thru

### 2.7.1 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. This 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 who are deploying storage area networks in their enterprise. The switches 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 by using an integrated management appliance. Alternatively, if you are using end-to-end IBM B-type SAN, you can use IBM management tools.

Figure 2-19 shows the Flex System FC5022 16Gb Scalable Switch.



Figure 2-19 Flex System FC5022 16Gb Scalable Switch

**Installation note:** On a compute node where LAN-on-motherboard is activated, the Fibre Channel adapter is installed on the I/O expansion adapter 2 (A2) of the compute node. Flex System FC5022 can be installed only on switch bay 3 and 4.

## 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, and offers scalability in external SAN size and complexity, and enhanced systems management capabilities. The IBM Flex System FC3171 8Gb Pass-thru supports a fully interoperable solution for seamless integration of the Fibre Channel initiators to an existing fabric. The pass-thru module uses industry-standard N\_Port ID virtualization (NPIV) technology to provide a cost-effective connectivity solution for the IBM Flex System chassis.

Figure 2-20 shows the Flex System FC3171 8Gb SAN Switch.



Figure 2-20 Flex System FC3171 8Gb SAN Switch

**Installation note:** On compute nodes where LAN-on-motherboard is activated, the Fibre Channel adapter is installed on the I/O expansion adapter 2 of the compute node. SAN switch can be installed only on switch bay 3 and 4.

### 2.7.2 Fibre Channel adapters

In this section, we describe the following Fibre Channel adapters:

- ▶ FC3172 2-port and FC3052 2-port 8Gb FC adapters
- ▶ FC5022 2-port and FC5054 4-port 16Gb 16Gb FC Adapters
- ▶ IBM Flex System FC5024D 4-port 16Gb FC Adapter
- ▶ IBM Flex System FC5172 2-port 16Gb FC Adapter

### **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 N\_Port ID Virtualization (NPIV) and virtual fabric.

### **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 and 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. In addition, it provides performance head room to support demanding SSD storage technologies.

The Flex System 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.
- ▶ Over 500,000 IOPS per port, which maximizes transaction performance and density of VMs per compute node.
- ▶ NPIV allows multiple host initiator N\_Ports to share a single physical N\_Port, which dramatically reduces SAN hardware requirements.
- ▶ Uses 16 Gbps bandwidth to eliminate internal oversubscription.

### **IBM 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 is based on the Brocade architecture and 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.
- ▶ Over 500,000 IOPS per port, which maximizes transaction performance and density of VMs per compute node.
- ▶ NPIV allows multiple host initiator N\_Ports to share a single physical N\_Port, which dramatically reduces SAN hardware requirements.
- ▶ Uses 16 Gbps bandwidth to eliminate internal oversubscription.
- ▶ Delivers considerable value by simplifying the deployment of server and SAN resources, which reduces infrastructure and operational costs.

### **IBM Flex System FC5172 2-port 16Gb FC Adapter**

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

## **2.8 IBM Flex System Manager functions and considerations**

Flex System Manager is a systems management appliance that drives efficiency and cost savings in the data center. It provides a pre-integrated and virtualized management environment for servers, storage, and networking that is managed easily from a single interface. Flex System Manager provides a focal point for seamless multichassis management that gives an instant and resource-oriented view of chassis and chassis resources for IBM System x and IBM Power Systems compute nodes.

Flex System Manager provides the following advantages:

- ▶ Reduces the number of interfaces, steps, and clicks 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.

Figure 2-21 shows the Flex System Manager management appliance.



Figure 2-21 The Flex System Manager management appliance

Flex System Manager is designed to help you get the most out of your IBM PureFlex System while automating repetitive tasks. Flex System Manager can reduce the number of manual navigational steps for typical management tasks. Flex System Manager provides core management functions with automation so you can focus your efforts on business innovation. These functions include simplified system setup procedures with wizards and built-in expertise to consolidate monitoring for all of your physical and virtual resources (compute, storage, and networking).

Flex System Manager has the following key features:

- ▶ Optimizing your workload management through built-in expertise
- ▶ Managing all of your resources with one solution: Compute, storage, networking, virtualization

Flex System Manager base feature set 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)
- ▶ Virtual machine lifecycle management (VMControl Express)
- ▶ I/O address management (IBM Fabric Manager)

The Flex System Manager advanced feature set upgrade offers the following advanced features:

- ▶ Image management (VMControl Standard)
- ▶ Pool management (VMControl Enterprise)

The Flex System Manager Node has the following fixed hardware specifications:

- ▶ One Intel Xeon processor E5-2650 8C 2.0 GHz 20 MB Cache 1600 MHz 95 W 32 GB of memory with eight 4 GB (1x4 GB, 1Rx4, 1.35 V) PC3L-10600 CL9 ECC DDR3 1333 MHz LP RDIMMs
- ▶ Integrated LSI SAS2004 RAID controller
- ▶ Two IBM 200 GB SATA 1.8-inch MLC SSD configured in a RAID 1
- ▶ One IBM 1 TB 7.2 K 6 Gbps NL SATA 2.5-inch SFF HS HDD
- ▶ Dual-port 10 Gb Ethernet Emulex BladeEngine 3 (BE3) network controller for data network connections
- ▶ Dual-port Broadcom 5718-based network adapter with integrated Broadcom 5389 8-port basic L2 switch for internal chassis management network connections
- ▶ Integrated Management Module II (IMM2)

Flex System Manager includes the following functions and features:

- ▶ Management network
- ▶ Chassis Management Module
- ▶ Integrated Management Module II
- ▶ Configuration Patterns
- ▶ Storage connectivity selection guidance

These functions and features are described next.

## 2.8.1 Management network

The management network is a private and secure Gigabit Ethernet network. It is used to complete management-related functions throughout the chassis, including management tasks that are related to the compute nodes, switches, and the chassis.

The management network is shown in Figure 2-22 on page 53 as the blue line. It connects the Chassis Management Module (CMM) to the compute nodes, the switches in the I/O bays, and the Flex System Manager.



The Flex System Manager connection to the management network is through a special Broadcom 5718-based management network adapter (Eth0). The management networks in multiple chassis can be connected through the external ports of the CMMs in each chassis by using a GbE top-of-rack switch.

The yellow line that is shown in Figure 2-22 indicates the production data network. Flex System Manager also connects to the production network (Eth1) so that it can access the Internet for product updates and other related information.

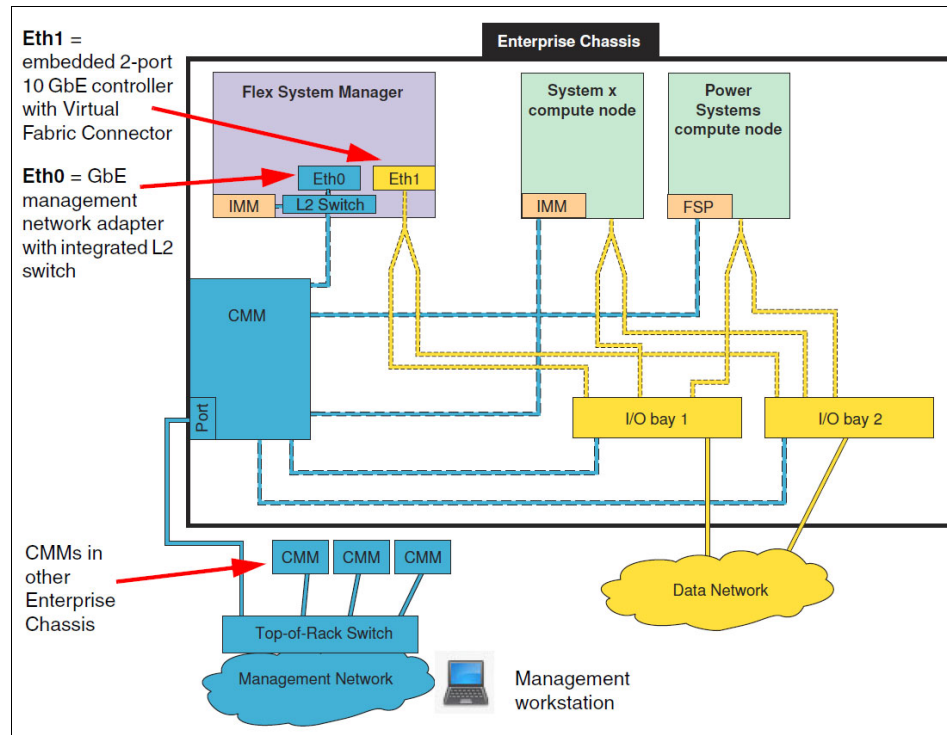


Figure 2-22 Separate management and production networks

One of the key functions that the data network supports is discovery of operating systems on the various network endpoints. Discovery of operating systems by Flex System Manager is required to support software updates on an endpoint, such as a compute node. Flex System Manager Checking and Updating Compute Nodes wizard assists you in discovering operating systems as part of the initial setup.

## 2.8.2 Chassis Management Module

The Chassis Management Module (CMM) provides single-chassis management and 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 itself, and take control of the system without any disruption. The CMM is central of 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 functions:

- ▶ 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 multi-chassis monitoring.
- ▶ Set power policies and view power consumption history for chassis components.

## 2.8.3 Integrated Management Module II

Integrated Management Module II (IMM2) is the next generation of the integrated service processors for the IBM x86-based server family. IMM2 enhancements include a more responsive user interface, faster power on, and increased remote presence performance. IMM2 incorporates a new web-based user interface that provides a common interface across all IBM System x software products.

IMM2 provides the following major features as standard:

- ▶ IPMI v2.0-compliance
- ▶ Remote configuration of IMM2 and 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 features on-demand enablement of server functions, option card features, and System x solutions and applications

## 2.8.4 Configuration Patterns

By using Configuration Patterns, you can provision or pre-provision multiple systems from a single pattern. Then, subsequent pattern changes are applied automatically to all associated systems.

Configuration Patterns also integrate support for IBM Fabric Manager so that you can virtualize server fabric connections and so that you can fail over or repurpose servers without disruption to the fabric. In addition, you can start fabric change requests through your change management process before your hardware arrives by preconfiguring host interconnect addresses.

By using Server Configuration Patterns, you can configure storage, I/O adapter, boot order, and other Integrated Management Module (IMM) and Extensible Firmware interface (UEFI settings).

Chassis Configuration Patterns allow you to configure CMM network management interface, users and security, power and acoustic settings, and basic I/O module and node IP address assignments.

Consider using these patterns to configure your Flex System infrastructure easily and quickly.

## 2.8.5 Storage connectivity selection guidance

vSphere supports many protocols (including Fibre Channel, iSCSI, Fibre Channel over Ethernet, and network-attached storage), with no preference given to any one protocol over another. However, many customers still want to know how these protocols compare to each other and to understand their respective pros and cons.

Because of the deployment and management differences within each protocol, determining which of these protocols to use is one of the key steps in designing a virtualized infrastructure. Knowing how each protocol performs in terms of throughput and CPU usage can be helpful in deciding about this important design consideration.

Recommending one or another protocol to use with VMware is challenging and depends on many factors, including the following factors:

- ▶ Customer requirements
- ▶ Customer bandwidth/performance expectations
- ▶ Existing infrastructure
- ▶ Implementation skills

In this section, we describe the protocols with corresponding Flex System components and provide guidance about selecting the correct protocol to meet your requirements that can be adopted easily in your existing infrastructure.

### Fibre Channel

Fibre Channel (FC) presents block devices that are similar to iSCSI. I/O operations are carried out over a network by using a block access protocol. In FC, remote blocks are accessed by encapsulating SCSI commands and data into FC frames. FC is commonly deployed in most mission-critical environments.

FC is implemented as a lossless network, which can run on a dedicated 1 Gb, 2 Gb, 4 Gb, 8 Gb, and 16 Gb, host bus adapter (HBAs) (typically two for redundancy and multipathing) but there is no support for full, end-to-end 16 Gb connectivity from host to array. To get full bandwidth, a number of 8 Gb connections can be created from the switch to the storage array.

Choosing FC protocol includes the following advantages:

- ▶ FC protocol typically affects a host's CPU the least because HBAs manage most of the processing (encapsulation of SCSI data into FC frames).
- ▶ Supports load balancing by distributing load across multiple paths to an FC target.
- ▶ Features a well-known and well-understood protocol. Also, it is mature and trusted and found most mission-critical environments.

FC also includes the following disadvantages:

- ▶ Requires dedicated HBA, FC switch, and FC-capable storage array, which makes an FC implementation more expensive. More management overhead (for example, switch zoning) is required.
- ▶ The configuration involves zoning at the FC switch level and LUN masking at the array level after the zoning is complete. It is more complex to configure than IP storage.
- ▶ Still runs only at 8 Gb, which is slower than other networks (16 Gb throttled to run at 8 Gb).

The following required Flex system components that support FC connectivity are available:

- ▶ IBM Flex system compute nodes, which are described in “IBM Flex System Compute Nodes” on page 20
- ▶ Flex System Fibre Channel switches, which are described in “Flex System Fibre Channel switches” on page 47
- ▶ Fibre Channel Adapter, which is described in “Fibre Channel adapters” on page 48
- ▶ IBM Flex System V7000, which is described in “IBM Flex System V7000” on page 27

## **Fibre Channel over Ethernet**

Fibre Channel over Ethernet (FCoE) also presents block devices, with I/O operations carried out over a network by using a block access protocol. In this protocol, SCSI commands and data are encapsulated into Ethernet frames.

Figure 2-23 shows compute node connectivity to the IBM Flex System V7000 Storage Node that uses an IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch, which provides FCF and DCB functionality.

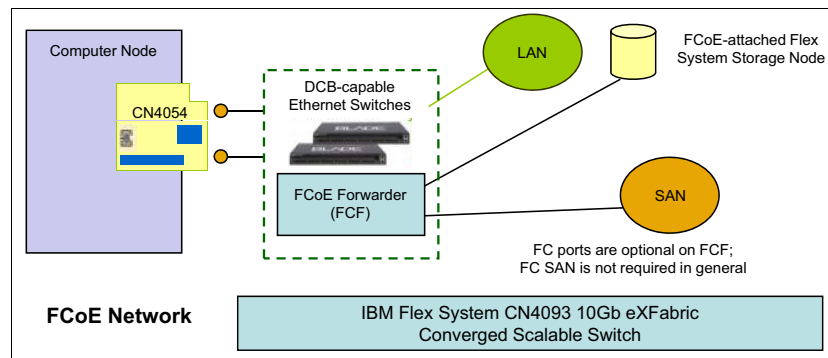


Figure 2-23 Compute node with CNA and FCF I/O module switch

FCoE has many of the same characteristics as FC, except that the transport is Ethernet. FCoE is implemented as a lossless network, which requires converged network adapter (CNA) or Network adapter with FCoE capabilities that uses software FCoE initiator. In addition, this protocol requires 10 Gb Ethernet. (FCoE is SCSI over Ethernet, not IP.) This protocol also requires jumbo frames because FC payloads are 2.2 KB and cannot be fragmented.

FCoE does have a number of benefits, especially when virtualized server space is considered. FCoE has the following advantages:

- ▶ The ability to unify I/O through host-linked CNAs and multiprotocol switches connections or ports and switch ports, which save on power and cooling through reduced cabling.
- ▶ The ability to hold onto the existing Fibre Channel storage and backup targets, which protects the existing investment.
- ▶ The ability to use advanced Ethernet networking QoS and other management practices inside the FC space.
- ▶ FCoE unifies I/O through host-linked CNAs and multiprotocol switches, which enables IT to considerably lower the total network devices and cables for interconnecting the clusters.
- ▶ Reduction in network management overhead.
- ▶ With FCoE, there is no IP encapsulation of the data as there is with NFS and iSCSI, which reduces some of the overhead or latency.
- ▶ Supports load balancing by distributing load across multiple paths to an FCoE target.

FCoE has the following disadvantages:

- ▶ FCoE configuration involves zoning at the FCoE switch level and LUN masking at the array level and it is more complex than IP-based storage.
- ▶ Requires a 10 Gb lossless network infrastructure, which can be expensive.
- ▶ Might be complex to troubleshoot or isolate issues with network and storage traffic using the same pipe.

The following Flex system components support FCoE connectivity:

- ▶ IBM Flex system compute nodes, which are described in “IBM Flex System Compute Nodes” on page 20.
- ▶ LAN-on-motherboard, which is described in “LAN-on-motherboard” on page 43 with the IBM Virtual Fabric Software Upgrade option to enable FCoE feature.
- ▶ IBM Flex system CN4054 Virtual Fabric Adapter, which is described in “IBM Flex system CN4054 10Gb Virtual Fabric Adapter” on page 44, with the IBM Flex System Virtual Fabric Adapter-SW Upgrade option to enable FCoE feature.

- ▶ IBM Flex System Fabric CN4093 Converged 10Gb Scalable Switches, which is described in “CN4093 10Gb Converged Scalable Switch” on page 39.

This switch offers FCoE Full Fabric and Fibre Channel NPV Gateway operations to deliver a truly converged integrated solution.

- ▶ IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches, which is described in “EN4093 and EN4093R 10Gb Scalable Switches” on page 37.

For FCoE implementations, the EN4093R acts as a transit switch forwarding FCoE traffic upstream to other devices, such as the IBM RackSwitch™ G8264CS, Brocade VDX, or Cisco Nexus 5548/5596 where the FC traffic is broken out.

- ▶ IBM Flex System Fabric SI4093 System Interconnect Module, which is described in “SI4093 System Interconnect Module” on page 40.

For FCoE implementations, the SI4093 acts as a transit switch forwarding FCoE traffic upstream to other devices, such as the IBM RackSwitch G8264CS, Brocade VDX, or Cisco Nexus 5548/5596 where the FC traffic is broken out.

- ▶ IBM Flex System V7000, which is described in “IBM Flex System V7000” on page 27.

## Internet Small Computer System Interface

Internet Small Computer System Interface (iSCSI) is a protocol that uses TCP to transport SCSI commands for a storage network, which enables existing TCP/IP infrastructure to be used as a SAN. iSCSI presents block devices to a VMware ESXi host. Rather than accessing blocks from a local disk, I/O operations are carried out over a network by using a block access protocol.

In the case of iSCSI, remote blocks are accessed by encapsulating SCSI commands and data into TCP/IP packets. You can mount block devices (disks) across an IP network to your local system, then use them as you do any other block device.

iSCSI can run over a 1 Gb or a 10 Gb TCP/IP network. Multiple connections can be multiplexed into a single session, which is established between the initiator and target. VMware supports jumbo frames for iSCSI traffic, which can improve performance.

iSCSI provides the following advantages:

- ▶ You can use existing networking hardware components and iSCSI driver from VMware, so it is inexpensive to implement.
- ▶ This protocol is well-known and well-understood, thus it is easy to implement.
- ▶ iSCSI supports authentication (CHAP) and encryption for security and multipathing for increased throughput and reliability.
- ▶ No special training and skills are needed to implement and manage iSCSI.
- ▶ Speed and performance is greatly increased with 10 Gbps Ethernet.
- ▶ Software initiators can be used for ease of use and lower cost.

iSCSI has the following disadvantages:

- ▶ Network latency and non-iSCSI network traffic can diminish performance.
- ▶ When an iSCSI path is overloaded, the TCP/IP protocol drops packets and requires them to be resent, which can cause latency.
- ▶ When a network path that is carrying iSCSI traffic is oversubscribed, the performance degrades because dropped packets must be resent.
- ▶ Possible security issues can occur because there is no built-in encryption to isolate traffic.
- ▶ Software iSCSI can cause more CPU overhead on the ESX host.



iSCSI supports the following Flex system components:

- ▶ IBM Flex system compute nodes, which are described in “IBM Flex System Compute Nodes” on page 20.
- ▶ LAN-on-motherboard, which is described in “LAN-on-motherboard” on page 43, with the IBM Virtual Fabric Software Upgrade option to enable iSCSI feature.
- ▶ IBM Flex system CN4054 Virtual Fabric Adapter, which is described in “IBM Flex system CN4054 10Gb Virtual Fabric Adapter” on page 44, with the IBM Flex System Virtual Fabric Adapter-SW Upgrade option to enable iSCSI feature.
- ▶ IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches, which are described in “EN4093 and EN4093R 10Gb Scalable Switches” on page 37.
- ▶ IBM Flex System Fabric SI4093 System Interconnect Module, which is described in “SI4093 System Interconnect Module” on page 40.
- ▶ IBM Flex System V7000, which is described in “IBM Flex System V7000” on page 27.

## Network-attached storage

Network-attached storage (NAS) uses a file-sharing protocol to communicate with the storage device that maintains the disk file system. NAS offloads the storage device functions that are responsible for writing data to the drives from the host server to the storage device. NAS encompasses the NFS and CIFS protocols and refers specifically to the use of file-based storage to store virtual guests. VMware ESXi supports only NFS for file-level access.

**Documentation note:** Because VMware supports only NFS, this documentation covers only NFS.

NFS presents file devices over a network to an ESXi host for mounting. The NFS array makes its local file systems available to ESXi hosts. ESXi hosts access the metadata and files on the NFS array by using an RPC-based protocol.

NFS can run over 1 Gb or 10 Gb TCP/IP networks. NFS also supports UDP, but the VMware implementation requires TCP.

Figure 2-24 shows files sharing with NFS on IBM Storwize V7000 Unified.

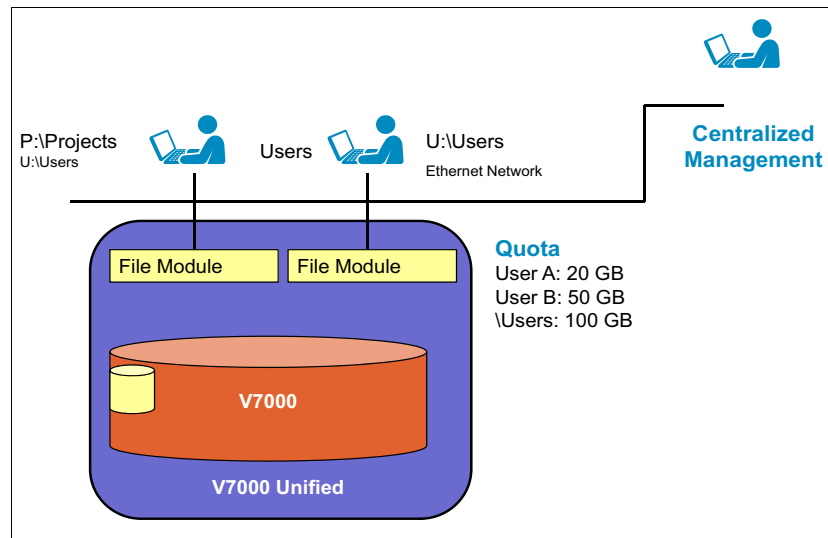


Figure 2-24 NFS on IBM Storwize V7000 Unified

NFS offers the following advantages:

- ▶ VMs are stored in directories on NFS shares, which makes them easy to access without the use of the hypervisor. This method is useful for taking VM backups, snapshots, or cloning an individual virtual guest. VMware configuration files also can be directly created or edited.
- ▶ Network shares can be expanded dynamically (if the storage filer supports it) without any affecting the ESXi.
- ▶ No extra server hardware is required to access NFS shares, which can be achieved over standard NICs.
- ▶ File locking and queuing are handled by file system, which can result in better performance where locking and queuing are handled by the host server.
- ▶ Virtual storage can easily be shared among multiple virtual servers. VMware uses a locking file on the share to ensure integrity in a clustered environment.
- ▶ Virtual guests can be thinly provisioned if the underlying storage hardware supports it.
- ▶ VMFS LUNs top out at approximately 2 TB in size, but NFS has no such limits (some arrays go as high as 16 TB).

NFS has the following disadvantages:

- ▶ Scalability is limited to eight NFS shares per VMware host, which can be expanded to 64 but requires the TCP/IP heap size to be increased.
- ▶ NFS does not support multipathing; therefore, high availability must be managed at the physical network layer with bonded networks on ESXi and virtual interfaces on the storage array.
- ▶ Although NFS shares can scale to the maximum size that is permitted by the storage filer, the share often is created from one group of disks with one performance characteristic; therefore, all guests on the share experience the same I/O performance profile.
- ▶ There is CPU overhead because the hypervisor must use a software client to communicate with the file system.
- ▶ Possible security issues exist because there is no built-in encryption.

NAS supports the following Flex system components:

- ▶ IBM Flex system compute nodes, which are described in “IBM Flex System Compute Nodes” on page 20.
- ▶ LAN-on-motherboard, which is described in “LAN-on-motherboard” on page 43
- ▶ IBM Flex System Fabric EN4093 and EN4093R 10Gb Scalable Switches, which are described in “EN4093 and EN4093R 10Gb Scalable Switches” on page 37.
- ▶ IBM Flex System Fabric SI4093 System Interconnect Module, which is described in “SI4093 System Interconnect Module” on page 40.
- ▶ IBM Storwize V7000 Unified system, which is described in “IBM Storwize V7000 Unified System” on page 29.

There are many factors to consider when you are choosing a storage device for your virtual environment; however, decisions ultimately come down to simple factors, such as budget, performance, and capacity. Among the many decisions IT managers face when they are deploying server virtualization is what protocol to use. For example, should you use block protocol, such as Fibre Channel, FCoE and iSCSI, or file-sharing protocol, such as NFS?

Block protocol is proven to work well in virtualized environments. Although it is highly reliable, provides excellent performance, and can scale to meet any capacity requirement, it can require more hardware. Directly accessing data is an issue for iSCSI/Fibre Channel/FCoE, which makes data cloning and backup more complex. Based on VMware, most VMware deployments rely on block-based protocol. However, file-sharing protocol, such as NFS, is an affordable alternative with many features, including ease of management and more flexible snapshot and replication capabilities.

NFS provides better out-of-band access to guest files without the need to use the hypervisor, large data stores, and cost-saving features, such as data deduplication.

Finally, iSCSI and NFS can use existing network infrastructures, can require less hardware, and are easy to implement, which can be an eligible protocol where the cost is the major factor. Alternatively, Fibre Channel and FCoE are qualified for environments where superior performance, reliability, and higher throughput are required.



## VMware vSphere design considerations

VMware vSphere is used for hosting, configuring, provisioning, and managing virtual machines and is a fundamental part of the Horizon View implementation, which provides the virtualization platform that is on top of the physical hardware. The virtual machines can be used as sources for virtual desktop pools and to host vSphere and Horizon View infrastructure components, vCenter Server, Active Directory, and Connection Servers.

This chapter presents some important design considerations for the use of VMware vSphere 5.1 infrastructure on IBM Flex System hardware as a part of VMware Horizon View deployment. We describe the five main layers of a vSphere infrastructure: datacenter management, compute servers, network, storage, and virtual machines.

This chapter includes the following topics:

- ▶ Compute servers layer
- ▶ Networking considerations
- ▶ Storage considerations

## 3.1 Compute servers layer

In this section, we describe the compute servers layer design of a vSphere infrastructure, which includes the vSphere Hypervisor, vSphere clusters, and cluster features.

### 3.1.1 ESXi hypervisor

The hypervisor in vSphere is ESXi. ESXi is installed directly on the compute nodes and provides a virtualization layer that abstracts the necessary processor, memory, storage, and networking resources. It also provides these resources to the virtual machines (VMs).

ESXi has a small disk footprint of less than 150 MB. This allows it to be on internal flash memory, such as a USB flash drive that is plugged into the motherboard of Flex System compute nodes. The IBM-customized version of ESXi is preinstalled on the flash drive and provides more drivers and CIM modules specific to IBM hardware.

The following design considerations are important when you are using the ESXi hypervisor:

- ▶ Selecting the server model:
  - Ensure that the server model and CPU are listed in the VMware HCL.
  - Consider the use of Flex System PCIe with a NVIDIA graphics card for environments where the virtual desktop users run graphics-intensive applications.
- ▶ Consider whether Lockdown mode should be enabled for a higher level of security. Remember that if Lockdown is enabled, only vCenter can authenticate remotely to the ESXi host.
- ▶ Use the latest stable version of ESXi that is compatible with the products that are used in the solution.
- ▶ Select hosts with high CPU core count per CPU socket to minimize VMware licensing costs.
- ▶ Use fewer, larger hosts in big environments, and more, smaller hosts in small environments. For Flex System, the server models that are recommended for desktop virtualization include x222, x240, and x440.
- ▶ Do not use memory overcommitting or, if you must use it, do so only for non-critical environments, as recommended by VMware. If the ESXi host does not run into memory contention issues, ballooning or memory compression and swapping does not occur.

- ESXi hosts should have fully redundant hardware components, including redundant network cards, host bus adapters (HBAs) for SAN access, and power supplies. Flex System compute nodes match perfectly with these requirements.

### 3.1.2 VMware vCenter Server

VMware vCenter Server is a mandatory component of a VMware View VDI that provides a centralized and extensible platform for managing a virtual infrastructure. Many advanced features, such as HA, DRS, vMotion, and dvSwitches are available only through vCenter Server.

vCenter Server is also a required component in a Horizon View environment due to its central role of managing all communication between View and vSphere. Each VMware cluster relies on vCenter to perform cluster management and other hosting infrastructure tasks; therefore, the delivery of desktops can be affected if vCenter becomes slow or unresponsive under high stress conditions, such as in a large View environment with many users logging on at the same time each morning or when workers' shifts change.

VMware recommends the use of vCenter as a VM, which allows for protection of vCenter with various high availability features and policies. VMware also recommends achieving high availability for the vCenter Server for View deployments due to the important role vCenter performs in provisioning virtual desktops.

Starting with version 5.1, the vCenter architecture changed. Some components, such as inventory services, were decoupled, and other new components were introduced, such as single sign-on, which can be installed on separate servers.

In addition to the classic vSphere client, VMware introduced a new web client. All operations of the classic client can be done by using the web client, and some web client-only operations were introduced.

The new vCenter architecture allows more flexibility in sizing and designing your system, but sometimes introduces more complexity and can require more compute resources than the previous version. For example, if you place all vCenter components on a single server, a minimum of 10 GB of RAM must be allocated.

### 3.1.3 vMotion and Storage vMotion

Live migration or vMotion technology allows you to move running VMs from one physical server to another with no downtime. This enables companies to perform hardware maintenance without disrupting business operations.

vMotion relies on the following mechanisms:

- ▶ Encapsulation of VM state in a file that is stored on shared storage.
- ▶ Transfer of the active memory of a VM over the network.
- ▶ Use of a virtualized network by the VM, which ensures 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 technology enables moving VM disks from one physical storage location to another without any 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 needs to expand its storage and there is not enough available space in the current physical container.

Before vSphere 5.1, vMotion required shared storage between hosts, and Storage vMotion required a host to have access to the source and destination datastores. vSphere 5.1 removes these requirements and allows combining vMotion and Storage vMotion into one process. This combined migration process 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 on page 69).

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

The following vMotion-specific design considerations are important:

- ▶ VMs should use virtual hardware version 9.
- ▶ Separate the vMotion network from the management and VM networks. Remember that vMotion traffic is not encrypted.
- ▶ If possible, leave some CPU resources available for vMotion operations. To ensure the ability to use full network bandwidth, ESXi reserves CPU resources on the source and destination hosts.






Figure 3-1 vMotion operations

### 3.1.4 Distributed Resource Scheduler

vSphere Distributed Resource Scheduler (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 that is based on CPU and memory activity.

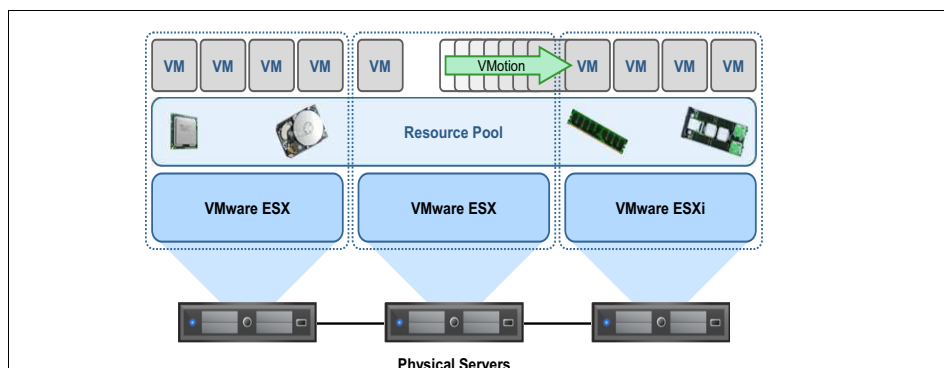


Figure 3-2 Distributed Resource Scheduler (DRS) operations

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

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

- ▶ Initial placement of VMs  
When a VM is powered on, DRS places it on an appropriate host or generates a recommendation depending on the specified 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 recommends VM migrations to maximize workload performance.
- ▶ Power management  
Distributed Power Management (DPM) works with DRS, and can place vSphere hosts in standby mode or power them back on as capacity demands. 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 when hosts are placed into maintenance mode.

The following DRS design considerations are important:

- ▶ Enable DRS on the entire cluster in fully automated mode, unless there are specific constraints.
- ▶ If needed, change the default DRS settings on specific VMs.
- ▶ Configure affinity and anti-affinity rules and DRS groups only when necessary (for example, if certain VMs must run on certain hosts).

### **3.1.5 High Availability considerations**

The vSphere High Availability (HA) provides an automated process for restarting VMs when a physical host becomes unavailable (see Figure 3-3 on page 71). When this situation occurs, the VMs are automatically registered and restarted on the remaining hosts in the cluster.

HA helps organizations meet their defined SLAs and reduces the potential for long outages by managing the risk that is associated with having aggressive consolidation ratios on physical hosts. When hardware failures occur, HA helps avoid incremental labor costs by providing automated recovery processes.




Figure 3-3 High Availability (HA) operations

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

The master host in an HA cluster has the following responsibilities:

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

The subordinate hosts primarily run VMs, monitors their runtime states, and reports state updates to the master host. A master host can also run and monitor VMs. Subordinate 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 VMware HA design considerations are important:

- ▶ HA should be enabled on all clusters with strict admission control. If you have a cluster that contains only non-persistent desktops, HA on the cluster should be disabled.
- ▶ Spare failover capacity should be determined based on specific customer requirements. If there is no specific requirement, the following general guideline can be applied:
  - Clusters with 12 hosts or fewer: Allow for the loss of at least one physical host
  - Clusters with more than 12 hosts: Allow for the loss of at least two physical hosts.
- ▶ Configure the Percentage of Cluster Resources Reserved policy to reserve failover capacity for at least one host. Use the Host Failures Cluster Tolerates policy if the virtual machine 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 only  $N + 10\%$ .
- ▶ HA works even if vCenter is down; however, vCenter is needed to configure HA.

### 3.1.6 vSphere licensing considerations

ESXi 5.1 is licensed per CPU socket. The vRAM entitlement that was introduced in ESXi 5.0 was ended with vSphere 5.1.

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

vSphere is available in three editions, ranging from the Standard edition with basic features to the Enterprise Plus edition with a full range of features. Table 3-1 on page 73 shows some of key vSphere features and the vSphere edition in which they are provided.

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
Distributed Resource Scheduler (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
Distributed switches	Enterprise Plus

### 3.1.7 Flex System integration with VMware

The Flex System Manager (FSM) accelerates the provisioning of compute node, networking, and storage resources that are used by the VMware environment. These capabilities decrease deployment time significantly.

VMware integration features make the following actions possible:

- ▶ Deploying hardware patterns from the FSM to new compute nodes to ensure that adapter interfaces are logically assigned to the compute resources.
- ▶ Installing IBM-customized ESXi 5.1 images to the new compute nodes from within the FSM interface.
- ▶ Providing VMware environment visibility and ESX resource inventory and topology views from within the FSM interface, including the ability to deploy new VM images.

- ▶ Providing extensibility from the native vCenter server to the Flex System hardware by using the specialized IBM Systems Director Upward Integration Module (UIM). Capabilities of the UIM include monitoring power and thermals of the Flex System components, viewing and updating firmware and software levels for various components in the chassis, and modifying the settings for predictive failure alerts in the chassis.

## 3.2 Networking considerations

Networking considerations apply to the physical network infrastructure and the VMware vSphere virtual network infrastructure.

From the physical network perspective, the host networking resources are shared by the virtual desktops the network supports. If there is insufficient bandwidth, users experience a reduced level of performance. As such, it is recommended to use fast network cards. IBM Flex System compute nodes offer 10 Gb Ethernet connectivity.

Also, performance might be improved by separating different types of network traffic. For example, traffic that is related to system management, VMs, storage, provisioning, and backups can all be isolated from each other. For more information about network design, see Chapter 5, “IBM Flex System and VMware View lab environment” on page 117.

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

### 3.2.1 Virtual switches

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

Although VSSs are defined at the host level, VDSs are defined at the data center level, which means that the switch configuration is pushed consistently to all hosts within the same data center. The VDS is also called a dvSwitch.

In addition, VDSs enable advanced features, such as Rx traffic shaping, improved monitoring through port mirroring (dvMirror), consistent network statistic monitoring, and Link Layer Discovery Protocol (LLDP), a vendor-neutral standard that is equivalent to Cisco Discovery Protocol (CDP).

The use of VDSs requires an Enterprise license. It is recommended that VDSs be used to enable their advantages.

When you are using VDSs, remember that they often can be controlled only from your vCenter server (unless you use a third-party VDS, such as IBM Distributed Virtual Switch 5000V, or Cisco 1000V). This means that 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 virtual switch 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 are called *dvPort groups*.

VMkernel Port is a specialized type of virtual switch port 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. vSphere 5.x includes ESXi hosts only, so 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. They provide a connection between a virtual network and a physical network.

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

The following networking design considerations are important:

- ▶ Ensure redundancy by using a single dvSwitch with redundant uplinks on all of the hosts in the cluster.
- ▶ Create separate, highly available port groups for each of management and vMotion traffic types. The Flex System platform positions Ethernet switch hardware inside the chassis, which provides improved network performance for activities that use network bandwidth (such as VMware vMotion) when compared to traditional top-of-rack network switching.
- ▶ Use VMware NetQueue to enable Intel Virtual Machine Device Queues (VMDq) support for the GbE ports.
- ▶ Use the TCP offload engine (TOE) capabilities of x240 network adapters to improve network performance by enabling stateless offload of the following elements:
  - Checksum offload
  - TCP segmentation offload (TSO)
  - Jumbo frames (JF)
  - Large receive offload (LRO)

## **3.3 Storage considerations**

Storage has a major effect on the performance, scalability, and availability of the Horizon View 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 non-persistent virtual desktops are hosted, 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. Instead, the use of tiered storage, where different storage technologies (such as solid-state drives and network-attached and Fibre Channel-attached storage systems) and drive access technologies (such as SAS and SATA) are grouped into storage tiers, which provide an effective way to offer a range of storage options that are based on needs that relate to performance, scalability, redundancy, and cost.



In this way, different virtual workloads with similar storage requirements can be grouped and a similar cost model can be applied.

### 3.3.3 Load balancing

VMware vSphere Storage DRS is a new feature that was introduced in vSphere 5.0. It provides load-balancing mechanisms that are based on I/O and space capacity and initial VM placement. VMware vSphere Storage DRS helps to decrease the operational effort that is associated with the provisioning of VMs and monitoring of the storage environment.

vSphere Storage DRS includes the following key features:

- ▶ Resource aggregation

This is the main feature of vSphere Storage DRS and is the one that all other features depend on. Datastores can be aggregated to a single-unit datastore cluster, and these datastore clusters form the basis of vSphere Storage DRS. By using this feature, you can manage the storage resources in a way that is similar to how vSphere DRS manages compute resources in a cluster. As with a cluster of hosts, a datastore cluster is used to aggregate storage resources, which enable smart initial placement of the virtual disk files and load balancing of existing workloads.

- ▶ Initial placement

During the manual provisioning of a VM, crucial provisioning factors, such as current space usage and I/O load, are often ignored. vSphere Storage DRS provides initial placement and ongoing balancing recommendations, which helps vSphere administrators make placement decisions that are based on space and I/O capacity. Initial placement simplifies and speeds up the provisioning process by automating the selection of a datastore.

- ▶ Load balancing

Load balancing can be thought of as a tool that proactively prevents high latencies and reactively prevents out-of-space scenarios that result from overloads on individual datastores. As load imbalances begin to occur, vSphere Storage DRS makes recommendations to correct them. Space-utilization load balancing is reactive to alleviate bottlenecks or extreme imbalances.

- ▶ Affinity rules

vSphere Storage DRS applies smart placement rules (in the form of affinity rules) on the VM files. Affinity rules help prevent placing VMs with similar tasks on the same datastore. These rules also help keep VMs together when required.

- Datastore maintenance mode

Datastore maintenance mode can be compared to host maintenance mode. When a datastore is placed in maintenance mode, all registered VMs on that datastore are migrated to other datastores in the datastore cluster. A typical use case involving this feature is a data migration to a new storage array.

### 3.3.4 Redundancy

vSphere datastores must be designed to meet the redundancy requirements of the components that they support, including RAID levels, storage adapters, and the back-end storage configuration. A leading 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) for Active/Active storage arrays and MRU (VMware) for Active/Passive storage arrays. Active/Active storage arrays can also use Round Robin (VMware) multipathing policy if the storage vendor supports it.

Flex System V7000 is an Active/Active storage solution that works well with Round Robin. Round Robin is the recommended multipathing policy because it provides more optimal use of the storage paths to every LUN. At any time, the LUN is accessed over a single path, but that path switches, by default, after every 1000 sent I/Os.



# VMware Horizon View design considerations

This chapter describes design considerations for the VMware Horizon View on IBM Flex System solution.

**Note:** This chapter is based on the information that is contained in the IBM Reference Architecture (RA) for VMware View. At the time of this writing, the IBM RA was based on VMware View 5.1; however, most testing and validation results can also be applied to VMware Horizon View 5.2, which is described in this book.

The IBM RA is updated regularly to include new features and components. For more information about the most recent IBM RA for VMware View, see this website:

<http://ibm.co/17c0yaN>

The chapter includes the following topics:

- ▶ VMware Horizon View components
- ▶ Choosing a desktop protocol
- ▶ VMware View provisioning
- ▶ Storage configuration
- ▶ Network configuration
- ▶ Choosing desktop and application delivery model
- ▶ Operational model and sizing guidelines

## 4.1 VMware Horizon View components

This section describes the functions of the various components of VMWare Horizon View (which is also known as *View*).

The IBM reference architecture for VMware View defines two types of server clusters:

- ▶ Compute clusters:
  - Hosts the virtual desktop workloads
  - Composed of multiple IBM compute nodes, the number of which varies based on the number of users that are hosted (for more information, see 4.7, “Operational model and sizing guidelines” on page 106)
  - Should not host workloads other than virtual desktops
  - Separate compute clusters for dedicated and stateless virtual desktops
- ▶ Management cluster:
  - Hosts the VMware Horizon View management components
  - Can be hosted on an existing or new vSphere environment
  - Contains VMware vCenter, vCenter SQL server, View Connection server, and other optional components
  - Can host more infrastructure services (AD, DNS, DHCP, and so on) if they do not exist in the environment

Figure 4-1 shows View components. These components are described next.

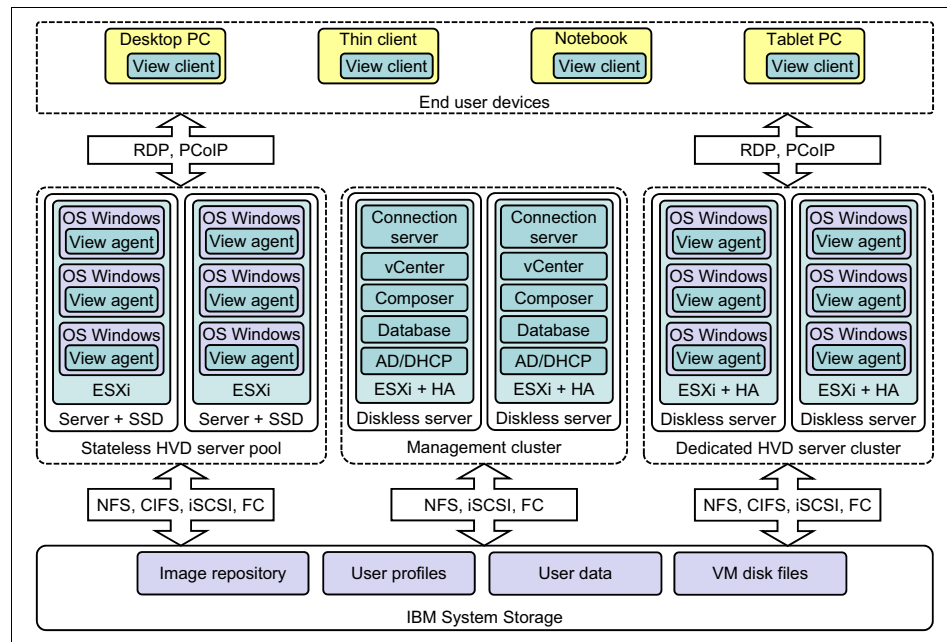


Figure 4-1 VMware Horizon View components

VMware Horizon View solution features the following components:

► VMware Horizon View Connection Server

VMware Horizon View Connection Server is a software service that acts as a broker for client connections by authenticating and then directing incoming user requests to the appropriate View desktop, thus ensuring that only valid users are allowed access.

If a virtual desktop is not available, the broker works with the management and provisioning layers to ensure that a virtual machine (VM) is ready and available.

View Connection Server must be installed on a dedicated physical or VM server, and the server must be a member of an Active Directory (AD) domain that is trusted by all View clients.

The View Administrator console (sometimes called *View Administrator*) must be installed on View Connection Server to manage the View environment and perform the following tasks:

- Deploy virtual desktops
- Create desktop pools
- Control access to desktop pools
- Examine View system events

View Administrator is a web-based application that is installed when you install View Connection Server.

The desktop on which you start View Administrator must trust the root and intermediate certificates of the server that hosts View Connection Server.

**Note:** The physical or virtual machine that host View Connection Server must use a static IP address.

The following operating systems support all View Connection Server types, including standard, replica, and security server installations:

- Windows Server 2008 R2 64-bit Standard and Enterprise
- Windows Server 2008 R2 SP1 64-bit Standard and Enterprise

View Connection Server requires specific versions of VMware virtualization software. For more information about which versions of Horizon View are compatible with which versions of vCenter Server and ESX/ESXi, see the VMware Product Interoperability Matrix at this website:

[http://www.vmware.com/resources/compatibility/sim/interop\\_matrix.php](http://www.vmware.com/resources/compatibility/sim/interop_matrix.php)

When you are installing replicated View Connection Server instances, you must configure the instances in the same physical location and connect them over a high-performance LAN. Do not use a WAN to connect replicated View Connection Server instances.

► VMware Horizon View Composer

VMware Horizon View Composer is a software service that can be used with VMware vCenter and View Connection Servers to deploy multiple linked-clone desktops from a single centralized base. View Composer can be installed directly on the vCenter Server or a dedicated server.

**Note:** View Composer is only required if linked-clone desktops are deployed.

View Composer can be installed on the same physical or virtual machine as vCenter Server or on a separate server. The following operating systems are supported:

- Windows Server 2008 R2 64-bit Standard and Enterprise
- Windows Server 2008 R2 SP1 64-bit Standard and Enterprise

View Composer requires an SQL database to store data. The View Composer database must be on, or be available to, the View Composer server host. Any of the following databases can be used:

- Microsoft SQL Server 2005 Express
- Microsoft SQL Server 2005 SP3 and later, Standard and Enterprise
- Microsoft SQL Server 2008 R2 Express
- Microsoft SQL Server 2008 SP1 and later, Standard and Enterprise
- Oracle 10g (Release 2)
- Oracle 11g (Release 1 and 2)

If a database server exists for vCenter Server, View Composer can use that existing database server.

**Note:** If you create the View Composer database on the same SQL Server instance as vCenter Server, make sure that you do not overwrite the vCenter Server database.

#### ► VMware Horizon View Transfer Server

VMware Horizon View Transfer Server is an optional software service that is used for offline desktops. It supports check in, check out, and replication of desktops that run in local mode. The View Client with Local mode is used where access to a virtual desktop is required during times where no network access is available. View Transfer Server is installed on a dedicated server or virtual machine.

View Transfer Server transfers static content to and from the View Transfer Server repository and transfers dynamic content between local desktops and remote desktops in the datacenter. View Transfer Server has the following storage considerations:

- The View Transfer Server repository must have enough space to store static image files.
- View Transfer Server supports 20 concurrent disk transfers.
- View Transfer Server must have access to the datastores that store the desktop disks to be transferred.



Install View Transfer Server on one of the following supported operating systems with at least 4 GB of RAM:

- Windows Server 2008 R2 64-bit Standard and Enterprise
- Windows Server 2008 R2 SP1 64-bit Standard and Enterprise

► VMware vCenter Server

VMware vCenter Server provides a central administration point for VMware vSphere hosts and other components of the vSphere suite. VMware vCenter Server creates and manages all virtual desktops that are based on instructions that are received from the View Connection Server and the View Composer Server.

VMware vCenter server can be installed on a dedicated physical or virtual machine.

► vCenter SQL Server

vCenter database is a data store that is used to centralize farm configuration information and transaction logs. Because the SQL server is a critical component of the View infrastructure, redundant servers must be available to provide fault tolerance. The following databases are supported:

- IBM DB2® 10 Enterprise and IBM DB2 Enterprise 9.7.2
- Microsoft SQL Server 2005 Standard, Enterprise, and Datacenter editions (SP4)
- Microsoft SQL Server 2008 Standard and Enterprise editions (SP2, SP3) and Microsoft SQL Server 2008 Datacenter edition (SP2)
- Microsoft SQL Server 2008 R2 Express (64-bit only), Standard, and Enterprise editions (SP1)
- Oracle 10g (Release 2) and Oracle 11g (Release 1 and 2)

► ESXi hypervisor

ESXi is a bare-metal hypervisor for the compute servers. The hypervisor provides a virtualized environment for running VMs with the desktop operating systems in them. These VMs are called *hosted virtual desktops*.

vSphere is the only hypervisor that is fully supported for hosting View virtual desktops as it fully integrates with View for full desktop lifecycle management.

► VMware Horizon View Agent

VMware Horizon View Agent is installed on the virtual desktops, physical desktops, and Windows Terminal Servers that are managed by View.

The View agent connects the virtual desktop to View's devices and services, such as client-attached USB devices, client connection monitoring, virtual printing, single sign-on, and View Persona Management.

VMware Horizon View Persona Management is an optional component of the View Agent that can be used as an alternative to Microsoft Windows roaming profiles for managing user Windows profile data and application settings.

View Persona Management has the following specific benefits:

- User profile data is loaded only as required, which speeds up the user desktop login process.
- Logoff times can be accelerated by syncing back the user profile updates to the remote persona management repository.

► VMware Horizon View Client

VMware Horizon View Client communicates with a View Connection Server and starts connections to desktops and Windows Terminal Servers. Users can access their virtual desktop from any device that is supported by the respective desktop virtualization solution.

The View Client is available for Microsoft Windows, Apple OS X, Android, iOS, and Ubuntu Linux.

► 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.

► VMware ThinApp

VMware ThinApp is an application virtualization product that integrates with View Manager to package conventional applications so that they become portable applications.

VMware ThinApp packages applications into executable files (in MSI or EXE format) that are encapsulated from other applications and from the underlying machine's operating system. The goal is to eliminate application conflicts and streamline application delivery and management.

ThinApp has the following capabilities:

- Eliminates application conflicts by isolating desktop applications from each other and from the underlying operating system.
- Enhances security policies by deploying ThinApp packages on PCs and by allowing users to run their favorite applications without compromising security.
- Increases users' mobility by deploying, maintaining, and updating virtualized applications on USB flash drives.
- Usable as a component of VMware View to reduce desktop storage costs and streamline updates to endpoints.

- Reduces the number of applications that must be installed on the master virtual desktop image, which reduces the need to deploy and maintain many images for different user bases.

The Figure 4-2 shows an example of ThinApp application delivery to devices.




Figure 4-2 Application delivery by using VMware ThinApp

## 4.2 Choosing a desktop protocol

When you are designing a View infrastructure, it is important to determine how much network bandwidth is required to support the View Client connections.

Desktop protocols and software provide access over a network connection to View desktops that are in the datacenter.

Depending on which type of client device you are using, you can choose between the PC-over-IP protocol (known as PCoIP) or Microsoft Remote Desktop Protocol (RDP). The preferred protocol for VMware Horizon View is PCoIP.

## PCoIP

PCoIP provides an optimized delivery of the entire desktop environment, including applications, images, audio, and video content for various users on the LAN or across the WAN. PCoIP can compensate for an increase in latency or a reduction in bandwidth, which helps ensure that users can remain productive regardless of network conditions.

The PCoIP protocol has the following features that make it ideal for connecting to View desktops:

- ▶ Supports image caching to store display data and minimize bandwidth usage.
- ▶ Achieves compression ratios of up to 100:1 for images and audio.
- ▶ Provides optimization controls for reducing bandwidth usage on the LAN and WAN.
- ▶ Enables more efficient encoding and decoding of content between the Virtual Desktop and the remote Client by using multiple codecs.
- ▶ Supports multiple monitors for some client types. For example, on Windows based clients, you can use up to four monitors and adjust the resolution for each monitor separately, up to a maximum of 2560 x 1600 per display. When the 3D feature is enabled, up to two monitors are supported by a resolution of up to 1920 x 1200.
- ▶ Supports 32-bit colors for virtual displays.
- ▶ Supports the advanced encryption standards AES-128, AES-192, or AES-256 (AES-128 is turned on by default).
- ▶ Supports USB redirection.
- ▶ Supports audio redirection with dynamic audio quality adjustment for LAN and WAN.
- ▶ Supports copy and paste text and images between the local system and the desktop is supported, up to 1 MB. Supported file formats include text, images, and Rich Text Format (RTF).
- ▶ Eliminates handshakes that are used in Transmission Control Protocol (TCP)-based display protocols.

**Note:** PCoIP is supported as the display protocol for View desktops with virtual machines and with physical machines that contain Teradici host cards.

The VMware Horizon View Architecture Planning guide provides estimates for PCoIP bandwidth usage that is based on the application workload of the client. A selection of these estimates is shown in Table 4-1.

*Table 4-1 Estimates of PCoIP bandwidth usage*

Workload characteristics	Bandwidth
2D display and single monitor with web and limited office applications	50 - 100 kbps
2D display and single monitor with office applications	100 - 150 kbps
3D display and single monitor with office applications	400 - 600 kbps
3D display and multiple monitors with office applications	500 kbps -1 Mbps
3D display and multiple monitors with 480p video and images and frequent window changes	2 Mbps

## RDP

Microsoft RDP is a TCP-based display protocol that does not have many of the WAN optimization and acceleration features that are found in PCoIP. Microsoft Remote Desktop Connection (RDC) uses RDP to transmit data.

Microsoft RDP provides the following features:

- ▶ 128-bit encryption
- ▶ 32-bit color for virtual displays
- ▶ Supports up to 16 monitors in spanning mode
- ▶ Supports copy-paste between the local system and the View desktop for text and system objects, such as folders and files

## 4.3 VMware View provisioning

VMware Horizon View offers the ability to create and provision pools of desktops as its basis of centralized management.

You can create a virtual desktop pool from one of the following sources:

- ▶ A physical system, such as a physical desktop PC or a Windows Terminal Services server
- ▶ A virtual machine that is hosted on an ESX or ESXi host and managed by vCenter Server

- A virtual machine that runs on VMware Server or some other virtualization platform that supports View Agent

VMware View supports floating and dedicated desktop assignment models. Provisioning for VMware View is a function of vCenter server and View Composer for linked clones.

vCenter Server allows for manually created pools and automatic pools. In addition, it allows for provisioning full clones and linked clones of a parent image for dedicated and stateless virtual desktops.

You can configure a desktop pool so that users have dedicated assignments or floating assignments to the desktops in the pool. You must choose a user assignment for automated pools that contain full virtual machines, automated linked-clone pools, and manual pools.

### 4.3.1 Dedicated and floating desktop pools

In this section, we describe the dedicated and floating desktop pool models.

#### **Dedicated desktop pools**

A virtual desktop from the Dedicated pool (which is also called *persistent pool* or *stateful pool*) is often assigned to the user upon its first logon. After that first logon, each user always connects to the same virtual desktop, which allows them to personalize the appearance of the desktop and have constant access to the data and documents they create there. It is also possible to manually pre-assign a user to a virtual desktop.

Figure 4-3 on page 91 shows the concept of dedicated user assignments.




Figure 4-3 Dedicated pool user assignments

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.

Dedicated desktops can be implemented by using full or linked clones provisioning models. For more information, see 4.3.2, “Provisioning by using full and linked clones” on page 93.

**Full and linked clones:** Dedicated desktops that are based on linked clones are difficult to back up and restore. Consider the use of dedicated desktops that are based on full clones instead.

Dedicated desktops that use full clones are independent copies of a parent virtual machine that share nothing with it. Ongoing operation of a full clone is entirely separate from the parent virtual machine.

Dedicated pools that use linked clones are copies of a parent virtual machine that shares virtual disks with it in an ongoing manner. This configuration conserves disk space and allows multiple virtual machines to use the same software installation. To achieve persistency, a persistent disk is used.

The persistent disk is created in addition to the base operating system disk image, and the VMware View agent instructs the guest OS to offload the user profile to this separate disk. The user profile consists of application data, registry entries, and all other user-specific folders.

When persistent disks are used, it is possible to replace the virtual desktop base image using recompose or refresh operations. This is a useful way to accommodate application upgrades and patches without losing user data. However, if persistent disks are used, it is important to design a good backup solution for the data that is on the disks. Backup agents from the guest OS can be used, but that approach increases the overall cost of the solution. You also can use VMware View storage tiering to dedicate a datastore for all persistent disks and then back up the entire datastore or LUN.

Another alternative to backing up the persistent disks is to enable active directory roaming profiles. However, you need to take into account the design of roaming profiles where the data is copied from the network during the user logon process and then copied back to the network at logoff.

Dedicated desktops have many drawbacks. The desktop images are large, grow quickly, and must be updated and patched individually because they have no common base image. There is often no separation between the operating system and user data. Additionally, because the image is unique to each user, data backup is critical and it typically involves large data sets.

More importantly, high availability must be considered. The user must always be able to connect to the same image, even if a VDI host fails, which is impossible if the desktop is hosted locally. Therefore, dedicated desktops require access to expensive shared storage.

## **Floating desktop pools**

With a *floating* desktop pool (which is also called a *non-persistent* pool or *stateless* pool), View Manager dynamically assigns desktops in the pool to entitled users. Users connect to a different desktop each time they log in. When a user logs off, the desktop is returned to the pool.

Floating desktops are allocated to users temporarily. After the user logs off, changes to the image often are discarded (that is, the image is reset) and the desktop becomes available for the next user or a new desktop is created for the next user session. To achieve a persistent user experience (that is, the ability to personalize the desktop and save data), developers rely on user profile management, folder redirection, difference data collection, and other approaches. Specific applications, if needed, can be provided to floating desktops by using application virtualization technologies, such as ThinApp.



Figure 4-4 shows the concept of floating user assignments.




Figure 4-4 Floating pool user assignments

This floating or stateless approach is based on a logical separation of the operating system, application, and user layers. The approach allows a common, centrally managed base image to be used for all users in the same pool. If the image is corrupted or becomes unavailable, the user connects to another image in the pool, relying on high-availability features that are provided by connection brokers rather than through a storage-hungry, VM failover approach. Backups are simplified as only a small subset of the overall data (such as, profile information and saved data) must be archived.

The stateless approach enables the use of local storage instead of shared storage, with only a fraction of the data on distributed storage (for example, profile and user data). This method directly reduces the cost per desktop. The only potential restriction to storing a desktop locally is that it cannot be moved from one server to another without restarting the VM. If a live migration of virtual desktops is required, a stateless desktop can still be used but all of the data is on a shared storage and there is a corresponding increase in the performance requirements of that shared storage.

### 4.3.2 Provisioning by using full and linked clones

Dedicated and floating desktop pools can be automatically provisioned by using the following primary provisioning models that are built into View:

- ▶ Full clones
- ▶ Linked clones

## Full clones

Full virtual machine desktops (which also are called *full clones*) are created by using a virtual desktop master image that was converted to the vSphere template format. A full virtual machine or full clone is an independent copy of the template and is managed separately from any other desktops and the template on which it was based. The full clone duplicates only the state of the virtual machine at the instant of the cloning operation. Because a full clone does not share virtual disks with the parent virtual machine, full clones generally perform better than linked clones. However, full clones take longer to create than linked clones. If the involved files are large, creating a full clone can take several minutes.

Use full clones for dedicated desktop pools, where the users expect to be connected to the same desktop virtual machine every time. Also, use full clones when specific software is required.

Figure 4-5 shows how a full clone environment operates.

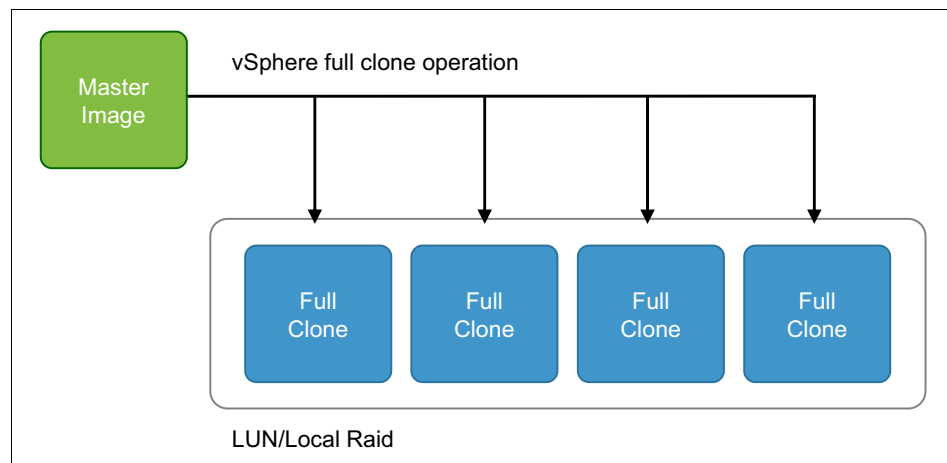


Figure 4-5 Full clone environment

To successfully deploy full virtual machine desktops, you must first complete the following preliminary tasks:

- ▶ Prepare a virtual machine template for View Manager to use to create the desktops. View Agent must be installed on the template.
- ▶ Ensure that any customizing specifications are accurate. Deploy and customize a virtual machine from your template by using the customizing specification.
- ▶ Verify that enough ports are available on the ESX virtual switch that is used for desktop virtual machines.

Full clones do not require an ongoing connection to the parent virtual machine. Overall performance of a full clone is the same as it is on a non-cloned virtual machine, but a linked clone trades potential performance reductions for a guaranteed conservation of disk space. If you are focused on performance, you should use a full clone over a linked clone.

## Linked clones

VMware View with View Composer uses the concept of linked clones to quickly provision virtual desktops. View Composer uses a parent image to create a pool of linked clone virtual machines. A parent image is a tuned desktop that is used to create new replica images.

Linked-clone desktop images optimize desktop storage space and improve image control. Changes to a master image apply to user desktops without affecting user settings, data, and applications. Each linked clone acts like an independent desktop, with a unique host name and IP address, yet the linked clone requires less storage. All files that are available on the parent at the time of the snapshot remain available to the linked clone. The operating system reads all of the common data from the read-only replica and the unique data that is created by the operating system or user is stored on the linked clone.

The linked-clone virtual machines each have unique identities and can be powered on, suspended, or reconfigured independently of the master image.

Figure 4-6 shows how the linked clone environment operates.




Figure 4-6 Linked Clone environment

Changes can be made to the snapshot of the virtual desktop master image while still retaining the ability to deploy more desktops that are based on the condition of the desktop when the snapshot was taken. When it is time to deploy the updated image, you take a second snapshot and recompose the desktops.

A linked clone desktop has the following advantages over a full clone desktop:

- ▶ Linked clone desktops share the parent virtual disk for read operations, so the amount of disk space they require is greatly reduced.
- ▶ Linked clone desktops can be recomposed, a process in which software updates or other changes are applied to the master image once and then propagated to the replica disks, which applies those changes to the entire desktop pool.
- ▶ Linked clone desktops can be refreshed, a process that deletes the modified contents of the linked clone operating system and disposable data disks. This action discards any changes that were made after the desktop was deployed, which allows for tight control over the user experience.
- ▶ A linked clone desktop pool can be rebalanced, which redistributes linked clone storage across datastores to prevent an imbalance in storage usage.
- ▶ A linked clone approach improves agility in Horizon View by reducing provisioning time, which provides near-instant provisioning of virtual machines.
- ▶ By using tiered storage with View Composer linked clones, you can redirect user data to a different datastore. This allows the linked-clone virtual machine OS to be refreshed (which is also referred to as being *rebased*) while preserving local user data because you can detach and attach the persistent disk to the linked clone virtual machine.

Although a full-clone desktop requires only one virtual hard disk, a linked-clone desktop requires up to four of the following virtual hard disks, and the replica disk that is shared among the desktops in the pool:

- ▶ **Replica disk:** When a desktop pool is created, a clone of the virtual desktop master image hard disk is created on each datastore that contains linked clones. These clones of the virtual desktop master image are referred to as *replica disks*. A replica disk can be created on a dedicated datastore, which results in only one replica disk being created rather than one for each linked-clone datastore. The replica disks are read-only; all changes are written to the individual linked-clone virtual hard disks.
- ▶ **OS disk:** This disk stores the system data that associates the linked clone with the base image and functions as a unique desktop.

- ▶ **Persistent disk:** This optional disk is used in dedicated assignment pools only. This disk can be used to store user profile data and the contents are retained during a refresh, recompose, or rebalance operation. If a persistent disk is not used, the user profile data is stored in the OS disk and is lost during refresh or recompose operations.
- ▶ **Disposable data disk:** This optional disk is used to store the OS paging and temporary files. The contents of this disk are discarded when the desktop is powered off and during refresh and recompose operations. If a disposable data disk is not used, the page file and temporary files are stored in the OS disk.
- ▶ **QuickPrep configuration data disk:** This disk stores QuickPrep and other OS-related data that must be preserved during refresh and recompose operations.

Because linked clones can grow in size over time, consider the use of the space-efficient sparse disk (SE Sparse disk) feature (which is enabled by default) to reclaim unused space in the linked-clone virtual machine. This feature was introduced in VMware Horizon View 5.2.

SE Sparse disks help optimize storage capacity in the VDI environments that use linked clones. Before Horizon View 5.2, clients needed to perform Recompose or Refresh operation to address the issue with the linked clone capacity growth. With SE Sparse disks, the space reclamation is automated through Horizon View, and there is no need to reclaim unused space manually through refresh operations.

Automated space reclamation by using SE Sparse disks generates a substantial amount of storage I/O and uses processor cycles; therefore, you should plan to run these operations during low or no activity by defining blackout windows when you are configuring SE sparse disk feature for the desktop pool.

## **Storage capacity for full and linked-clone desktop pools**

For sizing full and linked clone desktop pool's storage capacity, the following types of swap files are available:

- ▶ The virtual machine swap file (.vswp) that is stored with the virtual machine is equal to the amount of allocated, non-reserved vRAM, or 100% of the allocated vRAM if not using memory reservations.
- ▶ The secondary or overhead swap file is created to accommodate operations when the host is under memory pressure.

For the virtual machine swap file, consider reserving a portion of the allocated vRAM in the virtual machine to balance the capacity overhead that the swap file produces. For example, for a virtual machine with 2 GB of vRAM, consider a 1 GB reservation, which reduces the swap file size by 50%.

**Swap files:** It is common to size VM memory requirements in a way that avoids swapping, which helps improve overall VDI performance.

#### ***Full clone per virtual machine calculation***

The storage capacity that is required for a full clone virtual machine is simple to calculate by using the following formula:

Full clone + .vsmp + Overhead

#### ***Linked clone per virtual machine calculation***

To calculate storage capacity for a linked-clone virtual machine, use the following formula:

Replica (per LUN) + Linked Clone + Growth + .vsmp + Overhead

The replica size, which is equal to the master image size, is taken into account on a per-LUN basis. Capacity per linked clone virtual machine begins with the linked clone (50% of the replica size is a good estimate). After that, add the linked clone growth that occurs between refresh and rebase operations (20% of the linked clone size is a good estimate). Finally, add the amount that is needed for virtual machine swap and overhead, if enabled.

## **4.4 Storage configuration**

VDI workloads place huge demands on network shared storage, whether it is to support virtual desktop provisioning, VM loading across the network, or accessing user profiles and data files. In this section, we describe the storage considerations for non-persistent (stateless) and persistent (dedicated) virtual desktops that use the View deployment models with full and linked clones.

VMware datastores can be hosted on supported shared storage that uses FC, FCoE, iSCSI, or NFS storage protocols.

The sizes and IOPS for user data files and user profiles that are described in this section are based on the IBM Reference Architecture for VMware View, and they can vary depending on the customer environment. For example, power users might require more storage space and 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.

Many customers need a hybrid environment of stateless and dedicated desktops for their users. The IOPS for dedicated users outweighs those for stateless users; therefore, it is typical to use dedicated users in defining any storage controller configuration requirements.

The storage configurations that are presented in this section feature conservative assumptions about the VM size, changes to the VM, and user data sizes to ensure that the configurations can manage the most demanding user scenarios.

### **Non-persistent (stateless)**

For non-persistent (or stateless) virtual desktops, the following local storage components are available:

- ▶ USB flash drive for ESXi hypervisor

Each compute node runs the IBM ESXi custom image that is on a USB flash drive.

- ▶ Compute node's local drives

The replicas and linked clones are stored on local solid-state drives (SSDs).

Because of the stateless nature of the architecture, there is little added value in configuring reliable SSDs in more redundant RAID configurations.

Redundancy is not achieved on a host level. Rather, it is 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 for non-persistent (stateless) virtual desktops are available:

- ▶ Datastores

Stateless virtual desktops use datastores to store all virtual images for linked clones.

The paging file (or vSwap) is transient data that also can be redirected to the datastore. In general, it is recommended that swapping is disabled, which reduces storage use (shared or local). The designated desktop memory size should match the user workload and should not depend on a smaller image and swapping, which reduces overall desktop performance.

- ▶ User profiles (for Roaming Profiles or Persona Management):  
If you are using Microsoft Roaming Profiles (MSRP) or View Persona Management, user profiles are stored by using the Common Internet File System (CIFS).
- ▶ User data files  
In the stateless user model, you must redirect persistent user data, such as documents and other file repositories to user-specific file shares (CIFS-based) or network drives.

Figure 4-7 shows the storage allocation for stateless virtual desktops.




Figure 4-7 Stateless virtual desktop storage allocation

Table 4-2 summarizes the peak input/output operations per second (IOPS) and disk space requirements for stateless virtual desktops on a per-user basis as a starting point.

Table 4-2 Shared storage considerations for stateless desktops

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 (dedicated)

The following local storage components of persistent (dedicated) virtual desktops are available:

- ▶ USB flash drive for ESXi hypervisor  
Each compute node runs the IBM ESXi custom image that is on a USB flash drive.
- ▶ For dedicated hosts, no local storage is configured.

The following shared storage components for persistent (dedicated) virtual desktops are available:

- ▶ Datastores  
Datastores are used to store all virtual desktops' associated data, such as the master image, replicas, linked clones, persistent disks, and full clones.
- ▶ User profiles (for Roaming Profiles or View Persona Management, if used)  
User profiles are typically hosted on a CIFS-based file share.
- ▶ User data files  
CIFS and NFS-based file shares are used to redirect persistent user data (documents, other file repositories, and so on) to user-specific file shares or network drives.

Table 4-3 summarizes the peak IOPS and disk space requirements for dedicated virtual desktops on a per user basis. The last two rows in the table contain the same information as was shown for stateless desktops. It is a leading practice is to keep the AppData folder with the linked clones.

*Table 4-3 Shared storage considerations for dedicated desktops*

Data type	Protocol	Size	IOPS	% Write
Master image	Block or NFS	30 GB	18	85%
Linked clones	Block or NFS	10 GB	18	85%
User AppData folder			18	85%
User files	CIFS or NFS	5 GB	1	75%
User profiles (MSRP)	CIFS	100 MB	0.8	75%

**Storage IOPS:** Depending on the environment, maximum IOPS requirements for the dedicated virtual desktops can be as high as 160 - 190 IOPS. Use Table 4-3 on page 101 as a starting point and evaluate your actual requirements that are based on an existing or projected workload.

Figure 4-8 shows the required storage tiers. It also shows a View hybrid environment that consists of a floating model and a dedicated pool model that are connected to the same storage system.

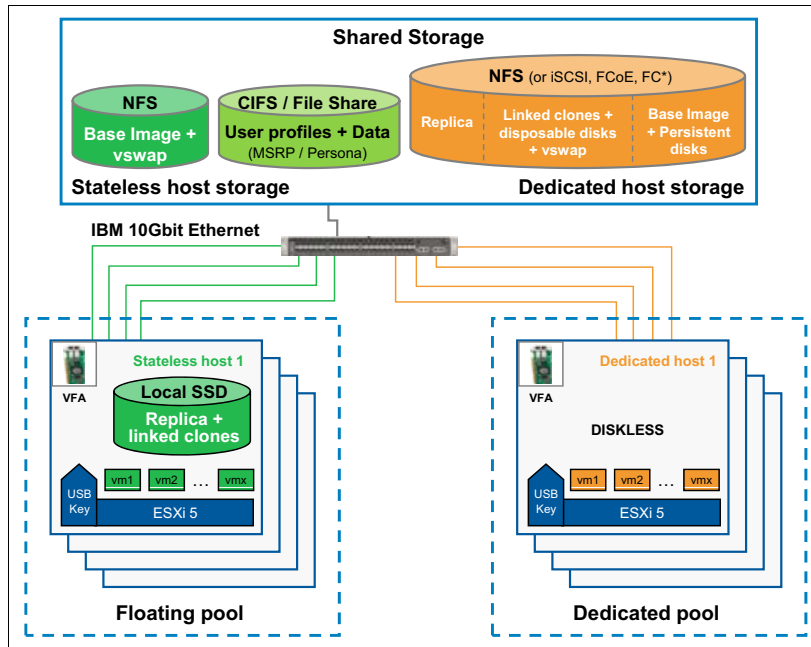


Figure 4-8 Storage layout for View: floating and dedicated models

With VDI infrastructures, storage IOPS performance takes precedence over storage capacity. This means that more drives are needed to achieve the required performance.

The large rate of IOPS (and the resulting need for many drives to support dedicated virtual desktops) can be lessened somewhat by caching read data in flash memory. This can be achieved by using the flash cache feature of some IBM System Storage N series controllers or, in IBM Flex System V7000 or Storwize V7000 storage systems, by implementing SSD storage with IBM EasyTier functionality.

VMware View 5.1 introduced the View Storage Accelerator feature that uses Content Based Read Cache (CBRC). Although CBRC is a vSphere feature, it is used in a unique way by VMware View to provide the host caching capability. The VSA feature provides a per-host RAM-based solution for View desktops that considerably reduces the read I/O requests that are issued to the storage layer and addresses boot storm issues.

With VSA, View indexes the contents of each virtual disk file when a virtual machine is created. The indexes are stored in a virtual machine digest file. At run time, the ESXi host reads the digest files and caches common blocks of data in memory. To keep the ESXi host cache up to date, View regenerates the digest files at specified intervals and when the virtual machine is recomposed.

Storage configurations should be based on your peak performance requirement, which often occurs during a so-called “logon storm”. This happens when all or most of the workers at a company arrive at work at the same time and try to start their virtual desktops simultaneously.

Storage configurations should also use conservative assumptions about the VM size, changes to the VM, and user data sizes. This approach helps ensure that the configurations can cope with the most demanding user scenarios.

## 4.5 Network configuration

A redundant 10 Gb network infrastructure is used to provide the network connectivity between all components of the VMware Horizon View architecture.

The following virtual local area networks (VLANs) are commonly deployed:

- ▶ Storage VLAN to provide storage connectivity (assuming that NFS, FCoE, or iSCSI storage is used). With Fibre Channel, no storage VLAN required.
- ▶ VM data VLAN for production (user) access.
- ▶ Management VLAN for dedicated access to the management interface of systems.
- ▶ VM control traffic VLAN for inter-VM communications such as vMotion.

On the server side, all networks are provided by a single dual port IBM 10GbE Virtual Fabric LOM. Each physical 10 Gbps port can be divided into four virtual ports with bandwidth that is allocated in 100 Mbps increments, up to the maximum 10 Gbps per physical port.

**Note:** The VLAN configuration and bandwidth allocation depends on your individual requirements. Ensure that you have adequate bandwidth available for each traffic type. For example, you might have another network that is dedicated to live migrations or a dedicated backup network.

The following starting points for bandwidth allocation can be useful:

- ▶ Management traffic: 0.5 Gbps
- ▶ VM control traffic: 1 Gbps
- ▶ VM data traffic: 1- 2 Gbps
- ▶ Storage traffic (if used): 1- 2 Gbps

Figure 4-9 shows logical network separation for the IBM SmartCloud VDI environment.

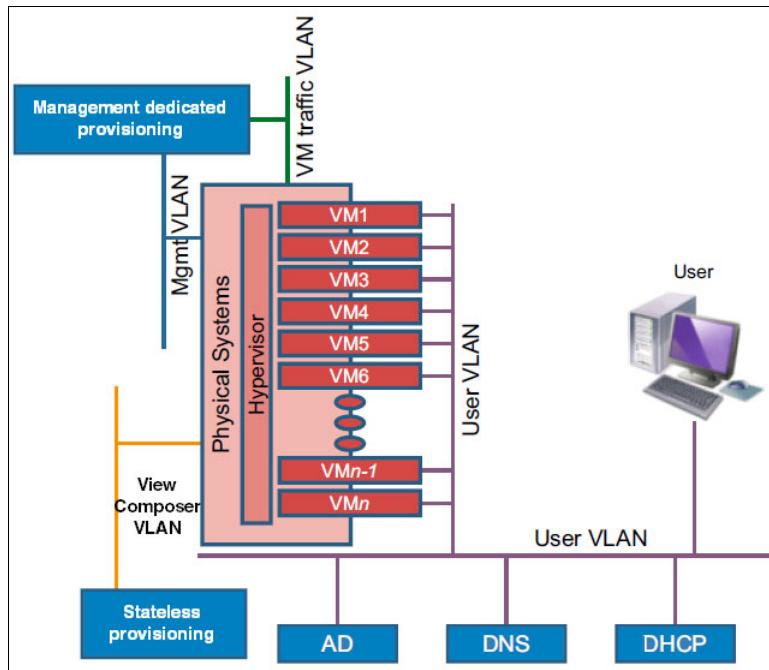


Figure 4-9 VDI logical network separation

## 4.6 Choosing desktop and application delivery model

VMware Horizon View-based virtual desktop infrastructure provides flexible desktop and application delivery by using the hosted applications and hosted virtual desktops.

The choice of a specific delivery model or combination of delivery models depends on user and application compatibility and application customization requirements, as shown in Table 4-4.

**Note:** The terms Low, Medium, and High that are used in Table 4-4 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 pooled desktops have better user density than dedicated desktops, and hosted applications have better user density than pooled desktops.

Table 4-4 Delivery model comparison, virtual application versus virtual desktop

Feature or requirement	Hosted virtual desktops		Hosted applications
	Dedicated	Pooled	
Provisioning model	Full clones or linked clones	Linked clones	Application virtualization
VDI component	vCenter VM Template	View Composer	ThinApp
Desktop OS compatibility	Yes	Yes	Yes
Server OS compatibility			Yes
User customization	Yes	Yes	
Application customization	Yes		
Professional graphics	Yes		
Management	Complex	Simplified	Simplified
Relative storage IOPS	High	Low	Low
Relative user density	Low	Medium	High
Relative cost	High	Medium	Low

If the application requires no user customization, the most cost-efficient way to deploy virtual desktop infrastructure is to use a hosted applications delivery model.

For highly customized user application environments, hosted virtual desktops provide a flexible and efficient way to deploy centralized desktop infrastructure. Within this method, a non-persistent model is more cost-optimized and a persistent model is more optimized for application customization.

## 4.7 Operational model and sizing guidelines

In this section, we describe operational models that cover stateless and dedicated environments. Stateless desktops that require live migration of a VM from one physical server to another are considered the same as dedicated desktops because they both require shared storage. In some client environments, both stateless and dedicated image models might be required, so a mixed operational model is needed.

**Sizing considerations:** Sizing considerations that are described in this section are based on the validated results that are obtained from the IBM Reference Architecture for VMware View. The most recent IBM RA for VMware View is available at this website:

<http://ibm.co/17c0yaN>

To show the operational model for different customer environments and size needs, four different configurations are described for supporting 600, 1,500, 4,500, and 10,000 users. Because the operational model for 10,000 users is approximately 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.

This section includes the following topics:

- ▶ 4.7.1, “Workload definition for the IBM RA test environment” on page 107
- ▶ 4.7.2, “VDI compute node configuration” on page 108
- ▶ 4.7.3, “Management services configuration” on page 111
- ▶ 4.7.4, “Shared storage configuration” on page 113

## 4.7.1 Workload definition for the IBM RA test environment

VDI is a performance-intensive workload that can stress all parts of the system including processors, memory, storage, networking, and the VDI software infrastructure. To successfully validate the performance, each of these attributes must be stressed in turn to determine its limits under defined workload. IBM Reference Architecture for VMware View uses Login VSI to generate user loads and monitor and measure the system performance under a particular load.

Login VSI is a VDI vendor-independent benchmarking tool to test and measure the performance and scalability of centralized Windows desktop environments, such as Server Based Computing and VDI. Login VSI measures the capacities of virtualized infrastructures by simulating typical user workloads and application usage.

IBM RA for VMware View uses Login VSI medium workload that simulates a medium-level knowledge worker, which uses Microsoft Office, Internet Explorer, and PDFs.

The medium workload is scripted in a 12 - 14-minute loop when a simulated Login VSI user is logged on. Each test loop performs the following operations:

- ▶ Microsoft Outlook 2007 and Outlook 2010: Browse 10 messages.
- ▶ Internet Explorer: Browse two websites.
- ▶ Flash application is run.
- ▶ Microsoft Word 2007 and Word 2010: Review and edit document.
- ▶ PDF Printer and Acrobat Reader: The Word document is printed to PDF and reviewed.
- ▶ Microsoft Excel 2007 and Excel 2010: A large randomized sheet is opened.
- ▶ Microsoft PowerPoint 2007 and PowerPoint 2010: A presentation is reviewed and edited.
- ▶ Archiving: The output of the session is archived.

After the loop finished, it is restarted automatically. Each loop takes approximately 14 minutes to run.

The following parameters and rules were used for Login VSI tests:

- ▶ User login interval: 30 seconds (some tests were run at 15-second intervals).
- ▶ Workload: Medium for most tests but more tests were performed by using the light, heavy, and multi-media workloads.
- ▶ All virtual desktops were pre-booted before the tests.

## 4.7.2 VDI compute node configuration

The VDI compute node is the base system unit that makes up the compute clusters. The compute clusters can consist of any IBM systems that are listed in 2.4, “IBM Flex System Compute Nodes” on page 20. Compute nodes run the VMware ESXi hypervisor and host user VMs.

An important consideration for compute servers is system memory. For stateless users, the typical range of memory that is required for each desktop is 1.5 GB - 4 GB; for dedicated users, the range of memory for each desktop is 2 GB - 6 GB. High-end computer-aided design (CAD) users that require 3D VDI technology might require 8 GB - 16 GB of RAM per desktop. In general, power users that require larger memory sizes also require more virtual processors.

The virtual desktop memory should be large enough so that swapping is not needed.

As a part of validating the reference architecture, IBM tested x222 and x240 compute nodes that were running VMs with different memory sizes of 1.5 GB, 2 GB, and 3 GB<sup>1</sup> and identified the maximum number of virtual desktops per compute node under specified workload (for more information, see 4.7.1, “Workload definition for the IBM RA test environment” on page 107). The results are summarized in Table 4-5.

Table 4-5 Number of virtual desktops per compute node

Feature (per node)	VM memory size		
	1.5 GB	2 GB	3 GB
<b>x222 compute node (dual-server)</b>			
System memory	384 GB (2x 192 GB)	384 GB (2x 192 GB)	384 GB (2x 192 GB)
Desktop VMs	204 (2x 102)	158 (2x 79)	104 (2x 52)
Desktop VMs (failover)	250 (2x 125)	190 (2x 90)	126 (2x 63)
<b>x240 compute node</b>			
System memory	256 GB	256 GB	384 GB
Desktop VMs	125	105	105
Desktop VMs (failover)	150	126	126

<sup>1</sup> For more information, see IBM Reference Architecture for VMware View at this website:  
<http://ibm.co/17c0yaN>



IBM testing shows that the number of users that are specified in Table 4-5 on page 108 is a good baseline and results in an average 75% usage of the processors in the server.

If a server fails, the 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.

The following configurations of the compute nodes are suggested:

- ▶ Non-persistent host:
  - Processor: Dual socket (8-core Intel Xeon processor E5-2680 or E5-2470)
  - Memory: 256 GB (16x 16 GB) or 384 GB (24x 16 GB)
  - Disks: IBM 2.5-inch MLC HS SSDs
  - Disk Controller: Standard integrated disk controller
  - Hypervisor: IBM USB Memory Key for VMware ESXi 5.1 (ESXi IBM Custom Image)
  - Network adapter: Integrated Dual Port 10 GbE Virtual Fabric LOM
- ▶ Persistent host
  - Processor: Dual socket (8-core Intel Xeon processor E5-2680 or E5-2470)
  - Memory: 256 GB (16x 16 GB) or 384 GB (24x 16 GB)
  - Disks: none
  - Disk Controller: None
  - Hypervisor: IBM USB Memory Key for VMware ESXi 5.1 (ESXi IBM Custom Image)
  - Network adapter: Integrated Dual Port 10 GbE Virtual Fabric LOM

If you intend to use the host in a dedicated user model that uses full virtual machines, you can remove the local SSDs because all VM data is on shared external storage.

Table 4-6, Table 4-7, and Table 4-8 show the number of compute nodes that are needed for different numbers of users size. The figures are based on the desktop VM quantity per server that is shown in Table 4-5 on page 108.

*Table 4-6 Compute nodes that are needed for different numbers of users (VM size of 1.5 GB)*

Description (VM size of 1.5 GB)	600 users	1500 users	4500 users	10000 users
<b>x222 compute node (dual-server)</b>				
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 to 1	4.5 to 1
<b>x240 compute node</b>				
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 4-7 Compute nodes that are needed for different numbers of users (VM size of 2 GB)*

Description (VM size of 2 GB)	600 users	1500 users	4500 users	10000 users
<b>x222 compute node (dual-server)</b>				
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
<b>x240 compute node</b>				
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 4-8 Compute nodes that are needed for different numbers of users (VM size of 3 GB)*

Description (VM size of 3 GB)	600 users	1500 users	4500 users	10000 users
<b>x222 compute node (dual-server)</b>				
Compute nodes @ 104 users	6	15	45	96
Compute nodes @ 126 users (failover)	5	12	36	80

Description (VM size of 3 GB)	600 users	1500 users	4500 users	10000 users
Failover ratio	5 - 1	4 - 1	4 - 1	5 - 1
<b>x240 compute node</b>				
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

### 4.7.3 Management services configuration

Management services are provided by the VDI solution for creating desktops, provisioning desktops, connecting to desktops, maintaining and managing desktops, and licensing. A typical VMware Horizon View environment requires several management components. In many cases, these management services can be installed as desktops and thus do not need separate stand-alone servers. In some cases, such as large-scale deployments, the use of so-called bare-metal management servers is required.

It is recommended 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 purposes, management components should be installed on a cluster that is different from the one that is used for VDI compute nodes.

In practice, a management cluster can be built on an existing vSphere environment that has spare capacity, or you can use more IBM systems to create a new management cluster

To optimize network traffic, keep the provisioning services close to the compute nodes that are running the target virtual machines.

For example, the following virtual machines are needed to host the management components that are on the management cluster (the IBM reference architecture assumes that you run each of these components in virtual machines):

- ▶ vCenter Server
- ▶ vCenter SQL Server
- ▶ View Connection Server
- ▶ View Composer

Management servers have the same hardware specification as VDI compute nodes (for more information, see 4.7.2, “VDI compute node configuration” on page 108) so they can be used interchangeably in a worst-case scenario. The management servers also use ESXi as the hypervisor but have management VMs instead of user VMs.

Table 4-9 summarizes the VM requirements and performance characteristics of each management service.

*Table 4-9 Requirements and performance characteristics for management services*

Management service	Virtual processors	Memory	Storage	Windows Server OS	HA needed	Performance characteristic
vCenter Server VM						
vCenter server	2	4 GB	40 GB	2008 R2	Yes	Up to 2,000 VMs.
vCenter server	16	48 GB	180 GB	2008 R2	Yes	Up to 10,000 VMs.
vCenter SQL Server VM						
vCenter SQL Server	4	4 GB	15 GB	2008 R2	Yes	Up to 2,500 VMs.
vCenter SQL Server	8	8 GB	15 GB	2008 R2	Yes	Up to 10,000 VMs.
View Connection Server VM						
View Connection Server	4	10 GB	70 GB	2008 R2	Yes	Up to 2,000 connections.
View Composer VM						
View Composer	2	4 GB	40 GB	2008 R2	Yes	Up to 2,000 VMs.
View Composer	4	10 GB	50 GB	2008 R2	Yes	Up to 10,000 VMs.

In a Horizon View environment, vSphere high availability (HA) clusters are used to protect from physical server failures. Horizon View 5.2 supports up to 32 ESXi nodes in a cluster. Each vCenter server can handle two clusters of up to 4,000 VMs each (*32 nodes x 125 users per node*), and each cluster exists on two vCenter servers.

Table 4-10 on page 113 lists the number of management VMs for different numbers of users that are based on the high-availability and performance characteristics that are listed Table 4-9.

Table 4-10 Management VMs needed

Management service	600 users	1500 users	4500 users	10000 users
vCenter server	2 (1 + 1)	2 (1 + 1)	3 (2 + 1)	4 (3 + 1)
vCenter SQL Server	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)
View Connection Server	2 (1 + 1)	2 (1 + 1)	4 (3 + 1)	7 (5 + 2)
View Composer	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)	2 (1 + 1)

It is assumed that common services, such as Microsoft Active Directory, DHCP, DNS, and Microsoft licensing servers, exist in the customer environment.

Typically, physical management servers have the same hardware configuration as compute servers. Based on the number and type of VMs that were shown in the previous tables, Table 4-11 lists the suggested number of physical management servers. In all cases, there is redundancy in the management servers and the management VMs.

Table 4-11 Physical management servers needed

Description	600 users	1500 users	4500 users	10000 users
Number of physical servers	2	2	3	4

#### 4.7.4 Shared storage configuration

Experimentation with VDI infrastructures shows that the IOPS performance takes precedence over storage capacity. This means that more of the slower speed drives are needed to get the required performance than higher speed drives. Even with the fastest drives available today (15,000 rpm), there can still be an excess capacity in the storage system because extra spindles are needed to provide the IOPS performance. Typically, 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 large rate of IOPS (and therefore, large number of drives that are needed for dedicated virtual desktops) can be ameliorated to some extent by caching data in flash memory or SSDs. The storage configurations are based on the peak performance requirement, which usually occurs during the so-called “logon storm”. This is when all workers at a company arrive at the same time and try to start their virtual desktops at the same time.

It is always recommended that user data files (shared folders) and user profile data are stored separately from the user image. By default, this must be done for stateless virtual desktops and should also be done for dedicated virtual desktops. It is assumed that 100% of the users at peak load times require concurrent access to user data and profiles.

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. Investigation into the performance shows that 600 GB 10,000 rpm drives in a RAID 10 array give the best ratio of input/output operation performance-to-disk space. It was 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 recommended to use a mixture of both drives for persistent desktops and shared folders/profile data.

If users need more than 5 GB, 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 helps with performance of provisioning virtual desktops that 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 4-12 shows an example scenario of calculating storage capacity for VM images.

Table 4-12 Storage capacity for storing VM images

Description	600 users	1500 users	4500 users	10000 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

In our example scenario, we describe IBM Flex System V7000 Storage Node as a shared storage.

For stateless desktops, the Flex System V7000 storage configuration is summarized in Table 4-13.

*Table 4-13 Flex System V7000 configuration for stateless desktops*

<b>Stateless desktops</b>	<b>600 users</b>	<b>1500 users</b>	<b>4500 users</b>	<b>10000 users</b>
400 GB SSDs in a RAID 1 for master images	2 (1x RAID 1)	2 (1x RAID 1)	4 (2x RAID 1)	8 (4x RAID 1)
Hot spare SSDs	2	2	4	4
600 GB 10,000 rpm HDDs in a RAID 10 for users	12	28	80	168
Hot spare 600 GB HDDs	2	2	4	12
V7000 Control Enclosure	1	1	1	1
V7000 Expansion Enclosure	0	1	3	7

For persistent desktops, the Flex System V7000 storage configuration is summarized in Table 4-14.


*Table 4-14 Flex System V7000 configuration for persistent desktops*

<b>Stateless desktops</b>	<b>600 users</b>	<b>1500 users</b>	<b>4500 users</b>	<b>10000 users</b>
400 GB SSDs in a RAID 1 for master images	2 (1x RAID 1)	2 (1x RAID 1)	4 (2x RAID 1)	8 (4x RAID 1)
Hot spare SSDs	2	2	4	4
600 GB 10,000 rpm HDDs in a RAID 10 for users	12	28	80	168
Hot spare 600 GB HDDs	2	2	4	12
300 GB 15,000 rpm 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
V7000 Control Enclosure	1	1	2	4
V7000 Expansion Enclosure	2	6	16 (2 x 8)	36 (4 x 9)

It is common to cluster multiple Flex System 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 such a case, two more physical management nodes are added to the solution and Windows Storage Server is deployed on them in a highly available cluster.





# IBM Flex System and VMware View lab environment

In this chapter, we describe how the environment is structured and the implementation plan for VMware Horizon View 5.2 in the ITSO lab.

The lab setup that is described in this chapter shows the main infrastructure components that are applied to the production VDI environments.

This chapter includes the following topics:

- ▶ Lab environment
- ▶ VMware View solution overview
- ▶ IBM Flex System chassis overview
- ▶ Storage configuration overview
- ▶ Network configuration overview
- ▶ VDI solution planning

## 5.1 Lab environment

In this section, we describe the physical environment and the software components that are used in the implementation of VMware View on IBM PureFlex Systems.

The IBM Flex Systems consist of the following components:

- ▶ An IBM Flex Enterprise Chassis
- ▶ Two Chassis Management Modules (CMMs)
- ▶ One IBM Flex System Manager (FSM) node for management purposes
- ▶ Four IBM Flex x240 compute nodes that are equipped with VMware ESXi 5.1 embedded
- ▶ An IBM Flex System V7000 Storage Node (FC based) that is used for shared storage
- ▶ An IBM Flex System Fabric EN4093 10Gb Ethernet Scalable Switch for Ethernet networking
- ▶ An IBM Flex System FC3171 8Gb SAN Switch for storage networking

Table 5-1 describes the software components that are used in the lab and their roles.

*Table 5-1 Software components*

Software Component	Description
VMware ESXi 5.1	Hypervisor that is used to build the VMware View solution.
VMware vCenter 5.1	Used to manage VMware ESXi servers and monitor the health of virtual machines.
Windows 2008 R2	Base operating system for the VMware View infrastructure.
Windows 7	Operating system for the virtual desktops.
SQL Server 2008 R2	Database server that is used to store View Composer DB.
VMware Horizon View 5.2 Composer	VMware View component that is used for linked-clone desktops from a centralized base image.
VMware Horizon View 5.2 Connection server	VMware View component that is used for user authentication. It acts as a connection broker to redirect connection to the user's appropriate virtual desktop.
VMware Horizon View 5.2 Agent	Agent that is used to permits user to access their virtual desktop.
VMware Horizon View 5.2 Client	Client that is used to connect to View desktop.

Software Component	Description
VMware Horizon View 5.2 Administrator	Web-based administration console that is used to manage virtual desktops

## 5.2 VMware View solution overview

The term *component model* describes the main components that are part of the lab environment.

Component model offers a high-level point of view that is useful to locate the exact position of each component that is related to its location in the infrastructure.

The two VMware clusters interact with the same network and storage stacks.

Management cluster contains all infrastructural server roles that are related to Active Directory, file services, vCenter, and VMware Horizon View main component.

VDI cluster contains all virtual desktops pools and the base OS images.

Client devices and infrastructure servers are on the same VLAN and can communicate with each other.

The component model of the virtual desktop infrastructure is shown in Figure 5-1.

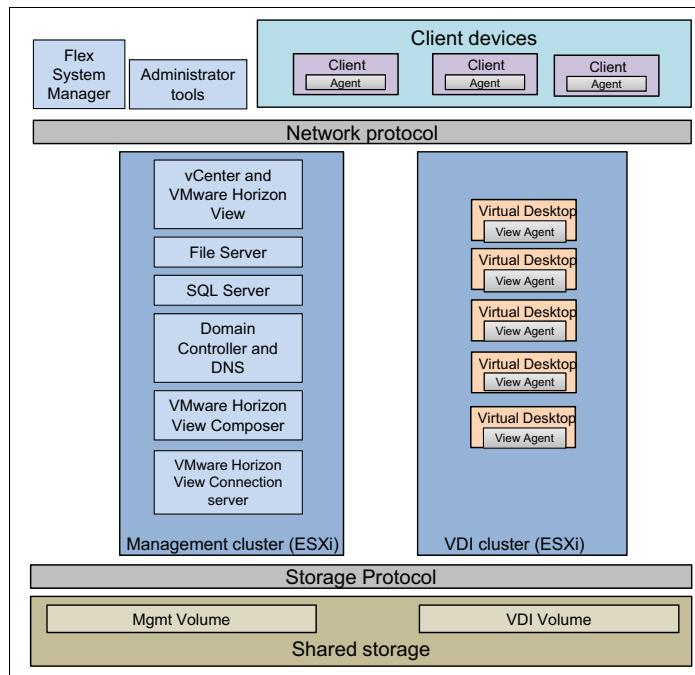



Figure 5-1 VDI component model

## 5.3 IBM Flex System chassis overview

Figure 5-2 shows all of the components on an IBM Flex System Enterprise Chassis. It integrates compute nodes, storage, Ethernet, and SAN switches in a single box.

This chassis is configured and managed by the Flex System Manager node.

Figure 5-2 also shows the IBM Flex System that is used in the lab environment.



*Figure 5-2 Front view of the IBM Flex Systems Chassis that is used for the lab*

The rear of the chassis shows two CMM modules: Ethernet and SAN switches.

Those modules are used to manage internal chassis communications and interact with the external network, SAN, or other external infrastructure services.

Figure 5-3 shows the rear view of the chassis and the exact position of each component.




Figure 5-3 Rear view of the IBM Flex Systems Chassis that is used for the lab

Figure 5-4 shows the physical view of the network that is between the components.

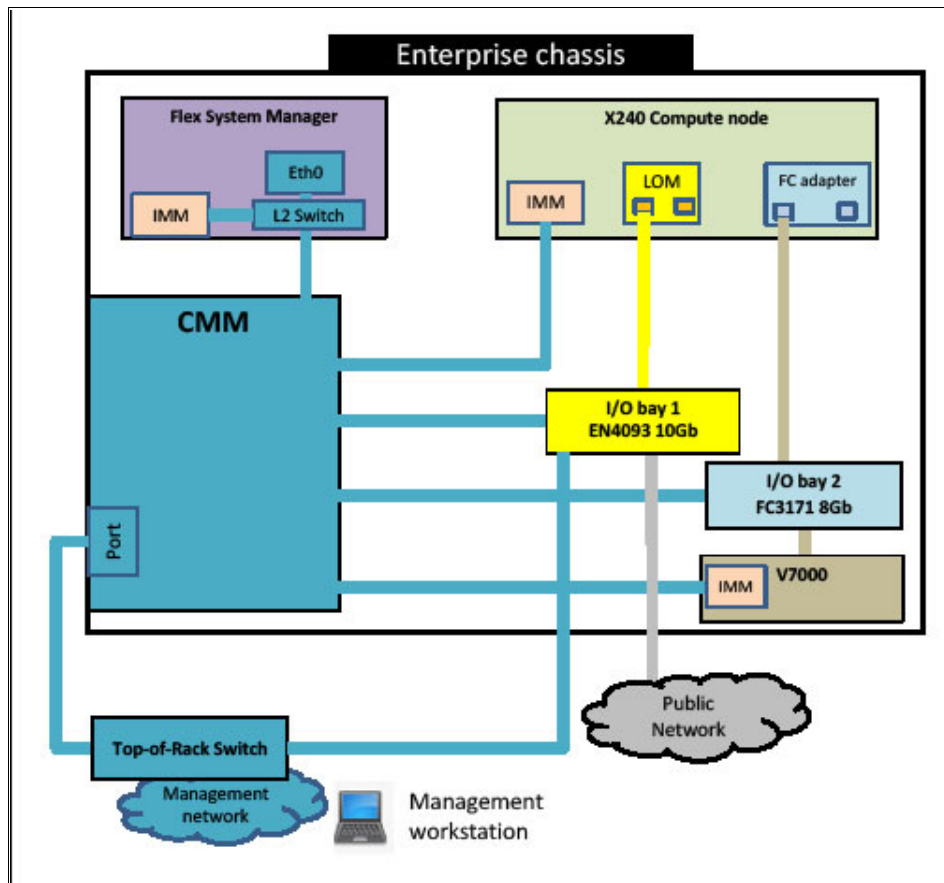


Figure 5-4 Physical view of network that is between components

## 5.4 Storage configuration overview

We create the MDisk on the IBM Flex System V7000 and it is then divided into two different volumes with the Thin Provision option enabled.

The volumes are created for the following purposes:

- ▶ Mgmt Volume: One volume to store all infrastructure servers and VMware View components.
- ▶ VDI Volume: One volume that is dedicated to store the provisioned desktops and the desktop pools.

The volumes are shown in Figure 5-5.

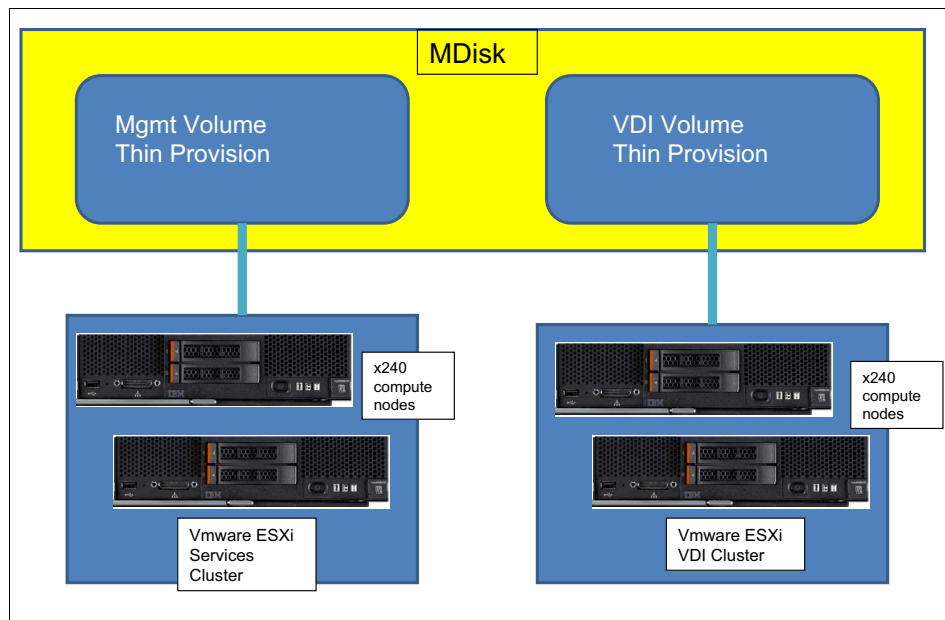


Figure 5-5 Graphical storage structure

From a VMware vSphere point of view, each volume represents a single, per cluster, VMFS formatted datastore.



## 5.5 Network configuration overview

In this section, we describe the different networks flows (internal and external) between Flex components and software components, administrators, and users.

An Ethernet segmentation in the meaning of VLANs tagging is configured to split the traffic according to the software component requirements by using the following perspective:

- ▶ Management VLAN (VLAN 42)

The management VLAN allows the technical support team to connect to the environment for management purposes. Externally from the switch, VLAN 42 has a VLAN ID of 42; internally, the vNIC Group to which it belongs has a private VLAN ID of 127. It connects all Flex components (FSM, compute nodes, storage, and switches) and VMware ESXi hosts. For security reasons, management traffic is not shared with the user access segment (Public/Access).

- ▶ Kernel/VMotion VLAN (VLAN 10)

The Kernel/VMotion VLAN is used for VMware ESXi operation; this VLAN is responsible for allowing the virtual machines to be transferred from one physical node to another in case of maintenance or hardware failure. Externally from the switch, VLAN 10 has a VLAN ID of 10; internally, the vNIC Group to which it belongs has a private VLAN ID of 128. This internal VLAN is defined only on the VMware ESXi internal DvSwitch; however, the vNIC Group to which it belongs has a physical link to one external switch port. If the IBM Flex System is placed in an existing environment, it allows VMs to be migrated on chassis-external hosts.

- ▶ Public/Access Network (VLAN 20)

The Public/Access Network segment is used for user access. This VLAN is available for clients access their virtual desktops, active directory authentication, database interactions, and access and for VMware View-specific transactions. Externally from the switch, VLAN 20 has a VLAN ID of 20; internally, the vNIC Group to which it belongs has a private VLAN ID of 128. As with the VLAN 10, this VLAN is defined only on the ESXi internal DvSwitch. The vNIC Group to which this connection belongs has a link to an external port trunk because of its heavy load.

Each compute node has a maximum of four vNICs.

Figure 5-6 shows how it is named and the switch internal connections.

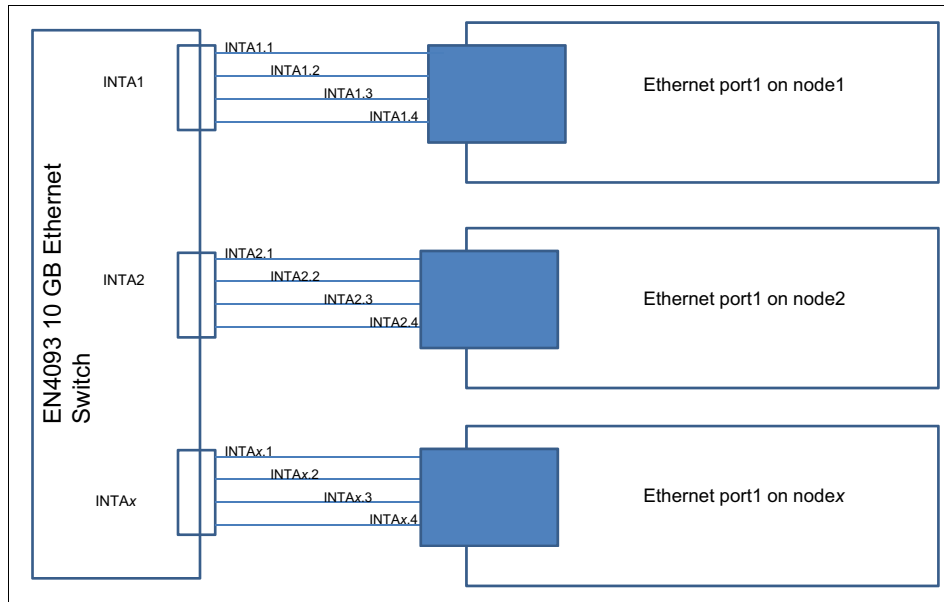


Figure 5-6 Compute nodes to switch internal vNIC names and connections

Figure 5-7 show the logical view of the connections.

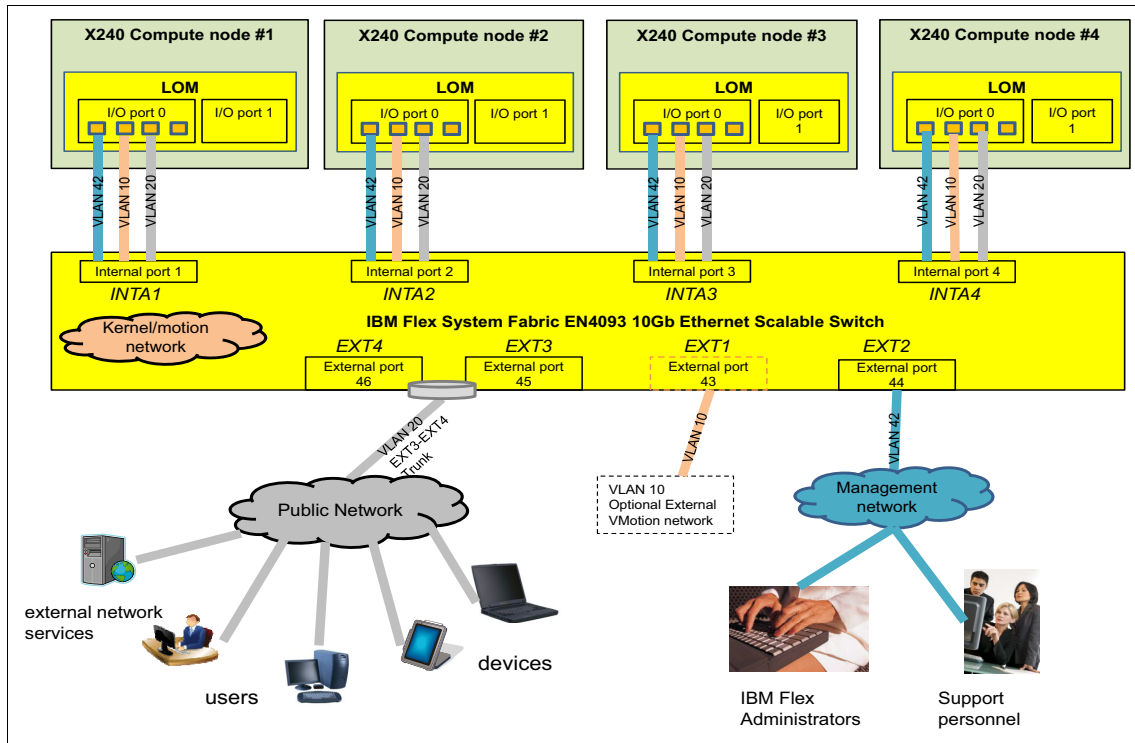


Figure 5-7 Connections logical view

Table 5-2 shows the bandwidth that is allocated on any port-related VLAN.

Table 5-2 Adapter bandwidth allocation

VLAN	vNIC	Bandwidth allocation
VLAN 42	INTAx.1	25%
VLAN 10	INTAx.2	25%
VLAN 20	INTAx.3	50%

Bandwidth allocation is part of the steps to virtualize the network adapters, where virtual network interface cards (vNICs) are created and then presented to the hosts as traditional adapters that are configured at your own VLAN.

Bandwidth allocation is configured at the IBM Flex System Fabric EN4093 Scalable Network Switch level to the specified vNICs that forms the related Port Group.

Because the Management Cluster is not expected to grow, it uses standard vSphere network switches. The standard switches are used for the management port group (VLAN 42), Public VLAN 20, and vMotion VLAN 10.

To have a consistent network configuration across all VDI hosts and to ease future scalability, a single distributed virtual switch (dvSwitch) is created for VDI Cluster VM traffic that contains all networks for VLAN 42, VLAN 20, and VLAN 10 for vMotion.

## 5.6 VDI solution planning

Flex System Manager (FSM) is used to manage the following Flex System components:

- ▶ Compute nodes
- ▶ Shared storage
- ▶ Network and storage switches

FSM is also used to create the patterns that are applied in the computer nodes to standardize the configuration characteristics and accelerate the deployment.

The virtual desktop infrastructure is distributed across two different VMware clusters within the same data center.

This allows for segmentation of the resource usage and to align with a standard pattern that is deployed in production environments in which the following clusters have a specific purpose:

- ▶ Management VMware cluster

This two-node cluster contains all of the infrastructure's servers that are used to deliver the essential VMware View services.

This cluster is based on two compute nodes and contains Active Directory server, SQL server, View Connection server, View Composer server, View Administrator Console, and VMware vCenter.

**Note:** In this lab, the vCenter is a virtual machine that is in the same cluster it manages.

In this two-node cluster configuration, if the node that hosts the vCenter server fails, some manual intervention must be done to migrate the vCenter virtual machine on the running instance of VMware ESX unless VMware HA Admission Control is enabled and configured at least 50% of cluster resources to be available.

► VDI VMware cluster

This two-node cluster contains all of the virtual desktops and all Desktop pools.

The shared storage that is used by the clusters is provided by the IBM Flex System V7000.

Each compute node has its local storage that is used to cache virtual machine stateless disks.

The following SAN volumes are presented to physical hosts:

- One volume that is to be used as a datastore for the Management cluster
- One volume that is to be used as a datastore for the VDI cluster

The network traffic is split on different VLANs, which are managed by the IBM Flex System Fabric EN4093 10Gb Ethernet Scalable Switch. It also provides external connectivity for client device connections.

The storage connections are managed by the IBM Flex System FC3171 8Gb SAN Switch. Different storage zones are created for each ESXi node.

Table 5-3 shows the relationship between the software components and the infrastructure servers.

*Table 5-3 Software components that are installed on each server*

<b>Server role</b>	<b>Installed operating system and software component</b>	<b>Dedicated server</b>
ESXi Server	VMware ESXi 5.1	Yes
vCenter and VMware Horizon View 5.2 Administrator	<ul style="list-style-type: none"> <li>▶ Windows 2008 R2</li> <li>▶ VMware vCenter 5.1</li> <li>▶ VMware Horizon View 5.2 Administrator</li> <li>▶ VMware Horizon View 5.2 Composer</li> </ul>	Yes
SQL Server	<ul style="list-style-type: none"> <li>▶ Windows 2008 R2</li> <li>▶ SQL 2008 R2 Server</li> </ul>	Yes
Active Directory Domain Controller DNS and DHCP	<ul style="list-style-type: none"> <li>▶ Windows 2008 R2</li> <li>▶ Domain Controller</li> <li>▶ DHCP and DNS AD-integrated zones</li> </ul>	Yes
File Server	<ul style="list-style-type: none"> <li>▶ Windows 2008 R2</li> <li>▶ File Server role enabled</li> </ul>	Yes
VMware Horizon View 5.2 Connection server	<ul style="list-style-type: none"> <li>▶ Windows 2008 R2</li> <li>▶ VMware Horizon View 5.2 Connection server</li> </ul>	Yes
VMware Horizon View 5.2 Client agent	<ul style="list-style-type: none"> <li>▶ Windows 7 x64 Professional</li> <li>▶ VMware Horizon View 5.2 Client agent</li> </ul>	N/A

## 5.6.1 Management Cluster component model

Management Cluster is formed by using x240\_node\_1 and x240\_node\_2 IBM Flex System compute nodes. It has access to a management shared volume that is defined on the IBM Flex System V7000 Storage Node and contains the following virtual machines:

- ▶ Domain Controller
- ▶ DNS DHCP
- ▶ File Server
- ▶ SQL Server
- ▶ vCenter and View Composer
- ▶ View Connection Server

You should consider adding more virtual disks to specific VMs (for example, the SQL Server) to separate DB files from the log files.

Table 5-4 shows the VM role and its functionality in the VMware View implementation.

*Table 5-4 VM roles and functionality*

VM role	Function	VLAN
Domain Controller DNS & DHCP	Provides authentication services to users, Group Policy settings, network name resolution, and IP addresses to View Desktops	Public VLAN 20
File Server	Stores user's profile that is used by VMware View Persona Manager	Public VLAN 20
SQL Server	Contains the vCenter and View Composer main databases	Public VLAN 20
VMware vCenter	Infrastructure's main management console	Public VLAN 20 Management VLAN 42
VMware View Composer	Installed on vCenter server, creates linked clones from a parent VM	Public VLAN 20
VMware View Connection server	Point of contact for client devices that are requesting virtual desktops. View Connection Server authenticates users and directs the request to the appropriate VM or desktop. After the authentication is complete, users are directed to their assigned VM or desktop.	Public VLAN 20

VM role	Function	VLAN
VMware View Administrator Console	Web-based application allows administrators to configure View Connection Server, deploy and manage View desktops, control user authentication and troubleshoot user issues. It is installed during the installation of View Connection Server.	Public VLAN 20

Active Directory represents a crucial part of the VMware View implementation; therefore, you should accurately configure domain controller replicas and provide a full redundancy of Active Directory domain controllers.

A specific root VDI organizational unit (OU) that contain two other sub-OU's Users and Computers are needed to organize users and virtual desktops.

Specific Group Policy Objects (GPOs) are needed to centrally assign users permission to connect to specified desktops.

For more information, see Chapter 7, "Deploying VMware Horizon View infrastructure" on page 277.

## 5.6.2 VDI Cluster component model

VDI Cluster is formed by using x240\_node\_3 and x240\_node\_4 IBM Flex System compute nodes. It has access to a VDI shared volume that is defined on IBM Flex System V7000 Storage Node and contains all Virtual Desktops and all Virtual Desktop Pools.

The following Virtual Desktop Pools are available:

- ▶ Full Virtual Machine (FVM) desktop pool
 

This pool generates a virtual desktop that is based on an existing VM template. FVM virtual desktops user assignments can be configured in the following distinct ways:

  - Dedicated-assignment pool: Each user is assigned a virtual desktop at the first-time login. The same desktop is assigned when the same user logs in.
  - Floating-assignment pool: Each user receives a different virtual desktop from the pool at login.
- ▶ Linked-Clone (LCVM) desktop pool: This pool generates a virtual desktop that is based on a running VM snapshot, which is also called *parent VM*. LCVM desktops share with the parent VM the base operating system disk. As a result, each linked-clone generated virtual desktop uses less hard disk drive disk space than the FVM.



LCVM virtual desktops user assignments can be configured in the following two distinct ways:

- Dedicated-assignment pool: Each user is assigned a virtual desktop at the first-time login. The same desktop is assigned when the same user logs in.
- Floating-assignment pool: Each user receives a different virtual desktop from the pool at login.

For each pool at creation time, the Virtual Desktop retention also can be defined and the action to perform when the user logs off; for example, to switch off the virtual desktop, delete the virtual desktop, or reset it to default to be ready for the next assignment.

Desktop Pools often are based on user-type.

Some users might require a full permanent desktop assignment where they can also install more applications. For other types of users, a linked-clone desktop in floating-assignment is enough for their daily work.

In our lab, we configure the following Desktop Pools:

- ▶ FVM with dedicated assignment for power users
- ▶ LCVM with floating assignment for standard users (task workers)

For LCVM desktops, local compute node’s SSDs can be used to store the linked clones desktops for improved performance. View Storage Accelerator manages of this and is enabled by default on each Desktop Pool.

Two replicas must be stored for each master image. Each LCVM virtual desktop requires a linked clone, which tends to grow over time until it is refreshed at log out. Because of the stateless nature of the architecture, there is no need to configure SSDs in a redundant RAID 1 configuration.

Each virtual desktop runs Microsoft Windows 7 x64 Professional and is configured as shown in Table 5-5.

Table 5-5 Virtual Desktop Configuration

User type	Virtual CPU	Amount of vRAM	Virtual Desktop type
Power user	4 vCPU	8 GB	Full Virtual Machine
Standard user	2 vCPU	4 GB	Linked-Clone Virtual Machine

Standard user's profiles are managed by VMware View Persona Manager, which is responsible to preserve user's profiles in a centralized File Server store.

Unlike classical Roaming Profiles, VMware View Persona Manager permits synchronizing the entire profile and provides files to the user only when the files are needed.

Power users do not need any profile management scenario because their virtual desktops are in dedicated assignment.

### **5.6.3 Desktop pool consideration**

As described in 5.6.2, "VDI Cluster component model" on page 132, with View Manager, we can create pools of desktops that deliver desktop access to users. In Chapter 9, "Operating VMware Horizon View infrastructure" on page 373, we configure two specific desktop pools for power users and for standard users.

Each Desktop Pool is configured to have a minimum of three VMs to a maximum of 20 VMs.

Because of the permanent assignment, no refresh is done on virtual desktops for the power users. For standard users, the VM is available at logoff to other users for login.

View Storage Accelerator is enabled by default on each desktop pool and enables ESXi hosts to locally cache virtual machine disk data.

This feature can reduce IOPS and improve performance during boot storms, or when many desktops are using hard disk drive-intensive applications at once.



## Deploying IBM Flex System

In this chapter, we describe how to fully deploy IBM Flex System and its components in the lab environment that is described in Chapter 5, “IBM Flex System and VMware View lab environment” on page 117.

The lab environment operates on an IBM Flex System Enterprise chassis. The four x240 compute nodes have the following components installed:

- ▶ Embedded 10Gb Virtual Fabric Ethernet Controller
- ▶ VMware ESXi 5.1 embedded hypervisor
- ▶ FC3052 8Gb Fibre Channel Adapter

The network is operated by the Flex System Fabric EN4093 10GB Ethernet switch. Compute nodes are connected to three different VLANs each. IBM Flex System V7000 Storage Node presents volumes to x240 compute nodes.

Each x240 compute node is part of a dedicated VMware cluster (Management or VDI) and has access to only one volume that is presented by IBM Flex System V7000 Storage Node.

The Management Cluster is based on VMware vSphere ESXi 5.1 and contains all of the infrastructure-related virtual servers on its own data store. The VDI Cluster is based on VMware ESXi 5.1 and contains all of the Virtual Desktop and Virtual desktop pools on its own data store.

This chapter includes the following topics:

- ▶ Initial configuration of Chassis Management Module
- ▶ IBM Flex System Manager Setup wizard
- ▶ Selecting a chassis to manage
- ▶ Discovery and inventory collection
- ▶ IBM Flex System Fabric EN4093 10Gb configuration
- ▶ IBM Flex System x240 compute node configuration
- ▶ IBM Flex System V7000 Storage Node configuration
- ▶ VMControl activation

## 6.1 Initial configuration of Chassis Management Module

In this section, we describe how to initially configure the Chassis Management Module (CMM) to enable chassis management tasks.

### 6.1.1 Connecting to the CMM


The CMM is configured with a static IP address but can respond to DHCP offers first before it uses its static pre-configured IP address. You can cable the CMM to support a management connection that best matches your site configuration. For the initial setup configuration, you connect an external client system to the CMM to configure and manage the IBM Flex System Enterprise Chassis.

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

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

Complete the following steps to connect to the CMM:

1. Ensure that the subnet of the client computer that is used to initially configure is set to the same value as in the CMM (the default CMM subnet is 255.255.255.0). The IP address of the CMM also must be in the same local network 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.  
You can also point-to-point connect the client computer to the CMM by using an Ethernet cable.
2. Open a web browser on the client computer and enter the CMM IP address. For the first connection to the CMM, use the default IP address of the CMM, as shown in Figure 6-1.



*Figure 6-1 Initial CMM IP address*

**Note:** The CMM has the following default settings:

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

3. In the CMM window that is shown in Figure 6-2, log in to the CMM by using the default credentials USERID/PASSWORD. Click **Log in**.



Figure 6-2 CMM login window

The CMM main window opens, as shown in Figure 6-3.

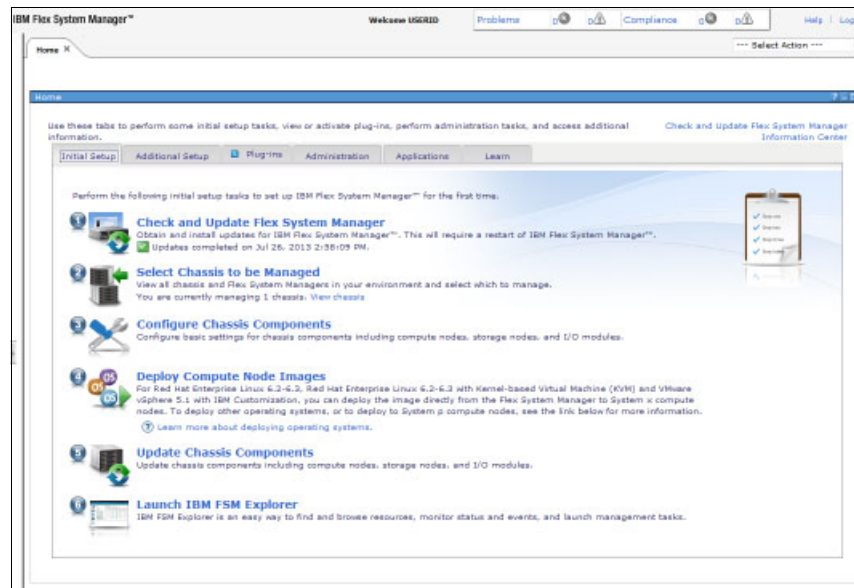


Figure 6-3 CMM main window

## 6.1.2 Using the initial setup wizard

For the initial configuration of the CMM, use the initial setup wizard. The initial setup wizard can help you configure the CMM through 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 start the initial setup wizard manually and perform the initial configuration:

1. From the CMM web interface home window, click **Mgt Module Management**, as shown in Figure 6-4.

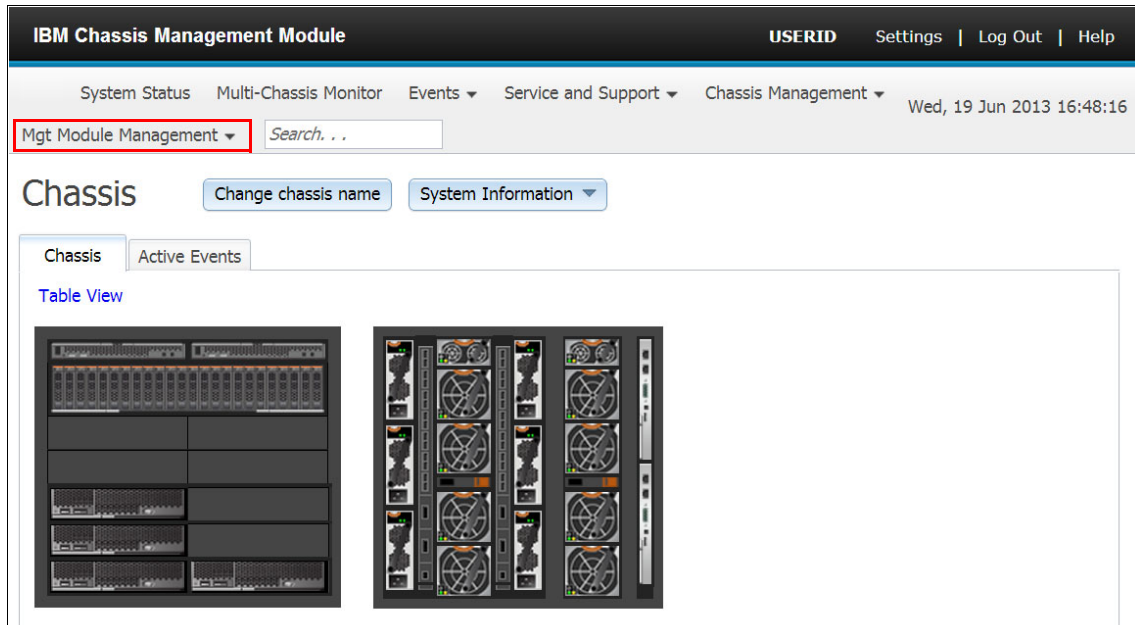


Figure 6-4 CMM main window: Mgt Module Management



The initial setup wizard is in the Configuration menu, as shown in Figure 6-5.

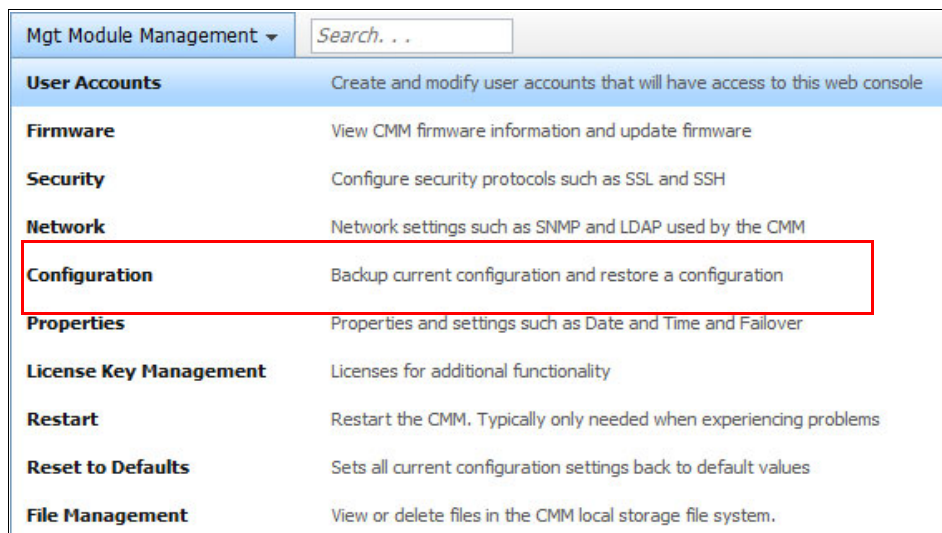


Figure 6-5 Mgt Module Management window

Several options are displayed for managing the CMM configuration.

2. For the first-time connection, click **Initial Setup Wizard**, as shown in Figure 6-6.




Figure 6-6 Manage Configuration window

- When the wizard starts, the first window displays the steps on the left side of the window that are performed. The basic description of the steps is displayed in the main part of the window.

Figure 6-7 shows the Welcome window of the setup wizard. This wizard is similar to other IBM wizards. Navigation buttons for the wizard are in the lower left corner of each window. Click **Next**.

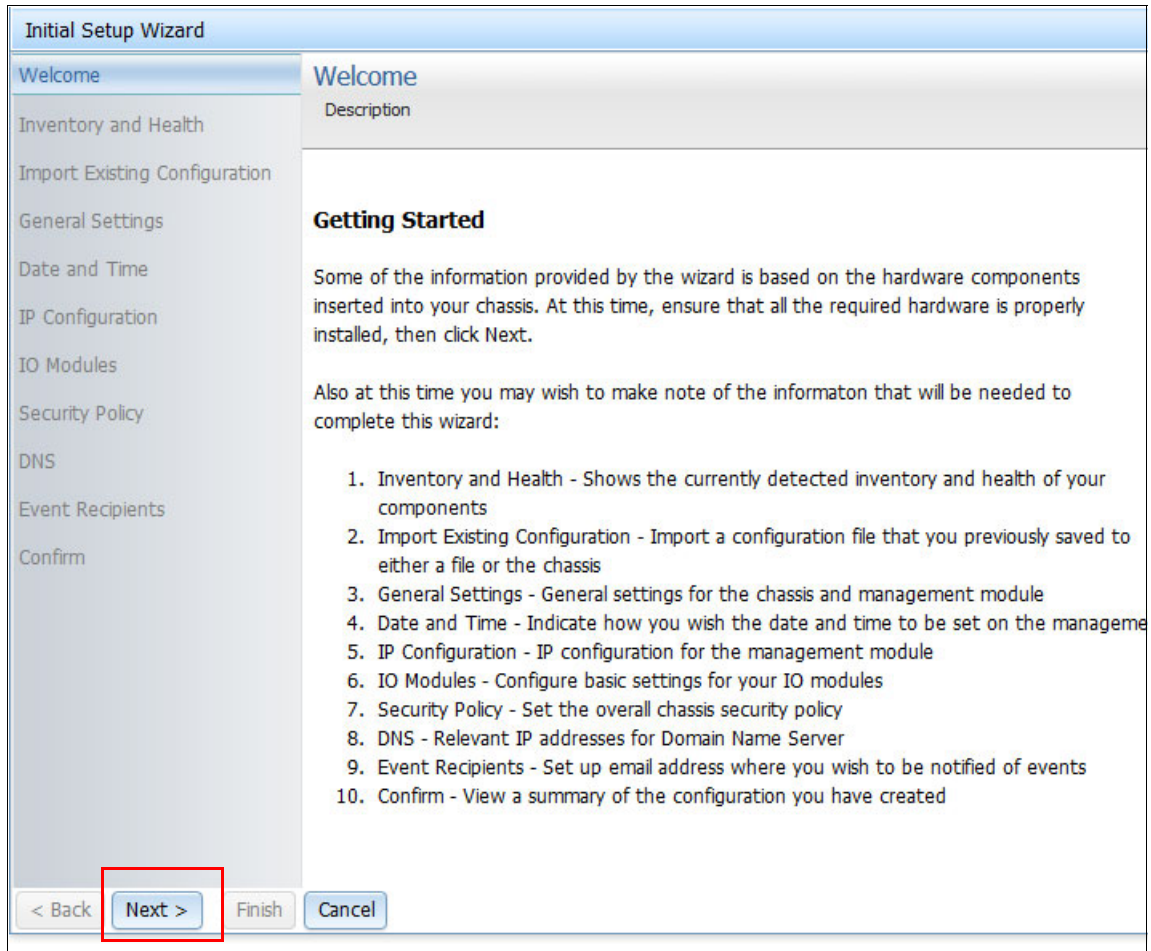
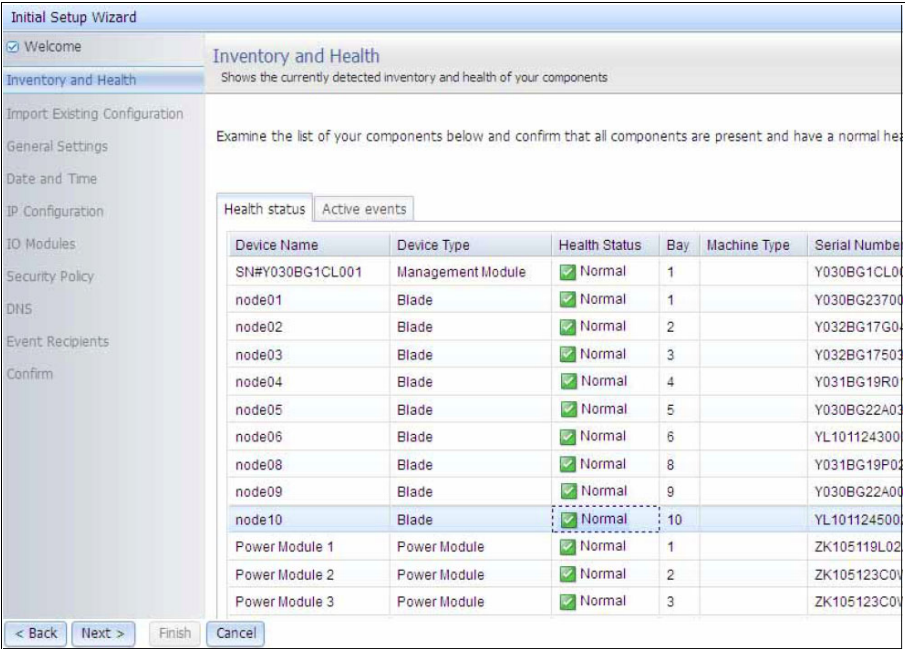


Figure 6-7 Welcome window

4. Select the **Health status** tab on the Inventory and Health window to view the detected components in the Chassis and their current health status, as shown in Figure 6-8. 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 and have a normal health status.

**Health status** | Active events

Device Name	Device Type	Health Status	Bay	Machine Type	Serial Number
SN#Y030BG1CL001	Management Module	✓ Normal	1		Y030BG1CL001
node01	Blade	✓ Normal	1		Y030BG23700
node02	Blade	✓ Normal	2		Y032BG17G04
node03	Blade	✓ Normal	3		Y032BG17503
node04	Blade	✓ Normal	4		Y031BG19R01
node05	Blade	✓ Normal	5		Y030BG22A03
node06	Blade	✓ Normal	6		YL101124300
node08	Blade	✓ Normal	8		Y031BG19P02
node09	Blade	✓ Normal	9		Y030BG22A00
node10	Blade	✓ Normal	10		YL101124500
Power Module 1	Power Module	✓ Normal	1		ZK105119L02
Power Module 2	Power Module	✓ Normal	2		ZK105123C0V
Power Module 3	Power Module	✓ Normal	3		ZK105123C0V

< Back | Next > | Finish | Cancel

Figure 6-8 Inventory and Health window

5. If you saved a configuration file, you can select the file by using the Import Existing Configuration window. It automatically enters the appropriate values in the fields of the wizard, as shown in Figure 6-9. This example shows a first configuration. So, you can ignore this window and click **Next**.

The screenshot shows a software window titled "Initial Setup Wizard". On the left is a vertical sidebar with a list of steps: "Welcome", "Inventory and Health", "Import Existing Configuration" (which is highlighted with a blue background), "General Settings", "Date and Time", "IP Configuration", "IO Modules", "Security Policy", "DNS", "Event Recipients", and "Confirm". The main area of the window has a title bar "Import Existing Configuration" and contains the following text: "To facilitate your task of setting up the management module, you can import a configuration file that you previously saved to a chassis. Importing a configuration will automatically fill in the fields of this wizard with the appropriate values." Below this, it says: "If this is your first time setting up a chassis, you will not have a configuration file to import. These files are useful for your management module settings, or for configuring multiple chassis. To create a configuration file, you can use the console under Mgt Module Management -> Configuration." There is a section labeled "Upload configuration file:" with a text input field and a "Browse for file" button. Below that is a checkbox labeled "Decode with passphrase" which is currently unchecked. At the bottom of the window are four buttons: "< Back", "Next >" (which is highlighted with a dashed border), "Finish", and "Cancel".

Figure 6-9 Import Existing Configuration window

6. In the General Settings window, enter some descriptive information about the Chassis, including location and contact person, as shown in Figure 6-10. Click **Next**.

The screenshot shows the 'Initial Setup Wizard' window. On the left is a sidebar with the following options: Welcome (checked), Inventory and Health (checked), Import Existing Configuration (checked), General Settings (selected), 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 the following fields:

Management module name	SN#Y030BG1CL001
Chassis description	
Contact person	No Contact Configured
Chassis location	No Location Configured
Room ID	
Rack ID	
Lowest U-position	0
Unit height of chassis	10

At the bottom of the window are four buttons: '< Back', 'Next >' (highlighted with a dotted border), 'Finish', and 'Cancel'.

Figure 6-10 General Settings window

7. Set the date and time for the CMM in the Date and Time window, as shown in Figure 6-11. There are two options to sync the time: by using NTP or setting it manually. 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

**Date and Time**  
Date and time settings for the management module

Indicate how you wish the date and time to be set on the management module. The management module date and time values the event log, for example.

Select method:

NTP server host name and/or IP address:

Synchronization frequency (minutes):

☒ Enable NTP v3 Authentication

NTP v3 Authentication key index:

NTP v3 Authentication key (M - MD5):

NTP last updated the clock on 06/12/2012 23:40:20 by 0 s.

GMT Offset:

☒ Automatically adjust for daylight savings time (DST)

< Back   **Next >**   Finish   Cancel

Figure 6-11 Date and Time window

8. IBM Flex System has two CMMs. Each CMM is configured with the same static IP address. Use the IP Configuration window that is shown in Figure 6-12 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 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' (MM5CF3FC25E3B7), 'Domain name', and a checkbox for 'Register this interface with DNS'. Below these are tabs for 'IPv4' and 'IPv6'. The 'IPv4' tab is active, showing 'Currently assigned IPv4 address information' with values: IP address: 9.27.20.56, Subnet mask: 255.255.252.0, and Default gateway: 9.27.20.1. It also shows 'IP address assignment methods:' set to 'Use static IP address'. Under 'Static IP Address Settings', there is a warning '\*Changing settings requires a CMM restart.' and fields for 'Static address: 9.27.20.56', 'Subnet mask: 255.255.252.0', and 'Default gateway: 9.27.20.1'. At the bottom are buttons for '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 6-12 IPv4 configuration window

9. If you need to set up IPv6, select the **IPv6** tab, as shown in Figure 6-13. 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

**IP Configuration**  
IP configuration for the management module

Host name: MM5CF3FC25E3B7  
Domain name:   
Register this interface with DNS: ☐

IPv4 | **IPv6**

☒ Enable IPv6

Link local address: fe80::5ef3:fcff:fe25:e3b7

Stateless address:	IP Address	Prefix Length
	fd55:faaf:e1ab:1015:5ef3:fcff:fe25:e3b7	64

Default gateway: 0::0  
Stateful address:

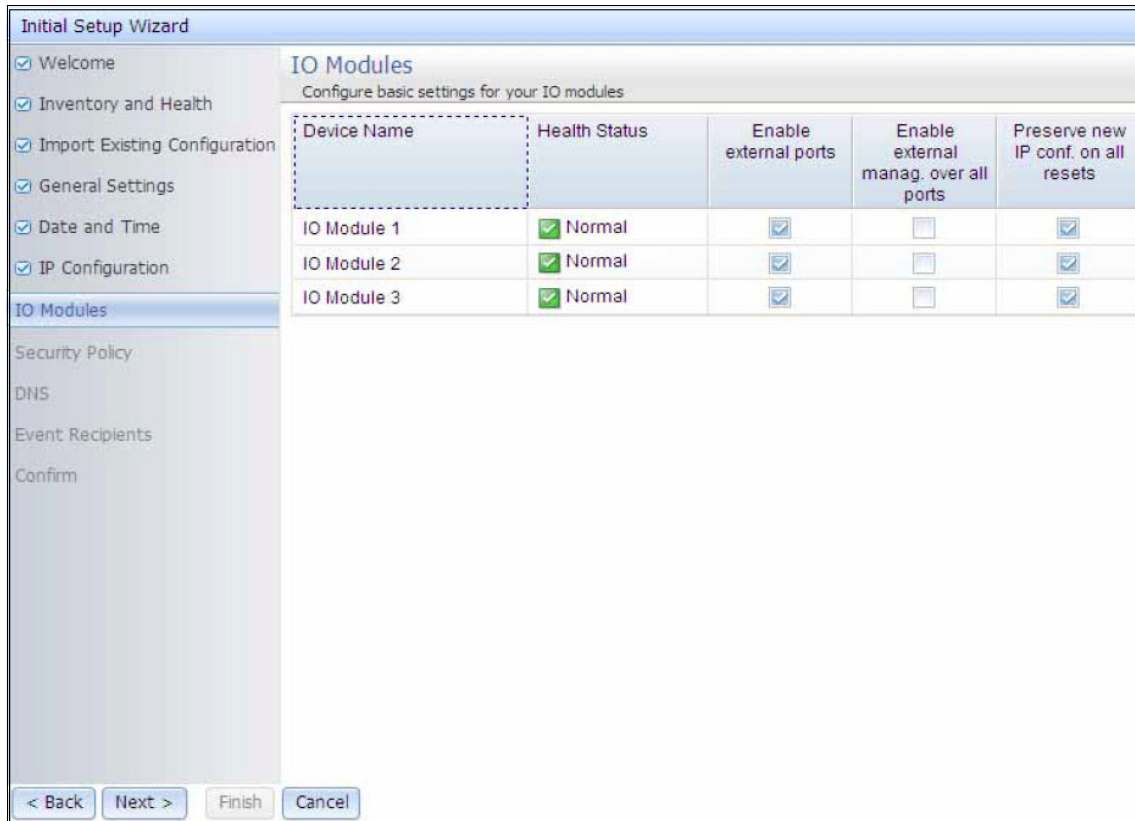
IP address assignment methods:  
☒ Use stateless address autoconfiguration  
☒ Use stateful address configuration (DHCPv6)  
☐ Use statically assigned IP address

< Back | **Next >** | Finish | Cancel

Figure 6-13 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 6-14. Click **Next**.



The screenshot shows the 'Initial Setup Wizard' window with the 'IO Modules' step selected. The left sidebar lists the steps: Welcome, Inventory and Health, Import Existing Configuration, General Settings, Date and Time, IP Configuration, IO Modules (selected), Security Policy, DNS, Event Recipients, and Confirm. The main area is titled 'IO Modules' and contains the instruction 'Configure basic settings for your IO modules'. Below this is a table with five columns: Device Name, Health Status, Enable external ports, Enable external manag. over all ports, and Preserve new IP conf. on all resets. The table lists three IO modules, all with a 'Normal' health status. The 'Enable external ports' column has checked checkboxes for all three modules. The 'Enable external manag. over all ports' column has unchecked checkboxes for all three modules. The 'Preserve new IP conf. on all resets' column has checked checkboxes for all three modules. At the bottom of the window are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Device Name	Health Status	Enable external ports	Enable external manag. over all ports	Preserve new IP conf. on all resets
IO Module 1	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
IO Module 2	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
IO Module 3	✓ Normal	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 6-14 I/O Modules window

11. Select the security policy for the CMM, as shown in Figure 6-15. Click **Next**.

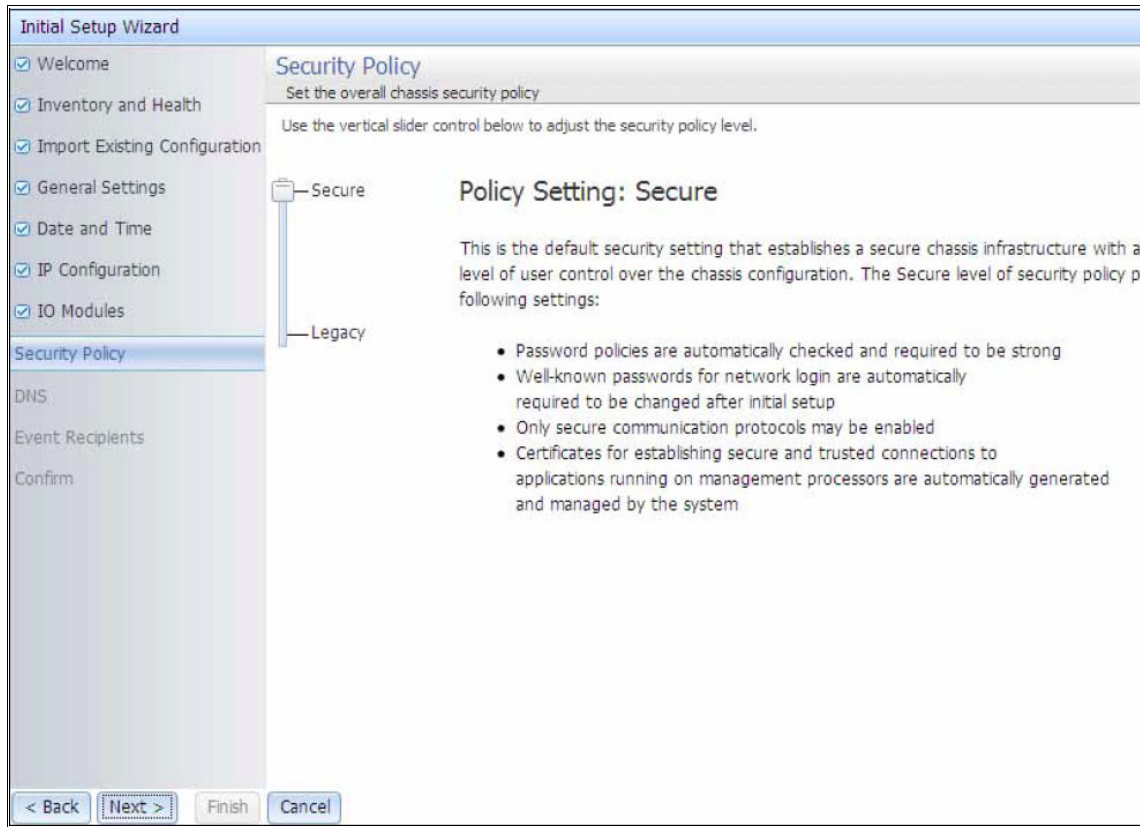
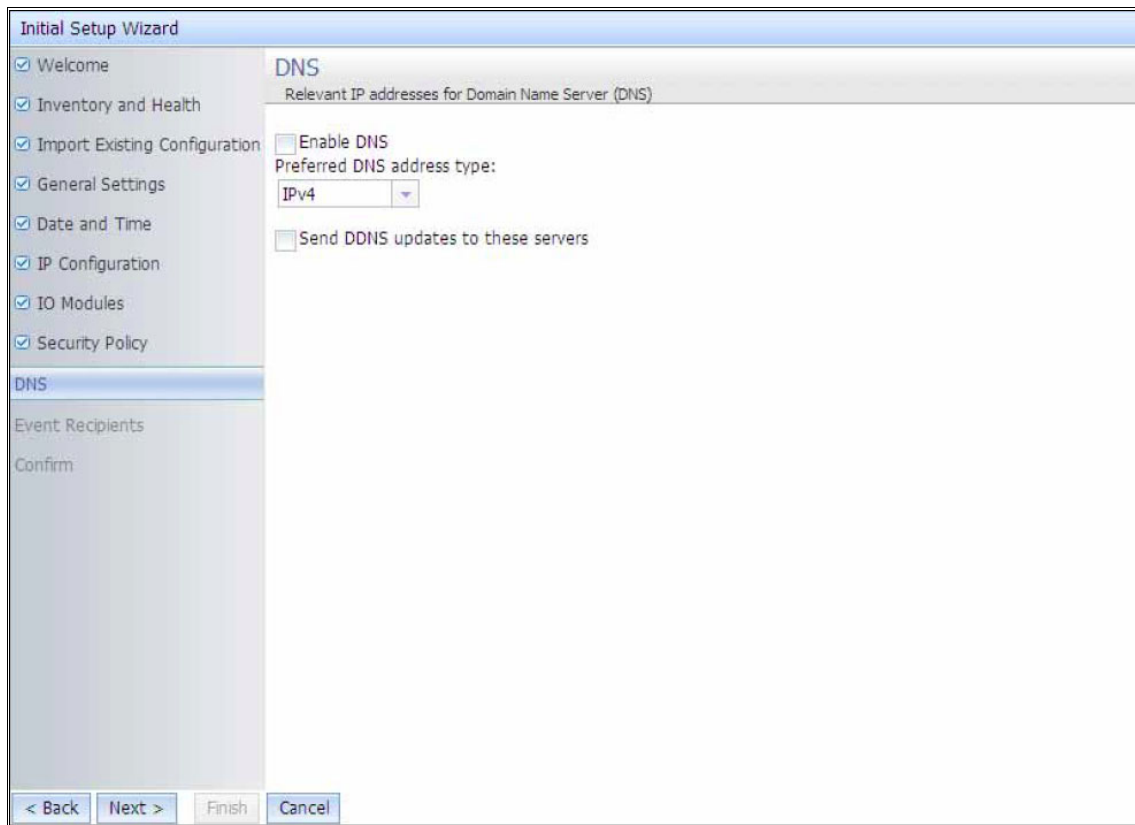


Figure 6-15 Security Policy window

**Important:** When the CMM is set to the Secure mode, you can use secure file transfer methods only (such as HTTPS and SFTP) 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, such as HTTP, FTP, and TFTP, are disabled when security is set to the Secure mode.

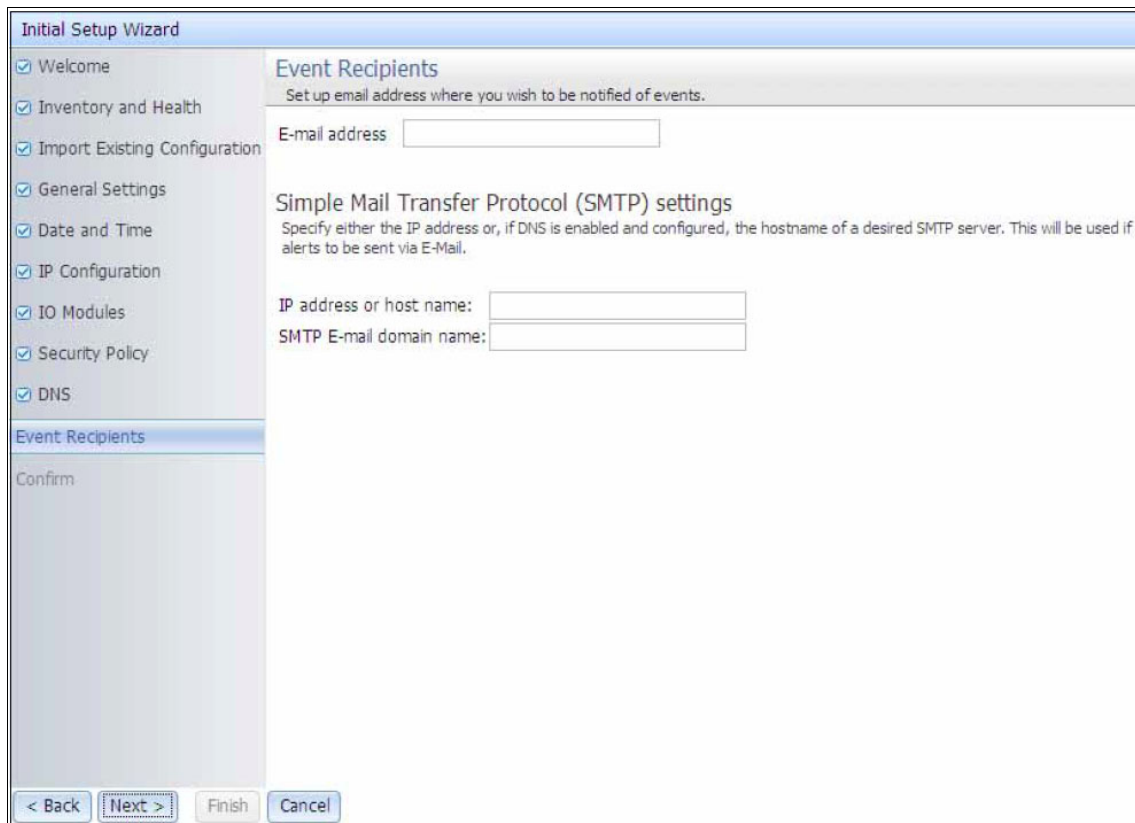
12. Select the appropriate DNS options for the CMM, as shown in Figure 6-16.  
Click **Next**.



The image shows a screenshot of the 'Initial Setup Wizard' window, specifically the 'DNS' configuration step. The window has a title bar 'Initial Setup Wizard' and a left-hand navigation pane. The navigation pane lists several steps: 'Welcome', 'Inventory and Health', 'Import Existing Configuration', 'General Settings', 'Date and Time', 'IP Configuration', 'IO Modules', 'Security Policy', 'DNS' (which is currently selected and highlighted in blue), 'Event Recipients', and 'Confirm'. The main content area is titled 'DNS' and contains the following options: 'Relevant IP addresses for Domain Name Server (DNS)' (a header), 'Enable DNS' (an unchecked checkbox), 'Preferred DNS address type:' (a label followed by a dropdown menu showing 'IPv4'), and 'Send DDNS updates to these servers' (an unchecked checkbox). At the bottom of the window, there are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 6-16 DNS setup window

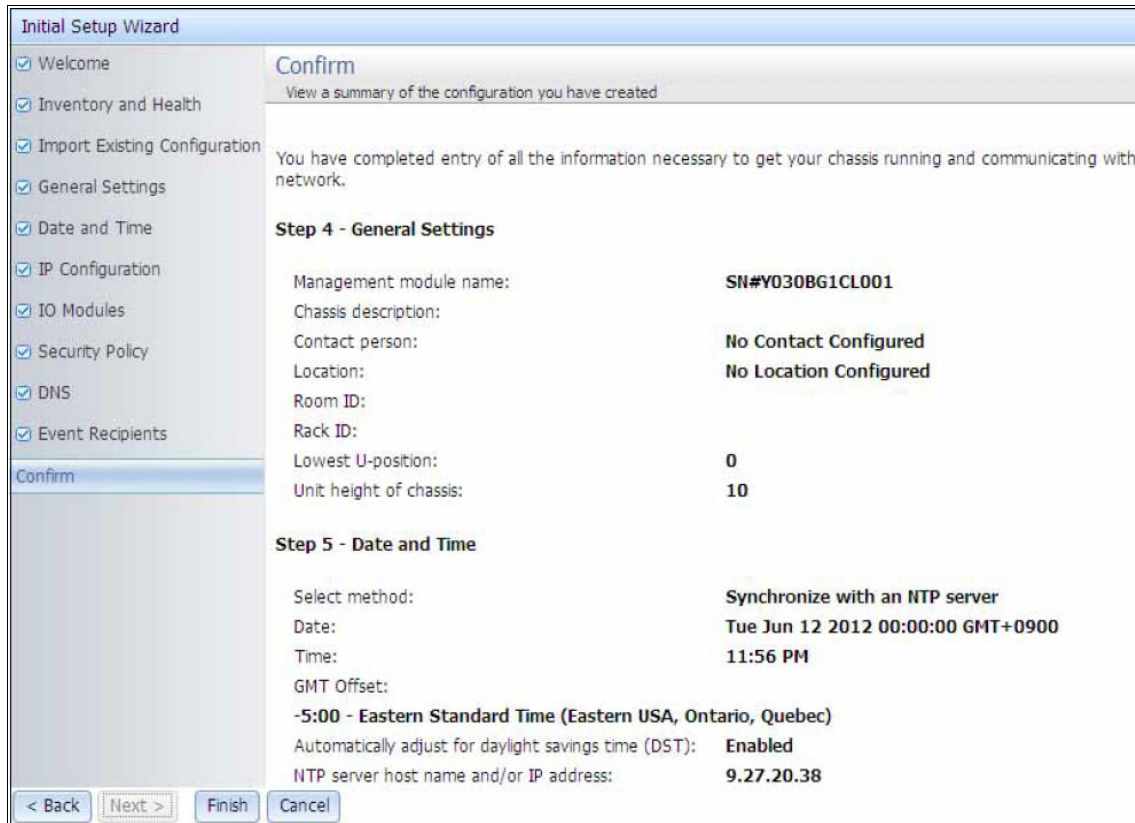
13. Enter the email addresses where notifications are to be sent when CMM events occur, as shown in Figure 6-17. Click **Next**.



The image shows a screenshot of the 'Initial Setup Wizard' window, specifically the 'Event Recipients' step. The window has a title bar 'Initial Setup Wizard' and a left sidebar with a list of steps: Welcome, Inventory and Health, Import Existing Configuration, General Settings, Date and Time, IP Configuration, IO Modules, Security Policy, DNS, Event Recipients (highlighted), and Confirm. The main area is titled 'Event Recipients' and contains the instruction 'Set up email address where you wish to be notified of events.' Below this is a text input field labeled 'E-mail address'. Further down, there is a section titled 'Simple Mail Transfer Protocol (SMTP) settings' with the instruction 'Specify either the IP address or, if DNS is enabled and configured, the hostname of a desired SMTP server. This will be used if alerts to be sent via E-Mail.' This section contains two text input fields: 'IP address or host name:' and 'SMTP E-mail domain name:'. At the bottom of the window are four buttons: '< Back', 'Next >' (highlighted with a dashed border), 'Finish', and 'Cancel'.

Figure 6-17 Event Recipients window

14. Confirm all of the information that you entered in the setup wizard, as shown in Figure 6-18. Click **Finish**.



The image shows a screenshot of the 'Initial Setup Wizard' window, specifically the 'Confirm' step. The left sidebar lists the steps of the wizard, with 'Confirm' selected. The main area displays a summary of the configuration and details for the next two steps.

**Initial Setup Wizard**

**Confirm**  
View a summary of the configuration you have created

You have completed entry of all the information necessary to get your chassis running and communicating with network.

**Step 4 - General Settings**

Management module name:	SN#Y030BG1CL001
Chassis description:	
Contact person:	No Contact Configured
Location:	No Location Configured
Room ID:	
Rack ID:	
Lowest U-position:	0
Unit height of chassis:	10

**Step 5 - Date and Time**

Select method:	Synchronize with an NTP server
Date:	Tue Jun 12 2012 00:00:00 GMT+0900
Time:	11:56 PM
GMT Offset:	-5:00 - Eastern Standard Time (Eastern USA, Ontario, Quebec)
Automatically adjust for daylight savings time (DST):	Enabled
NTP server host name and/or IP address:	9.27.20.38

< Back   Next >   Finish   Cancel

Figure 6-18 Confirm window

### 6.1.3 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 6-19.

#### Component IP Configuration

Configure IPv4 and IPv6 address information for the components below.

##### I/O Modules

Bay	Device Name	IPv4 Enabled	IP Address
1	EN4093 10Gb Ethernet Switch	Yes	<a href="#">View</a>
2	CN4093 10Gb Converged Switch	Yes	<a href="#">View</a>
3	FC3171 8Gb SAN Switch	Yes	<a href="#">View</a>

##### Compute Nodes

Bay	Device Name	IPv4 Enabled	IP Address
1	<span style="border: 1px solid red;">node01-x240</span>	Yes	<a href="#">View</a>
2	node02-x240	Yes	<a href="#">View</a>
3	node03-x240	Yes	<a href="#">View</a>
4	node04-x240	Yes	<a href="#">View</a>
5	node05-FSM	Yes	<a href="#">View</a>
10	node06-p270	Yes	<a href="#">View</a>

##### Storage Nodes

Bay	Device Name	IPv4 Enabled	IP Address
11-14:1	node01	Yes	<a href="#">View</a>
11-14:2	node02	Yes	<a href="#">View</a>

Figure 6-19 Component IP Configuration window

Click the I/O module or compute node link to open its IP properties window, as shown in Figure 6-20.

IP Address Configuration node01-x240

General Setting IPv4 IPv6

**Current IP Configuration**

Network Interface eth1

Configuration Method Use Static IP Address

IP Address 9.42.171.16

Subnet Mask 255.255.254.0

Gateway Address 9.42.170.1

**Change IP Configuration**

Configuration Method Use Static IP Address

**New Static Address Information**

IP Address

Subnet Mask

Gateway Address

Apply

Close

Figure 6-20 IP Address Configuration node01 window

## 6.2 IBM Flex System Manager Setup wizard

IBM Flex System Manager is an appliance that has all of the required software preinstalled. When this software stack is started for the first time, a startup wizard opens. This wizard guides you through the required configuration process, such as licensing agreements and TCP/IP configuration for the appliance.

When configuration is complete, Flex System Manager is ready to manage the chassis it is installed in and up to 16 more chassis. After the chassis is managed, individual components, such as compute nodes and switches, also can be managed.

Flex System Manager is based on an x86 compute node and has the same options for obtaining an initial console. You can use the IMM2 remote console or use the supplied dongle and front port on the Flex System Manager node to connect directly to a keyboard, display, and mouse or a console manager unit.

To monitor the Flex System Manager startup process, connect a console by using one of the methods that are described in this section before the Flex System Manager node is powered on. The following steps use the IMM2 remote console method:

1. Start a browser session as shown in Figure 6-21 and browse to the IP address of the Flex System Manager IMM2.

**Tip:** The IP address of the IMM2 of x86 compute nodes can be determined from the CMM or command-line interface (CLI). By default, the interface is set to use DHCP. However, it can be changed to a static address by using the CMM, a CLI, or a console that is connected directly to the VGA port on the front of Flex System Manager. The console is accessible with the use of the console breakout cable.

Figure 6-21 IMM2 login



2. After you log in to the IMM2, click **Server Management** → **Remote Control**, as shown in Figure 6-22.

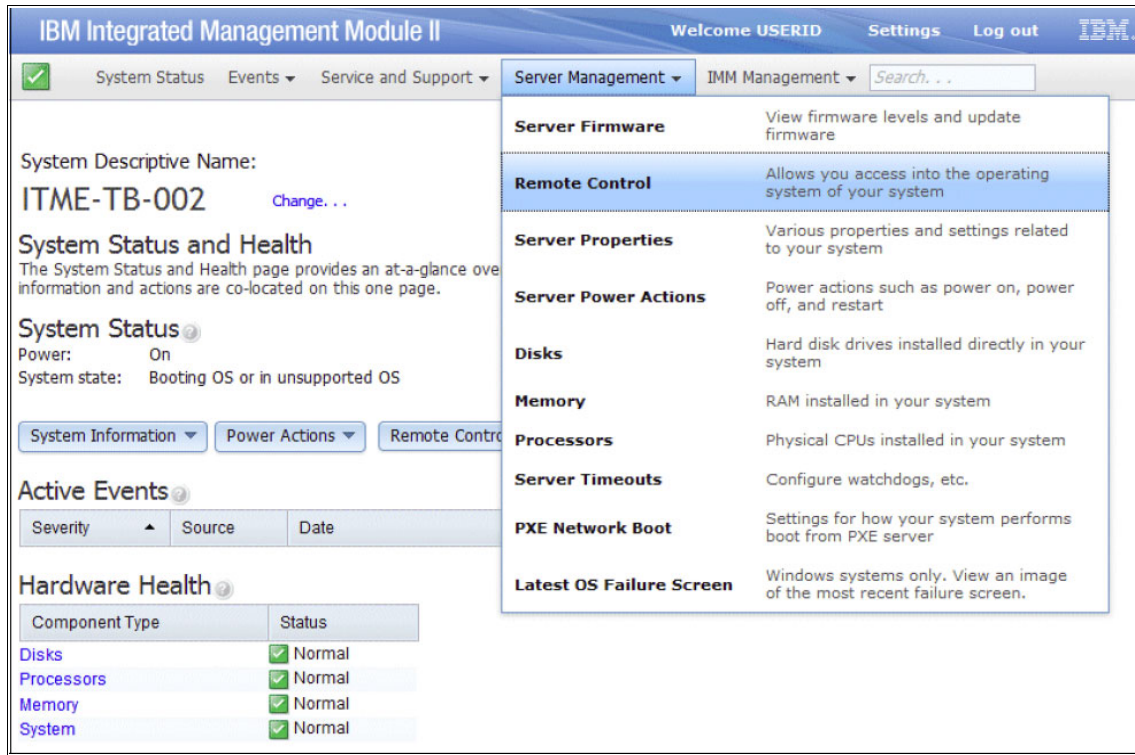


Figure 6-22 Remote control option in IMM2

3. In the Remote Control window, click **Start remote control in single-user mode**, as shown in Figure 6-23. A Java applet starts on the local desktop that is a console session to Flex System Manager.

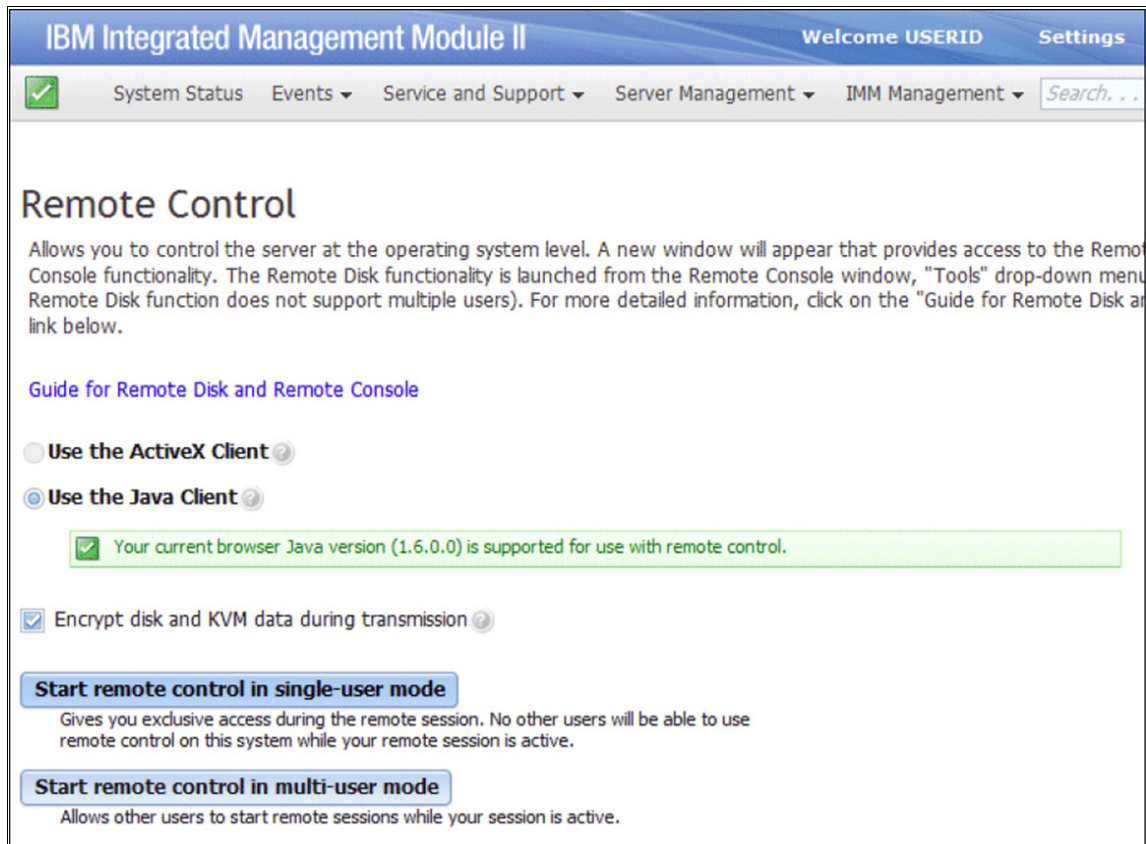


Figure 6-23 Starting remote console from IMM2

4. Flex System Manager can be powered on from several locations, including the physical power button on Flex System Manager or from the CMM. For this example, selecting the **Tools** → **Power** → **On** option from the remote console menu, as shown in Figure 6-24, is the most convenient method.

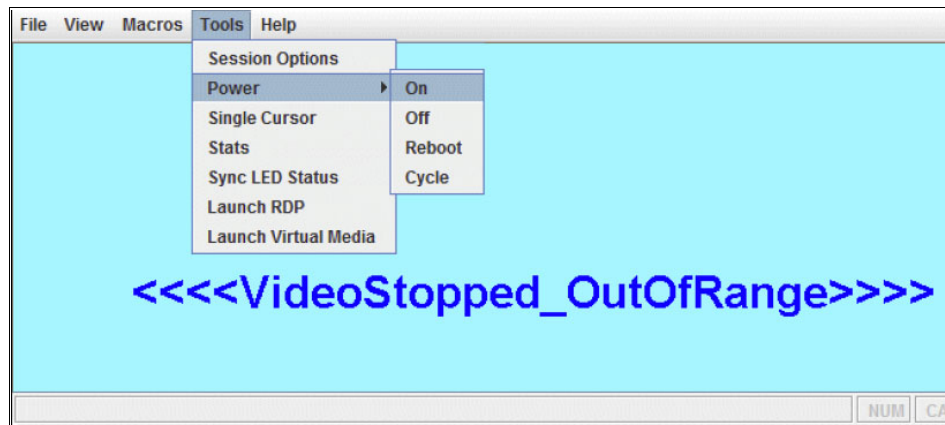


Figure 6-24 Powering on Flex System Manager from the remote console session

5. As Flex System Manager powers up and boots, you can monitor the process. No input is accepted until the License Agreement window (as shown in Figure 6-25) displays. Click **I agree** to continue.

**IBM Flex System Manager License Agreement**

By clicking on **I agree**, you agree that (1) you have had the opportunity to review the terms of all agreements presented below and (2) you govern this transaction. If you do not agree, click **I do not agree**.

Agreements:

Language: English (en)

IBM Programs:  
[IBM Flex System Manager Separately Licensed Code](#)

Third Party Licenses:  
[Red Hat EULA](#)

Agreement text:

International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON AN "ACCEPT" BUTTON OTHERWISE USING THE PROGRAM, LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ACCEPTING THESE TERMS ON BEHALF OF LICENSEE, YOU REPRESENT AND WARRANT THAT YOU 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 PROGRAM; AND
- PROMPTLY RETURN THE UNUSED MEDIA, DOCUMENTATION, AND PROOF OF ENTITLEMENT TO PARTY FROM WHOM IT WAS OBTAINED FOR A REFUND OF THE AMOUNT PAID. IF THE PROGRAM DOWNLOADED, DESTROY ALL COPIES OF THE PROGRAM.

You can view and print copies of the above referenced agreements by selecting an agreement and clicking Print.

Figure 6-25 IBM Flex System Manager license agreement

The startup wizard Welcome window displays, as shown in Figure 6-26.

**Welcome**

Use this wizard to complete the following set up tasks to install and configure t

**Setup Prerequisites**

- Date and Time
- Set the system level User ID and password
- Configure Local Area Network (LAN) adapters
- Configure network settings
- Configure Domain Name System (DNS)

Figure 6-26 Flex System Manager Welcome window

6. To continue, click **Data and Time** from the wizard menu to display the window that is shown in Figure 6-27. Set the time, date, time zone, and Network Time Protocol server, as needed. Click **Next**.

Welcome

➔ **Date and Time**

Password

Network Topology

LAN Adapters

Host and Gateway

Advanced Routing

DNS

Summary

### Date and Time

Set the date and time and select the correct time zone for the system, if neces

Date: 12/05/2011

Time: 9:22:14 PM

Time zone: America/New\_York

☒ Automatically adjust clock for Daylight Saving Time (DST)

#### Network Time Protocol (NTP) Server

Specify an NTP server to automatically synchronize the system clock periodically.

Time server hostname or IP address:

Add >

NTP version not specified

Remove

☐ Use NTP authentication

Key index:

Key type: M - MD5

Key:

[Learn more about a network time protocol server](#)

< Back Next >

Figure 6-27 Setting the Flex System Manager date and time

7. Create a user ID and password for accessing the GUI and CLI. User ID and password maintenance, including creating more user IDs, is available in Flex System Manager after the startup wizard completes. Figure 6-28 shows the creation of the USERID user ID and entering a password. Click **Next** to continue.

**System-Level User ID and Password**

Enter a user ID and password for the system-level access user. The default user ID is 'USERID', which matches the CMM user ID. This password will be applied to all local administrative accounts, including 'pe' (product engineer) and 'root'.

\*User ID:

\*New password:

\*Confirm password:

Group:

Note: You can change this password and add additional users after setup is complete.

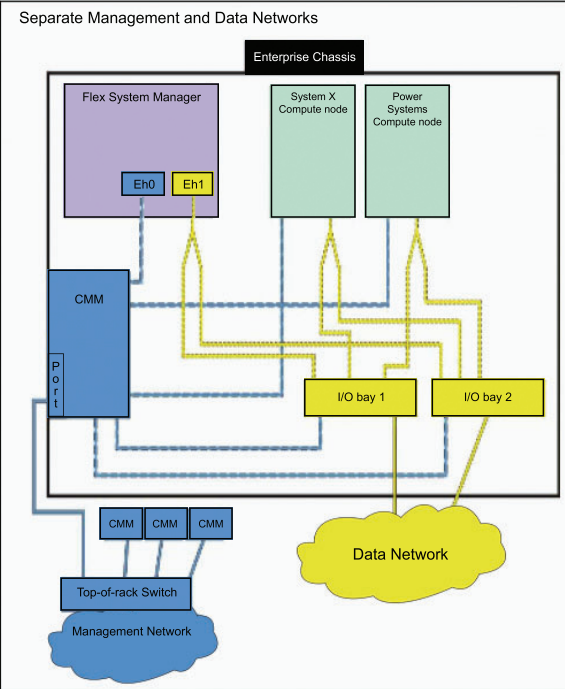
*Figure 6-28 Flex System Manager system level user ID and password step*

8. Network topology options include separate networks for management and data or a single network for both data and management traffic from the chassis. Often, it is best to have separate management and data networks. To simplify this example, a combined network is configured by using the topology that is shown on the right side of Figure 6-29 on page 163. Click **Next** to continue to the actual network configuration.

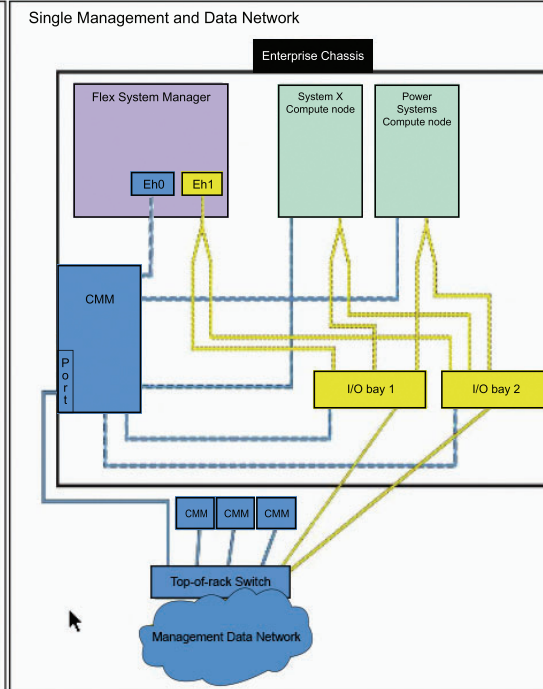
## Network Topology

There are two possible network topologies that can be configured for Flex System Manager (refer to network diagrams below). Note: Your network topology determines Ethernet adapter(eth0 and eth1 settings that will be made on later pages in this wizard.

Separate networks for data and management traffic.



One network for both data and management traffic.



### Advanced Routing

☐ I want to set up advanced routing in this wizard. Advanced routing provides settings for multiple gateways or subnets within your network. (Advanced routing involves defining specific network paths for data traffic. If you are unfamiliar with network paths, do this after the management server is running.)

Figure 6-29 Flex System Manager network topology options







10. In the Configure IP Address window that is shown in Figure 6-31, select the DHCP or static IP options for IPv4 and IPv6 addressing. Select the options that you want, enter the information as required, and then click **Next**.

### Configure IP Address

Configure the IP addresses for the specified LAN adapter. If the adapter is configured DHCP and is unable to get an IP address, the management server will not start.

LAN interface address: 5C:F3:FC:5F:40:E5 eth0 (Management Network)

☒ IPv4 address:

☐ Obtain an IP address automatically

☒ Use the following IPv4 address:

Static IP address:

Network mask:

☐ IPv6 address:

☐ Use DHCPv6 or Stateless Auto Configuration to configure IP settings

☒ Use the following IPv6 address:

Specify new static IPv6 address information and click Add:

IPv6 address:

Prefix length:

Configured static IPv6 addresses:

IPv6 address	Prefix length	Remove
fe80:0:0:0:5ef3:fcff:fe5f:40e5	64	<input type="button" value="Remove"/>

Figure 6-31 Flex System Manager IP address assignment

The wizard returns to the Initial LAN Adapter window and preselects the next adapter in the list, as shown in Figure 6-32. This example uses a combined network topology and a single adapter, so it does not need more IP addresses.

**Attention:** Click to clear the perform network validation option (which is selected by default) if DNS is not available or not configured.

Select **No** for “Do you want to configure another LAN adapter”, as shown in Figure 6-32. Click **Next** to continue.

### Configure Local Area Network (LAN) Adapters

Configure a LAN adapter for network access to the system.

Do you want to configure another LAN adapter?  
Select the LAN adapter to configure and click Next.

☒ No

☐ Yes, I want to configure another LAN adapter

Select	Adapter	Description	IP address	Configured this session
<input type="radio"/>	eth0 5CF3FC5F	Management Network	9.27.20.199	Yes <input type="button" value="Clear Configuration"/>
<input checked="" type="radio"/>	eth1 5CF3FC5F	Data Network	0.0.0.0	No

☒ Perform network validation and recovery when the setup wizard is complete. If the specified configuration produces errors, the management server will not start, the system will be reset to factory defaults, and this wizard will restart to allow you to change your network settings.

Figure 6-32 Flex System Manager LAN adapter configuration continue option

11. After IP address assignment, the host name and gateway are configured, as shown in Figure 6-33. Enter the host name, domain name, and default gateway address. Ensure that the IP address and the default gateway adapter are correct. Click **Next** to continue.

**Tip:** The host name of the Flex System Manager should be available on the domain name server.

### Configure Host and Gateway

Specify host name information, verify the domain name for the host, and specify the default gateway address and device.

*Host name IP address:	<input type="text" value="9.27.20.199"/>
*Short name:	<input type="text" value="r2-c1-ch1-itme1"/>
*Domain name:	<input type="text" value="stglabs.ibm.com"/>
*Default Gateway address:	<input type="text" value="9.27.20.1"/>
*Default Gateway device:	<input type="text" value="eth0"/>

Figure 6-33 Flex System Manager host name and gateway configuration

12. You can enable the use of a DNS service and to add the address of one or more servers and a domain suffix search order. Enter the information as shown in Figure 6-34 and then click **Next** to continue.

### Configure Domain Name System (DNS)

Enable DNS services and configure the search order for DNS servers and domain suffixes. An incorrect DNS can prevent the management server from starting.

☒ Enable DNS services

DNS server search order:

Domain suffix search order:

Figure 6-34 Flex System Manager DNS services configuration

13. A summary window of your configured options opens, as shown in Figure 6-35. To change a selection, click **Back**. If you do not need to make any changes, click **Finish**.

### Summary

Review the following settings, then click Finish. To change any other settings, click Back.

#### Date and Time

Date: 12/05/2011  
Time: 9:28:00 PM  
Time zone: America/New\_York  
Time servers:

#### LAN Adapters

LAN interface address: eth0 5CF3FC5F40E5  
IP address: 9.27.20.199  
Network mask: 255.255.252.0  
Static IPv6 address: Prefix length

#### Host and Gateway

\*Short name: r2-c1-ch1-itme1  
\*Domain name: stglabs.ibm.com  
\*Default Gateway address: 9.27.20.1  
IPv6 gateway address:  
\*Default Gateway device: eth0  
\*Host name IP address: 9.27.20.199

Figure 6-35 Flex System Manager startup wizard summary window

When you click **Finish**, the final configuration and setup continues automatically without the need of further input. Figure 6-36 shows the processing status display.

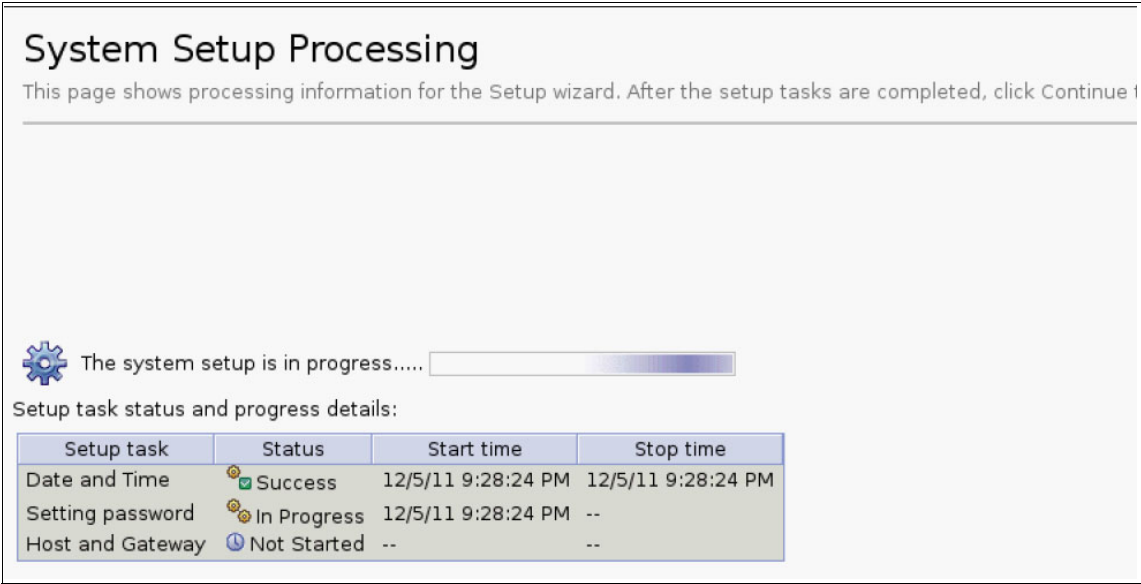



Figure 6-36 Flex System Manager system setup processing status

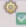

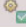
Figure 6-37 shows the message when the processing is complete.

## System Setup Processing

This page shows processing information for the Setup wizard. After the setup tasks are completed, click Continue!



Setup task status and progress details:


Setup task	Status	Start time	Stop time
Date and Time	 Success	12/5/11 9:28:24 PM	12/5/11 9:28:24 PM
Setting password	 Success	12/5/11 9:28:24 PM	12/5/11 9:28:36 PM
Host and Gateway	 Success	12/5/11 9:28:36 PM	12/5/11 9:28:36 PM

Congratulations. All setup tasks completed successfully.

[Continue](#)

Figure 6-37 Flex System Manager system setup processing completed

Figure 6-38 shows the message when server is started.

 **Attention:** The web server is being restarted as part of the setup process. Network setup can take up to 5 minutes after which the setup process will continue for approximately 45 minutes. If there are network errors, you will receive a notification within 5 minutes, after which the setup process can continue unattended. Do not close this page or your browser window.


 Please wait while the network settings are being applied

Figure 6-38 Flex System Manager startup

Figure 6-39 shows the startup process display.

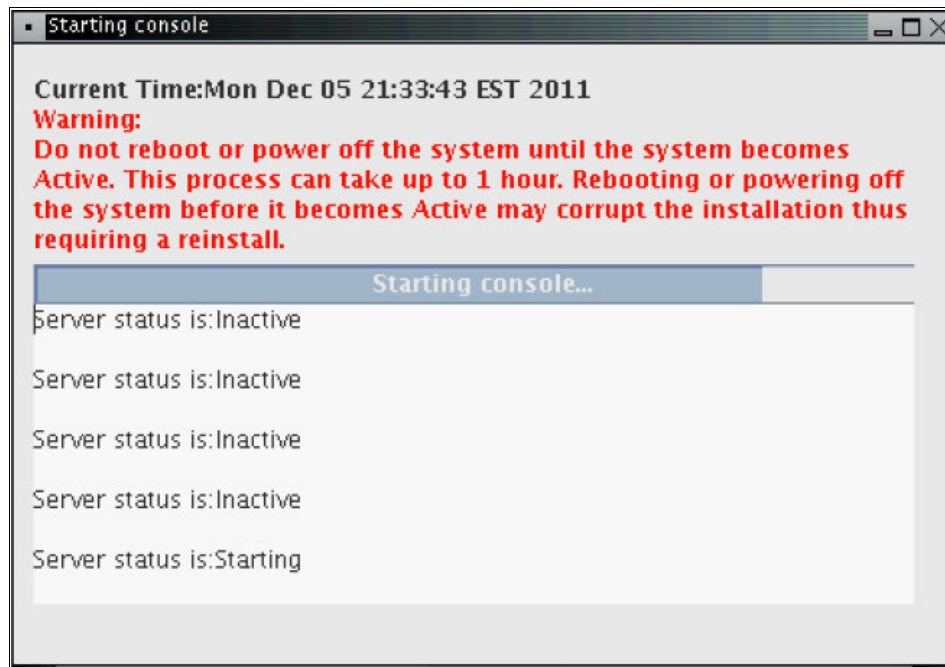


Figure 6-39 Flex System Manager startup status



14. When the startup is completed, the local browser on the Flex System Manager also starts. A list of untrusted connection challenges displays. Click **I Understand the Risks**, as shown in Figure 6-40. Accept these challenges by clicking **Add Exception**.

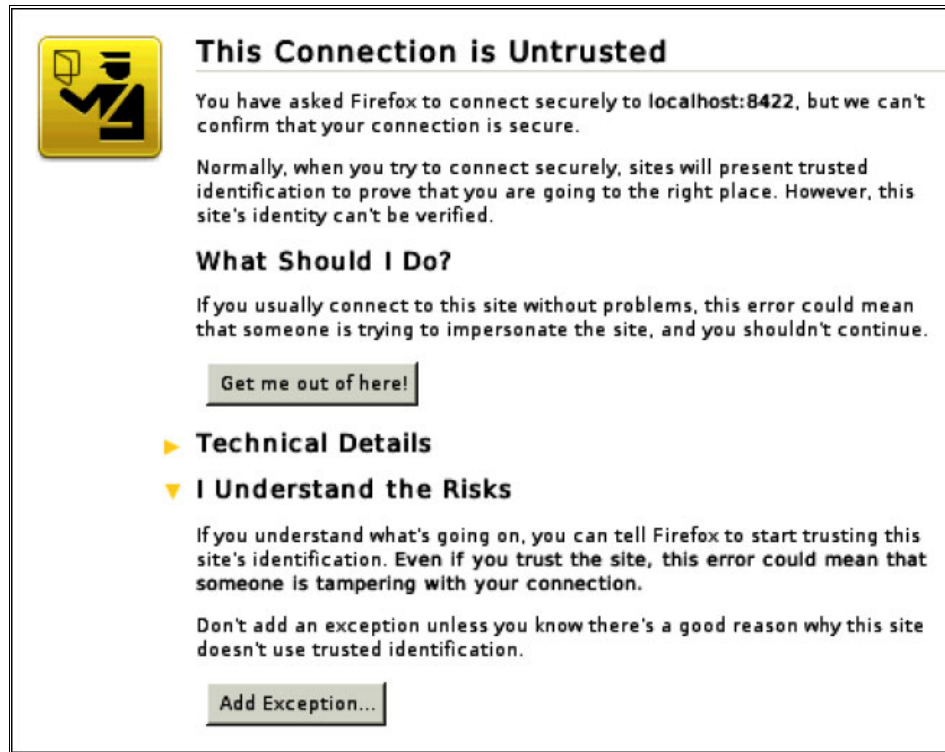


Figure 6-40 Flex System Manager browser add exception

15. With the security exceptions cleared, the Login window of the IBM Flex System Manager GUI opens. Enter the User ID and credentials that you entered in the startup wizard, and then click **Log in**, as shown in Figure 6-41.



*Figure 6-41 Flex System Manager Login window*

The startup wizard and initial login are now complete. Flex System Manager is ready for further configuration and use. This example uses a console from the remote console function of the IMM2. A secure browser session can now be started to Flex System Manager.

## 6.3 Selecting a chassis to manage

Most tasks in Flex System Manager can be performed with more than one method when you are using the GUI. In this example, the most common methods are shown.

After the initial setup of Flex System Manager, it discovers any available chassis. Selections can then be made as to which chassis are managed by the current Flex System Manager. Complete the following steps to select chassis:

1. From the Home tab, go to the Initial Setup tab to display the Initial Setup window. Click **Configure Chassis Components**, as shown in Figure 6-42.



Figure 6-42 Flex System Manager initial setup window

2. A list of available chassis opens. Select the chassis that you want to manage, as shown in Figure 6-43. Click **Manage**.

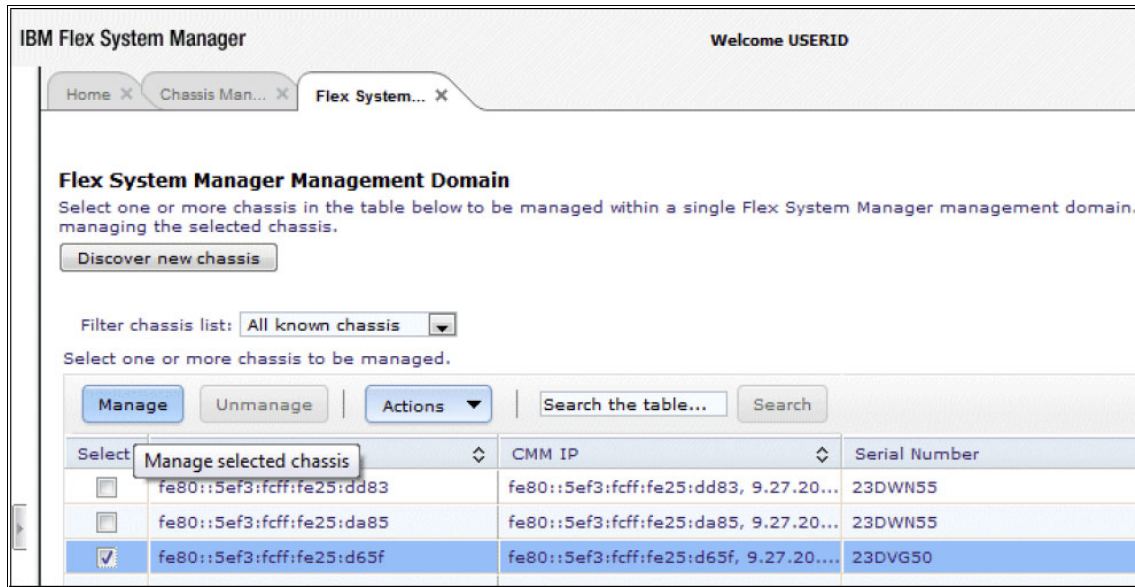


Figure 6-43 Flex System Manager chassis selection for management

The Manage Chassis window shows a list of the selected chassis, as shown in Figure 6-44. A drop-down menu shows the available Flex System Manager systems. Ensure that the chassis and Flex System Manager selections are correct.

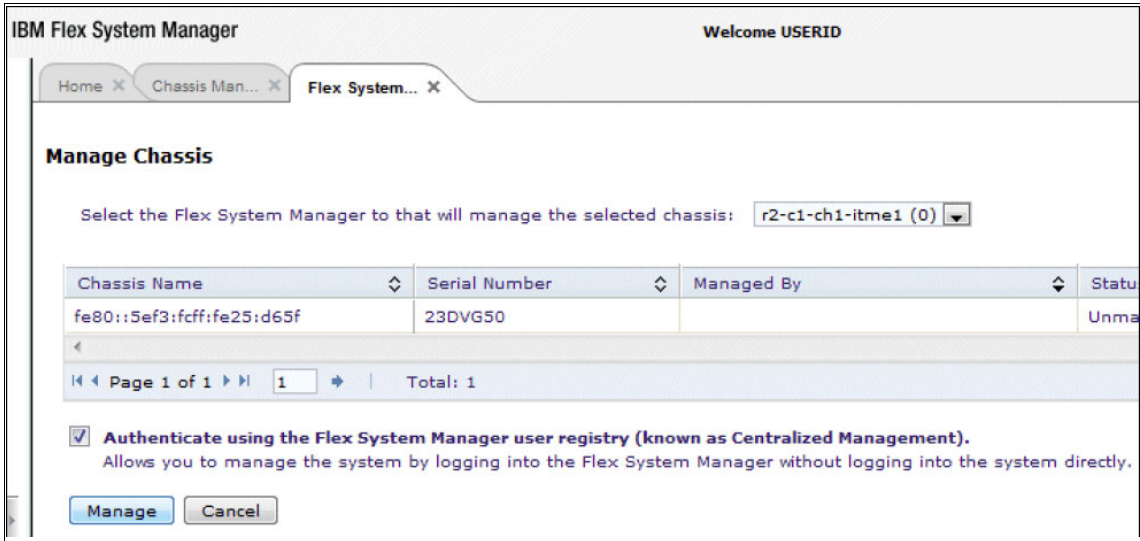


Figure 6-44 Flex System Manager Manage Chassis options

3. Click **Manage** to update the Message column from Waiting to Finalizing, as shown in Figure 6-45.

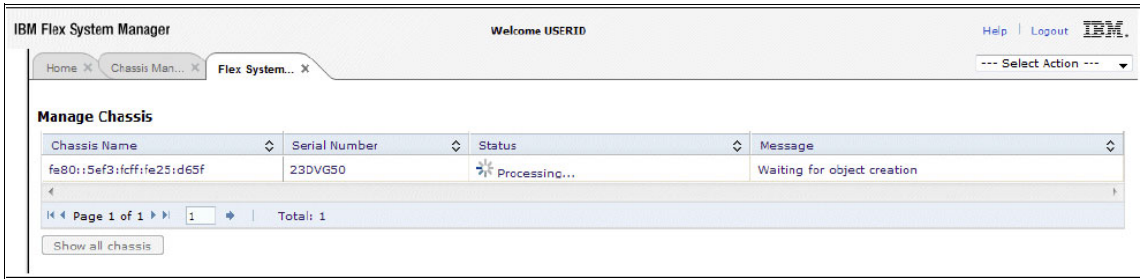


Figure 6-45 Flex System Manager manage chassis: Step 1

The Message column changes to Managed, as shown in Figure 6-46.

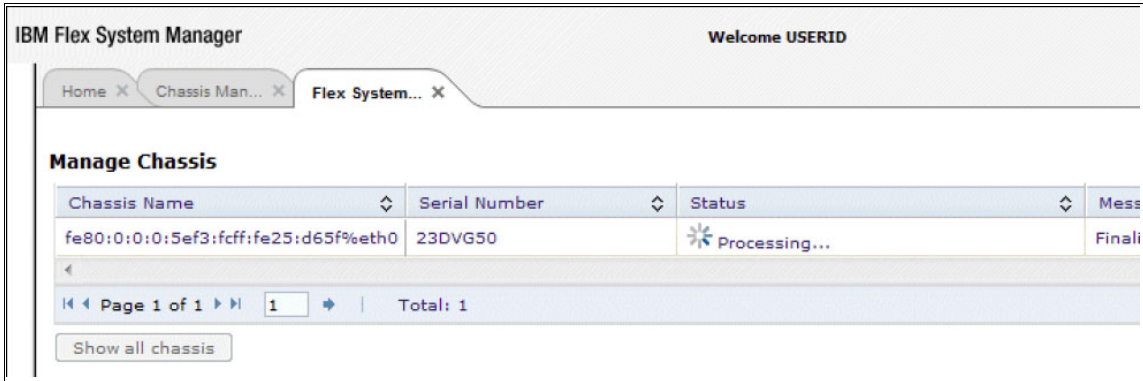


Figure 6-46 Flex System Manager manage chassis: Step 2

4. After the successful completion of the manage chassis process, click **Show all chassis**, as shown in Figure 6-47.

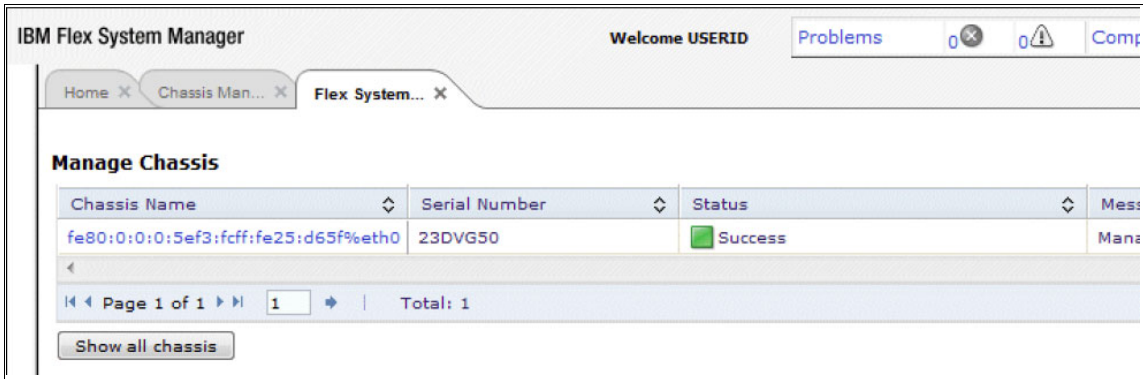


Figure 6-47 Flex System Manager manage chassis steps completed

The original Flex System Manager Management Domain window opens with the target chassis as the managing Flex System Manager, as shown in Figure 6-48.

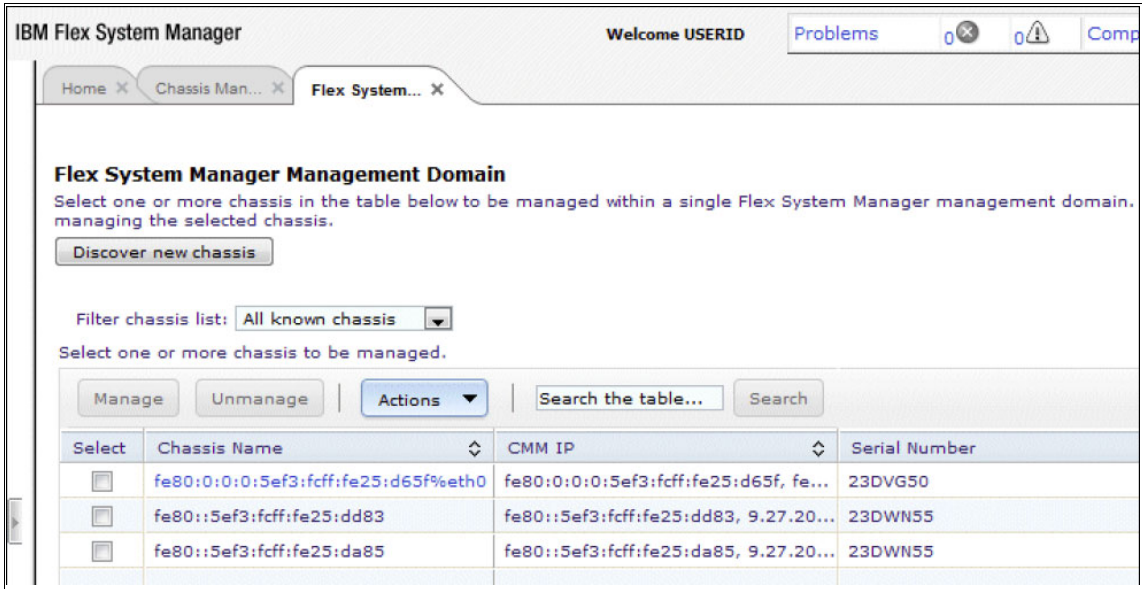


Figure 6-48 Flex System Manager with management domain updated

The Enterprise Chassis is now managed by Flex System Manager.

## 6.4 Discovery and inventory collection

To manage a resource within an environment or view inventory data about it, that resource must first be discovered. After access is granted, an inventory must be collected. The resource is recognized and added to the comprehensive list of native resources and native attributes for the system. Discovery and inventory collection are the two primary tasks that are used to connect to supported network resources and collect information about them.



## 6.4.1 Discovery

*Discovery* is the process by which Flex System Manager identifies and establishes connections with network-level resources that Flex System Manager can manage. These resources include compute nodes, switches, storage devices, operating systems, hypervisors, and virtual machines. Use system discovery to identify resources within your environment, collect data about those resources, and establish connections with them.

### Choosing which discovery method to use

Discovering your resources in the most efficient manner means deciding which method best suits your needs. Each method has advantages and disadvantages to consider.

Collecting the inventory of a chassis component requires the following overall steps:

1. Discovery
2. Grant Access
3. Collect Inventory

There are several paths to discover and collect the inventory on Flex System components. In this section, we describe the method that uses Discovery Manager. The next sections describe different paths to discover the three main components in an Enterprise Chassis (CMM, compute nodes, and I/O modules).

### Discovery protocols

A *discovery protocol* is any network communication protocol that Flex System Manager uses during the discovery process to discover a resource. The default discovery profile uses a predetermined list of protocols. When you specify a single IP address, a single host name, or a single range of IP addresses, system discovery uses one or more protocols. These protocols are based on the selected target resource type. With a discovery profile, you can refine the target resource type and configure specific protocols that you want to use.

The communication protocols that Flex System Manager uses during discovery depend on the protocols that are used by the target resource type. You must decide about the different protocols only when you create or edit a discovery profile. The Discovery Profile wizard helps you select and configure the correct protocol for the type of resource that you want to discover.



When you are discovering many resources, network traffic that is associated with the discovery process might cause timeouts. These timeouts might result in some discoverable resources remaining undiscovered. To help prevent this problem, use one or more discovery profiles. With a discovery profile, you can target specific resources and limit the number of communication protocols that are used during discovery.

By default, Flex System Manager supports the following discovery protocols:

- ▶ **Agent manager discovery**  
Agent manager discovery specifically targets the discovery of Tivoli common agents. In the Tivoli paradigm, Service Location Protocol (SLP) is not supported. Management nodes must contact an agent manager that knows about the agents in their environment. You can select the agent managers that you want to use in discovery.
- ▶ **Common Agent Services discovery**  
This discovery uses SLP discovery with which clients can locate servers and other services in the network.
- ▶ **Common Information Model (CIM) discovery**  
CIM discovery uses the SLP for discovery. With CIM discovery, clients can locate servers and other services in the network.
- ▶ **Interprocess communication (IPC) discovery**  
IPC uses services that Flex System Manager provides that components use to communicate with each other. By using these services, a server task can communicate with an agent task that is running on a target.
- ▶ **Secure Shell (SSH) discovery**  
Secure Shell is a command interface and protocol that is based on UNIX for securely accessing a remote computer. With SSH discovery, you can specify a single IP address or a range of IP addresses upon which to run discovery.
- ▶ **Simple Network Management Protocol (SNMP) discovery**  
SNMP is a network management standard that is widely used in Internet Protocol networks. SNMP runs management services by using a distributed architecture of management systems and agents. SNMP provides a method of managing network hosts, such as workstation and server computers, routers, bridges, and hubs from a centrally located computer that runs the network-management software.

- ▶ **Storage Management Initiative Specification (SMI-S) discovery**

With SMI-S discovery, clients can locate servers and other services in the network. This design specification was developed by the Storage Networking Industry Association (SNIA). It specifies a secure and reliable interface with which storage management systems can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools that are used to manage them.

- ▶ **Windows distributed component object model (DCOM) discovery**

Use Windows DCOM, which is an extension of the Microsoft Component Object Model (COM), to support objects that are distributed across a network configuration. Use DCOM to specify a single IP address or a range of IP addresses on which to run discovery.

## **Discovery Manager**

With Discovery Manager, you can discover and connect to the systems at your site. This window displays an overview of all discovered systems, the systems you have access to, and the systems from which you collected inventory. It has options to explore all discovered resources, in order by category, as shown in Figure 6-49 on page 183.

Complete the following steps to use Discover Manager:

1. Select **System Discovery** under Common tasks, as shown in Figure 6-49.

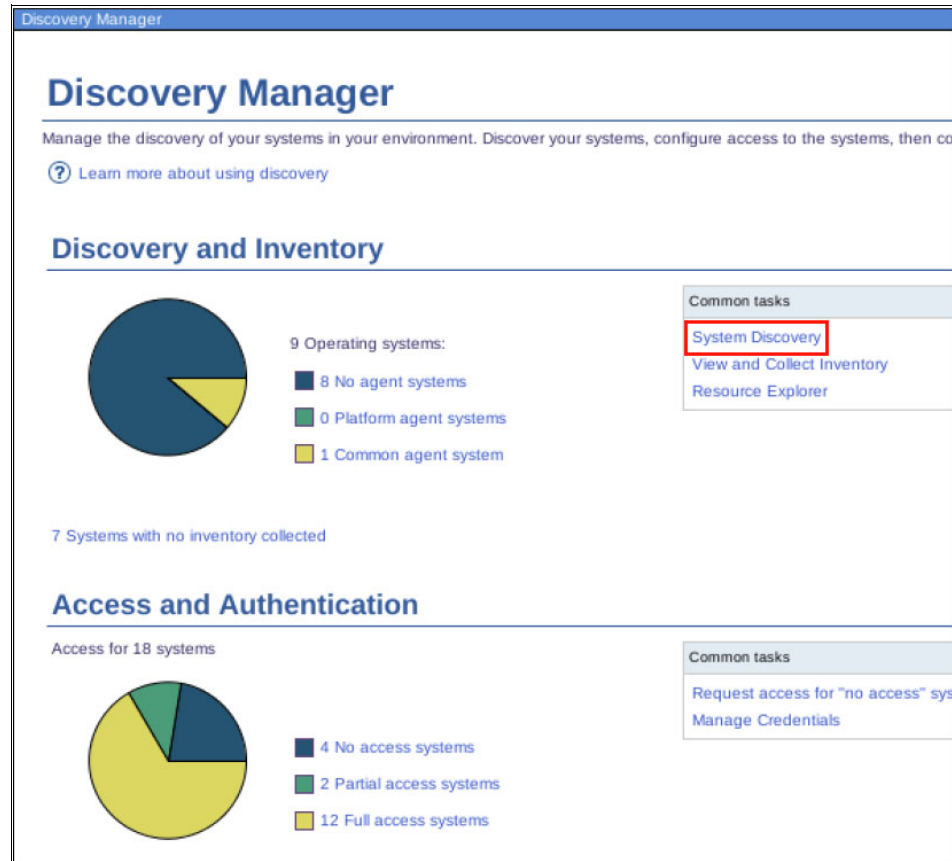
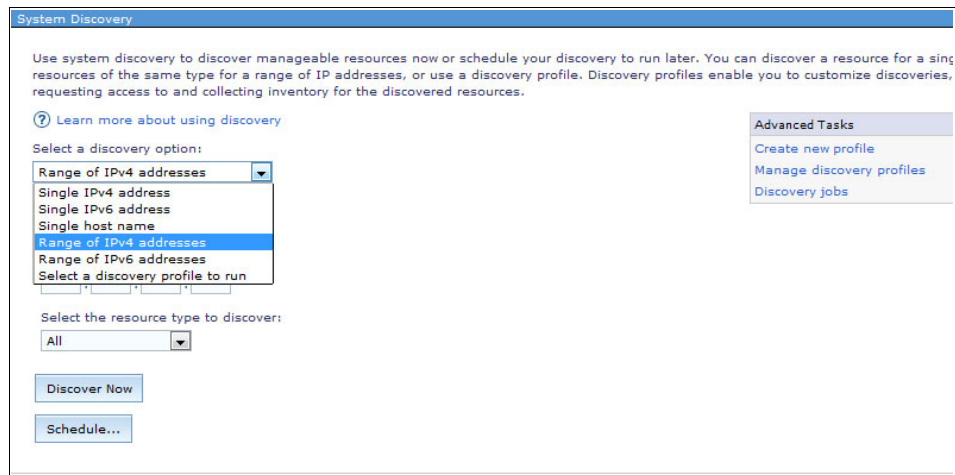


Figure 6-49 Discovery Manager window

**Tip:** You can run a discovery on a single IP address or on an IP address range.

2. In this example, you run a discovery on an IP range. Select **Range IPv4 address**, as shown in Figure 6-50.



The screenshot shows the 'System Discovery' window. At the top, there is a title bar and a descriptive paragraph about system discovery. Below this, there is a link to 'Learn more about using discovery'. The main section is titled 'Select a discovery option:' and contains a dropdown menu. The dropdown menu is open, showing several options: 'Range of IPv4 addresses' (which is highlighted in blue), 'Single IPv4 address', 'Single IPv6 address', 'Single host name', 'Range of IPv6 addresses', and 'Select a discovery profile to run'. To the right of the main content area, there is a sidebar titled 'Advanced Tasks' with three links: 'Create new profile', 'Manage discovery profiles', and 'Discovery jobs'. Below the dropdown menu, there is a section titled 'Select the resource type to discover:' with a dropdown menu set to 'All'. At the bottom of the window, there are two buttons: 'Discover Now' and 'Schedule...'.

Figure 6-50 IP address range selection

**Tip:** You can also choose to schedule your discovery, if required.

3. Enter your IP address range and click **Discover Now**, as shown in Figure 6-51.

The screenshot shows a 'System Discovery' dialog box. At the top, there is a blue header bar with the text 'System Discovery'. Below the header, there is a paragraph of text: 'Use system discovery to discover manageable resources now or host name, discover resources of the same type for a range of discoveries, including importing IP addresses, and requesting a'. Below this text is a link: '? Learn more about using discovery'. Underneath the link is the text 'Select a discovery option:'. Below this text is a dropdown menu with 'Range of IPv4 addresses' selected. Below the dropdown menu are two groups of input fields. The first group is labeled 'Starting IP address:' and contains four input boxes separated by dots. The second group is labeled 'Ending IP address:' and also contains four input boxes separated by dots. Below these input fields is the text 'Select the resource type to discover:'. Below this text is a dropdown menu with 'All' selected. At the bottom of the dialog box are two buttons: 'Discover Now' and 'Schedule...'. The 'Discover Now' button is highlighted with a red rectangle.

Figure 6-51 Enter IP address range for discovery

4. A blue information square displays and indicates that the job is started, as shown Figure 6-52. Click **Display Properties** to check the job status.

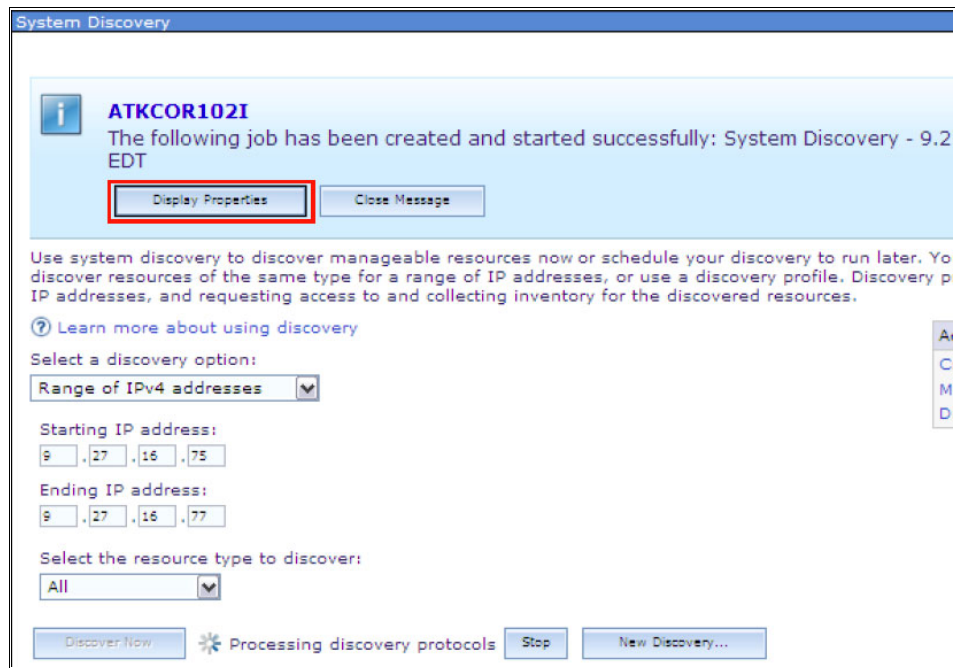


Figure 6-52 Discovery job information

Wait until the progress bar reaches 100%, which indicates that the discovery is complete, as shown in Figure 6-53.

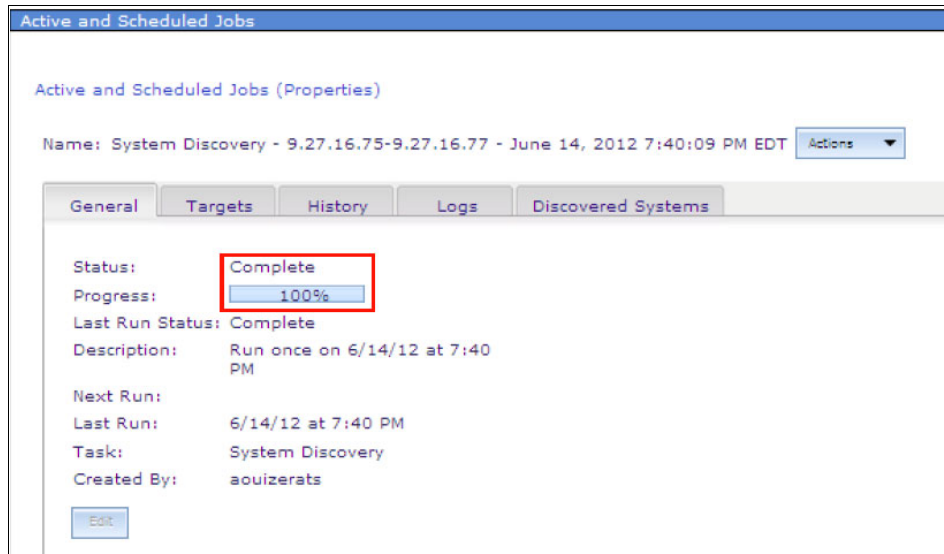


Figure 6-53 Discovery completed

If granting access to your object is required, the system displays No access in the Access field, as shown in Figure 6-54.

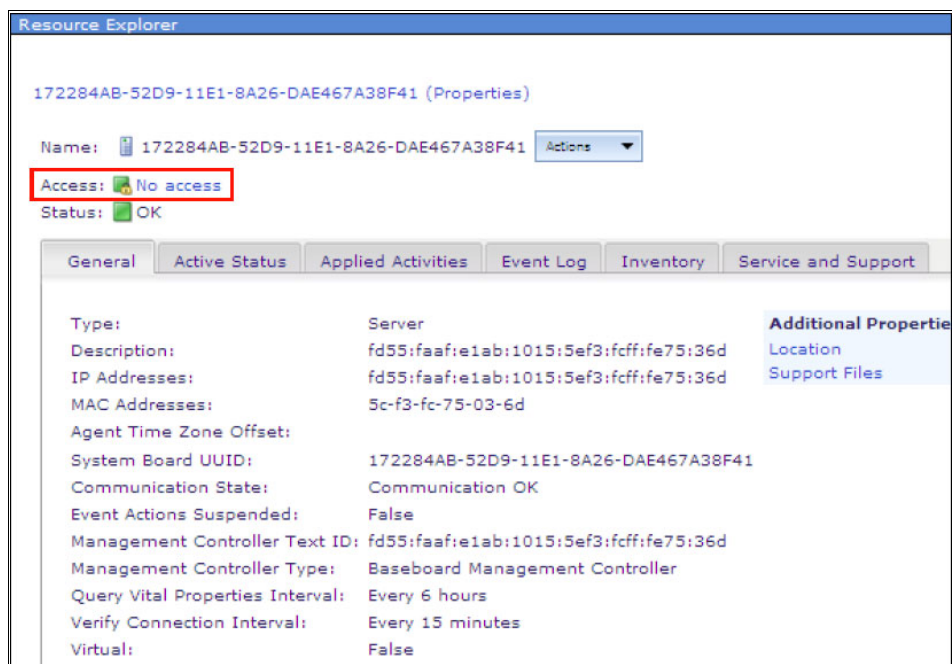


Figure 6-54 Object with no access

To manage the object, you must request access. Complete the following steps to request access:

- Click the **No access** link. A new window opens in which you enter the UserID and password to request access.
- Complete the information, as shown on Figure 6-55 on page 189.





**Request Access**

Specify the user ID and password to authenticate Flex System Manager to one or more target systems. Then click **Request Access** to request access to the target system(s).

\*User ID:

\*Password:

Selected targets:

Name	Access	Trust State
 ESXi_Node_4	 No access	<input type="checkbox"/> Not applicable

Page 1 of 1 | 1 | Total: 1

Figure 6-55 Authenticating to request access

- c. Click **Request Access**, and then click **Close**, as shown in Figure 6-56.



**Request Access**

Specify the user ID and password to authenticate Flex System Manager to one or more target systems. Then click **Request Access** to request access to the target system(s).

User ID:

Password:

Selected targets:

Name	Access	Trust State
 ESXi_Node_4	 OK	<input type="checkbox"/> Not applicable

Page 1 of 1 | 1 | Total: 1

Figure 6-56 Success on granting access

5. After you request access to the object, ensure that access is granted by clicking the **General** tab, as shown in Figure 6-57.



Figure 6-57 Access is granted

**Remember:** Do not forget to grant access for every object on which you want to collect inventory.

6. Go to the Resource Explorer tab, and click **Action** → **Inventory** → **View and Collect Inventory**, as shown in Figure 6-58.

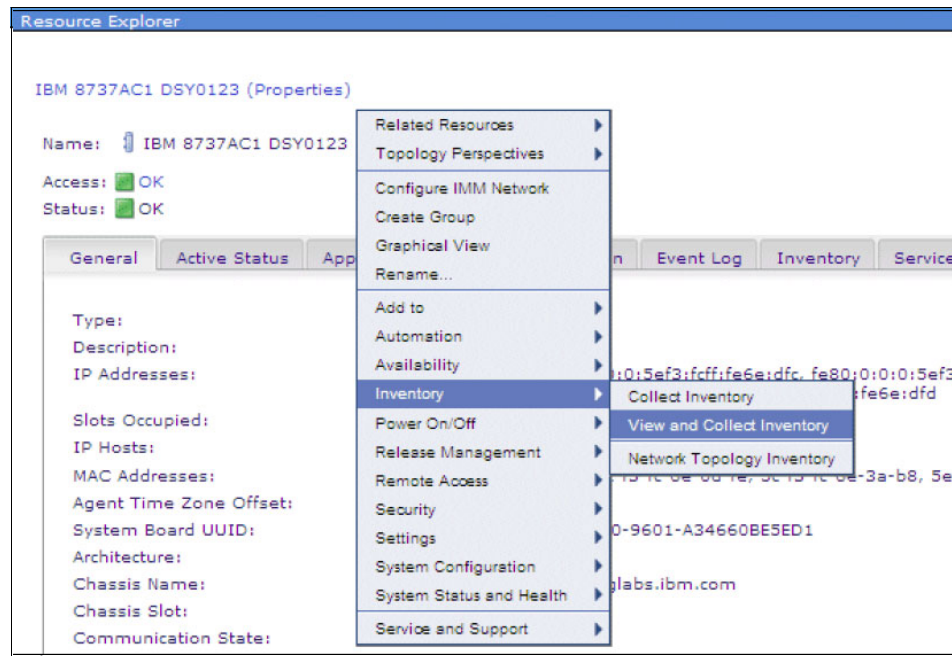


Figure 6-58 Inventory collection

7. Click **Run Now**. Your object is selected as the Target system. Click **Collect Inventory**, as shown in Figure 6-59.

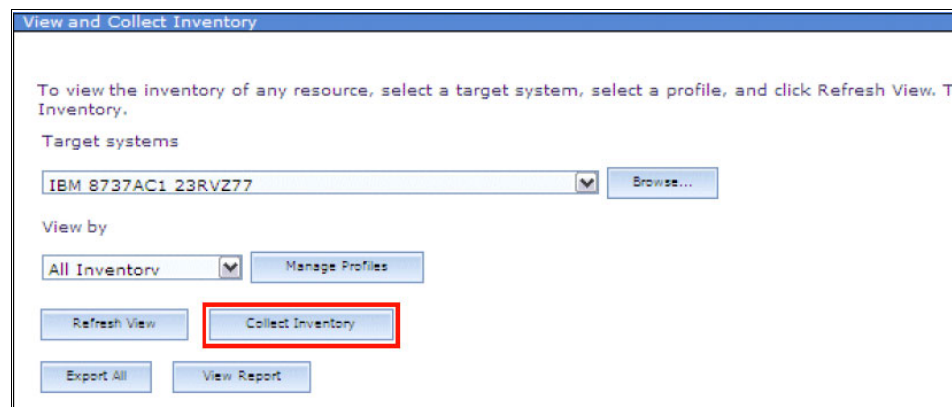


Figure 6-59 Collect Inventory button

8. To begin the inventory collection, select **Run Now**, and click **OK**, as shown in Figure 6-60.

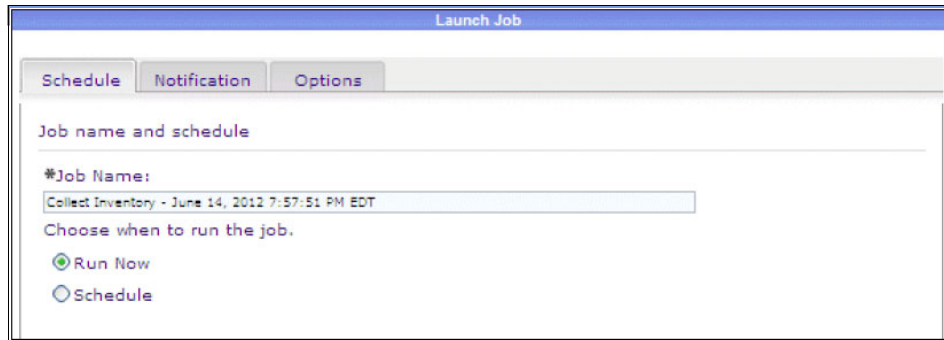


Figure 6-60 Run collect inventory

9. A blue information square indicates that the job is started, as shown in Figure 6-61. Click **Display Properties** to check the job status.

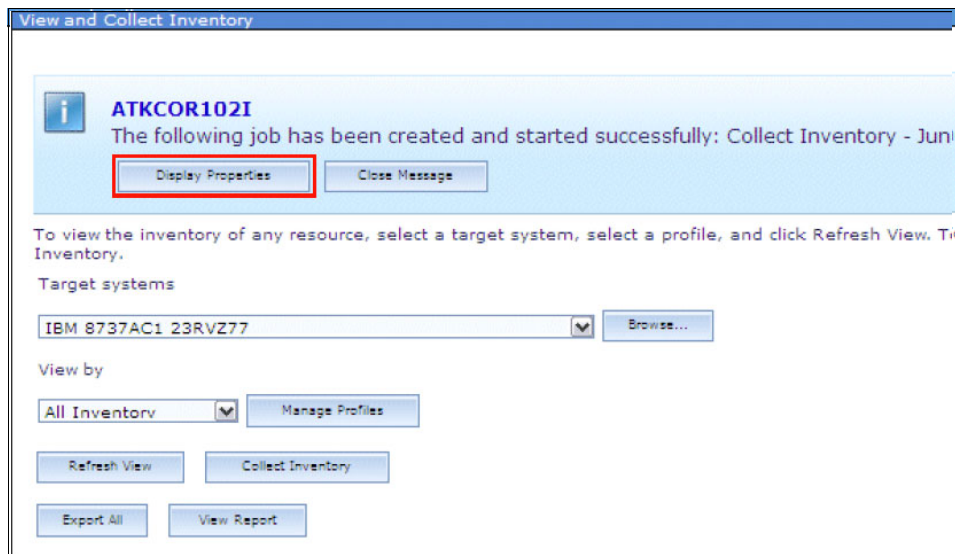


Figure 6-61 Collect inventory information

Wait until the job completes, as shown Figure 6-62.

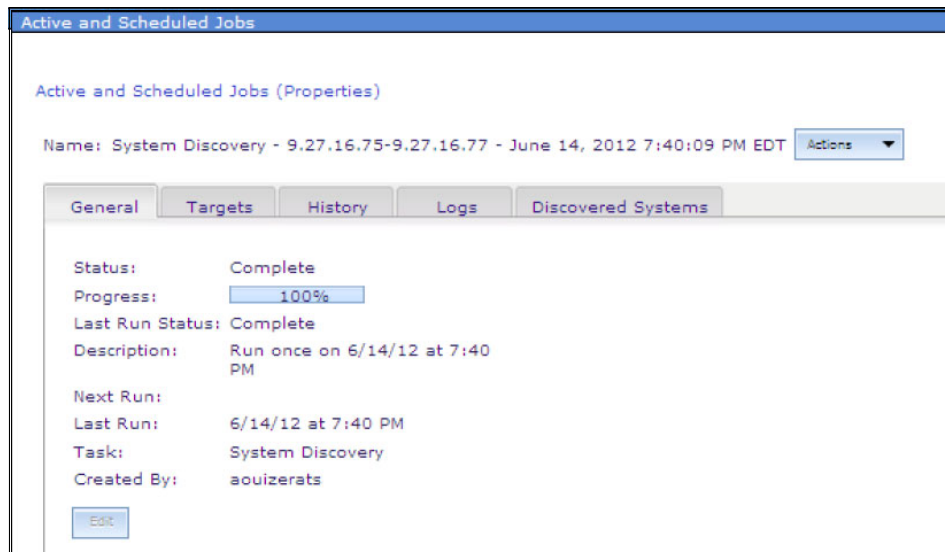


Figure 6-62 Collect inventory is completed

## 6.4.2 I/O modules

Complete the following steps to discover I/O modules:

1. Select your chassis from the Chassis Manager view, as shown Figure 6-63.

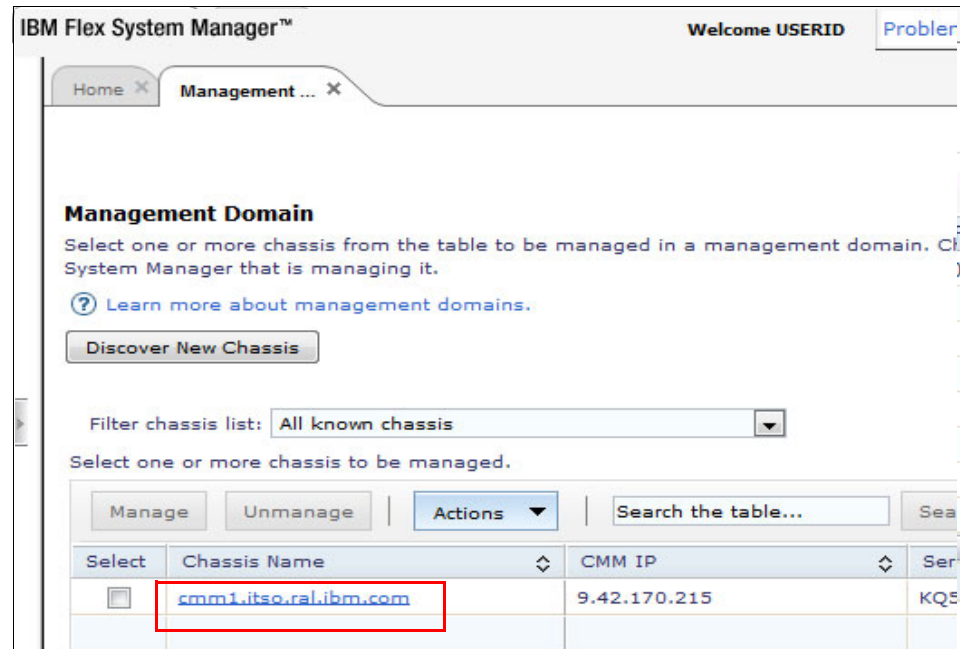


Figure 6-63 Chassis selection

A graphical front view of the chassis is shown, as shown in Figure 6-64. You can get information about chassis components by positioning the cursor over them.




Figure 6-64 Front chassis view (left part of the window)

You also see a rear view that functions in the same way, as shown in Figure 6-65.



Figure 6-65 Rear chassis view (right part of the window)



2. Select your chassis component (in this case, the switch module in bay 1), as shown in Figure 6-66.



*Figure 6-66 Select IO module component from the physical view*

3. Scroll down to the Action menu and select **Security** → **Configure Access**, as shown in Figure 6-67.



Figure 6-67 Configure access

The general access status is shown as Partial access, as shown in Figure 6-68.



Figure 6-68 I/O switch module partial access

4. Different protocols are available, and most have no access. Click **Request Access**, as shown in Figure 6-69.

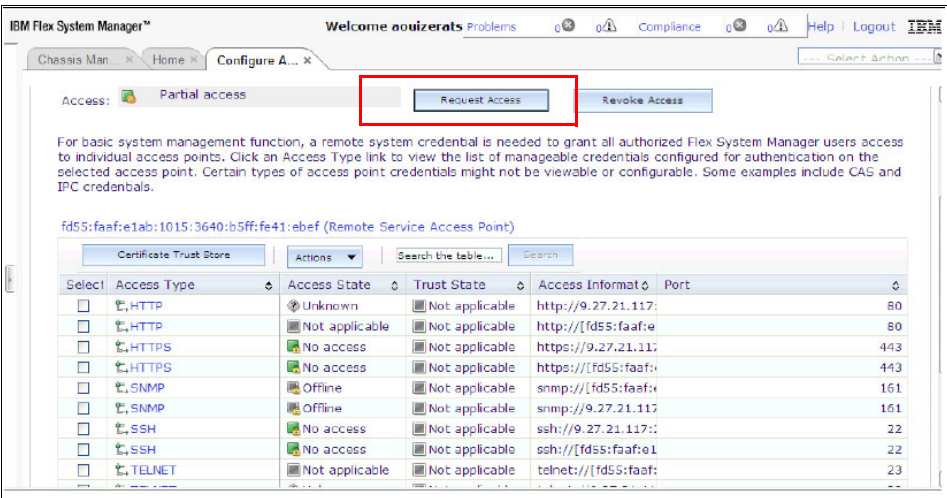


Figure 6-69 IBM switch protocols

5. Enter the credentials for your I/O module and click **Request Access**, as shown in Figure 6-70.

The 'Request Access' dialog box contains the following elements:

- Title Bar:** Request Access
- Instructions:** Specify the user ID and password to authenticate Flex System Manager to one or more target systems, grant all authorized Flex System Manager users access to the target system(s).
- \*User ID:** A text input field containing 'USERID'.
- \*Password:** A password input field with masked characters '\*\*\*\*\*'.
- Buttons:** 'Request Access' (highlighted with a red rectangle) and 'Close'.
- Selected targets:** A table with the following data:

Name	Access	Trust State
fd55:faaf:e1ab:1015:36	Partial access	Trusted
- Page Info:** Page 1 of 1, Total: 1

Figure 6-70 Partial access on an Ethernet I/O module

You might receive a message that not all of the protocols are enabled on the managed component. This error message indicates that not all discovery protocols are supported by the switch, as shown in Figure 6-71.

The 'Request Access' dialog box displays an error message with the following details:

- Title Bar:** Request Access
- Error Message:** The Request Access attempt was not successful on one or more target systems within the Access column for the most current access level. Resources with OK access have full access. Resources that do not display OK access might have a request access failure. For more handling request access failures, see the product documentation.
- Buttons:** 'Product Documentation' and 'Close Message'.
- Instructions:** Specify the user ID and password to authenticate Flex System Manager to one or more target systems, grant all authorized Flex System Manager users access to the target system(s).
- \*User ID:** A text input field containing 'USERID'.
- \*Password:** An empty password input field.
- Buttons:** 'Retry on Failed' and 'Close'.

Figure 6-71 Partial access message

6. Scroll down to note that more protocols are enabled now, as shown in Figure 6-72.

IBM Flex System Manager™

Welcome **anouizerats** Problems Compliance Help | Logout

Chassis Man... Home Configure A... Request Acc... Select Action

For basic system management function, a remote system credential is needed to grant all authorized Flex System Manager users access to individual access points. Click an Access Type link to view the list of manageable credentials configured for authentication on the selected access point. Certain types of access point credentials might not be viewable or configurable. Some examples include CAS and IPC credentials.

fd55:faaf:e1ab:1015:3640:b5ff:fe41:ebef (Remote Service Access Point)

Select	Access Type	Access State	Trust State	Access Informat	Port
<input type="checkbox"/>	HTTP	Unknown	Not applicable	http://9.27.21.117:	80
<input type="checkbox"/>	HTTP	Not applicable	Not applicable	http://[fd55:faaf:e1ab:1015:3640:b5ff:fe41:ebef]:	80
<input type="checkbox"/>	HTTPS	OK	Not applicable	https://9.27.21.117:	443
<input type="checkbox"/>	HTTPS	OK	Not applicable	https://[fd55:faaf:e1ab:1015:3640:b5ff:fe41:ebef]:	443
<input type="checkbox"/>	SNMP	Offline	Not applicable	snmp://[fd55:faaf:e1ab:1015:3640:b5ff:fe41:ebef]:	161
<input type="checkbox"/>	SNMP	Offline	Not applicable	snmp://9.27.21.117:	161
<input type="checkbox"/>	SSH	OK	Not applicable	ssh://9.27.21.117:	22
<input type="checkbox"/>	SSH	No access	Not applicable	ssh://[fd55:faaf:e1ab:1015:3640:b5ff:fe41:ebef]:	22
<input type="checkbox"/>	TELNET	Not applicable	Not applicable	telnet://[fd55:faaf:e1ab:1015:3640:b5ff:fe41:ebef]:	23
<input type="checkbox"/>	TELNET	Unknown	Not applicable	telnet://9.27.21.117:	23

Figure 6-72 Partial access with more protocols

**Remember:** Some protocols must be directly enabled on the I/O module. For example, if SSH was not enabled, you must enable it on the switch before you enable access for this protocol from Flex System Manager. To enable SNMP on the switch, you must configure SNMP credentials.

- To collect inventory for the I/O modules select **Action** → **Your I/O module name** → **Inventory** → **View and Collect Inventory**, as shown in Figure 6-73.

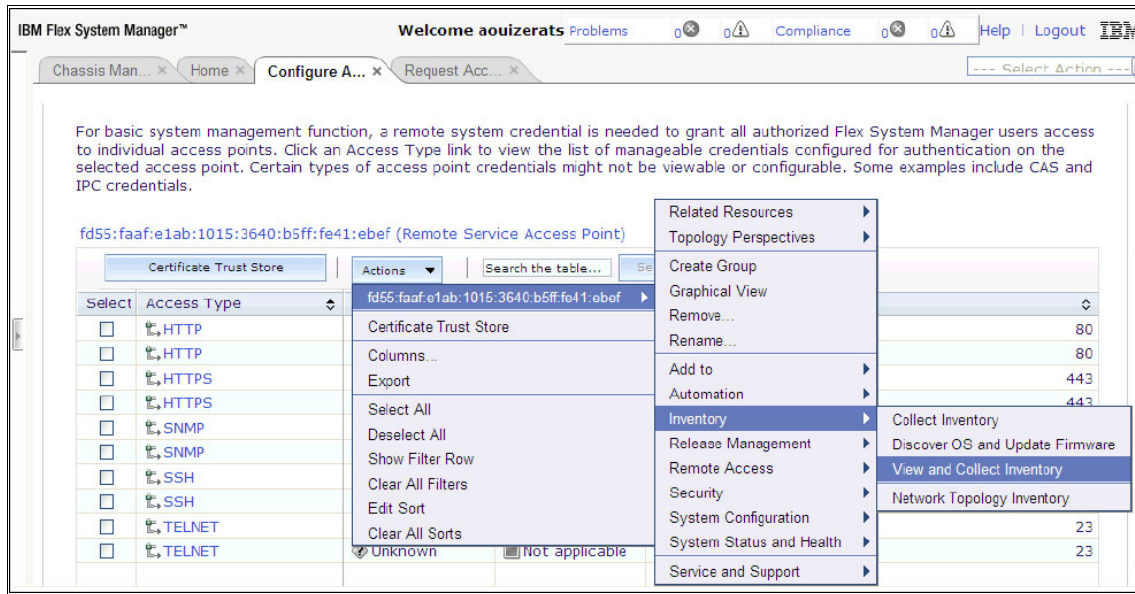


Figure 6-73 I/O module inventory collection

- Click **Collect Inventory**, as shown in Figure 6-74.

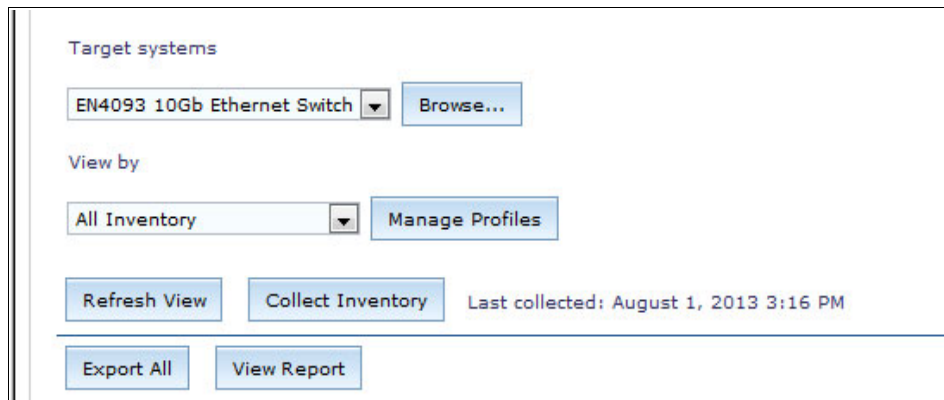


Figure 6-74 Collect Inventory



9. Click **OK** to run your collection task, as shown in Figure 6-75.

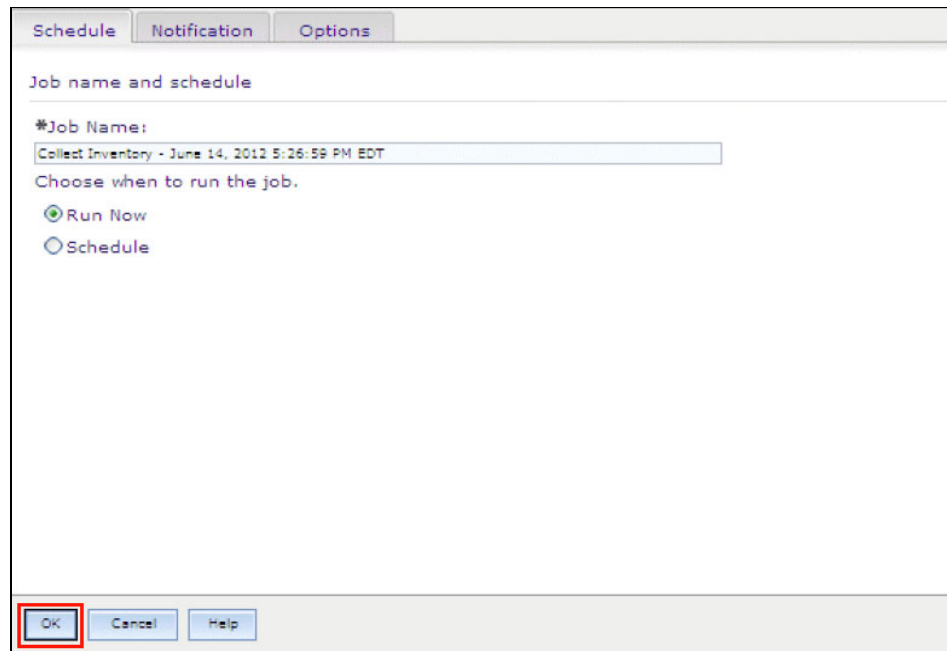


Figure 6-75 Run job

10. Click **Display Properties** to see the job status, as shown in Figure 6-76.

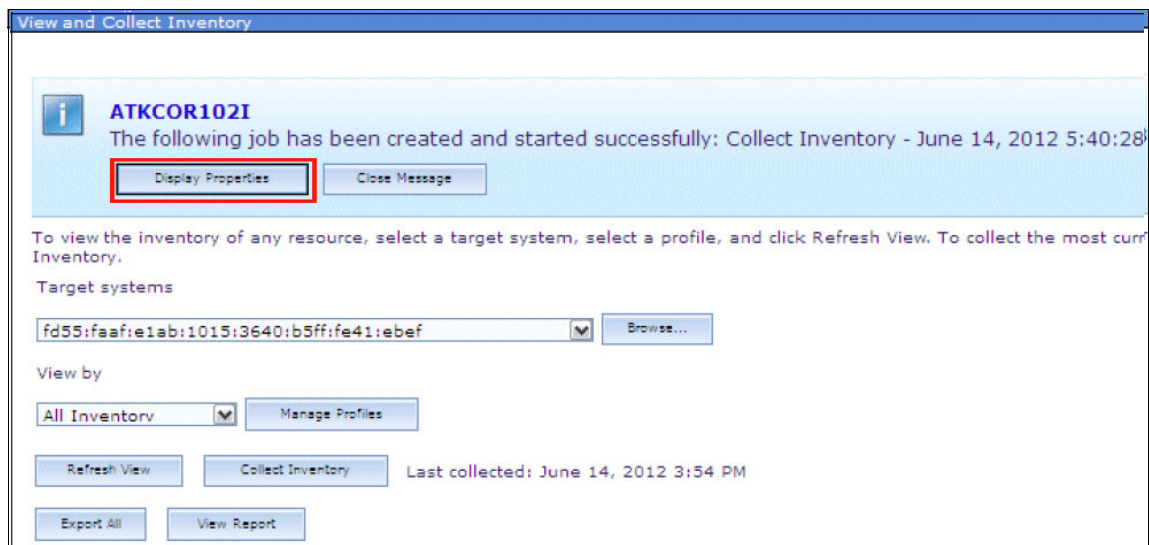


Figure 6-76 Job information

It takes a few minutes for the job to complete. You can check the status, as shown in Figure 6-77.

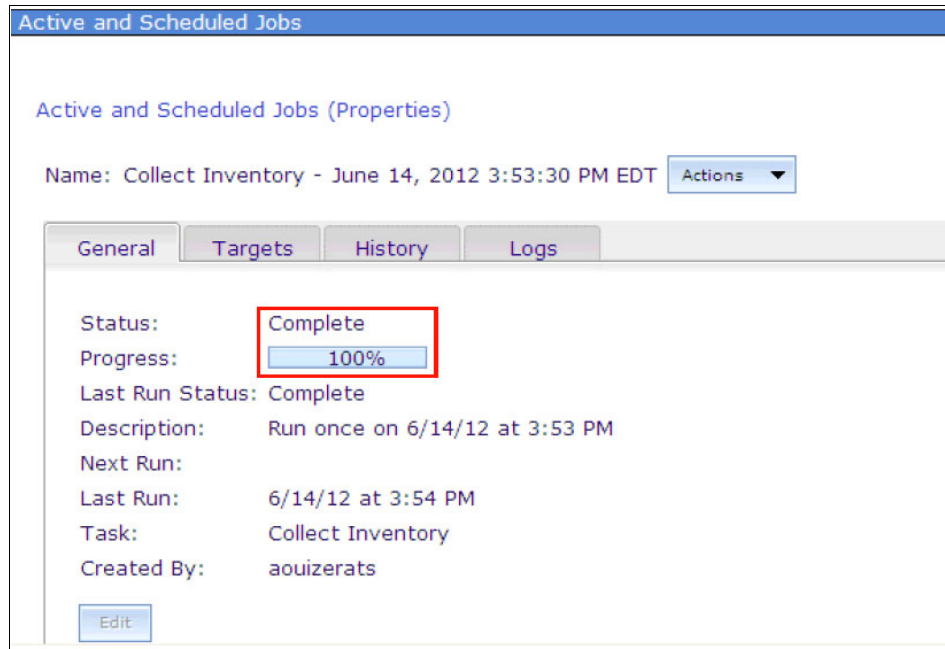


Figure 6-77 Job completed

## 6.5 IBM Flex System Fabric EN4093 10Gb configuration

In this section, we describe the configuration of the Flex System Fabric EN4093 10Gb, as a preparatory step for the configuration that was designed previously.

The virtual network interface card (vNIC) is configured on top of a system's physical network adapter. Flex System can have a maximum of four vNIC for each physical network adapter port. To gain the maximum benefit from Flex System, the examples in this book use a virtual fabric adapter vNIC that also offers flexibility for bandwidth configuration.

Virtual fabric adapters are switch-dependent, which means that no manual intervention is needed on the Ethernet adapter to change bandwidth allocation or to expand vNIC group membership. This type of configuration is made centrally on the switch from Flex System Manager by using server patterns.

To connect specific vNICs each other, you create vNIC *groups*. You also use vNIC group to bridge an internal switch port to an external switch port or trunk.



For each vNIC group that exists in the Flex System Fabric EN4093, you configure an internal, unique, and independent VLAN. Then, that VLAN is used only by the vNIC group members. All packet inside the VLAN is tagged with the internal VLAN tag. When the packet leaves the vNIC group, the tag is deleted automatically and the resulting frame is untagged.

The Flex System Fabric EN4093 10Gb has the following configuration:

- ▶ 14 internal ports activated (from INTA1 to INTA14)
- ▶ 10 external ports activated (from EXT1 to EXT10)

Because vNIC was distributed within configuration patterns, as described in 6.6, “IBM Flex System x240 compute node configuration” on page 222, compute nodes 1 - 4 are connected on the internal port from INTA1 to INTA4 of the Flex System Fabric EN4093 10Gb Ethernet switch.

Table 6-1 shows the internal vNIC to physical compute node ports association.

Table 6-1 vNIC-to physical association

Physical switch interface	vNIC internal name	Function
Port1	INTA1.1	vNIC1 for the first compute node
	INTA1.2	vNIC2 for the first compute node
	INTA1.3	vNIC3 for the first compute node
	INTA1.4	vNIC4 for the first compute node
Port2	INTA2.1	vNIC1 for the second compute node
	INTA2.2	vNIC2 for the second compute node
	INTA2.3	vNIC3 for the second compute node
	INTA2.4	vNIC4 for the second compute node
Port...	....	....
Portx	INTAx.x	vNIC x for the x compute node

Figure 6-78 shows the vNIC to switch connections. EXT1 (port 43), EXT2 (port 44), EXT3 (port 45), and EXT4 (port 46) are external ports of the network switch.

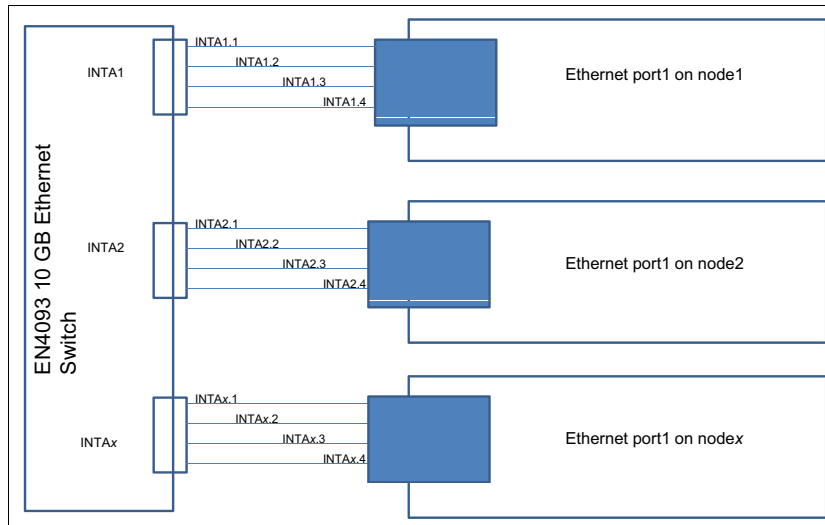


Figure 6-78 vNIC to switch interconnections

The following procedure shows how to configure the switch internal and external ports to reach the following connections:

- ▶ Port EXT2 connected to VLAN 42 and internal vNIC Group 1
- ▶ Port EXT1 connected to VLAN 10 and internal vNIC Group 2
- ▶ Ports EXT3 and EXT4 in Trunk Group for VLAN 20 and internal vNIC Group 3
- ▶ Relevant bandwidth settings for all switch ports

Figure 6-79 shows the relationship between internal interfaces, switch port groups, and external interfaces.

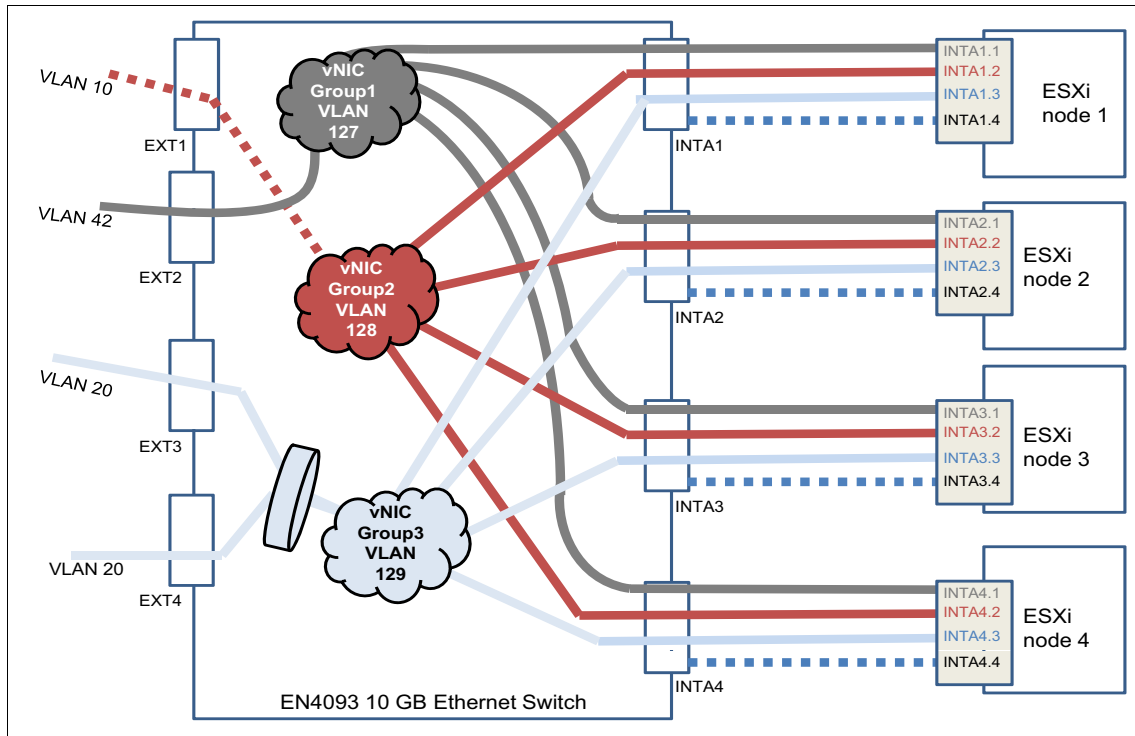


Figure 6-79 Switch configuration

Figure 6-79 shows the following network characteristics:

- ▶ vNIC Group 1 has an outer tag for VLAN 127. The group consists of vNIC pipes INTA1.1, INTA2.1, INTA3.1, INTA4.1 and external uplink port EXT2. This vNIC Group represents the Management network.
- ▶ vNIC Group 2 has an outer tag for VLAN 128. The group consists of vNIC pipes INTA1.2, INTA2.2, INTA3.2, INTA4.2 and external uplink port EXT1 (dotted to show a possible external chassis link to other vMotion networks).
- ▶ vNIC Group 3 has an outer tag for VLAN 129. The group consists vNIC pipes INTA1.3, INTA2.3, INTA3.3, INTA4.3 and external uplink trunk of port EXT3 and EXT4. This vNIC Group represents the Public network.
- ▶ vNIC bandwidth on ports INTA1.3, INTA2.3, INTA3.3 and INTA4.3 is set to 50%.
- ▶ Interfaces INTAx.4 are disabled and not connected to any vNIC Group.

Complete the following steps to configure the EN4093 switch:

1. Connect to the EN4093 switch from a web browser by using HTTPS. At the login window, enter the default credentials and click **Submit**, as shown on Figure 6-80.

The image shows a web browser window with a light beige background. At the top, the text "Login to" is centered in a dark blue font. Below it, the text "IBM Flex System Fabric EN4093 10Gb Scalable Switch(Upgrade1)" is centered in a smaller, dark blue font. In the center of the page, there is a login form. The form has two input fields: "Username:" with the text "USERID" entered, and "Password:" with seven black dots entered. Below the input fields are two buttons: "Submit" and "Reset".

Username:	USERID
Password:	•••••••
<input type="button" value="Submit"/> <input type="button" value="Reset"/>	

*Figure 6-80 EN4093 login window*

**Note:** The default login and password for EN4093 are USERID/PASSW0RD (where the sixth character is the number zero, not the letter O).

2. Select **Configure** in the upper frame. From the left menu, click **Virtualization** → **VNIC** → **General** and select **On** from the Global VNIC On/Off drop-down menu. Click **Submit**, as shown on Figure 6-81.

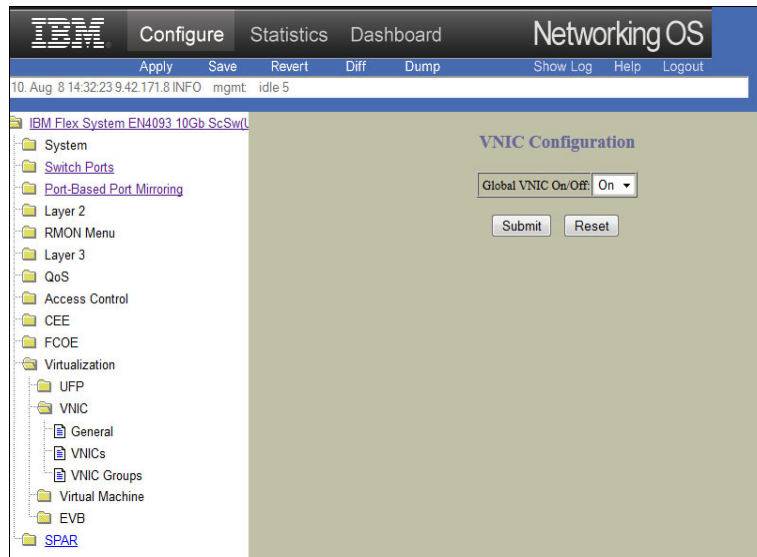


Figure 6-81 Enable vNIC feature

3. Under the Configure menu, click **Apply**.

**Tips:** No feedback is received after the Apply or Save action menu is selected. Also, the Show Log function is useful to monitor switch activities and configuration changes.

4. From the left frame, select **Layer 2** → **Trunk Groups**, as shown in Figure 6-82. Click **Trunk Group 1**.

The screenshot shows the IBM Networking OS configuration interface. The top navigation bar includes 'Configure', 'Statistics', and 'Dashboard'. Below this is a status bar with 'Apply', 'Save', 'Revert', 'Diff', 'Dump', 'Show Log', 'Help', and 'Logout'. The main content area is titled 'Trunk Groups Configuration'. On the left, a tree view shows the following structure:

- System
  - Switch Ports
  - Port-Based Port Mirroring
  - Layer 2
    - 802.1x
    - FDB
    - Virtual LANs
    - Spanning Tree Groups
    - MSTP/RSTP/PVRST
    - LLDP
    - ECP
    - Failover
    - Hot Links
    - Trunk Groups
    - Trunk Hash
    - LACP
    - PVST+ compatibility
    - VLAN Auto STG Assignment
    - MAC Address Notification
  - Layer 3
  - QoS

The 'Trunk Groups Configuration' table is as follows:

Trunk Group	State
<u>1</u>	disabled
<u>2</u>	disabled
<u>3</u>	disabled
<u>4</u>	disabled
<u>5</u>	disabled
<u>6</u>	disabled
<u>7</u>	disabled
<u>8</u>	disabled
<u>9</u>	disabled
<u>10</u>	disabled
<u>11</u>	disabled
<u>12</u>	disabled
<u>13</u>	disabled
<u>14</u>	disabled
<u>15</u>	disabled
<u>16</u>	disabled
<u>17</u>	disabled
<u>18</u>	disabled

Figure 6-82 Trunk Groups configuration

5. Configure Trunk Group 1 as it is shown in Figure 6-83. Click **Submit**.

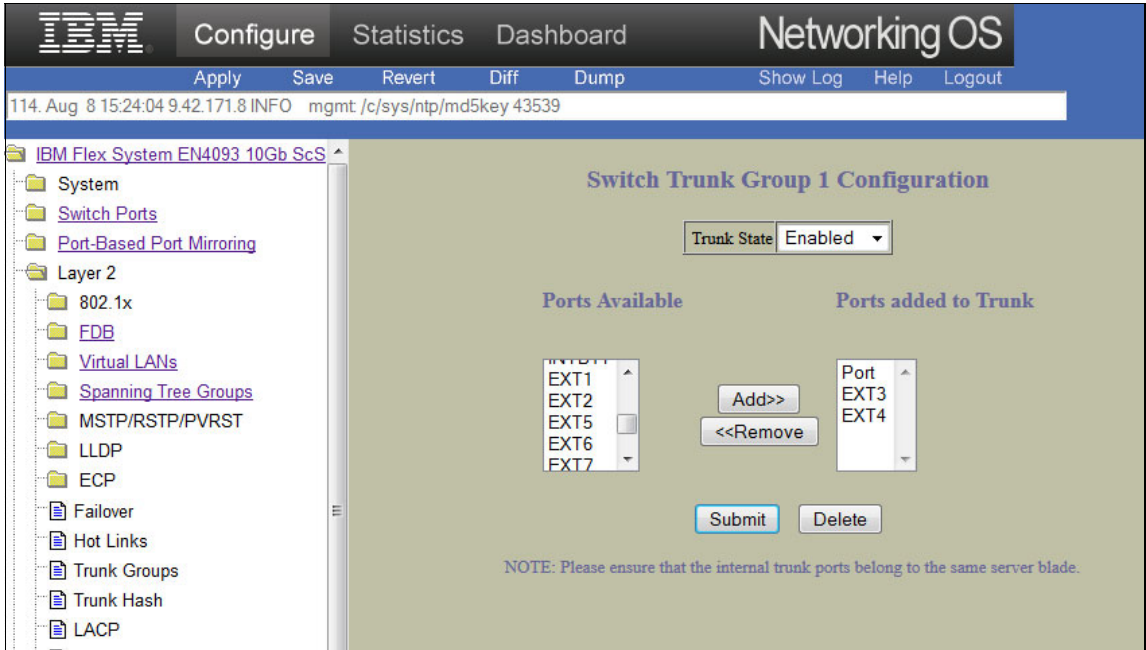


Figure 6-83 Trunk Group 1 configuration

6. Go to **Virtualization** → **VNIC** → **VNIC Groups** and select **vNIC Group 1**, as shown in Figure 6-84.

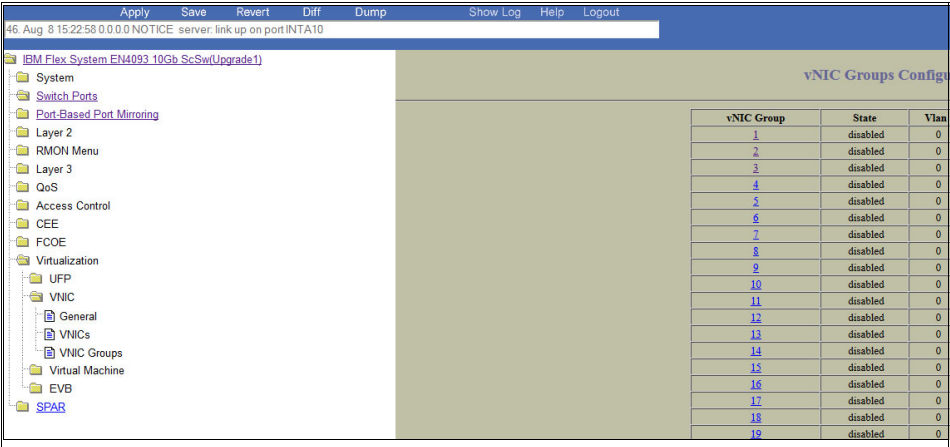


Figure 6-84 vNIC Groups

7. Configure the vNIC Group 1 as it is shown on Figure 6-85. Click **Submit**.

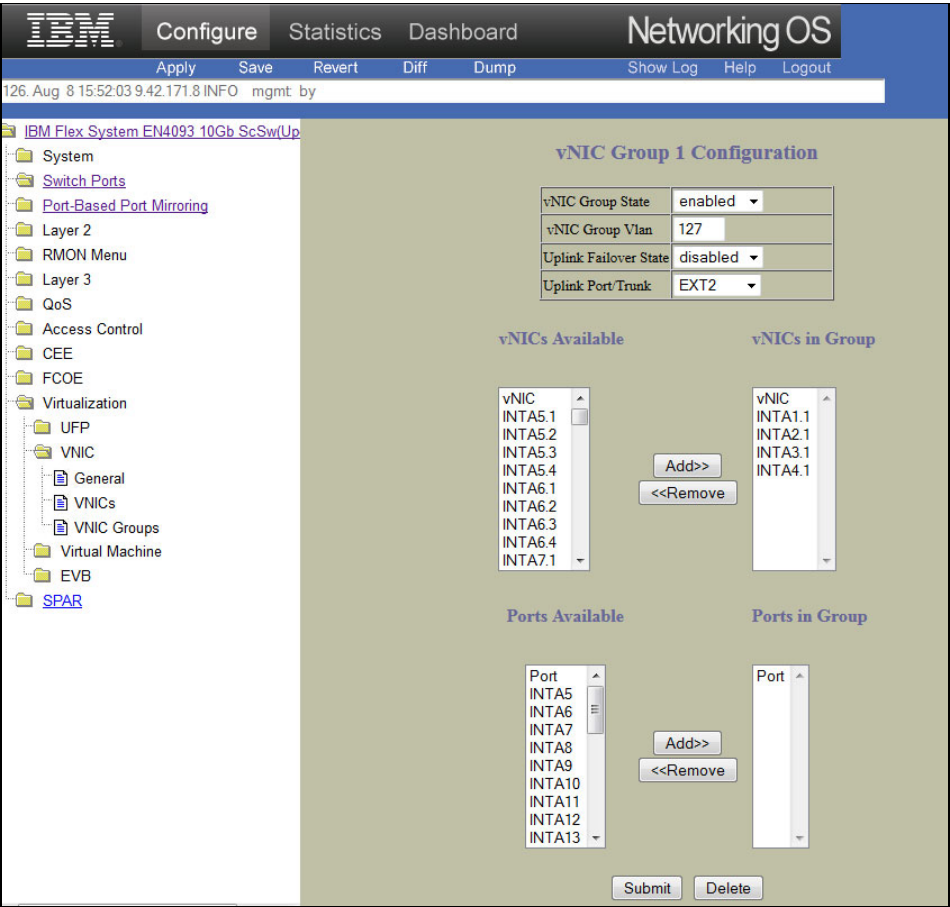


Figure 6-85 vNIC Group 1 configuration

8. Select **OK** when you are prompted to enable the vNIC interfaces, as shown in Figure 6-86.




Figure 6-86 Interface enabling window



9. Select **VNIC Groups** in the right menu and click to select **vNIC Group 2**. Complete the information as shown in Figure 6-87. Click **Submit**.

IBM Configure Statistics Dashboard Networking OS

Apply Save Revert Diff Dump Show Log Help Logout

200 Aug 8 16:54:26 9.42.171.8 INFO mgmt by

Port-Based Port Mirroring

Layer 2

802.1x

FDB

Virtual LANs

Spanning Tree Groups

MSTP/RSTP/PVRST

LLDP

ECP

Failover

Hot Links

Trunk Groups

Trunk Hash

LACP

PVST+ compatibility

VLAN Auto STG Assignment

MAC Address Notification

RMON Menu

Layer 3

QoS

Access Control

CEE

FCOE

Virtualization

UFP

VNIC

General

VNICs

VNIC Groups

Virtual Machine

EVB

SPAR

vNIC Group 2 Configuration

vNIC Group State	enabled
vNIC Group Vlan	128
Uplink Failover State	disabled
Uplink Port/Trunk	EXT1

vNICs Available

vNICs in Group

Port

Submit Delete

Figure 6-87 vNIC Group 2 configuration

10. Select **OK** when you are prompted to enable the vNIC interfaces, as shown in Figure 6-88.

Confirm enabling VNICs INTA1.2, INTA2.2, INTA3.2, INTA4.2?

OK Cancel

Figure 6-88 Interface enabling window

11. Select **VNIC Groups** in the right menu, and click to select **vNIC Group 3**. Complete the information as shown in Figure 6-89. Click **Submit**.

IBM Networking OS

Configure Statistics Dashboard

Apply Save Revert Diff Dump Show Log Help Logout

168 Aug 8 16:54:20 9.42.171.8 INFO mgmt admin

Layer 2

- 802.1x
- FDB
- Virtual LANs
- Spanning Tree Groups
- MSTP/RSTP/PVRST
- LLDP
- ECP
- Failover
- Hot Links
- Trunk Groups
- Trunk Hash
- LACP
- PVST+ compatibility
- VLAN Auto STG Assignment
- MAC Address Notification

Layer 3

- QoS
- Access Control
- CEE
- FCOE
- Virtualization
- UFP
- VNIC
- General
- VNICs
- VNIC Groups
- Virtual Machine
- EVB

SPAR

vNIC Group 3 Configuration

vNIC Group State	enabled
vNIC Group Vlan	129
Uplink Failover State	disabled
Uplink Port/Trunk	Trunk 1

vNICs Available

vNICs in Group

Ports Available

Ports in Group

Submit Delete

Figure 6-89 vNIC Group 3 configuration

12. Select **OK** when you are prompted to enable the vNIC interfaces, as shown in Figure 6-90.

Confirm enabling VNICs INTA1.3, INTA2.3, INTA3.3, INTA4.3?

OK Cancel

Figure 6-90 vNIC Group 3 configuration

13. Select **VNICs** in the right menu and select **INTA1.3**. Configure the bandwidth as show in Figure 6-91. Repeat this step for INTA2.3, INTA3.3, and INTA4.3.

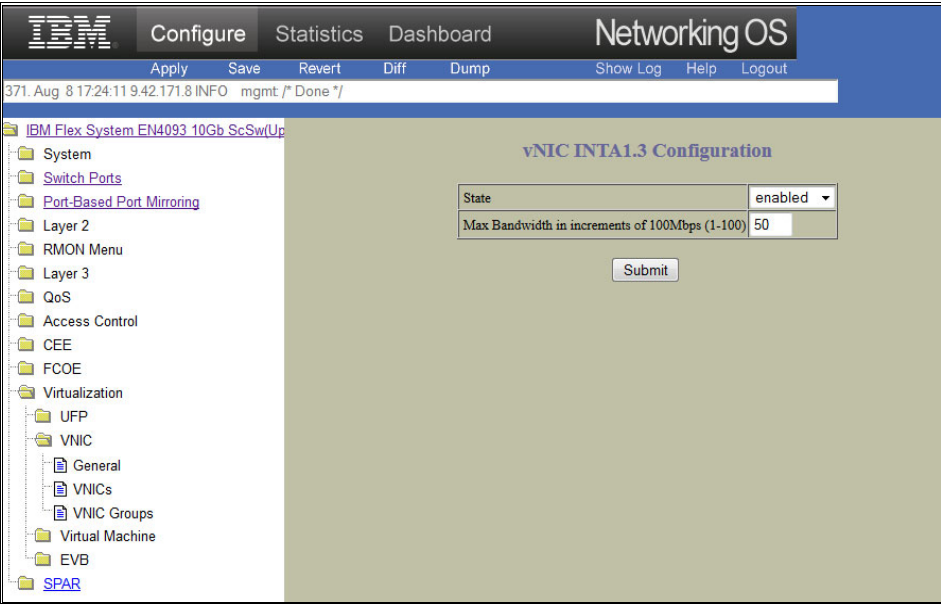


Figure 6-91 Bandwidth for INTA1.3

14. Click **Apply** in the upper menu and then click **Save**. The final vNIC Groups configuration displays as shown in Figure 6-92.

vNIC Groups Dashboard					
vNIC Group	State	Vlan	Failover State	Uplink Port/Trunk	vNICs
1	enabled	127	disabled	EXT2	INTA1.1 INTA2.1 INTA3.1 INTA4.1
2	enabled	128	disabled	EXT1	INTA1.2 INTA2.2 INTA3.2 INTA4.2
3	enabled	129	disabled	Trunk 1	INTA1.3 INTA2.3 INTA3.3 INTA4.3
4	disabled	0	disabled	none	empty
5	disabled	0	disabled	none	empty
6	disabled	0	disabled	none	empty
7	disabled	0	disabled	none	empty
8	disabled	0	disabled	none	empty
9	disabled	0	disabled	none	empty
10	disabled	0	disabled	none	empty

Figure 6-92 vNIC Groups final configuration

15. In the top menu, select **Dashboard**. From the left menu, select **Layer 2** → **Virtual LANs**. Figure 6-93 shows the final configuration.

IBM

Configure

Statistics

Dashboard

Networking OS

ApplySaveRevertDiffDumpShow LogHelpLogout

258, Aug 8 16:58:45 9.42.171.8 NOTICE: dcbx: Detected DCBX peer on port INTA10

IBM Flex System EN4093 10Gb ScSw(Upgrade1)

System

Switch Ports

Port-Based Port Mirroring

Layer 2

802.1x

FDB

Virtual LANs

Add VLAN

Spanning Tree Groups

MSTP/RSTP/PVRST

LLDP

ECP

Failover

Hot Links

Trunk Groups

Trunk Hash

LACP

PVST+ compatibility

VLAN Auto STG Assignment

MAC Address Notification

RMON Menu

Layer 3

VLANs Dashboard

1. Search Range

VLAN ID (1 - 4095) From 1 To 4095

2. Search Options

VLAN Name

VLAN State any

Search Operation or

Search

Number of existing VLANs: 5 Enabled: 5 Disabled: 0

VLAN ID	VLAN Name	VLAN Ports	VLAN VPorts	VLAN Type	Private VLAN Type	Private VLAN Map	Private VLAN State	Management VLAN State	
1	Default VLAN	INTA1-INTB14 EXT5-EXT10 EXT15-EXT22	empty	Port based	empty		disabled	disabled	en
127	VLAN 127	INTA1-INTA4 EXT2	empty	Port based	empty		disabled	disabled	en
128	VLAN 128	INTA1-INTA4 EXT1	empty	Port based	empty		disabled	disabled	en
129	VLAN 129	INTA1-INTA4 EXT3 EXT4	empty	Port based	empty		disabled	disabled	en
4095	Mgmt VLAN	EXTM MGT1	empty	Port based	empty		disabled	enabled	en

Figure 6-93 Final VLANs configuration

16. Select **Virtualization** → **VNIC** → **VNICs**. Figure 6-94 shows the final vNICs configuration.

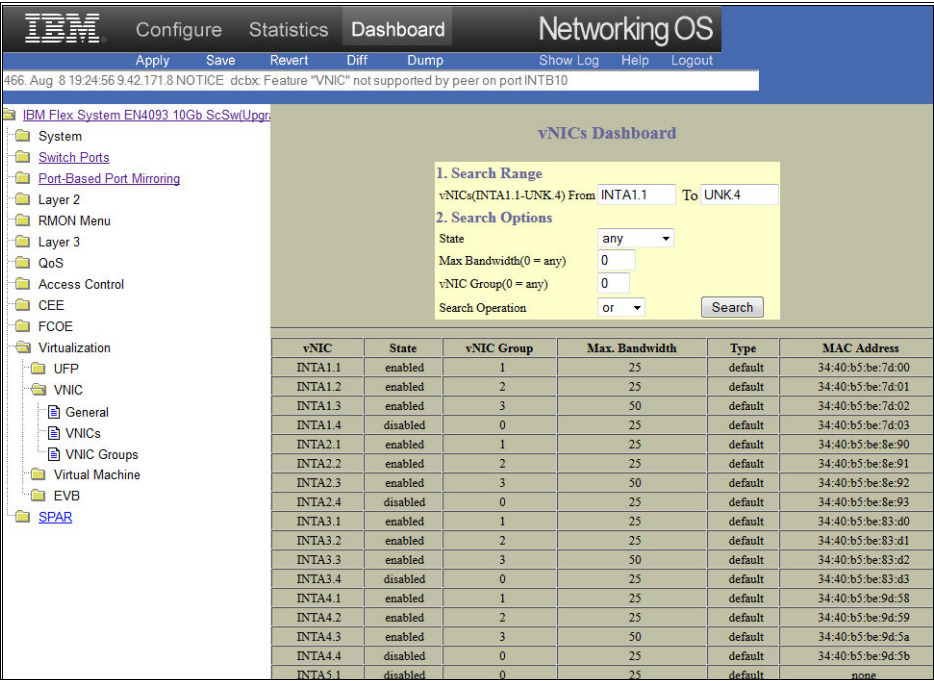


Figure 6-94 vNICs final configuration

You can also connect to the EN4093 Switch by using the CLI that uses SSH.

The following examples show the configuration from a CLI perspective.  
Example 6-1 shows all information that is related to the switch physical ports.

Example 6-1 Output of /info/port

```
>> Main# /info/port
```

Alias	Port	Tag	Type	RMON	Ln	Fld	PVID	NAME	VLAN(s)
INTA1	1	y	Internal	d	e	e	1	INTA1	1 127 128 129
INTA2	2	y	Internal	d	e	e	1	INTA2	1 127 128 129
INTA3	3	y	Internal	d	e	e	1	INTA3	1 127 128 129
INTA4	4	y	Internal	d	e	e	1	INTA4	1 127 128 129
INTA5	5	n	Internal	d	e	e	1	INTA5	1
INTA6	6	n	Internal	d	e	e	1	INTA6	1
INTA7	7	n	Internal	d	e	e	1	INTA7	1
INTA8	8	n	Internal	d	e	e	1	INTA8	1
INTA9	9	n	Internal	d	e	e	1	INTA9	1

INTA10	10	n	Internal	d	e	e	1	INTA10	1
INTA11	11	n	Internal	d	e	e	1	INTA11	1
INTA12	12	n	Internal	d	e	e	1	INTA12	1
INTA13	13	n	Internal	d	e	e	1	INTA13	1
INTA14	14	n	Internal	d	e	e	1	INTA14	1
INTB1	15	n	Internal	d	e	e	1	INTB1	1
INTB2	16	n	Internal	d	e	e	1	INTB2	1
INTB3	17	n	Internal	d	e	e	1	INTB3	1
INTB4	18	n	Internal	d	e	e	1	INTB4	1
INTB5	19	n	Internal	d	e	e	1	INTB5	1
INTB6	20	n	Internal	d	e	e	1	INTB6	1
INTB7	21	n	Internal	d	e	e	1	INTB7	1
INTB8	22	n	Internal	d	e	e	1	INTB8	1
INTB9	23	n	Internal	d	e	e	1	INTB9	1
INTB10	24	n	Internal	d	e	e	1	INTB10	1
INTB11	25	n	Internal	d	e	e	1	INTB11	1
INTB12	26	n	Internal	d	e	e	1	INTB12	1
INTB13	27	n	Internal	d	e	e	1	INTB13	1
INTB14	28	n	Internal	d	e	e	1	INTB14	1
EXT1	43	n	External	d	e	e	128	EXT1	128
EXT2	44	n	External	d	e	e	127	EXT2	127
EXT3	45	n	External	d	e	e	129	EXT3	129
EXT4	46	n	External	d	e	e	129	EXT4	129
EXT5	47	n	External	d	e	e	1	EXT5	1
EXT6	48	n	External	d	e	e	1	EXT6	1
EXT7	49	n	External	d	e	e	1	EXT7	1
EXT8	50	n	External	d	e	e	1	EXT8	1
EXT9	51	n	External	d	e	e	1	EXT9	1
EXT10	52	n	External	d	e	e	1	EXT10	1
EXT15	57	n	External	d	e	e	1	EXT15	1
EXT16	58	n	External	d	e	e	1	EXT16	1
EXT17	59	n	External	d	e	e	1	EXT17	1
EXT18	60	n	External	d	e	e	1	EXT18	1
EXT19	61	n	External	d	e	e	1	EXT19	1
EXT20	62	n	External	d	e	e	1	EXT20	1
EXT21	63	n	External	d	e	e	1	EXT21	1
EXT22	64	n	External	d	e	e	1	EXT22	1
EXTM	65	n	Mgmt	d	e	e	4095	EXTM	4095
MGT1	66	y	Mgmt	d	e	e	4095	MGT1	4095

\* = PVID is tagged.

# = PVID is ingress tagged.

Example 6-2 shows all defined VLANs and their related ports.

*Example 6-2 Output for command /info/l2/vlan*

---

```
>> Information# /info/l2/vlan
```

VLAN	Name	Status	MGT	Ports
----	-----	-----	---	-----
1	Default VLAN	ena	dis	INTA1-INTB14 EXT5-EXT10 EXT15-EXT22
127	VLAN 127	ena	dis	INTA1-INTA4 EXT2
128	VLAN 128	ena	dis	INTA1-INTA4 EXT1
129	VLAN 129	ena	dis	INTA1-INTA4 EXT3 EXT4
4095	Mgmt VLAN	ena	ena	EXTM MGT1

---

Example 6-3 shows the vNICs, bandwidth, VLANs, and the physical MAC address that is connected.

*Example 6-3 Output for command /info/virt/vnic/vnic*

---

```
>> Layer 2# /info/virt/vnic/vnic
```

vNIC	vNICGroup	Vlan	MaxBandwidth	Type	MACAddress	Link
-----	-----	-----	-----	-----	-----	-----
INTA1.1	1	127	25	Default	34:40:b5:be:7d:00	up
INTA1.2	2	128	25	Default	34:40:b5:be:7d:01	up
INTA1.3	3	129	50	Default	34:40:b5:be:7d:02	up
INTA2.1	1	127	25	Default	34:40:b5:be:8e:90	up
INTA2.2	2	128	25	Default	34:40:b5:be:8e:91	up
INTA2.3	3	129	50	Default	34:40:b5:be:8e:92	up
INTA3.1	1	127	25	Default	34:40:b5:be:83:d0	up
INTA3.2	2	128	25	Default	34:40:b5:be:83:d1	up
INTA3.3	3	129	50	Default	34:40:b5:be:83:d2	up
INTA4.1	1	127	25	Default	34:40:b5:be:9d:58	up
INTA4.2	2	128	25	Default	34:40:b5:be:9d:59	up
INTA4.3	3	129	50	Default	34:40:b5:be:9d:5a	up

---

Example 6-4 shows vNIC groups. Uplink port EXT1, EXT3, and EXT4 appear to be down because there is nothing that is connected to these ports.

*Example 6-4 Output for command /info/virt/vnic/vnicgrp*

---

```
>> vNIC Information# /info/virt/vnic/vnicgrp
```

```
-----
----
vNIC Group 1: enabled
```

```

-----
----
VLAN          : 127
Failover      : disabled

vNIC          Link
-----
INTA1.1       up
INTA2.1       up
INTA3.1       up
INTA4.1       up

Port          Link
-----

UplinkPort    Link
-----
EXT2          up

-----
----
vNIC Group 2: enabled
-----
----
VLAN          : 128
Failover      : disabled

vNIC          Link
-----
INTA1.2       up
INTA2.2       up
INTA3.2       up
INTA4.2       up

Port          Link
-----

UplinkPort    Link
-----
EXT1          down

-----
----

```



```

vNIC Group 3: enabled
-----
----
VLAN      : 129
Failover   : disabled

vNIC      Link
-----
INTA1.3    up
INTA2.3    up
INTA3.3    up
INTA4.3    up

Port      Link
-----

UplinkPort Link
-----
EXT3*     down
EXT4*     down
* = The uplink port is in trunk 1

```

---

The switch is now configured to allow vNIC on the internal ports INTA1, INTA2, INTA3, and INTA4.

## 6.6 IBM Flex System x240 compute node configuration

By using *configuration patterns*, you can quickly provision or pre-provision multiple systems from a single pattern. Then, subsequent pattern changes are applied automatically to all associated systems. Server configuration patterns give you the ability to configure local storage, I/O adapter, boot order, and other integrated management module (IMM) and Unified Extensible Firmware Interface (UEFI) settings.

Complete the following steps to configure and deploy the configuration patterns on x240 Compute nodes:

1. In the Initial Setup tab of the Flex System Manager web interface, click **Launch IBM FSM Explorer**, as shown in Figure 6-95.

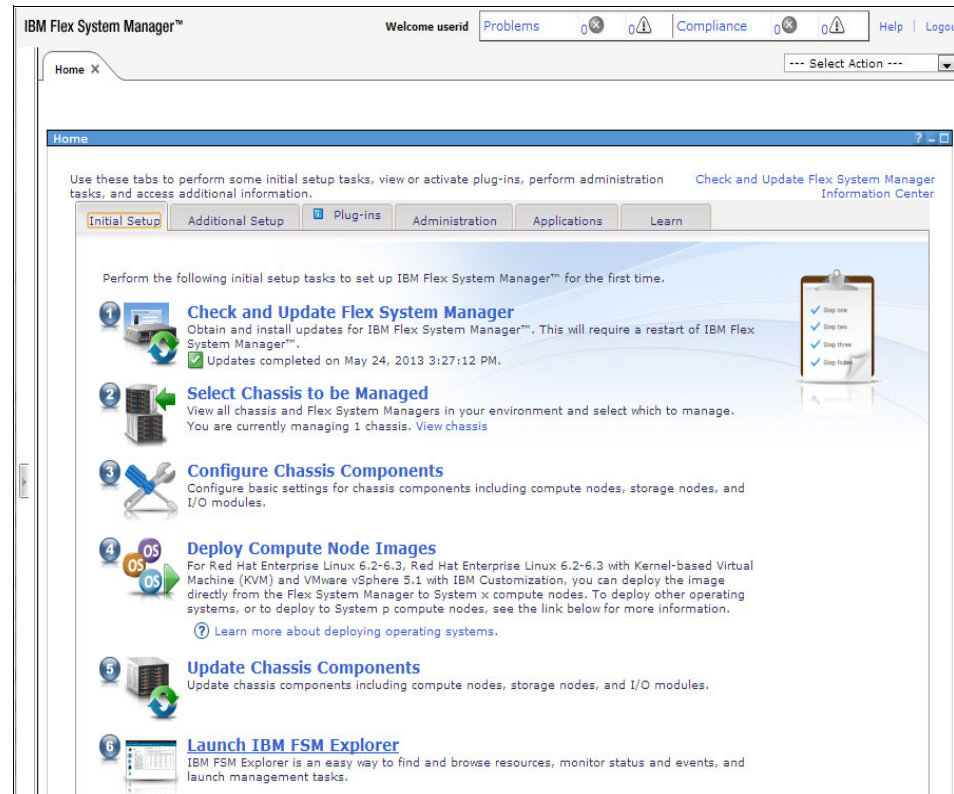


Figure 6-95 Launch Flex System Manager web interface

2. In the Flex System Manager Explorer interface, from the Systems tab, select **Configuration Patterns**, as shown in Figure 6-96.



Figure 6-96 Select Configuration Patterns

3. In the Configuration Patterns: Getting Started window, select **Patterns** under the Servers menu in the left column, as shown in Figure 6-97.

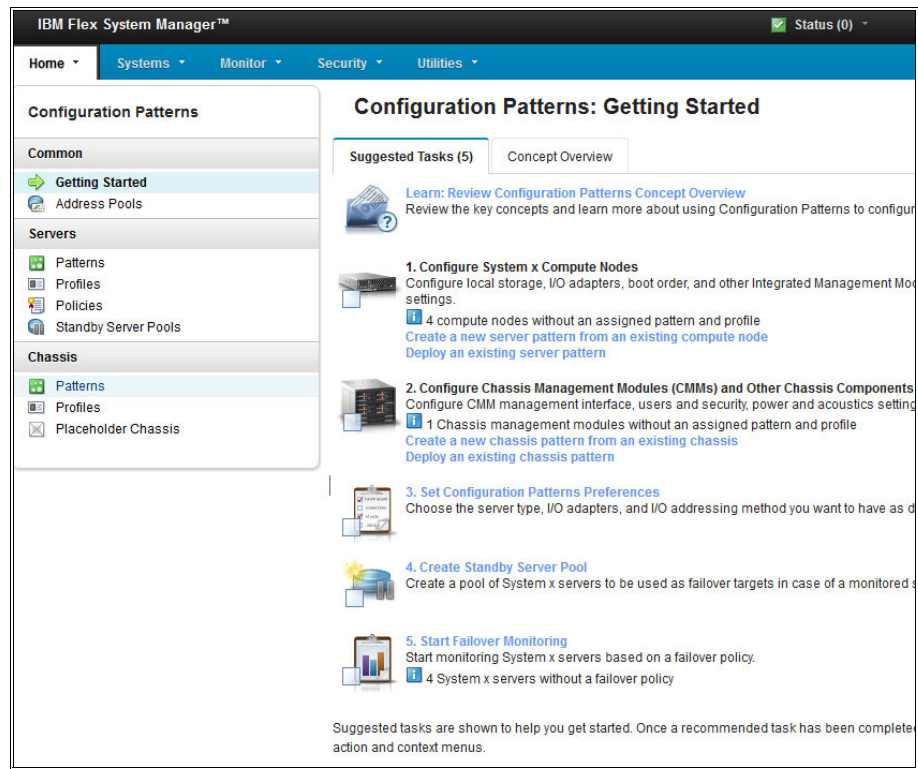


Figure 6-97 Patterns: Getting Started window

4. Select **New Server Pattern**, as shown in Figure 6-98.

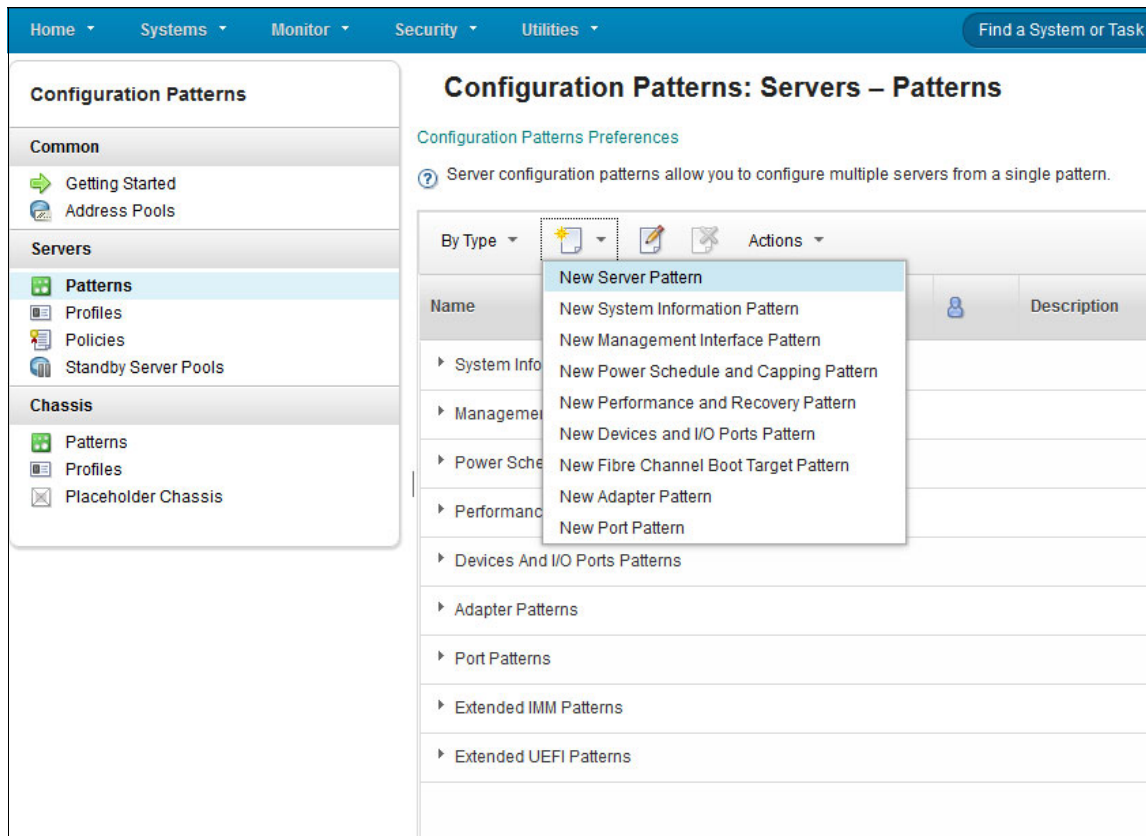


Figure 6-98 Create a server pattern

5. Select **Create a new pattern from scratch**, as shown in Figure 6-99.

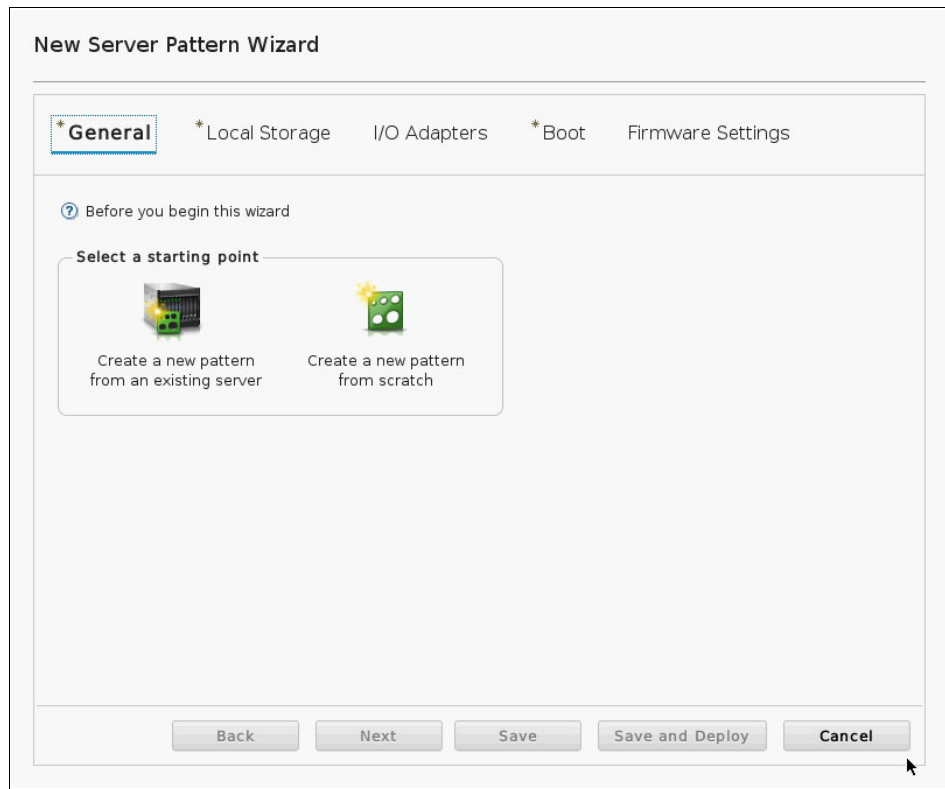


Figure 6-99 New Server Pattern Wizard: General configuration

6. Select the pattern form factor and specify the name of the pattern, as shown in Figure 6-100. Click **Next**.

The screenshot displays the 'New Server Pattern Wizard' interface, specifically the 'General' tab. The wizard is titled 'New Server Pattern Wizard' and has five tabs: 'General', 'Local Storage', 'I/O Adapters', 'Boot', and 'Firmware Settings'. The 'General' tab is active and contains the following sections:

- Before you begin this wizard**: A section with a question mark icon and a 'Select a starting point' button with a pencil icon. Below it is a button labeled 'Create a new pattern from scratch'.
- Specify pattern form factor**: A section with a 'Form Factor' dropdown menu set to '1 Bay Compute Node'.
- Specify pattern name and description**: A section with a '+Name:' field containing 'ESXi VFA Server Pattern' and a 'Description (limit of 500 characters)' text area.

At the bottom of the wizard, there are five buttons: 'Back', 'Next', 'Save', 'Save and Deploy', and 'Cancel'. The 'Next' button is highlighted in a darker shade, indicating it is the recommended action.

Figure 6-100 New Server pattern Wizard: General tab




7. Select **Specify storage configuration** to configure RAID1 automatically for the internal compute nodes disks, as shown in Figure 6-101. Click **Next**.

**Edit Server Pattern Wizard**

\* General    **\* Local Storage**    I/O Adapters    \* Boot    Firmware Settings

Define the storage configuration that will be applied to target servers when this pattern is deployed.

**Select local storage configuration**

 Specify storage configuration     Keep existing storage configuration on target     Disable local disk

This option provides basic RAID configuration for the local boot device.

**Specify storage configuration settings**

Disk Type: SAS Hard disk drive (HDD) ▼

Raid Level: RAID 1 (Mirroring) ▼

Number of drives: 2 ▲ ▼

A single volume will be created using the available array capacity.

Save    Save As...    Cancel

Figure 6-101 New Server Pattern Wizard: Local Storage



8. Configure I/O Adapters by selecting **Add I/O Adapter 1 or LOM**, as shown in Figure 6-102.

**New Server Pattern Wizard**

\*General \*Local Storage **I/O Adapters** \*Boot Firmware Settings

❓ If desired you can modify adapter addressing and define additional adapters to match the hardware you expect to configure with this pattern.

Graphic view [?](#) I/O adapter addressing: [?](#) **Burned in** Virtual [?](#)

[-](#) [+](#) [+](#) ☐ Advanced Settings | [✎](#) Assign Pattern | More ▾

Location	Type	I/O Bay	Configuration Pattern	I/O Add
▼ Compute Node				
<a href="#">+ Add I/O Adapter 1 or LOM</a>		1-2		
<a href="#">+ Add I/O Adapter 2</a>		3-4		

Figure 6-102 New Server Pattern Wizard: I/O Adapters

9. Select **Embedded 10Gb Virtual Fabric Ethernet Controller**, as shown in Figure 6-103. Click **Add**.

New Server Pattern Wizard

Add I/O Adapter 1 or LOM

Select type of adapter to add


Default ▾	Adapter Description	Physical Ports	Type
	Embedded 1Gb Ethernet Controller (LOM)	2	Ethernet
	Embedded 10Gb Virtual Fabric Ethernet Controller (LOM)	2	Fabric Connector
	IBM Flex System CN4054 10Gb Virtual Fabric Adapter	4	Virtual Fabric
	IBM Flex System EN4132 2-port 10Gb Ethernet Adapter	2	Ethernet
	IBM Flex System EN2024 4-port 1Gb Ethernet Adapter	4	Ethernet
	IBM Flex System FC5022 2-port 16Gb FC Adapter	2	Fibre Channel
	IBM Flex System FC3172 2-port 8Gb FC Adapter	2	Fibre Channel
	IBM Flex System FC3052 2-port 8Gb FC Adapter	2	Fibre Channel

Add Cancel

Figure 6-103 New Server Pattern Wizard, Add I/O Adapter 1, or LOM



10. Select the IBM provided patterns for the Initial Adapter and Initial Port patterns, as shown in Figure 6-104. Click **Add**.



Add I/O Adapter 1 or LOM

Select type of adapter to add 

Embedded 10Gb Virtual Fabric Ethernet Controller (LOM)

Select initial patterns ?

Initial adapter pattern: ? IBM VFA-LOM Virtual Fabric Mode - Ethernet  

Initial port pattern: ? IBM Virtual Fabric Balanced Ethernet  

Add Cancel

Figure 6-104 I/O Adapter specification

11. Click **Next**, as shown in Figure 6-105.

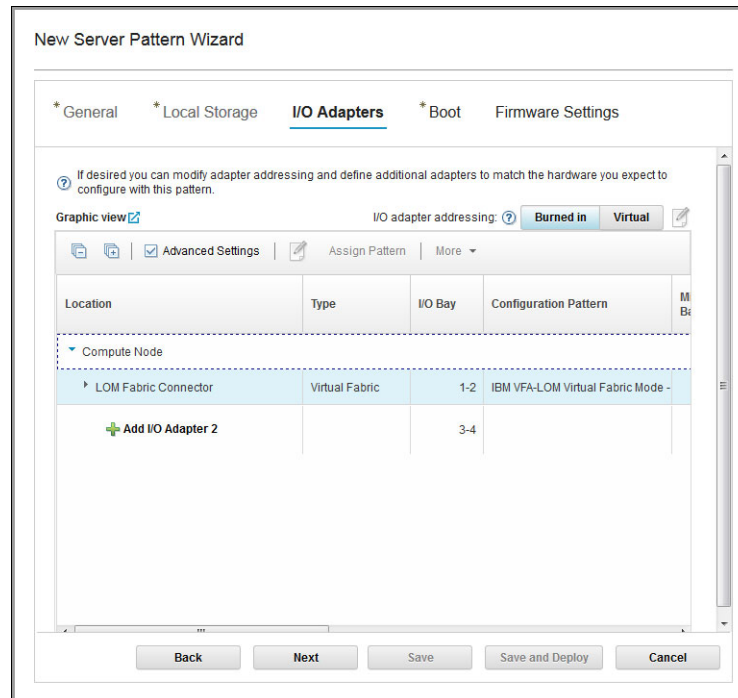


Figure 6-105 The I/O pattern defined

12. Configure the Normal Boot Order to Embedded Hypervisor in the first position of the boot sequence, as shown in Figure 6-106. Click **Next**.

New Server Pattern Wizard

\*General\*Local StorageI/O Adapters\*BootFirmware Settings

Review the boot configuration and modify as desired for the new pattern.

☐ ?UEFI Only Boot☒UEFI First, Then Legacy☐Legacy Only Boot

Normal Boot OrderWake on LAN (WoL) Boot Order

? Specify boot options and boot device sequences

Order	Boot Option	Boot Device	Boot Target Pattern		
1	Embedded Hypervisor			<div><div>+</div><div>×</div></div>	<div><div>↑</div><div>↓</div></div>

BackNextSaveSave and DeployCancel

Figure 6-106 New Server Pattern Wizard: Boot

13. Leave the default firmware settings shown in Figure 6-107. Click **Save and Deploy**.

**New Server Pattern Wizard**

\*General \*Local Storage I/O Adapters \*Boot **Firmware Settings**

**Integrated Management Module (IMM) and Server Firmware Settings (UEFI)**  
Select existing or create new category patterns as desired to include in this server pattern.

Category	Pattern
System Info: ?	— No Pattern Selected —
Management Interface: ?	— No Pattern Selected —
Power Schedule And Capping: ?	— No Pattern Selected —
Performance And Recovery: ?	— No Pattern Selected —
Device And IO Ports: ?	— No Pattern Selected —
Extended IMM: ?	— No Pattern Selected —
Extended UEFI: ?	— No Pattern Selected —

Back Next Save Save and Deploy Cancel

*Figure 6-107 New Server Pattern Wizard: Firmware Settings*

14. Select all of the compute nodes in the left column and import them into the right column, as shown in Figure 6-108.

Deploy Server Pattern - ESXi VFA Server Pattern

Deploy the server pattern to one or more individual servers, or groups of servers (e.g. chassis). On deploy, one server profile is created for each individual server.

\*Pattern To Deploy: ESXi VFA Server Pattern

Profile Activation: ? Full — start/restart server now

Available Servers

Name	Bay	Access
cmm1.itso.ral.ibm.com	...	OK
x240_Node_1	1	OK
x240_Node_2	2	OK
x240_Node_3	3	OK
x240_Node_4	4	OK

Selected Servers

Name	Bay	Access	Deploy Status
No data to display			

Deploy Cancel

Figure 6-108 Select compute nodes for pattern

15. Make sure to select **Full- Start/Restart server now** from the Profile Activation drop-down menu. Click **Deploy**, as shown in Figure 6-109.

Deploy Server Pattern - ESXi VFA Server Pattern

Deploy the server pattern to one or more individual servers, or groups of servers (e.g. chassis). On deploy, one server profile is created for each individual server.

\*Pattern To Deploy: ESXi VFA Server Pattern

Profile Activation: ? Full — start/restart server now

Available Servers

Name	Bay	Access
No data to display		

Selected Servers

Name	Bay	Access	Deploy Status
cmm1.itso.ral.ibm.com			
x240_Node_1	1	OK	Unknown
x240_Node_2	2	OK	Unknown
x240_Node_3	3	OK	Ready

Deploy Cancel

Figure 6-109 Deploy Server Pattern

16. Because the servers are online, a warning prompts you as to whether you want to deploy the patterns and reboot the online compute nodes. Click **Deploy** in the warning message that is displayed, as shown in Figure 6-110.

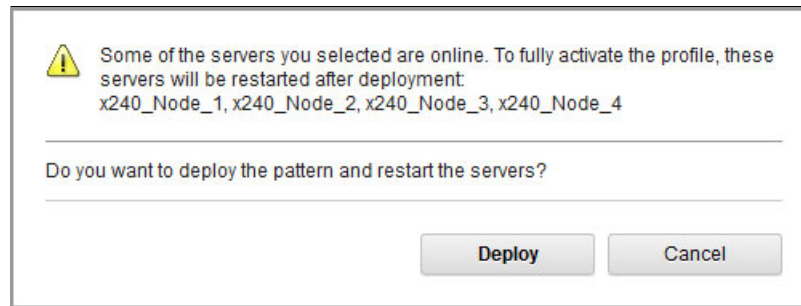


Figure 6-110 Deploy patterns confirmation

When the deploy request is submitted successfully, a message displays, as shown in Figure 6-111.



Figure 6-111 Deploy submission confirmation

The compute nodes are restarted and configured according to the server patterns.

## 6.7 IBM Flex System V7000 Storage Node configuration

In this section, we describe how to configure the Flex System V7000 Storage Node, which has 10 internal drives available (136.23 GB, SAS, 15000 rpm).

For disk space purposes, create a single MDisk that consists of a RAID6 array with nine disks member and one Hot Spare disk. The physical capacity is 942 GB.

This example configures one pool and implements the following data stores:

- ▶ A thin-provisioned volume with a capacity of 400 GB that is used for ESXi management cluster data store.
- ▶ A thin-provisioned volume with a capacity of 600 GB is used as ESXi VDI cluster data store.

The virtual capacity (600 GB) is higher than the physical capacity.

### 6.7.1 IBM Flex System V7000 Storage Node initial configuration

The following procedure guides you through the necessary steps when you are using the Flex System Manager web user interface:

1. Open a web browser and point it to the IP address of the Flex System Manager and log in. The menu panel that is shown in Figure 6-112 on page 237 opens, which features several selections.  
Select **Launch IBM FSM Explorer** from the menu list.



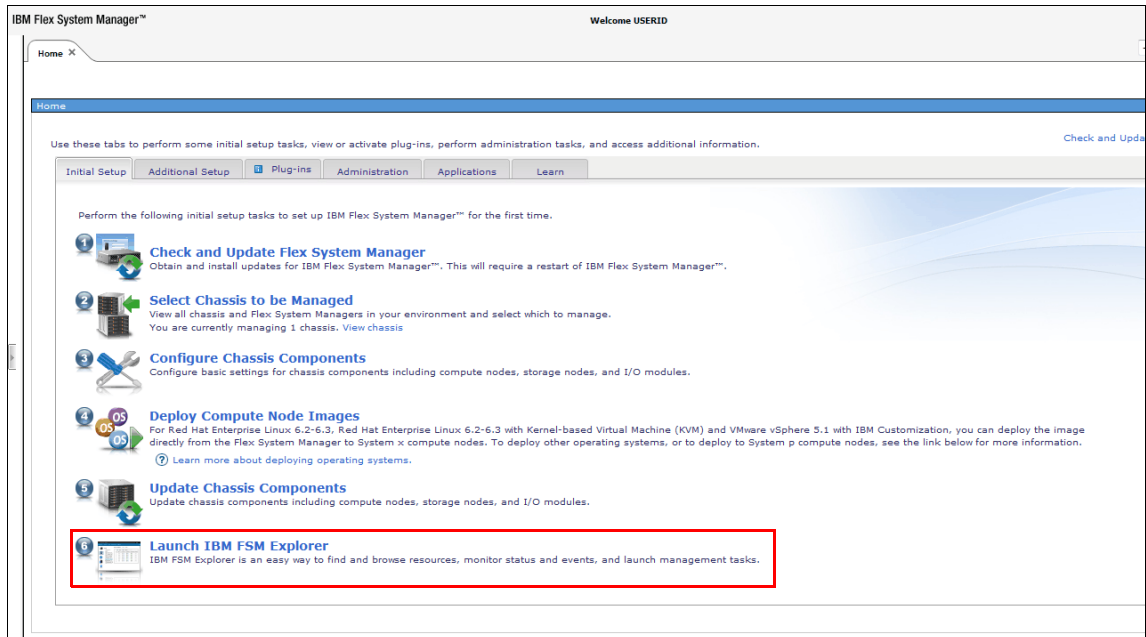


Figure 6-112 Launch IBM FSM Explorer

Notice that a new browser tab opens in which you can select the applicable enclosure from the Chassis Map, as shown in Figure 6-113.

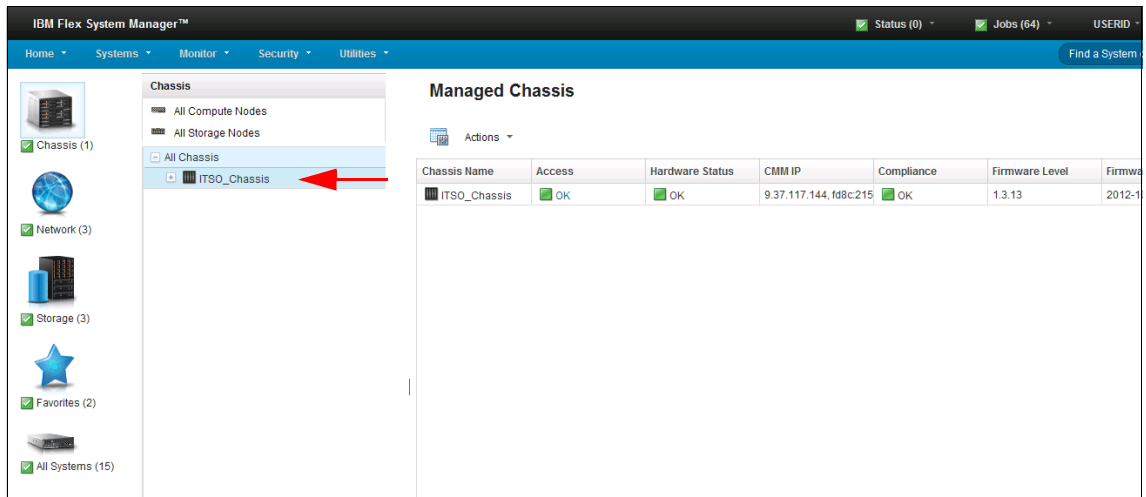


Figure 6-113 Select and launch the chassis in the Chassis Manager

2. In the Chassis Manager, select the applicable chassis that starts the chassis map for that chassis, as shown in Figure 6-114.

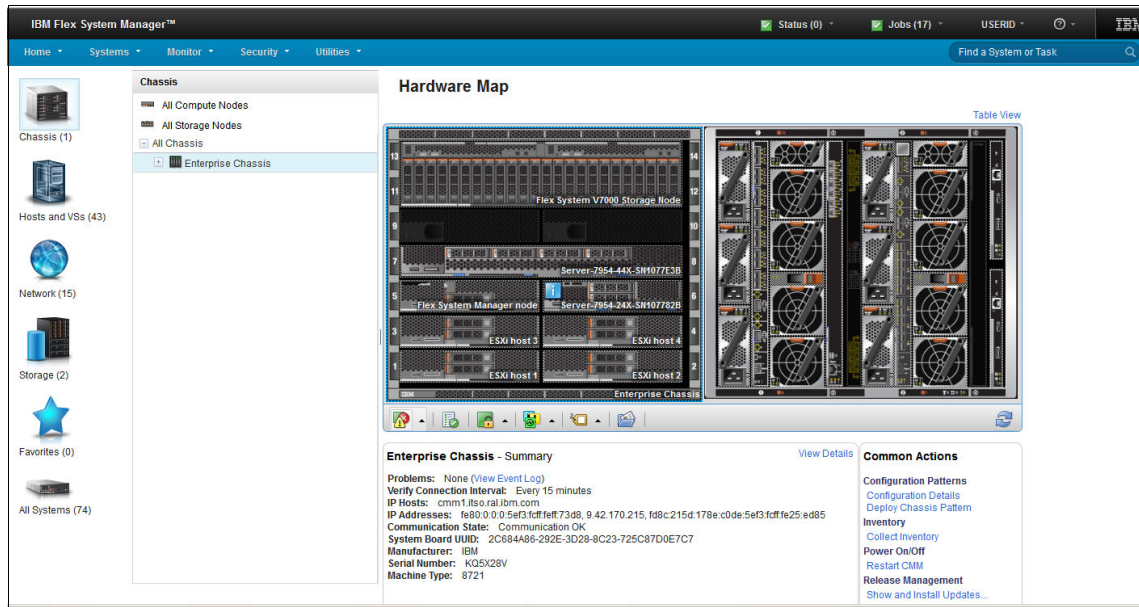


Figure 6-114 IBM Flex System Manager: Hardware Map

3. Right-click IBM Flex System V7000 Storage Node in the chassis map. Select **Remote Access** and then click **Launch IBM Flex System V7000** (as shown in Figure 6-115) to start the Initial Setup wizard.

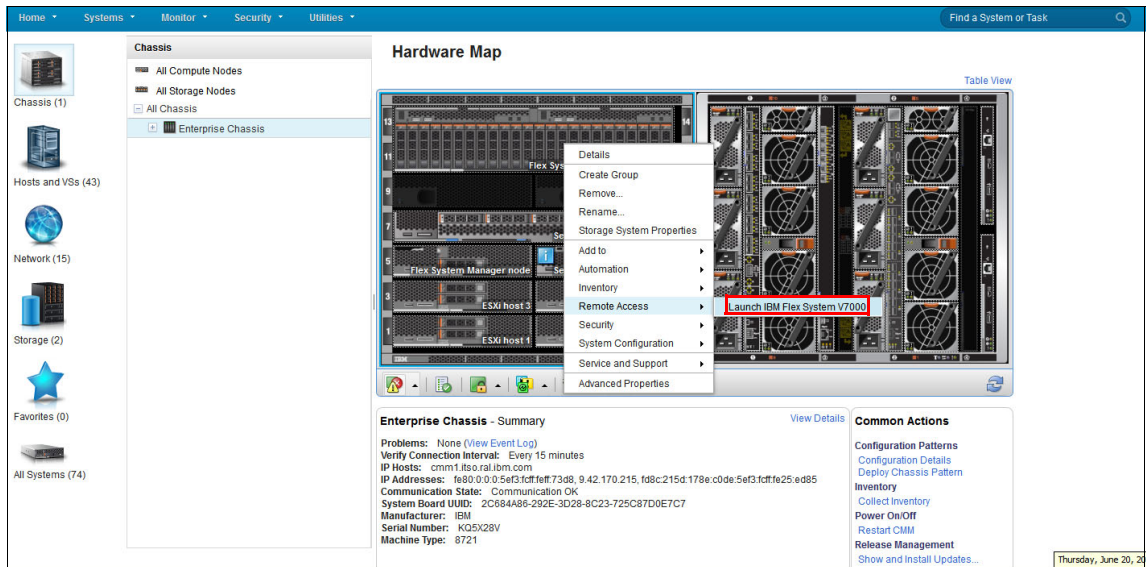


Figure 6-115 Launch Storage Manager

4. The next window is a welcome window from the IBM Flex System V7000 Storage Node interface. You can create a system (cluster) or add to an existing system, as shown in Figure 6-116. This example creates a system. Select **Create a new system** and then click **Next**.



Figure 6-116 First-time setup welcome window

5. Select whether you are using an IPv4 or IPv6 management IP address (as shown in Figure 6-117) and enter the IP address. (You can use DHCP or the static address that was assigned.) The subnet mask and gateway already have defaults that are listed, which you can edit.

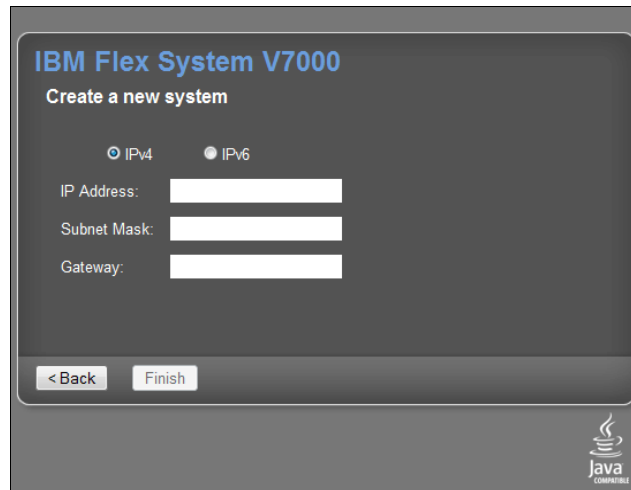


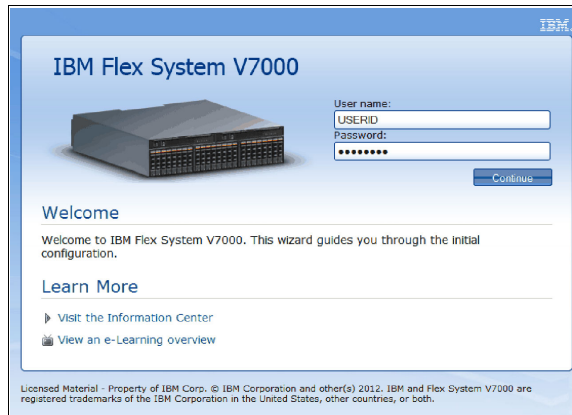
Figure 6-117 Creating a Storage Cluster

6. 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 starts 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.

## 6.7.2 IBM Flex System V7000 Storage Node Setup Wizard

After the initial configuration that is described in 6.7.1, “IBM Flex System V7000 Storage Node initial configuration” on page 236 completes, the IBM Flex System V7000 Storage Node Welcome window opens, as shown in Figure 6-118 on page 241.



*Figure 6-118 IBM Flex System V7000 Storage Node Welcome window*

**Tip:** During the initial setup of IBM Flex System V7000 Storage Node, the installation wizard asks for various information that you need available during the installation process. If you do not have this information available 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 6-119. Click **Next** after you accept the license agreement.

**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

Form 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ADAPTING, MODIFYING OR AN "ACCEPTED" BUTTON, OR OTHERWISE USING THE PROGRAM, LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON BEHALF OF YOURSELF, YOU REPRESENT AND WARRANT THAT YOU HAVE THE AUTHORITY TO BOND LICENSEE TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS, ...

- DO NOT DOWNLOAD, INSTALL, COPY, ADAPT, MODIFY OR AN "ACCEPTED" BUTTON, OR USE THE PROGRAM; OR
- PROMPTLY RETURN THE PHYSICAL MEDIA, DOCUMENTATION, AND PROOF OF SETTLEMENT 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.

**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 copies or other unit ("MSU"), Processor Value Units ("PVU"), 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 6-119 Setup wizard: License Agreement

2. Specify a System Name and Superuser Password, as shown in Figure 6-120. Click **Next**.

**System Name and Superuser Password (Step 2 of 7)**

**System Name**

Name: FlexSystem V2000

**Superuser Password**

New Superuser Password

Verify New Superuser Password

**Next >**

Figure 6-120 Setup wizard: Set system name and superuser password

3. Set up the system date and time, as shown in Figure 6-121. Click **Next**.

**Date and Time (Step 3 of 7)**

Current Date and Time

03/07/2012 11:06:37

Time Zone

(GMT-5:00) US Eastern Time

Use Router Settings

Next >

Figure 6-121 Setup wizard: Set date and time

4. Optionally, you can enter System licenses (as shown in Figure 6-122) 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. You do not need to add them here.

**System License (Step 4 of 7)**

The enclosure license already includes virtualization of internal Serial Attached SCSI (SAS) drives on your system. You can use this panel to set any additional options. If you are sharing the total authorized capacities across multiple clusters, enter only the capacities you wish to use on this system. The sum of the capacities across all systems must not exceed your authorized capacities.

Set License Options

External Virtualization Limit

10 enclosures

Remote-Copy Limit

30 enclosures

Real-time Compression Limit

24 enclosures

Next >

Figure 6-122 System license

5. Configure support notifications, as shown in Figure 6-123. Click **Next**.



Figure 6-123 Configure support notifications

6. Define company contact information as shown in Figure 6-124. 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**

<b>* Contact Name</b>	<b>* Email Reply Address</b>	
<input type="text" value="John Doe"/>	<input type="text" value="jd@ibm.com"/>	
<b>* Machine Location</b>	<b>* Telephone (Primary)</b>	<b>Telephone (Alternate)</b>
<input type="text" value="305"/>	<input type="text" value="9091234567"/>	<input type="text"/>

\* Required

**Next >** **Cancel**

Figure 6-124 Define company contact



7. Verify that all hardware was detected by the system correctly as shown in Figure 6-125. Click **Next**.

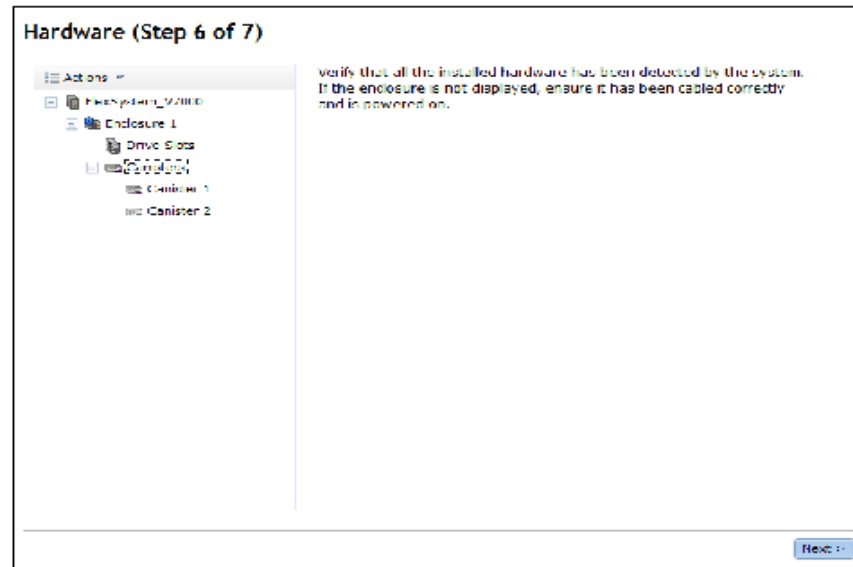


Figure 6-125 Verify hardware

8. Do not select **Yes** to automatically configure internal storage now because you are creating a customized storage layout.
9. Click **Finish** to complete the Setup wizard task and log in to IBM Flex System V7000 Storage Node, as shown in Figure 6-126. You log in as a Superuser with your newly defined password. If you did not change the password, the default is passw0rd.

**Remember:** The password includes a zero not the letter O.

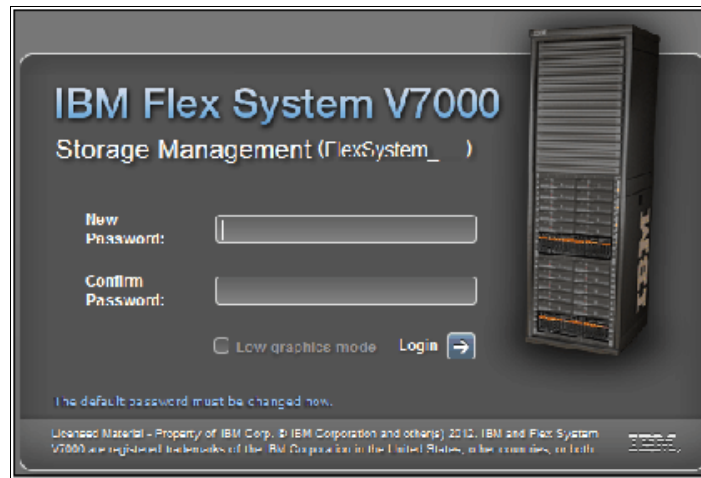


Figure 6-126 Setup wizard task complete

After a successful login, the IBM Flex System V7000 Storage Node Home Overview window looks similar to Figure 6-127.

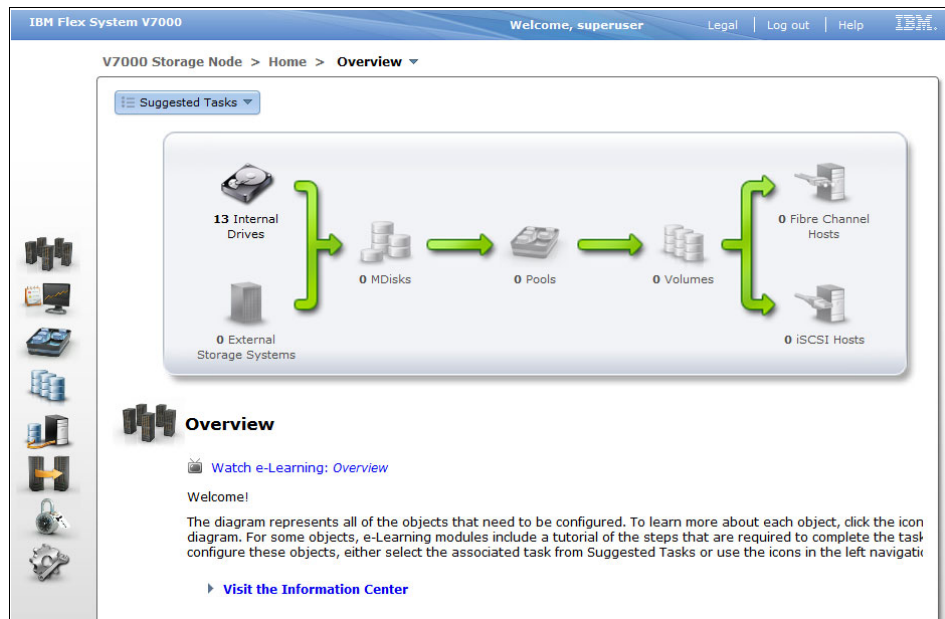


Figure 6-127 IBM Flex System V7000 Storage Node Home Overview window

IBM Flex System V7000 Storage Node initial configuration is complete and the cluster is up and running, as shown in Figure 6-128.

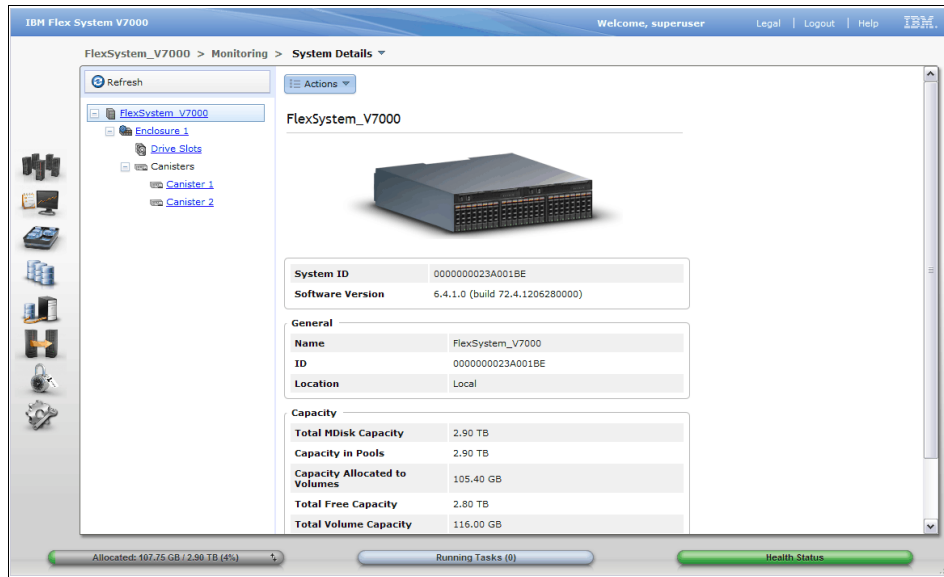


Figure 6-128 System details view in management GUI

You can continue to configure other functions and features for your environment to meet your implementation requirements.

### 6.7.3 MDisk configuration

Complete the following steps to configure the MDisk:

1. Return to the Overview window (as shown in Figure 6-129) and browse to the **Pools** menu. Select **Internal Storage**.

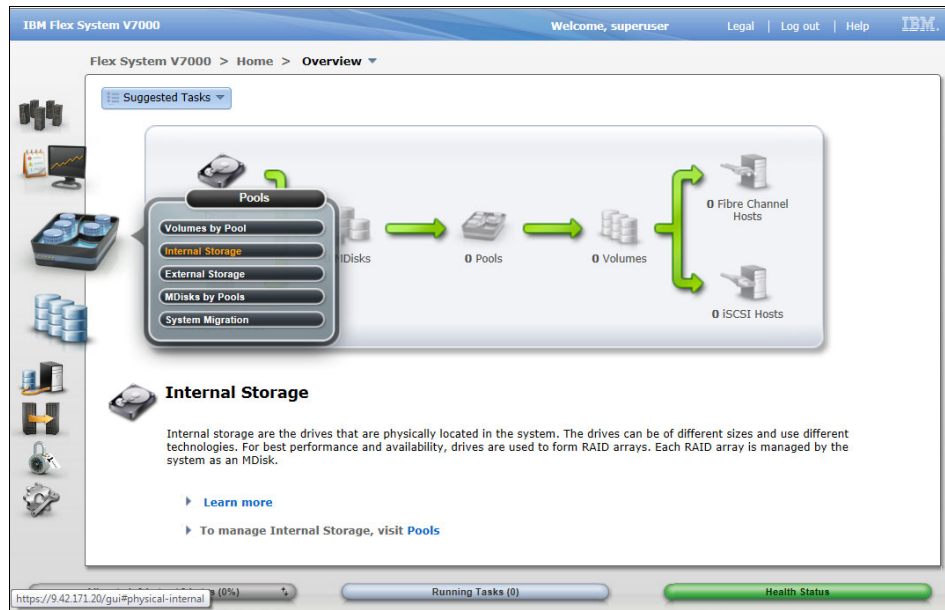


Figure 6-129 IBM Flex System V7000: Overview

2. Click **Configure Storage**, as shown in Figure 6-130.

IBM Flex System V7000

Welcome, superuser   Legal   Log out   Help

V7000 Storage Node > Pools > Internal Storage

Drive Class Filter

All Internal

136.23 GB, SAS  
15000 rpm  
io\_grp0

558.41 GB, SAS  
10000 rpm  
io\_grp0

Configure Storage

All Internal

Capacity Allocation

0%

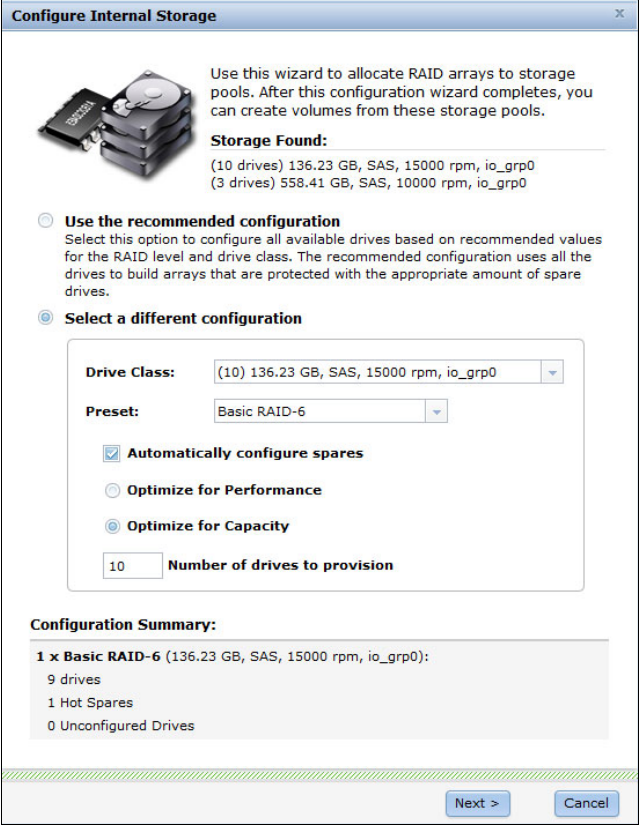
MDisk Capacity 0 bytes  
Spare Capacity 0 bytes  
Total Capacity 2.97 TB

Actions

Drive ID	Capacity	Use	Status	MDisk Name	Enclosure ID
12	136.23 GB	Candidate	Online		1
11	136.23 GB	Candidate	Online		1
10	136.23 GB	Candidate	Online		1
9	136.23 GB	Candidate	Online		1
8	136.23 GB	Candidate	Online		1
7	136.23 GB	Candidate	Online		1
6	136.23 GB	Candidate	Online		1
5	136.23 GB	Candidate	Online		1
4	136.23 GB	Candidate	Online		1
3	136.23 GB	Candidate	Online		1

Figure 6-130 IBM Flex System V7000: Internal Storage

3. Choose **Select a different configuration** and **Basic RAID 6**. Select **Optimize for Capacity** and specify 10 in the **Number of drives to provision** field, as shown in Figure 6-131. Click **Next**. The RAID 6 consists of nine drives and one drive is Hot Spare.



The screenshot shows the 'Configure Internal Storage' wizard window. It includes an icon of a storage pool and a description: 'Use this wizard to allocate RAID arrays to storage pools. After this configuration wizard completes, you can create volumes from these storage pools.' Below this, the 'Storage Found' section lists two storage pools: '(10 drives) 136.23 GB, SAS, 15000 rpm, io\_grp0' and '(3 drives) 558.41 GB, SAS, 10000 rpm, io\_grp0'. The 'Use the recommended configuration' option is unselected, while 'Select a different configuration' is selected. In the configuration section, 'Drive Class' is set to '(10) 136.23 GB, SAS, 15000 rpm, io\_grp0' and 'Preset' is set to 'Basic RAID-6'. The 'Automatically configure spares' checkbox is checked. Under 'Optimize for Capacity', the 'Number of drives to provision' is set to 10. The 'Configuration Summary' section shows '1 x Basic RAID-6 (136.23 GB, SAS, 15000 rpm, io\_grp0):' with '9 drives', '1 Hot Spares', and '0 Unconfigured Drives'. At the bottom, there are 'Next >' and 'Cancel' buttons.

**Configure Internal Storage**

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  
(3 drives) 558.41 GB, SAS, 10000 rpm, io\_grp0

☐ Use the recommended configuration  
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

**Drive Class:** (10) 136.23 GB, SAS, 15000 rpm, io\_grp0

**Preset:** Basic RAID-6

☒ Automatically configure spares

☐ Optimize for Performance

☒ Optimize for Capacity

10 **Number of drives to provision**

**Configuration Summary:**

1 x Basic RAID-6 (136.23 GB, SAS, 15000 rpm, io\_grp0):  
9 drives  
1 Hot Spares  
0 Unconfigured Drives

Next > Cancel

Figure 6-131 Configure Internal Storage: RAID configuration

4. Select **Create one or more new pools** and specify ESXiPool as the Pool Name or Prefix, as shown in Figure 6-132. Click **Finish**.

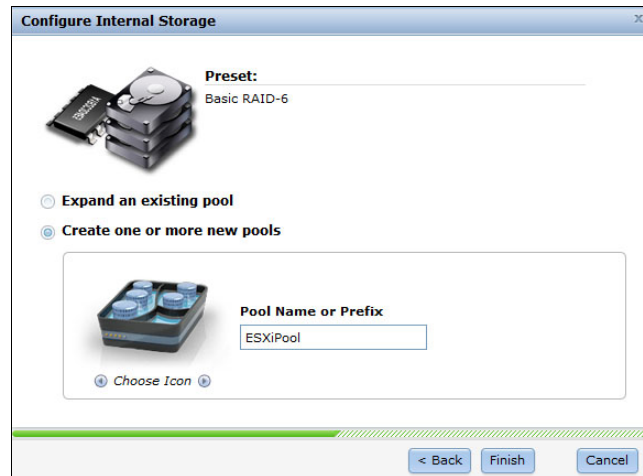


Figure 6-132 Configure Internal Storage: Pool creation

5. When the task completes, click **Close**, as shown in Figure 6-133.




Figure 6-133 Create RAID Arrays task

You created the ESXiPool, based on a RAID6 + 1 Hot Spare drive.

**Remember:** Because the option to automatically configure spares was selected as shown in Figure 6-131 on page 251, one of the 10 disks is configured for hot spare.



## 6.7.4 Zoning configuration

When the Flex System FC3171 8Gb SAN Switch is powered on, it includes a preconfigured switch zoning that automatically includes all of the HBAs that are connected to it. Figure 6-134 shows a conceptual representation of the new zoning configuration that was implemented in this chapter.

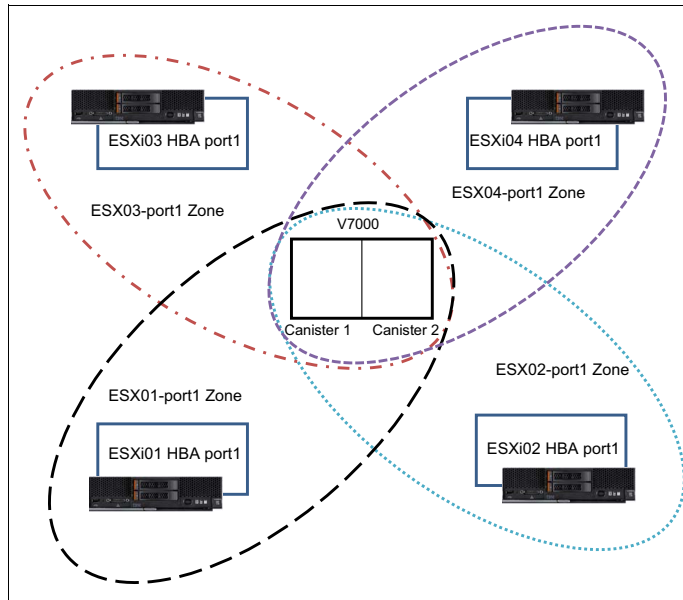


Figure 6-134 Zoning final layout

To configure the appropriate zoning on the Flex System FC3171 8Gb SAN Switch to support the storage configuration that is shown on Figure 5-5 on page 124, complete the following steps:

1. From the Chassis Manager on Flex System Manager, right-click the image of the FC3171 8GB SAN Switch, select **Remote Access**, and then click **Launch Web Browser**, as shown on Figure 6-135.

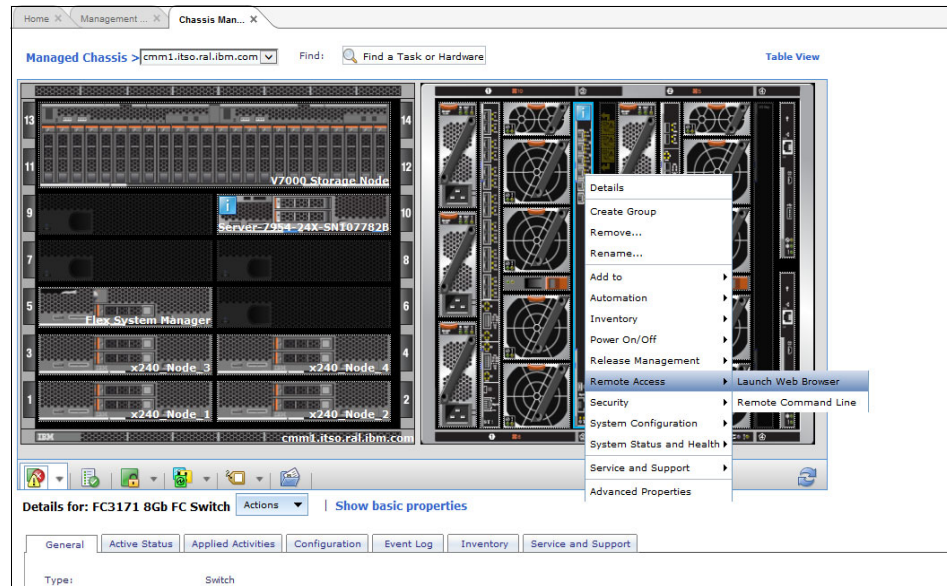


Figure 6-135 SAN Switch web access console start

2. A new Web Browser window opens. If you are prompted to accept a Java ActiveX installation, trust the signature and click **No** when you are prompted for the security warning, as shown in Figure 6-136.

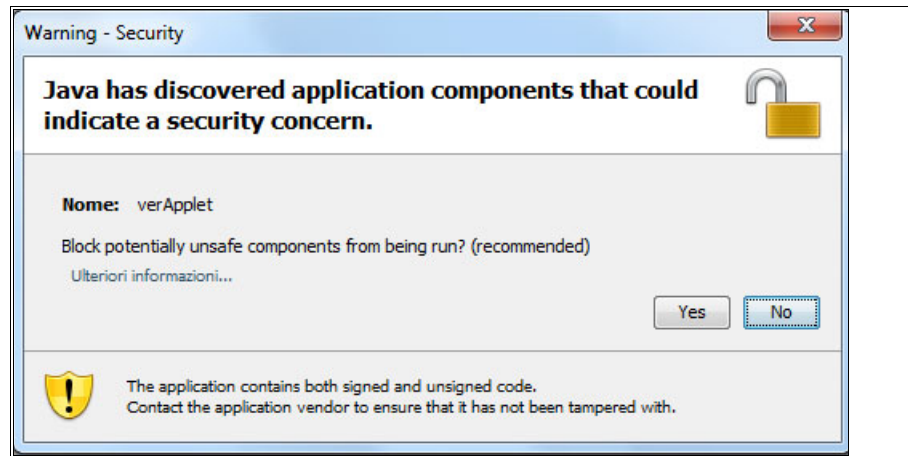


Figure 6-136 Java Security Warning connecting to Flex SAN Switch

3. Log in with the default user name and password (USERID and PASSWORD) as shown in Figure 6-137. Click **Add Fabric**.

**Remember:** The password uses a zero and not the letter O.

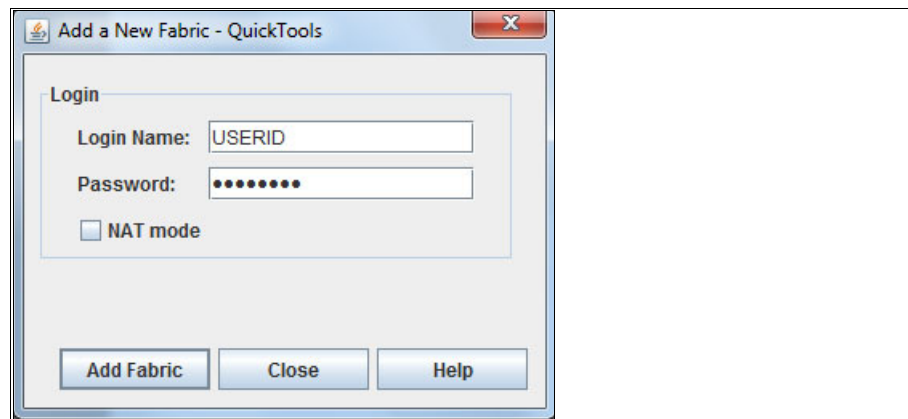


Figure 6-137 IBM Flex SAN Switch login windows

- Click **OK** when prompted that your password was not changed, as shown in Figure 6-138.



Figure 6-138 Password change reminder

- From the main console menu, select **Zoning** → **Edit Zoning**, as shown in Figure 6-139.

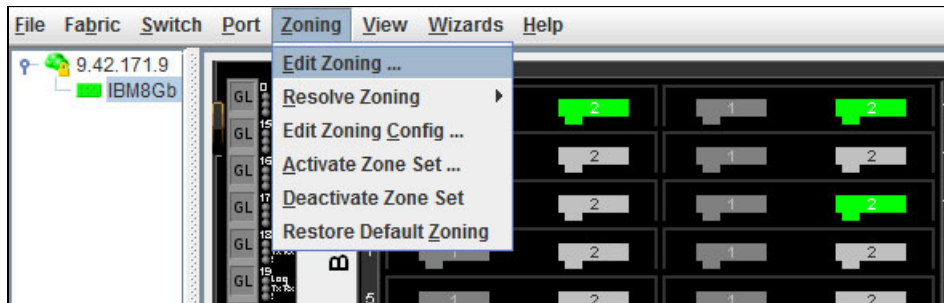


Figure 6-139 Zoning menu

- Click to highlight **Zone Sets** at the root, and then click **Zone Set** to add a new zone set. Enter ESXi to identify the zone set, as shown in Figure 6-140.




Figure 6-140 Create a zone set

- Right-click the ESXi zone set and select **Create a Zone**, as shown in Figure 6-141.

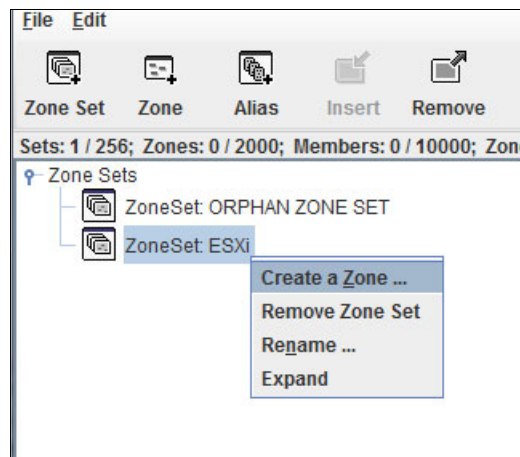


Figure 6-141 Create a zone

8. Name the zone ESXi01-port1, as shown in Figure 6-142.




Figure 6-142 Zone name window

9. Repeat steps 1 - 8 to create zones that are named ESXi02-port1, ESXi03-port1, and ESXi04-port1 under the ESXi root.
10. Right-click **Zone:ESXi01-port1** and select **Create Members**, as shown in Figure 6-143.

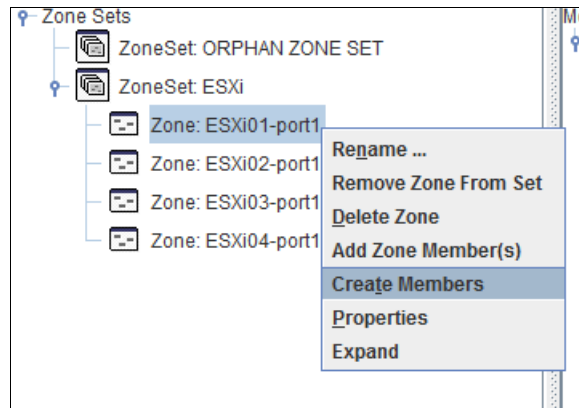


Figure 6-143 Create Zone Members

11. Each created zone must contain a single ESXi host port. Domain1-Bay1 and Domain1-Bay2 are the first and the second ESXi compute nodes host because Domain1-Bay3 and Domain1-Bay4 are the third and the fourth ESXi hosts. Complete the Enter hex value field and select the **FC Address** option for the Domain1-Bay1 host, as shown in Figure 6-144. Click **OK**.




Figure 6-144 Enter hex value

12. On the same ESXi01-port1 zone, add the two IBM Flex System V7000 Storage Node canisters Domain1-Bay13 and Domain1-Bay14, as shown in Figure 6-145.

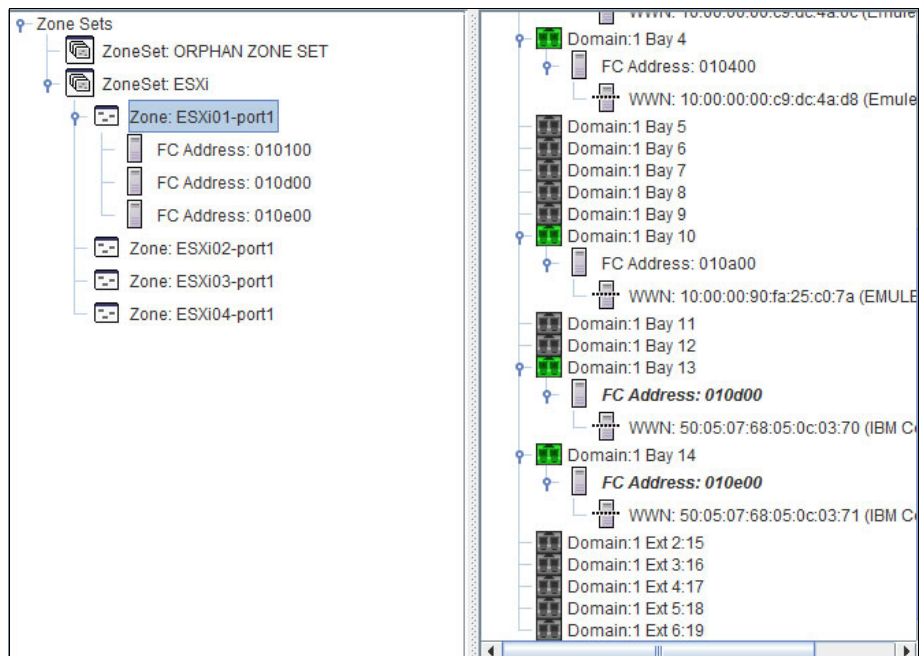


Figure 6-145 Single zone membership

13. Repeat this procedure on the other ESXi zones for all the ESXi compute nodes in Domain1-Bay2, Domain1-Bay3, and Domain1-Bay4.

**Note:** Each zone must contain a single ESXi host and the two IBM Flex Canister on Bay13 and Bay14.



Figure 6-146 shows the final configuration.

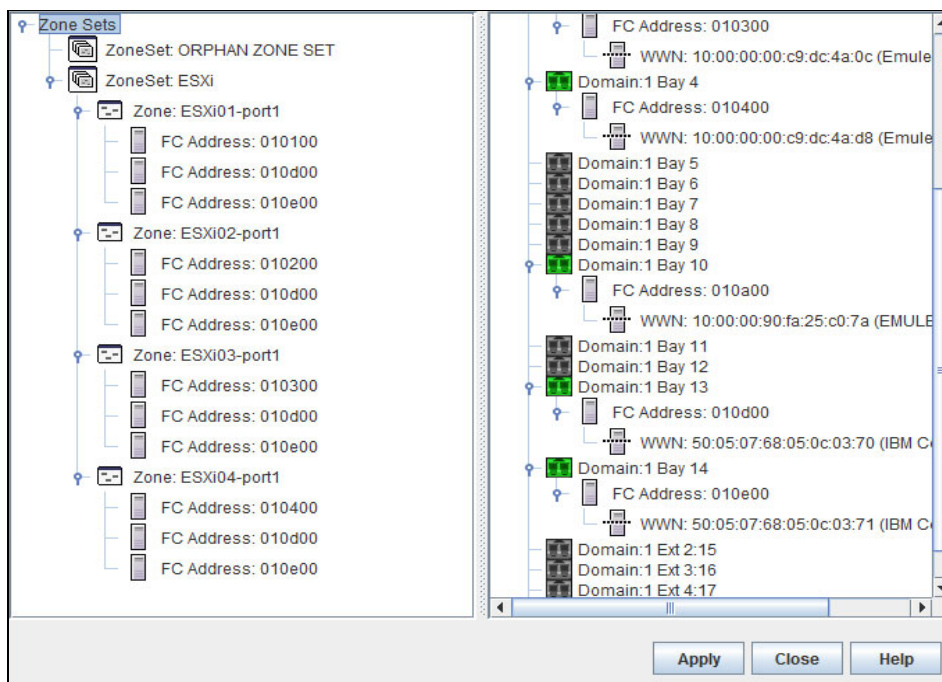


Figure 6-146 Final zoning configuration

14. Click **Apply** to write the configuration. In the Save Zoning & Error Check window (see Figure 6-147), click **Save Zoning**.




Figure 6-147 Save Zoning & Error Check window

15. When you are prompted to activate one zone after the zone save, select **No**, as shown in Figure 6-148.

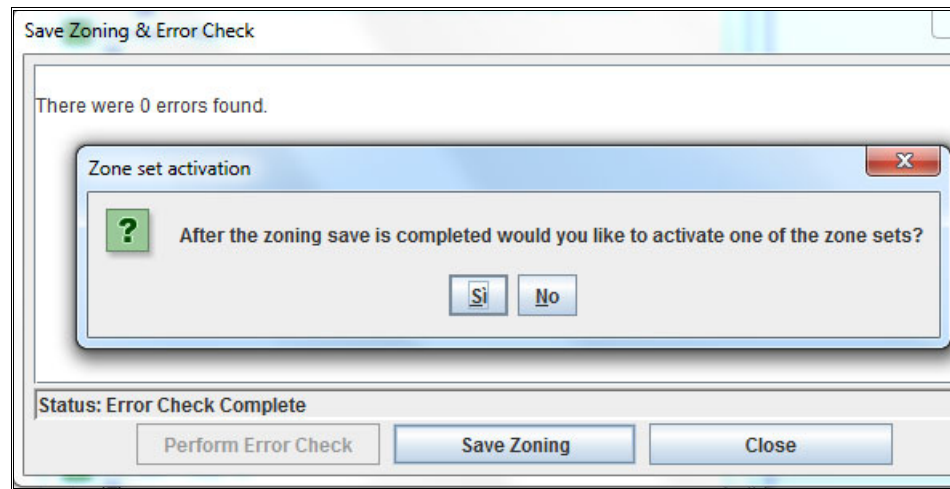


Figure 6-148 Activate zone set

16. When the zone save completes, click **Close**, as shown in Figure 6-149.

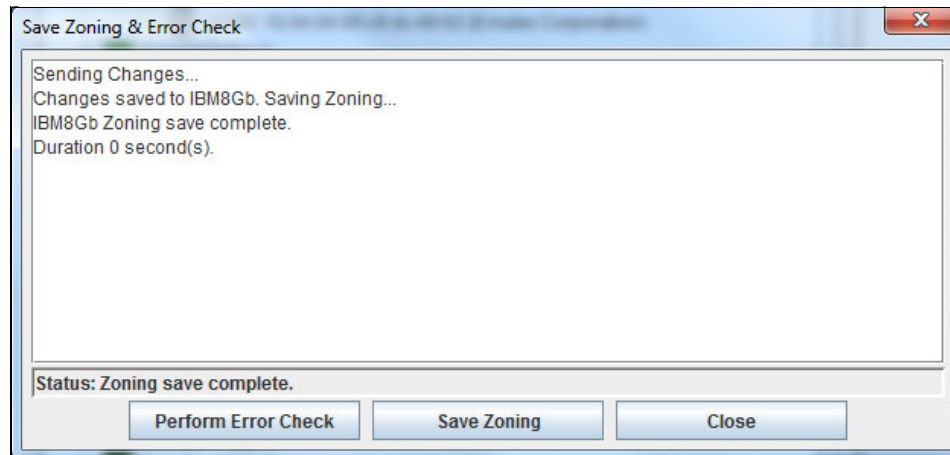


Figure 6-149 Zoning save complete

17.To return to the main window, on the Edit Zoning window, click **Close** to exit, as shown in Figure 6-150.

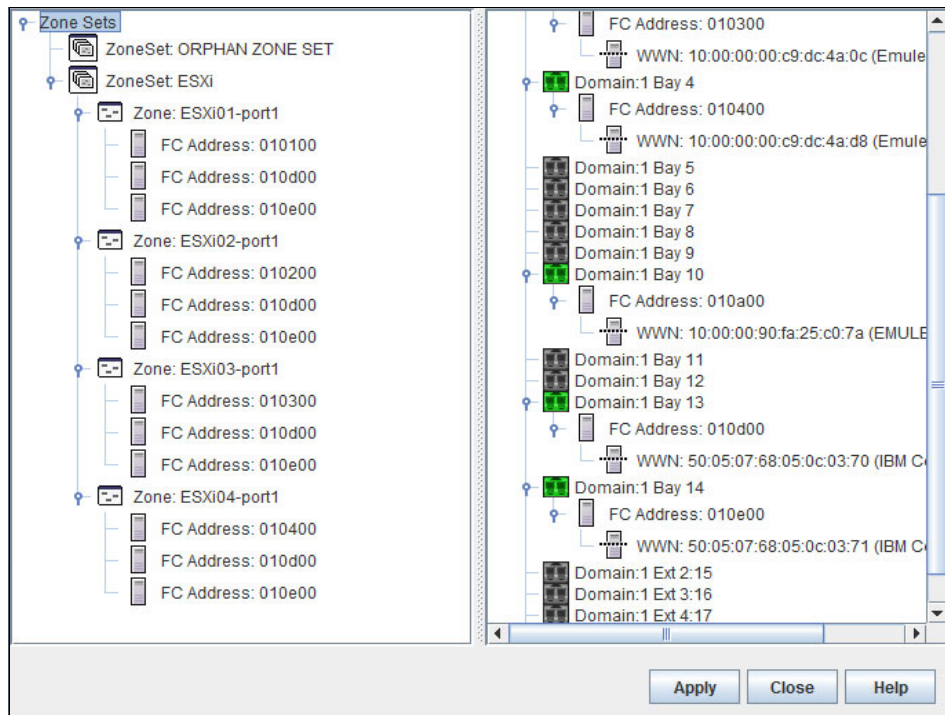


Figure 6-150 Edit Zoning window

18.Select **Zoning** → **Activate Zone Set**, as shown in Figure 6-151.

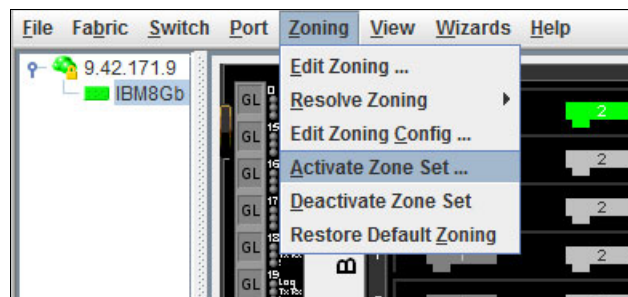
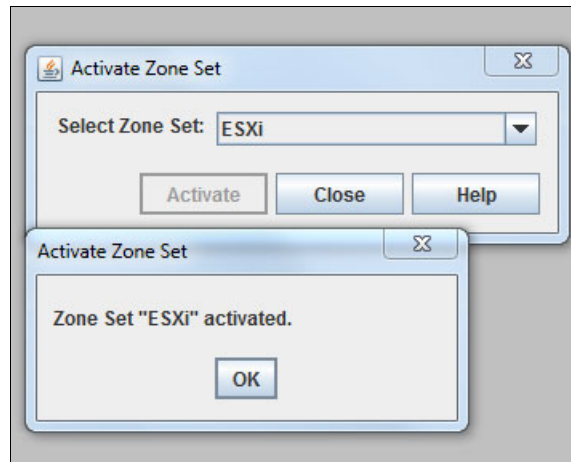


Figure 6-151 Activate Zone Set menu

19. Select the **ESXi** zone set and click **Activate**. Confirm the activation, as shown in Figure 6-152.



*Figure 6-152 Confirm the activation*

# 6.7.5 Configuring volumes

Complete the following steps to configure the volumes:

1. In the right side menu, browse to Volumes, as shown in Figure 6-153.



Figure 6-153 Volumes creation

2. Select **New Volume**, as shown in Figure 6-154.

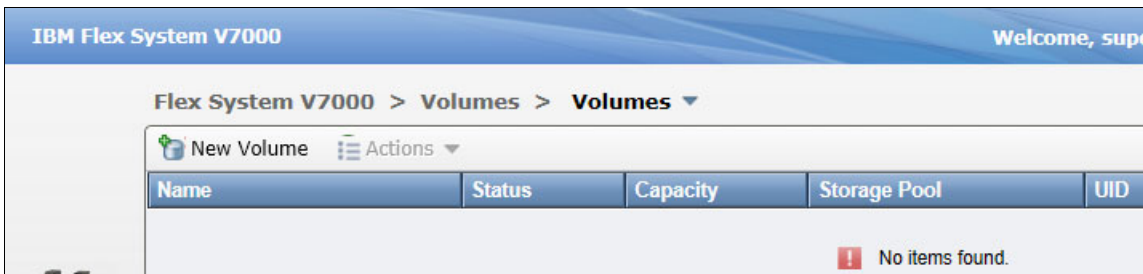


Figure 6-154 New volume

3. Select **Thin-Provision**, as shown in Figure 6-155.




Figure 6-155 Preset selection

4. Select the pool that is named ESXiPool, as shown in Figure 6-156.

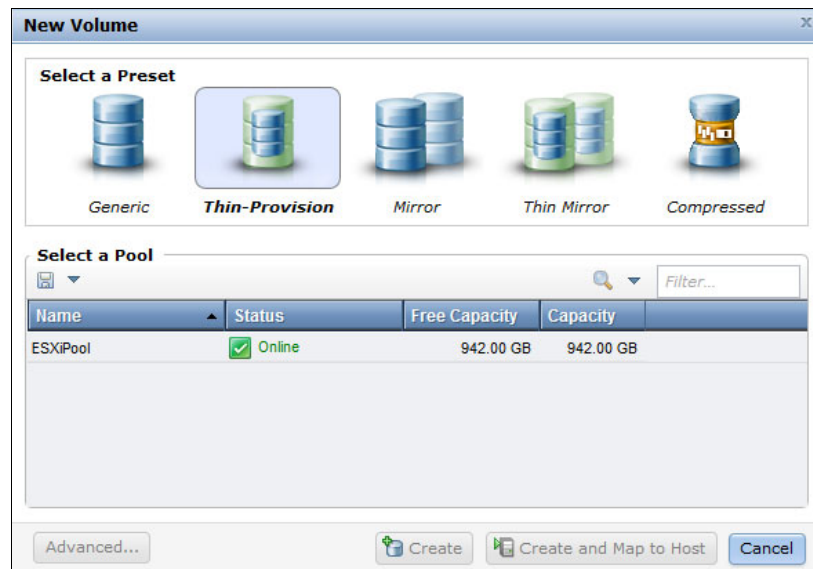


Figure 6-156 Pool selection

5. Create a volume that is named **ESXi Mgmt Volume** and set the size to 400GB, as shown in Figure 6-157. Click **Create**.

The 'New Volume' dialog box is shown with the following configuration:

- Select a Preset:** Five icons are displayed: Generic, Thin-Provision (selected), Mirror, Thin Mirror, and Compressed.
- Select a Pool:** The 'Primary Pool' is set to 'ESXiPool' with an 'Edit' button.
- Volume Details:**
  - Quantity: 1
  - Capacity: 400 GB
  - Name: ESXi Mgmt Volume
- Summary:** 1 thin-provisioned volume, 400.00 GB virtual capacity, 8.00 GB real capacity, 934.00 GB free in pool
- Buttons:** 'Advanced...', 'Create' (highlighted), 'Create and Map to Host', and 'Cancel'.

Figure 6-157 ESXi Mgmt Volume creation

6. Click **Close** in the Create Volumes window when the task is completed. Create a volume that is named **VDI Volume** and set the size to 600GB, as shown in Figure 6-158. Click **Create**.

The 'New Volume' dialog box is shown with the following configuration:

- Select a Preset:** Five icons are displayed: Generic, Thin-Provision (selected), Mirror, Thin Mirror, and Compressed.
- Select a Pool:** The 'Primary Pool' is set to 'ESXiPool' with an 'Edit' button.
- Volume Details:**
  - Quantity: 1
  - Capacity: 600 GB
  - Name: VDI Volume
- Summary:** 1 thin-provisioned volume, 600.00 GB virtual capacity, 12.00 GB real capacity, 921.00 GB free in pool
- Buttons:** 'Advanced...', 'Create' (highlighted), 'Create and Map to Host', and 'Cancel'.

Figure 6-158 VDI Volume creation

## 6.7.6 Configuring hosts

For this part of the procedure, you must know the host's HBA WWPN to successfully map the new volumes to the correct hosts. To retrieve the HBA WWPN, complete the following steps and make note of the hosts-to-WWPN association:

1. From the main page of Flex System Manager, click **Select Chassis to be Managed**, as shown in Figure 6-159.

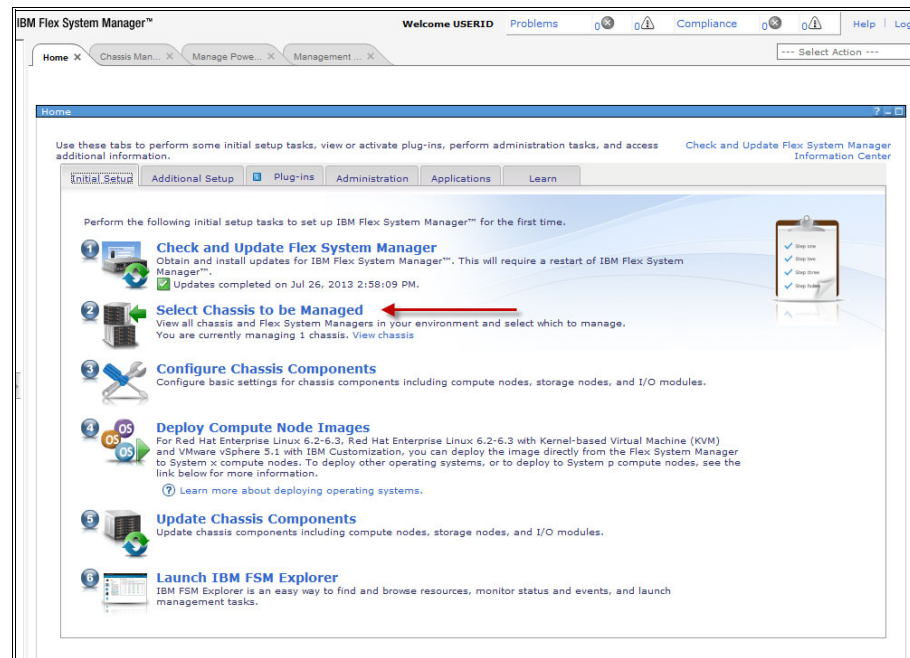


Figure 6-159 Flex System Manager main page



2. From the Management Domain tab, select the chassis, as shown in Figure 6-160.

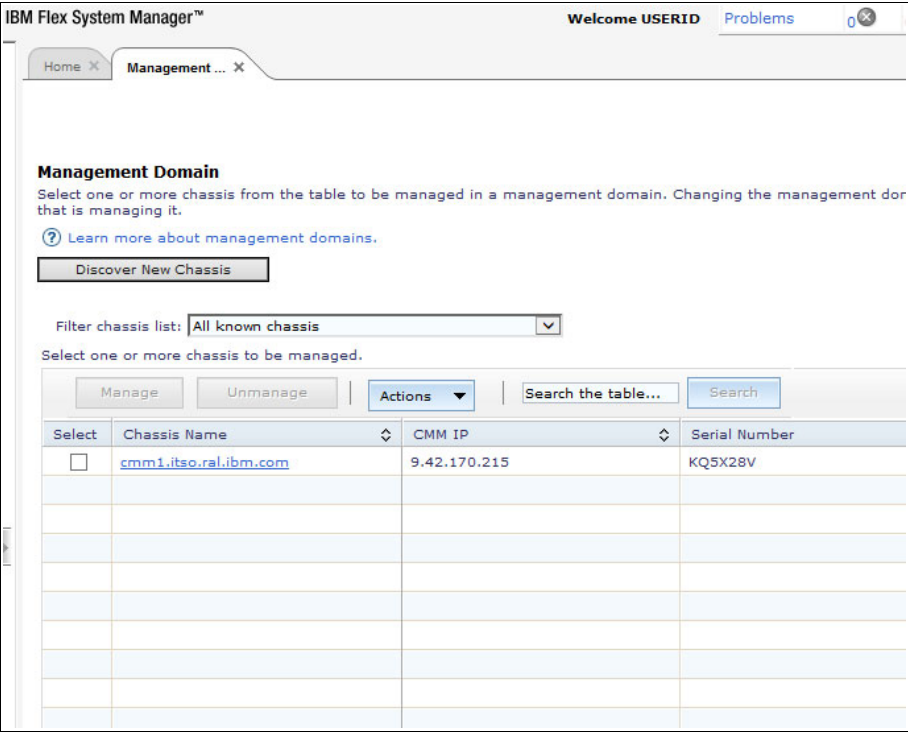


Figure 6-160 Management Domain tab

3. Select the x240\_Node\_1 compute node and go to the Inventory tab, as shown in Figure 6-161.

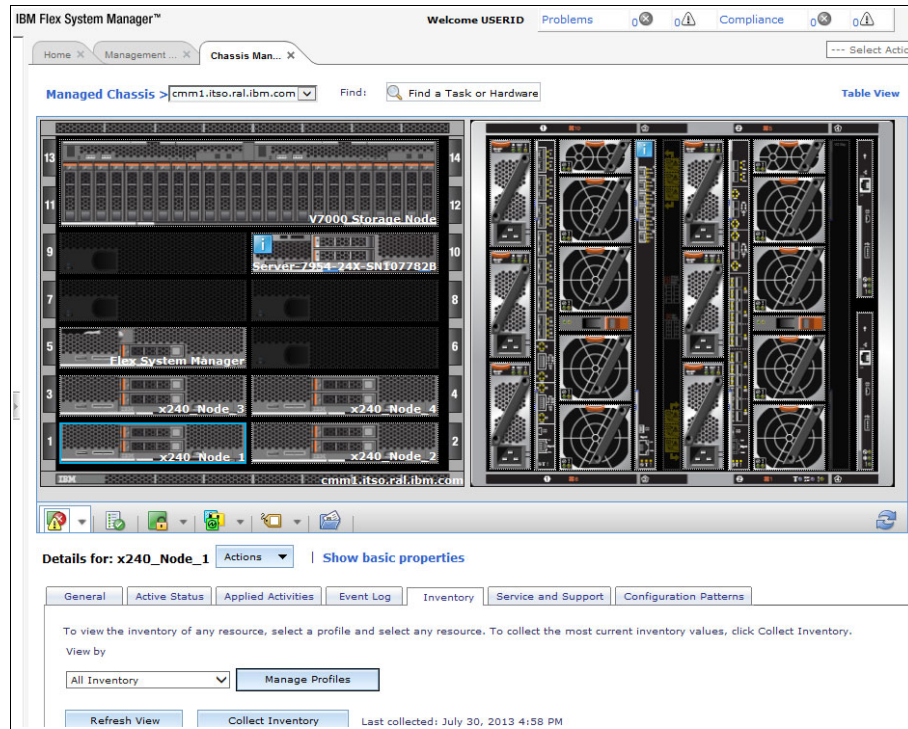


Figure 6-161 Managed Chassis main window

4. Scroll down the page and in the Collected Items section, expand **Network Configuration** and the **SCSI Interface** subsection. The HBA WWPN information is here, as shown in Figure 6-162.

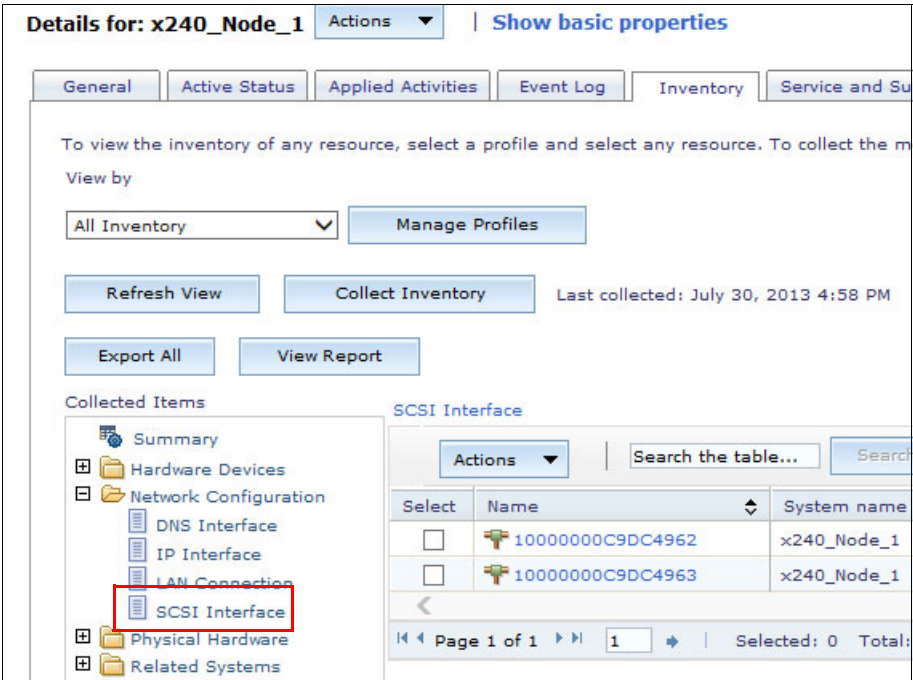


Figure 6-162 SCSI Interface and WWPN section

5. Repeat steps 1 - 4 for the remaining compute nodes.

Complete the following steps to configure the hosts:

1. In the right side menu, browse to Hosts and click **New Host**, as shown in Figure 6-163.




Figure 6-163 New host

2. Select **Fibre Channel Host**, as shown in Figure 6-164.



Figure 6-164 Choose the Host Type

3. Specify the Host Name, select a port from the list in Fibre Channel Ports field, and click **Add Port to List**. Click **Create Host**, as shown in Figure 6-165. Repeat this step for all the ESXi hosts.

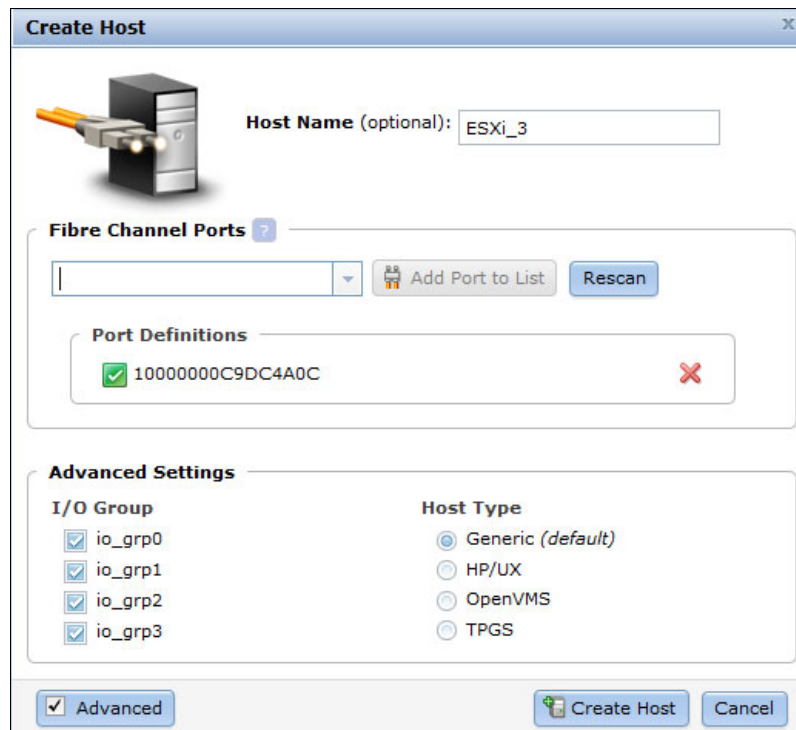


Figure 6-165 Create Host

- To modify the host mappings, select a host, right-click to display the menu, and select **Modify Mappings**, as shown in Figure 6-166.

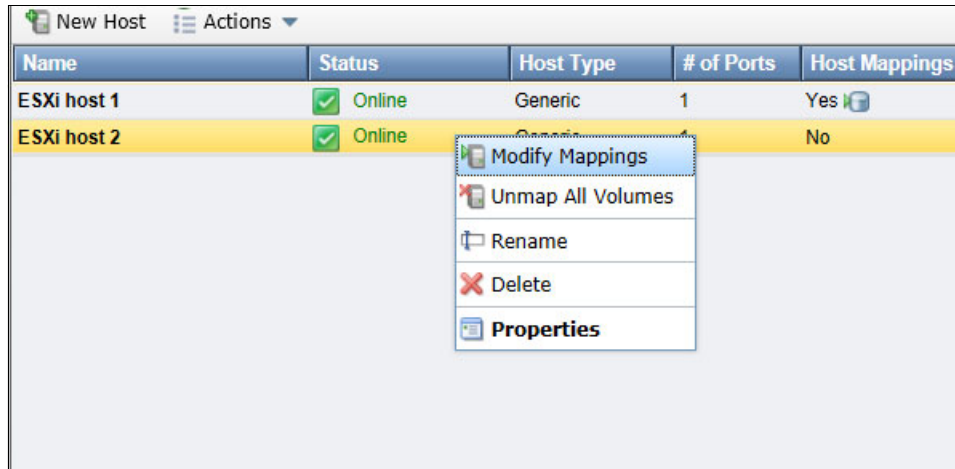


Figure 6-166 Modify Mappings

- Assign the needed volumes to each host, as described in 6.7.5, “Configuring volumes” on page 265. Click **Map Volume**, as shown in Figure 6-167.

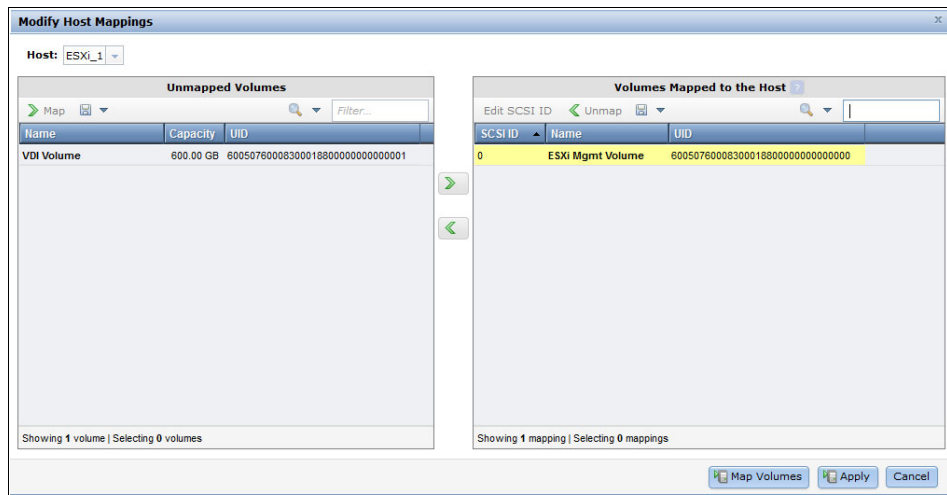


Figure 6-167 Modify Host mappings

Every volume is the shared disk volume for a two-node VMware Cluster. Therefore, you must assign the same volume to two hosts.

When you try to add a pre-assigned volume to another node, the message that is shown in Figure 6-168 opens. You can safely click **Map All Volumes** to map the volumes to the second node.

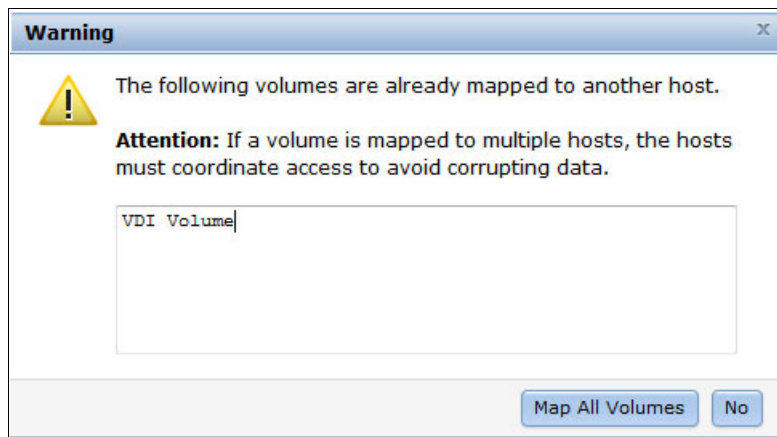


Figure 6-168 Multiple nodes mapping warning

## 6.8 VMControl activation

VMControl includes an evaluation license, which enables the use of the optional chargeable (fee-based) management function. The evaluation period begins after you activate VMControl and then restart the IBM Flex System Manager management node.

You must activate IBM Flex System Manager VMControl before it can be used.

**Deactivation note:** The evaluation period continues to run even if you deactivate VMControl.

Complete the following steps to activate VMControl:

1. From the plug-ins page of Flex System Manager, in the Additional Plug-ins to activate section, click **Activate 90-day evaluation now**, as shown in Figure 6-169.

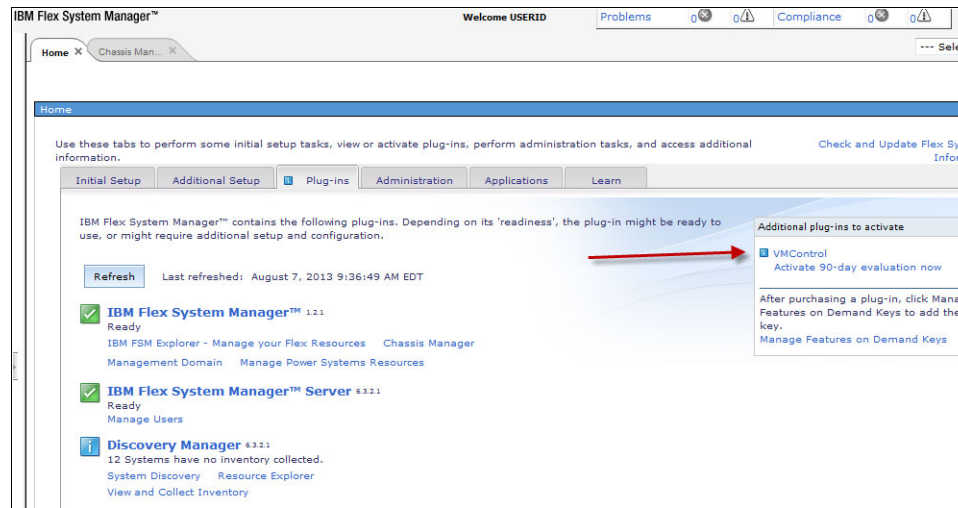


Figure 6-169 VMControl 90-day trial activation

2. A message displays with the successful activation. Click **Restart IBM Flex System Manager Server** to activate the VMControl plug-in, as shown in Figure 6-170.




Figure 6-170 Restart for the VMControl plug-in to be effective

VMControl is now active.

For more information about how to manage VMware environment on IBM Flex System to deploy and configure VMs by using VMControl, see Chapter 9, “Operating VMware Horizon View infrastructure” on page 373.







# Deploying VMware Horizon View infrastructure

This chapter provides the necessary steps to perform the installation and initial configuration of VMware Horizon View components.

**Important:** To complete the steps in this chapter, IBM Flex System must be deployed as described in Chapter 6, “Deploying IBM Flex System” on page 135.

This chapter includes the following topics:

- ▶ Installing vSphere components and infrastructure services
- ▶ Configuring vSphere
- ▶ Installing View Composer
- ▶ Installing View Connection Server
- ▶ Configuring View Connection Server initially

## 7.1 Installing vSphere components and infrastructure services

This section guides you through the installation and initial configuration of the base vSphere components and necessary infrastructure services for a Horizon View deployment on the Flex System that is deployed, as described in Chapter 6, “Deploying IBM Flex System” on page 135.

### 7.1.1 Configuring ESXi

The x240 compute nodes are configured with an IBM Customized ESXi 5.1 USB key that is installed on the node motherboard. The first mandatory configuration that must be done on each host is that of the management network-static IP, subnet, gateway, host name, and DNS. You also must configure the ESXi management password. You can complete this configuration by using the direct console user interface.

Complete the following steps to continue with the configuration:

1. Connect to an ESXi host by using vSphere Client.
2. Create local and shared data stores.
3. Upload the Windows 2008 R2 installation ISO to the Management shared data store.
4. Create a virtual machine (VM) and install Microsoft Windows 2008 R2. Then, shut down the VM.
5. Copy the VM or export and import by using the Open Virtualization Format Archive (OVA) template. Then, deploy the VMs that are needed for the infrastructure. For more information, see the VM configurations that are described in Chapter 5, “IBM Flex System and VMware View lab environment” on page 117.
6. Change the computer security identifier (SID) of each guest OS by using a supported method.

**Tip:** You can change the computer SID by using the generalize feature of the built-in Sysprep tool in Windows 2008 R2. For more information, see this website:

<http://technet.microsoft.com/en-us/library/hh824938.aspx>

## 7.1.2 Installing infrastructure services

The infrastructure services include the following components:

- ▶ Active Directory
- ▶ DNS
- ▶ DHCP
- ▶ File Server
- ▶ SQL server

Active Directory, DNS, and DHCP are vital for the functioning of virtual desktops, user authentication, and IP allocation of VMs. The example that is described here also deploys MS SQL Server 2008 R2, which is required by several of the Horizon View components. A file server is also commonly used in a View deployment. The example in this book uses it for Persona management.

For more information, see the VM configurations that are described in Chapter 5, “IBM Flex System and VMware View lab environment” on page 117.

Add the following server roles to the DC01 server:

- ▶ Active Directory Domain Services
- ▶ DHCP Server
- ▶ DNS Server

### Active Directory

Perform a standard Active Directory configuration and use the following settings:

- ▶ Domain name: companyA.local
- ▶ Domain name (pre Windows 2000): COMPANYA
- ▶ Domain functional level: Windows Server 2008 R2
- ▶ Forest functional level: Windows Server 2008 R2

Create a View Admins Global security group in the Users organization unit (OU), which is used for Horizon View administrators.

Create a VDI OU with the following sub OUs, as shown in Figure 7-1 on page 280:

- ▶ Computers

This OU contains the following sub OUs to store the virtual desktop computer accounts:

- FVM: Stores full VMs computer accounts
- LCVM: Stores linked clone VMs computer accounts

► Users

This OU contains the following objects:

- Standard Users: Security group for standard virtual desktop users
- VIP Users: Security group for VIP virtual desktop users
- Standard User: Standard virtual desktop user account
- VIP User: VIP virtual desktop user account

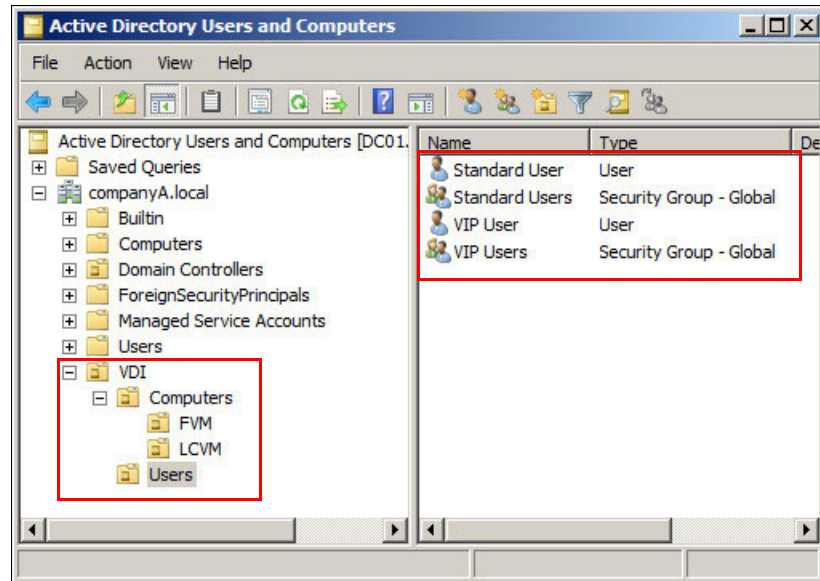


Figure 7-1 Active Directory OU configuration

## DHCP

Configure a DHCP scope in 10.20.20.x network, as shown in Figure 7-2.

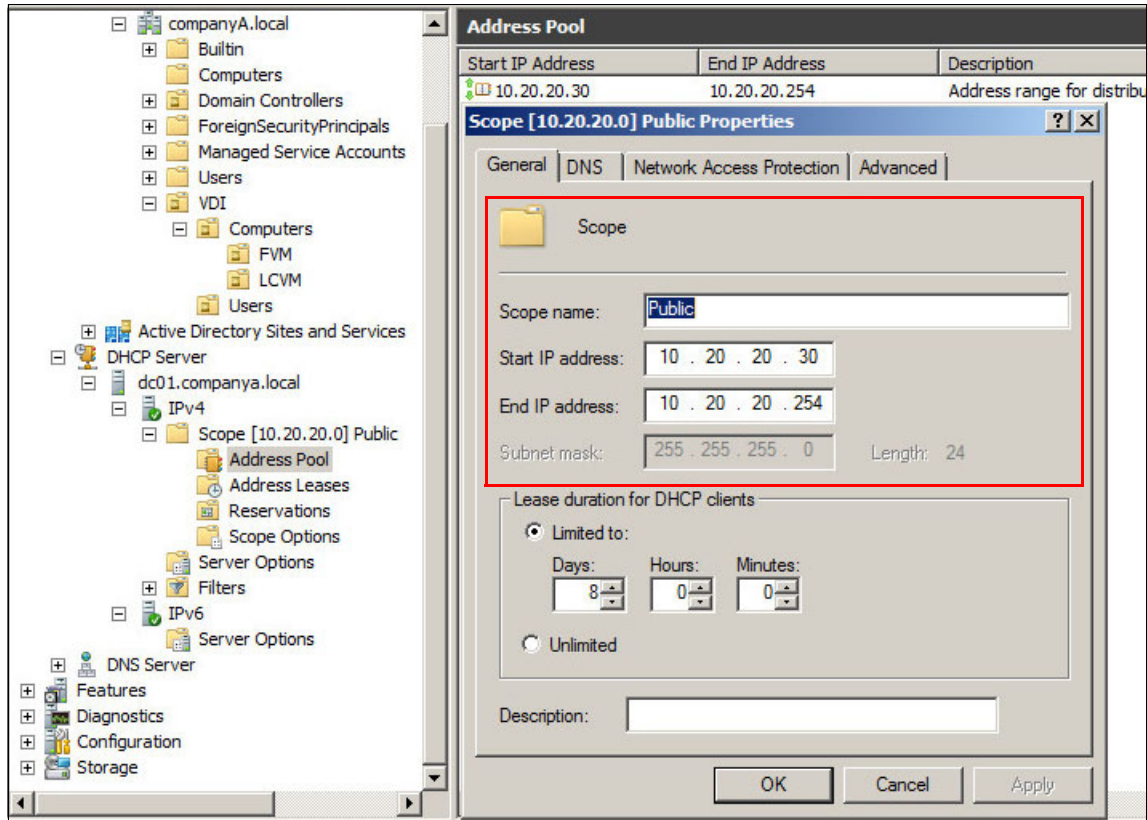


Figure 7-2 DHCP configuration window

## DNS

DNS is automatically configured with the Active Directory Service when Active Directory integrated DNS installation is performed. No other customization is required now.

## MS SQL

Create the necessary DB instances by using the standard procedure that is described in the following official VMware documentation:

- ▶ vCenter Server and SSO DB:  
<http://pubs.vmware.com/vsphere-51/topic/com.vmware.vsphere.install.doc/GUID-C6AF2766-1AD0-41FD-B591-75D37DDB281F.html>
- ▶ View Composer DB  
<http://pubs.vmware.com/view-52/topic/com.vmware.view.installation.doc/GUID-84F18501-3CF8-4584-9874-0243253786C3.html>

Create the MS SQL database instances that are listed in Table 7-1.

*Table 7-1 MS SQL DB instances*

DB name	Purpose	Owner
VCDB	vCenter Server	vpxuser
RSA	vCenter Single Sign-On	RSA_DBA
ViewCMPDB	View Composer	vpxuser
ViewEvent	View Event Database	vpxuser

### 7.1.3 Creating vCenter data source name

Complete the following steps to create the vCenter data source name:

1. Log in to the vCenter server OS. Browse to **Start** → **Administrative Tools** → **Data Sources (ODBC)**. Go to the System DSN tab, as shown in Figure 7-3. Click **Add**.

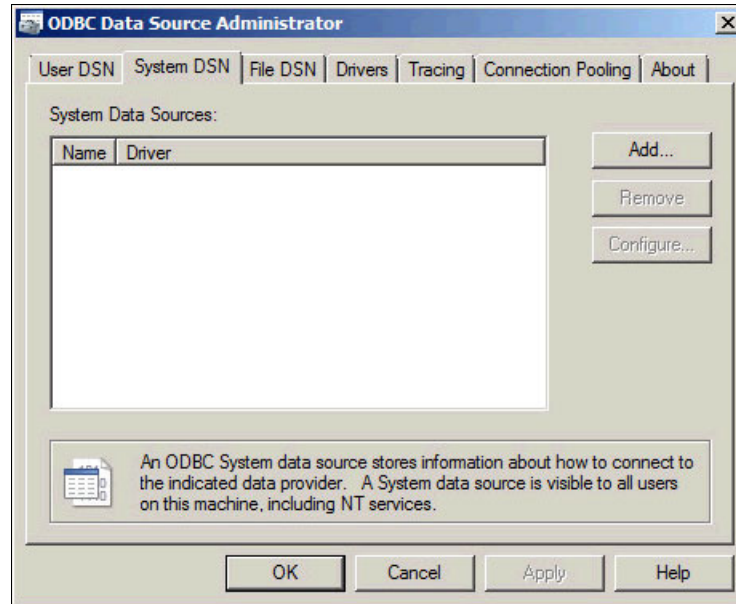


Figure 7-3 ODBC Data Source Administrator window

2. Select **SQL Server Native Client 10.0** and click **Finish**, as shown in Figure 7-4.

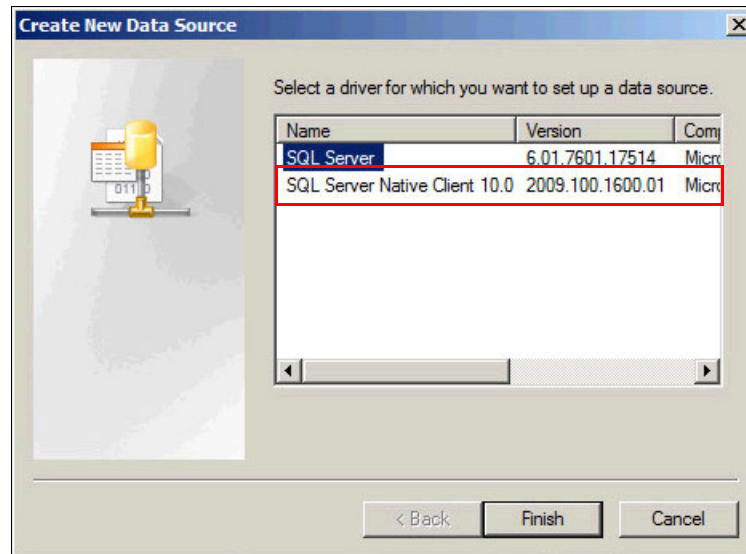
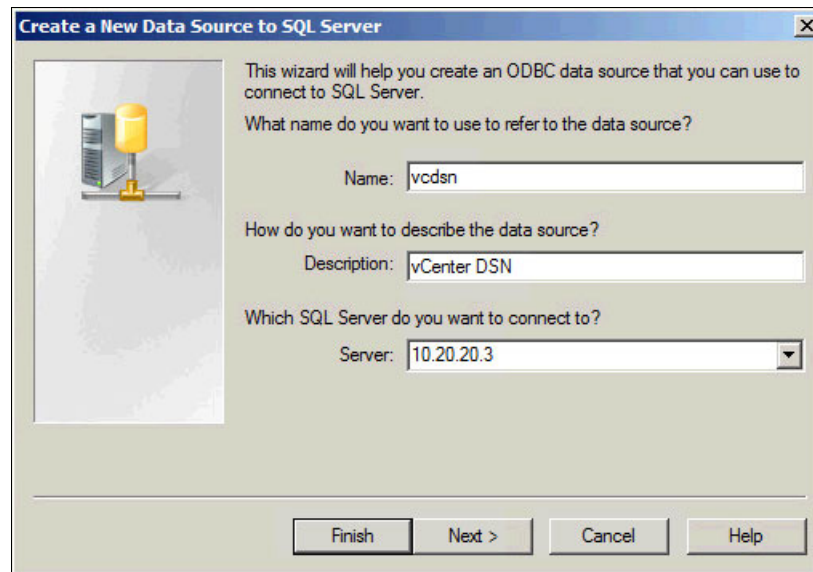


Figure 7-4 Create New Data Source driver selection window



3. Enter a name and description to which the data source and SQL Server connect. Enter `vcdsn` as the name and `10.20.20.3` as the SQL Server address, as shown in Figure 7-5.



*Figure 7-5 Data Source name and target server window*

4. Select **SQL Server authentication** and enter the Login ID and Password that has sufficient access rights to the vCenter DB instance, as shown in Figure 7-6.

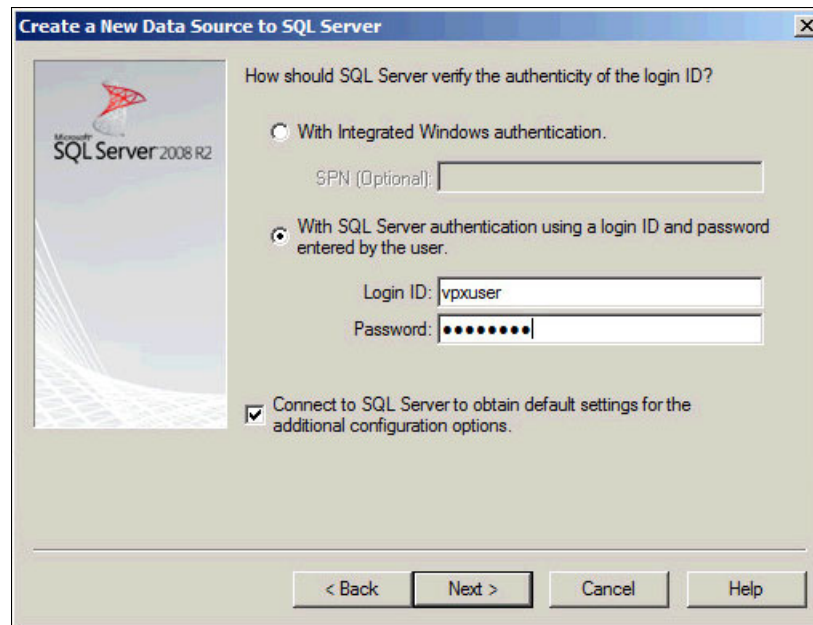


Figure 7-6 Data Source authentication details window

5. Ensure that the default database is the pre-created vCenter database instance; in this case, VCDB, as shown in Figure 7-7. If the default database is not the vCenter DB instance, select **Change the default database to** and choose the correct DB instance.

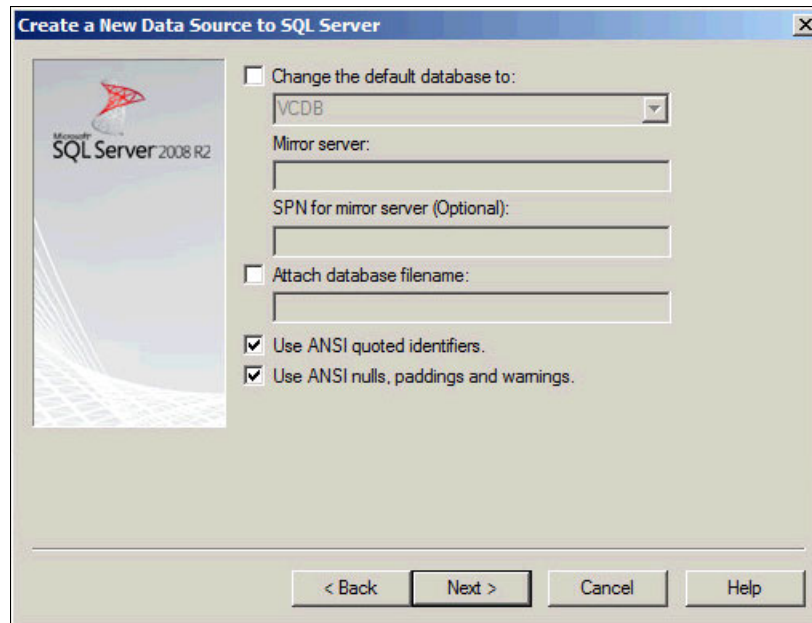


Figure 7-7 Data Source database instance selection

6. In the next window, leave the default settings, and click **Finish**, as shown in Figure 7-8.

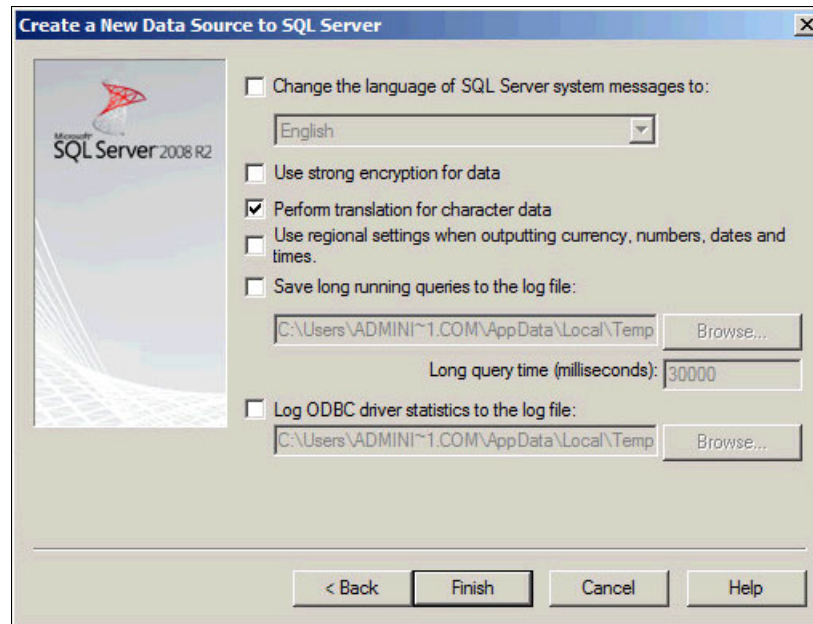


Figure 7-8 Data Source other configuration window

7. Review the data source creation summary window and click **Test Data Source** to ensure that the configuration is correct, as shown in Figure 7-9. The test should return a “TESTS COMPLETED SUCCESSFULLY!” message. Click **OK**.

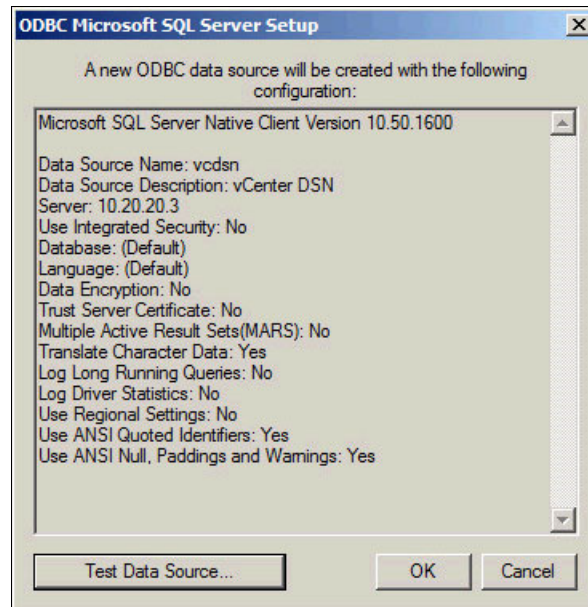
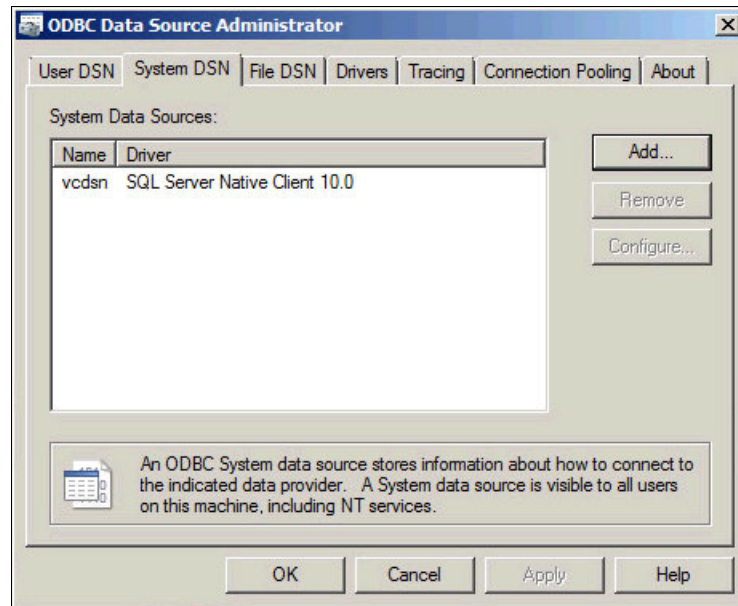


Figure 7-9 ODBC data source summary window

8. The new data source name is shown in the list of System Data Sources, as shown in Figure 7-10.



*Figure 7-10 ODBC Data Source Administrator window*

## 7.1.4 Installing vCenter Server

Complete the following steps to install vCenter Server:

1. Insert the VMware vCenter Installer media and ensure that autorun starts the main installation menu. From the list of products, select **VMware vCenter Simple Install**. Click the Simple Install option, as shown in Figure 7-11.

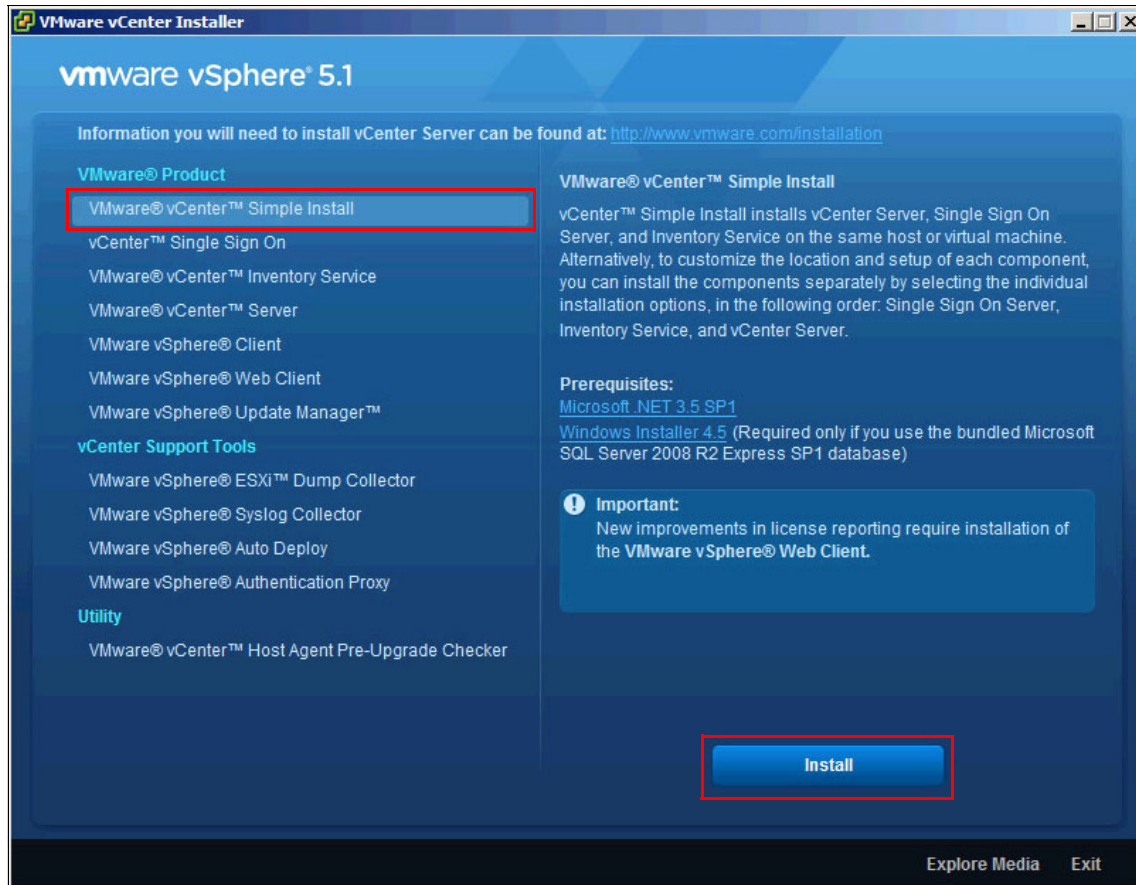
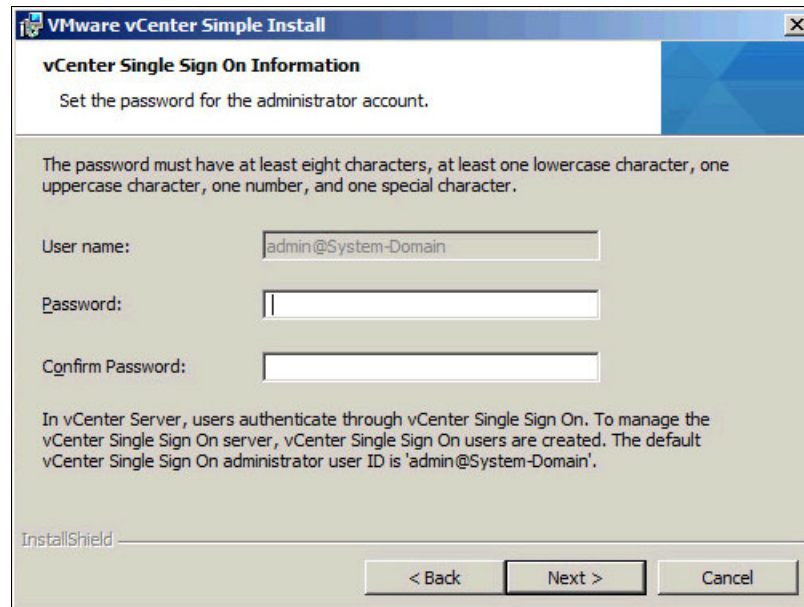


Figure 7-11 VMware vCenter installation options

2. The installation wizard opens. Click **Next** to continue. Then, read and accept the license agreement to proceed to the next step. Specify a strong password for the default vCenter Single Sign-On administrator, as shown in Figure 7-12. Click **Next**.



**VMware vCenter Simple Install**

**vCenter Single Sign On Information**

Set the password for the administrator account.

The password must have at least eight characters, at least one lowercase character, one uppercase character, one number, and one special character.

User name:

Password:

Confirm Password:

In vCenter Server, users authenticate through vCenter Single Sign On. To manage the vCenter Single Sign On server, vCenter Single Sign On users are created. The default vCenter Single Sign On administrator user ID is 'admin@System-Domain'.

InstallShield

< Back   Next >   Cancel

*Figure 7-12 Single Sign On administrator password*



3. SSO requires a database. Select **Use an existing supported database**, as shown in Figure 7-13. Click **Next**.

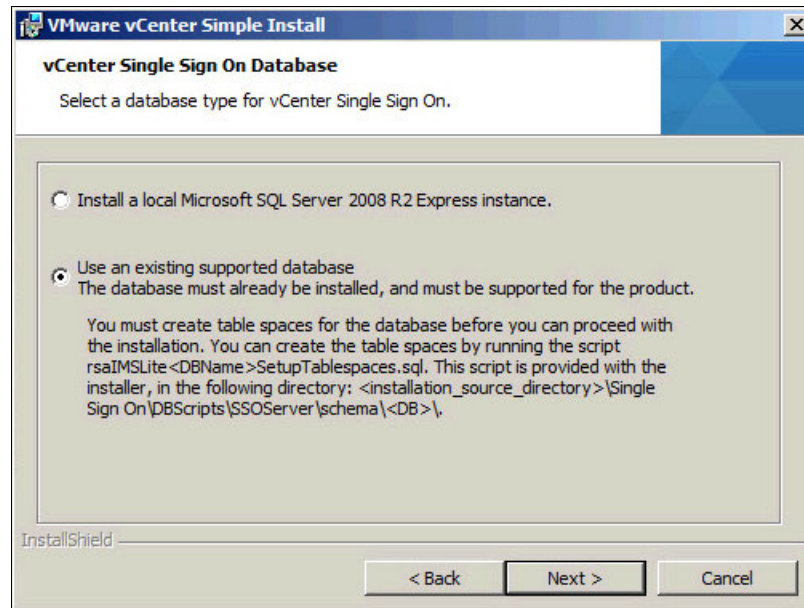


Figure 7-13 Single Sign On database selection

4. Enter the database type, IP address, database name, the database user name and password, and the database DBA user name and password, as shown in Figure 7-14. Click **Next**.

The screenshot shows the 'Database Information' window in the VMware vCenter Simple Install wizard. The window title is 'VMware vCenter Simple Install'. Below the title bar, the text 'Database Information' is displayed, followed by 'JDBC connection information for vCenter Single Sign On.' The form contains the following fields and options:

- Database Type:** A dropdown menu set to 'Mssql'.
- Database Name:** A text box containing 'RSA'.
- Host name or IP address:** A text box containing '10.20.20.3'.
- Port:** A text box containing '1433'.
- Database user name:** A text box containing 'RSA\_USER'.
- Database password:** A password field with 10 dots.
- Database DBA user name:** A text box containing 'RSA\_DBA'.
- Database DBA password:** A password field with 10 dots.
- Set MSSQL instance with dynamic port:** An unchecked checkbox.
- I will enter the JDBC URL myself:** An unchecked checkbox.
- JDBC URL:** An empty text box.

At the bottom of the window, there are three buttons: '< Back', 'Next', and 'Cancel'. The 'Next' button is highlighted. The 'InstallShield' logo is visible in the bottom left corner.

Figure 7-14 Single Sign On database configuration

5. Enter the Fully Qualified Domain Name of the system where SSO is installed. With vCenter Simple Install, SSO is installed on the same system as the vCenter Server. Enter vCenter01.CompanyA.local, as shown in Figure 7-15. Click **Next**.

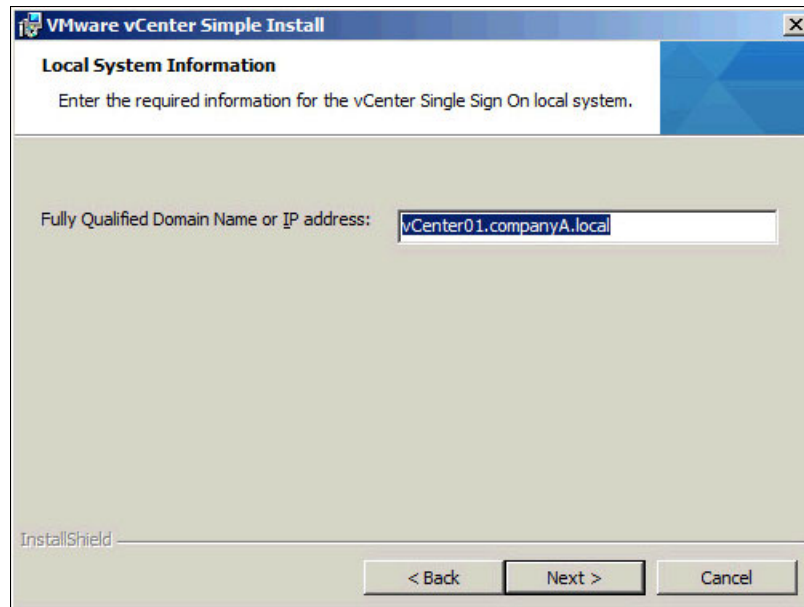


Figure 7-15 Single Sign On system FQDN

6. Configure the Security Support Provider Interface service to run in the default Windows NetworkService account, as shown in Figure 7-16. Click **Next**.

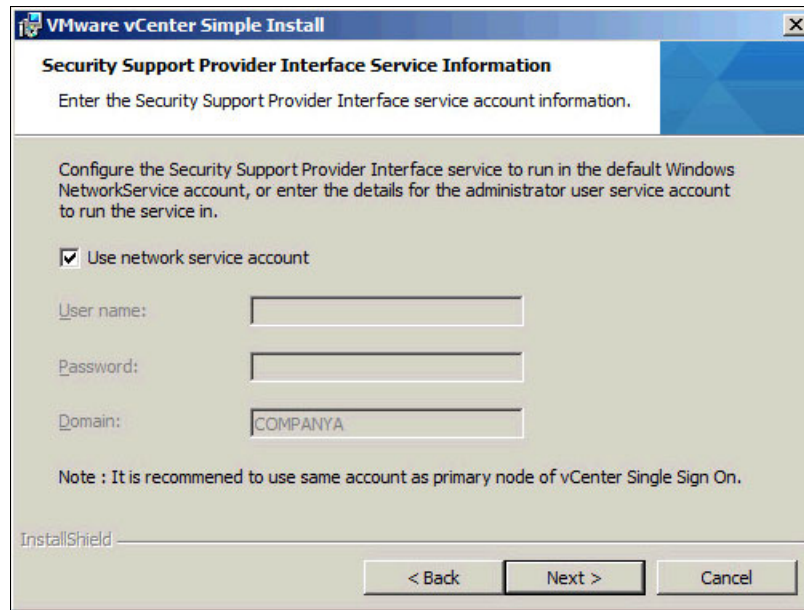
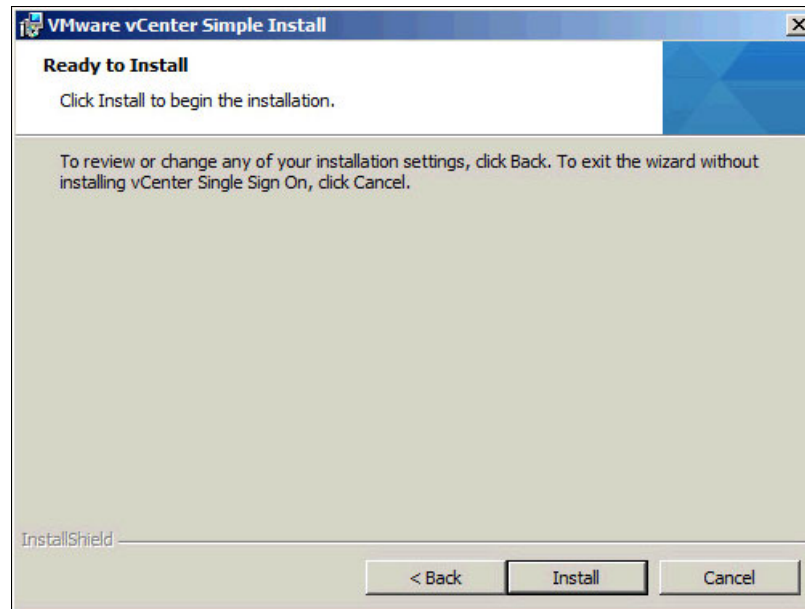


Figure 7-16 Single Sign On Security Support Provider Interface

7. Accept the default settings for the vCenter Single Sign On installation folder by clicking **Next**. Also, accept the default settings for the HTTPS port settings by clicking **Next**.

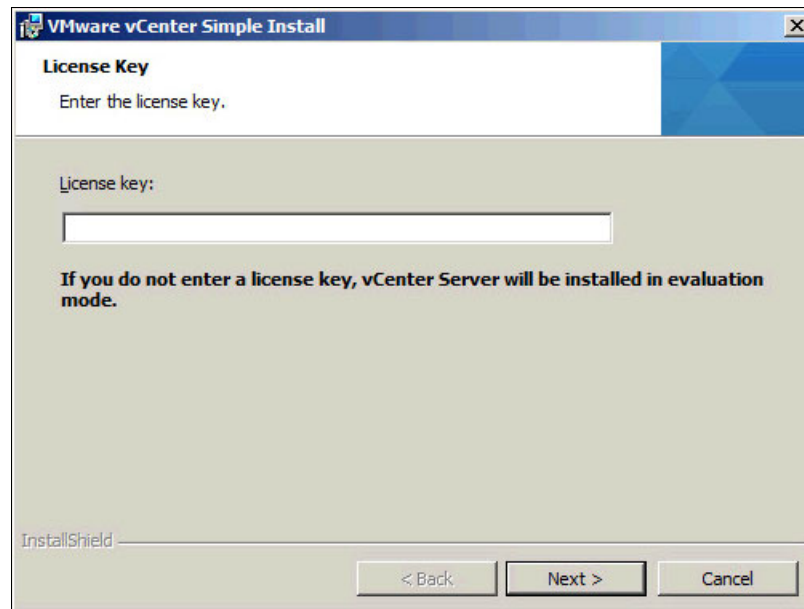
**Firewall:** The setup process opens the HTTPS port (7444 by default) automatically if the Windows operating system firewall service is running on the system.

8. Click **Install** to start the Single Sign On installation, as shown in Figure 7-17.



*Figure 7-17 Single Sign On Install completion window*

9. SSO and VMware vCenter Inventory Service are installed after SSO. The Simple Install setup starts the vCenter Server installation. Enter the license key (as shown in Figure 7-18) and then click **Next**.



*Figure 7-18 vCenter Simple Install License key window*

10. Select **Use an existing supported database** and enter the DSN that you created earlier (vcdsn), as shown in Figure 7-19. Click **Next**.

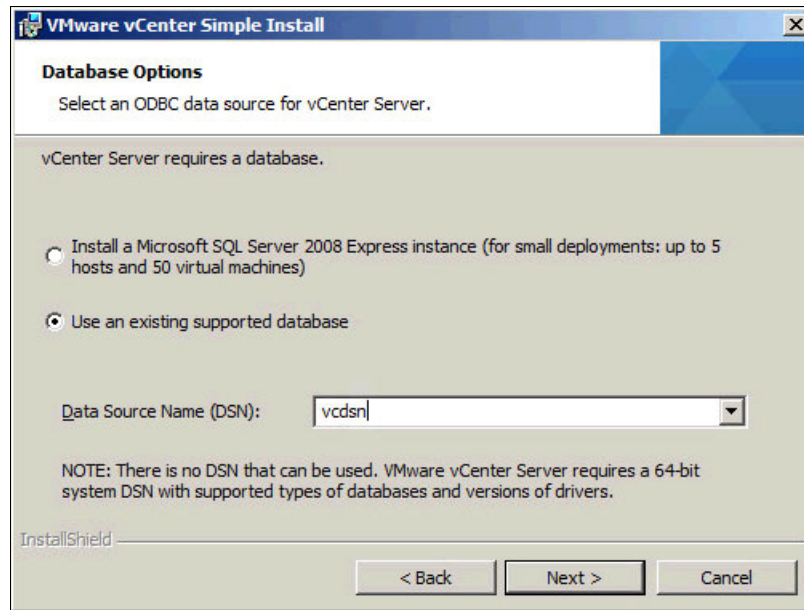
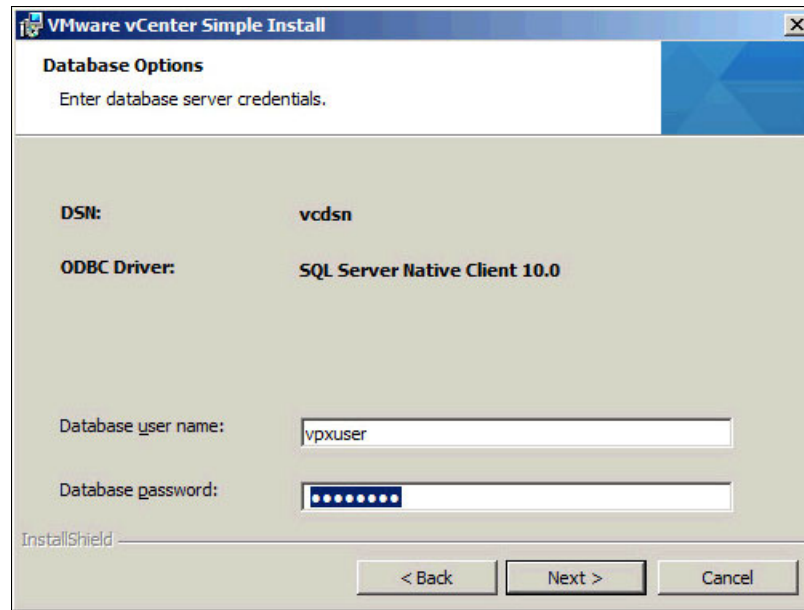


Figure 7-19 vCenter Install Database Options window

11. Enter the database user name and password for the user with sufficient access rights to the vCenter DB, as shown in Figure 7-20. Click **Next**.



The image shows a screenshot of the 'VMware vCenter Simple Install' window, specifically the 'Database Options' tab. The window has a title bar with the VMware logo and the text 'VMware vCenter Simple Install'. Below the title bar, the 'Database Options' section is highlighted. The instructions 'Enter database server credentials.' are displayed. The 'DSN:' field is set to 'vcdsn' and the 'ODBC Driver:' field is set to 'SQL Server Native Client 10.0'. The 'Database user name:' field contains 'vpxuser' and the 'Database password:' field is masked with dots. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted. The 'InstallShield' logo is visible in the bottom left corner.

DSN:	vcdsn
ODBC Driver:	SQL Server Native Client 10.0
Database user name:	vpxuser
Database password:	.....

Figure 7-20 vCenter Install Database server credentials window



12. Leave the “Use SYSTEM Account” option enabled, and enter the Fully Qualified Domain Name of the vCenter server machine, as shown in Figure 7-21. Click **Next**.

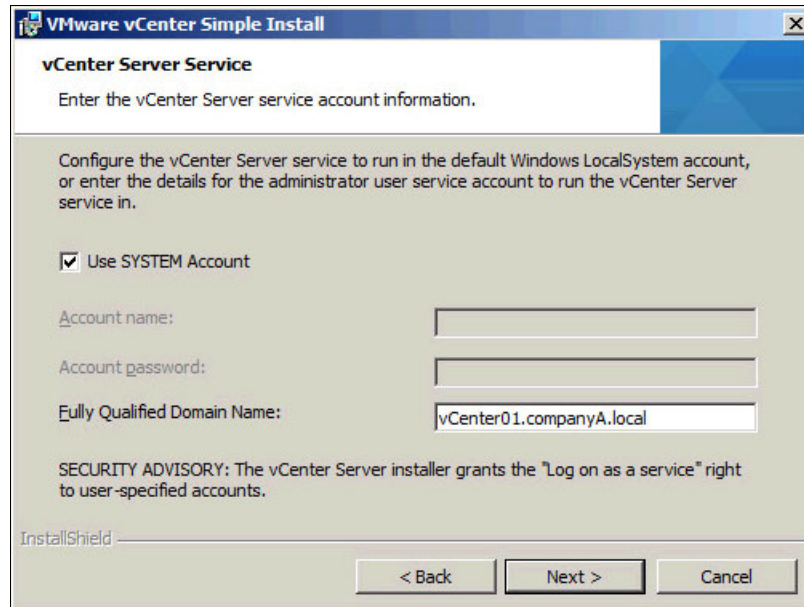
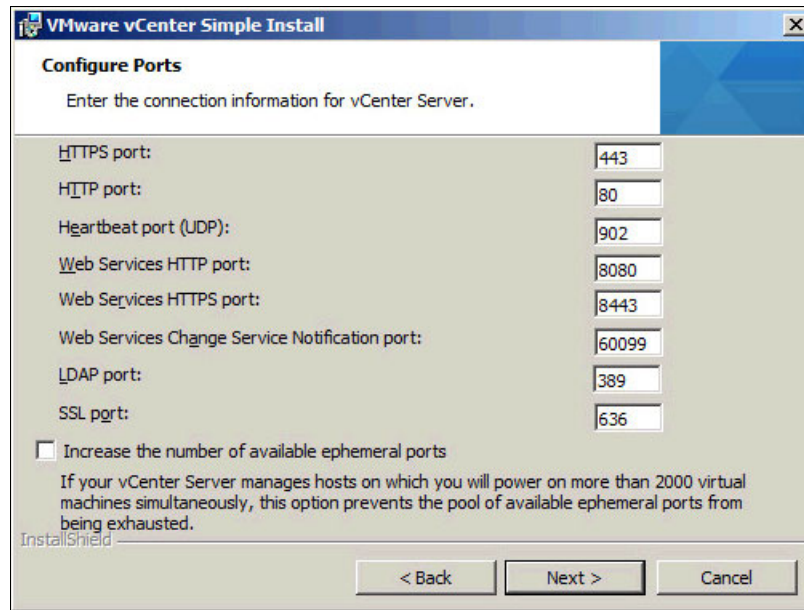


Figure 7-21 vCenter Server Service account information window

13. In the Configure Ports window, review the vCenter ports and leave the default settings by clicking **Next**, as shown in Figure 7-22.



The image shows the 'Configure Ports' window from the VMware vCenter Simple Install wizard. The window title is 'VMware vCenter Simple Install'. The main heading is 'Configure Ports' with a sub-instruction: 'Enter the connection information for vCenter Server.' The window contains several port configuration fields, each with a label and a text input box containing a default value:

Port Label	Default Value
HTTPS port:	443
HTTP port:	80
Heartbeat port (UDP):	902
Web Services HTTP port:	8080
Web Services HTTPS port:	8443
Web Services Change Service Notification port:	60099
LDAP port:	389
SSL port:	636

Below the port fields, there is a checkbox labeled 'Increase the number of available ephemeral ports'. Below this checkbox is a descriptive text: 'If your vCenter Server manages hosts on which you will power on more than 2000 virtual machines simultaneously, this option prevents the pool of available ephemeral ports from being exhausted.' At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a blue border. The 'InstallShield' logo is visible in the bottom left corner.

Figure 7-22 vCenter Server ports configuration window

14. In the vCenter Server JVM Memory configuration window, select **Small Inventory Size** and click **Next**, as shown in Figure 7-23.

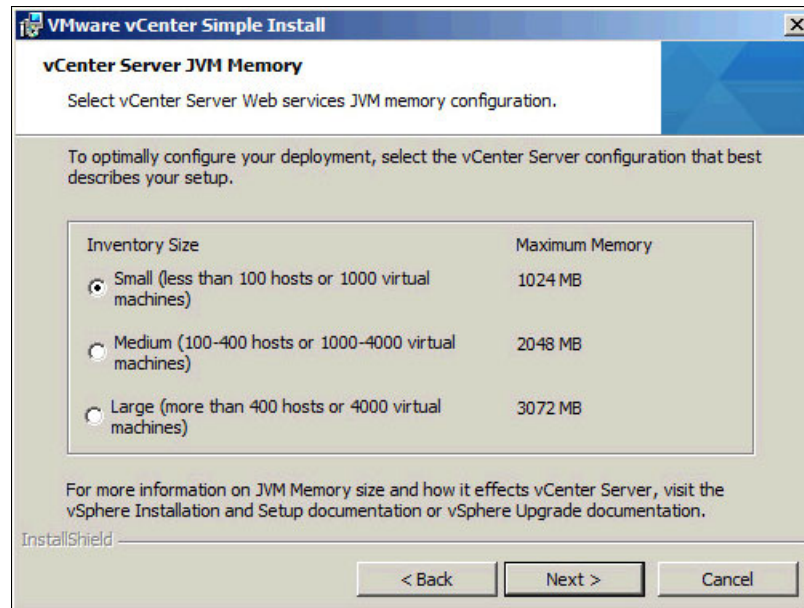


Figure 7-23 vCenter Server JVM Memory configuration

15. The vCenter Server installation settings are complete. Click **Install** to proceed with the installation, as shown in Figure 7-24.

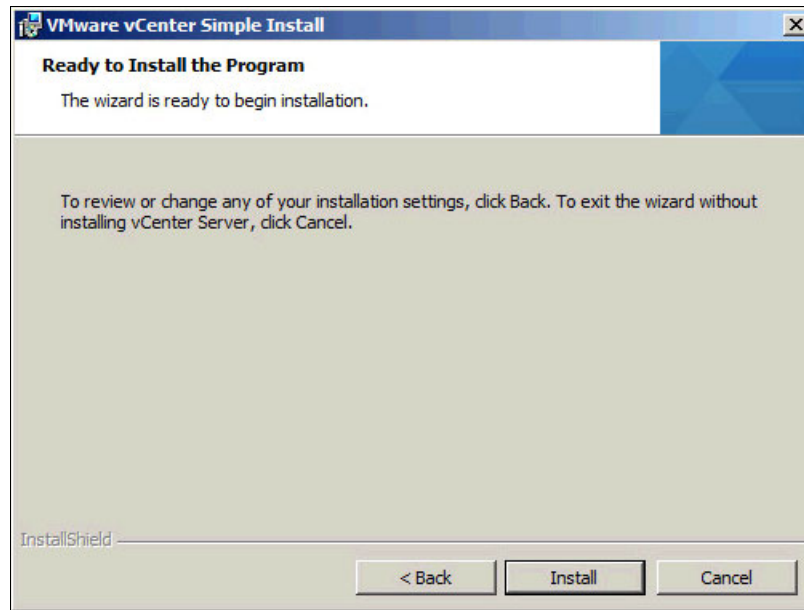


Figure 7-24 vCenter Server Ready to Install window

vCenter Server installation is now complete. All of the required components (SSO, Inventory Service, and vCenter Service) are installed. To connect to and manage the vCenter Server, you can use the Windows based vSphere Client. You can download the installation package from the VMware page. It is also available on the vCenter Server installation media and on each ESXi host. By using vSphere client, you also can connect and manage ESXi hosts directly.

vSphere 5.1 includes a vSphere Web Client with which you manage vCenter Server through a browser. The vSphere Web Client is the core administrative interface for vSphere and is platform-independent. The use of the vSphere Web Client provides the following advantages:

- ▶ You do not need a local Windows vSphere client application to access vCenter.
- ▶ The Tags feature provides user-defined labels or metadata with which you organize vCenter inventory.
- ▶ By using the Work In Progress feature, you can pause a configuration wizard and complete it later after you perform other tasks.
- ▶ An enhanced search feature provides a more granular search.

vSphere Web Client often is deployed in a vSphere 5.1 environment.

## 7.1.5 Installing vSphere Web Client

vSphere Web Client acts as a web server that connects to one or more vCenter Servers and accepts client browser connections. It can be installed separately from vCenter. The examples in this book install it on the same server as vCenter.

Complete the following steps to install vSphere Web Client:

1. Return to the autorun installation menu of the VMware vCenter Installer media. Select **VMware vSphere Web Client** and click **Install**, as shown in Figure 7-25.

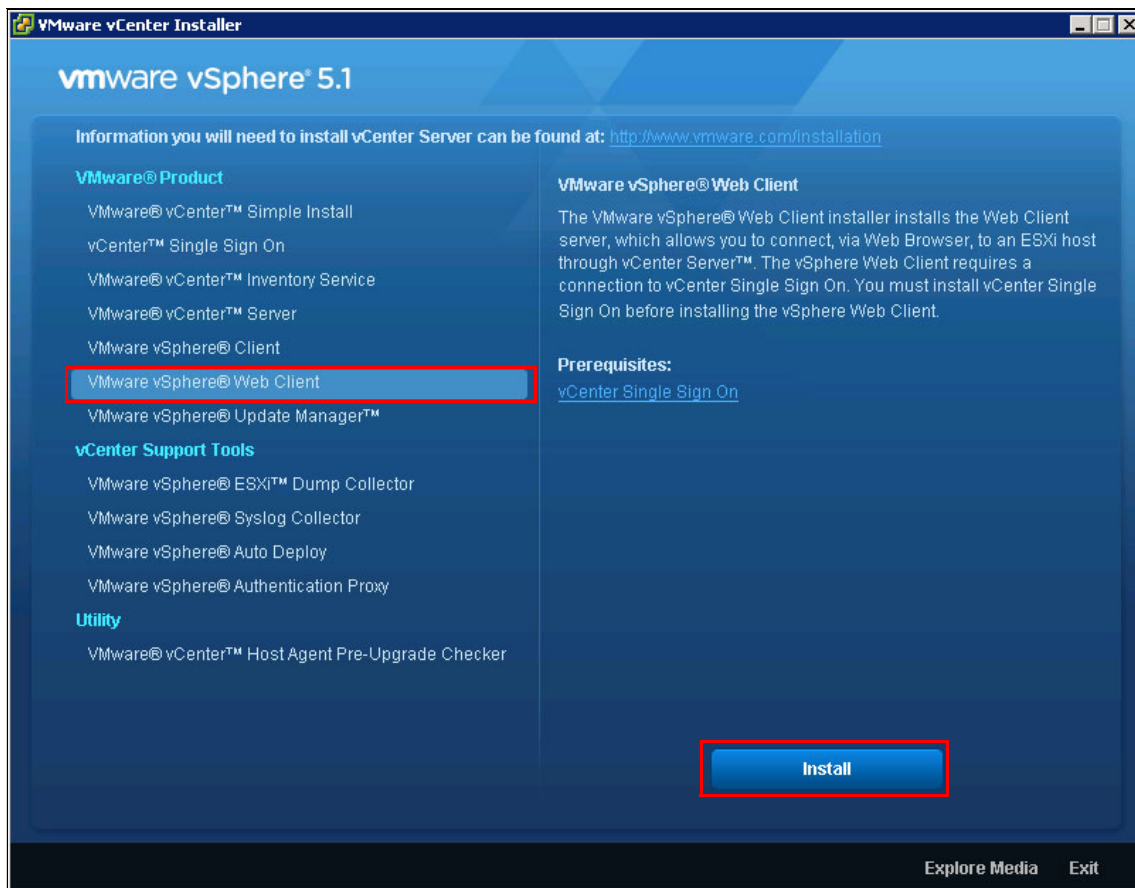


Figure 7-25 VMware vCenter Installer window

2. The wizard opens. Click **Next** to continue. Read and agree to the license agreement and then click **Next** to proceed. Specify the Web Client installation folder and click **Next**.
3. Leave the default Web Client port settings (as shown in Figure 7-26) by clicking **Next**.

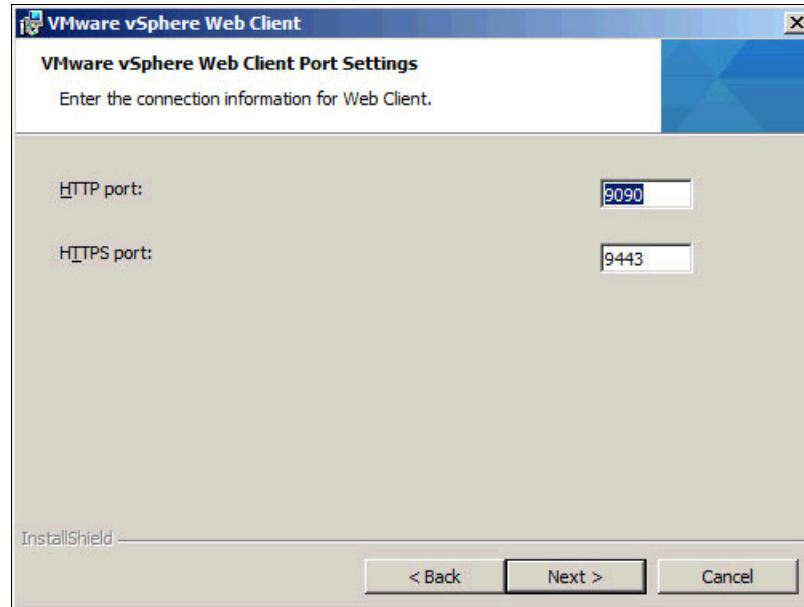
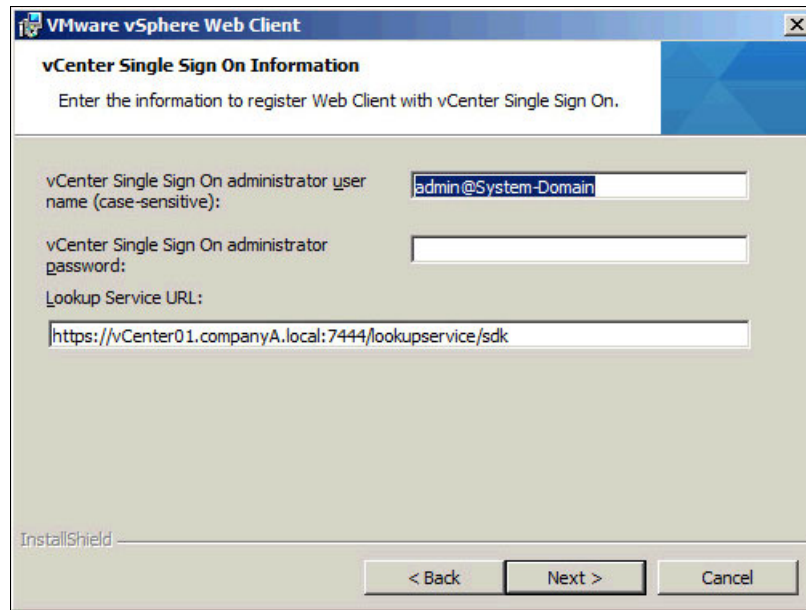


Figure 7-26 VMware vSphere Web Client Port Settings

4. Complete the Single Sign On information. Provide the information that was used during the SSO installation, as shown in Figure 7-27. Click **Next**.



The screenshot shows a window titled "VMware vSphere Web Client" with a sub-header "vCenter Single Sign On Information". Below the sub-header is the instruction "Enter the information to register Web Client with vCenter Single Sign On." The form contains three input fields: "vCenter Single Sign On administrator user name (case-sensitive):" with the value "admin@System-Domain", "vCenter Single Sign On administrator password:" which is empty, and "Lookup Service URL:" with the value "https://vCenter01.companyA.local:7444/lookupservice/sdk". At the bottom left is the "InstallShield" logo, and at the bottom right are three buttons: "< Back", "Next >" (which is highlighted), and "Cancel".

Figure 7-27 vCenter Single Sign On Information

5. The installation configuration is complete. Click **Install** to proceed with the installation process, as shown in Figure 7-28.

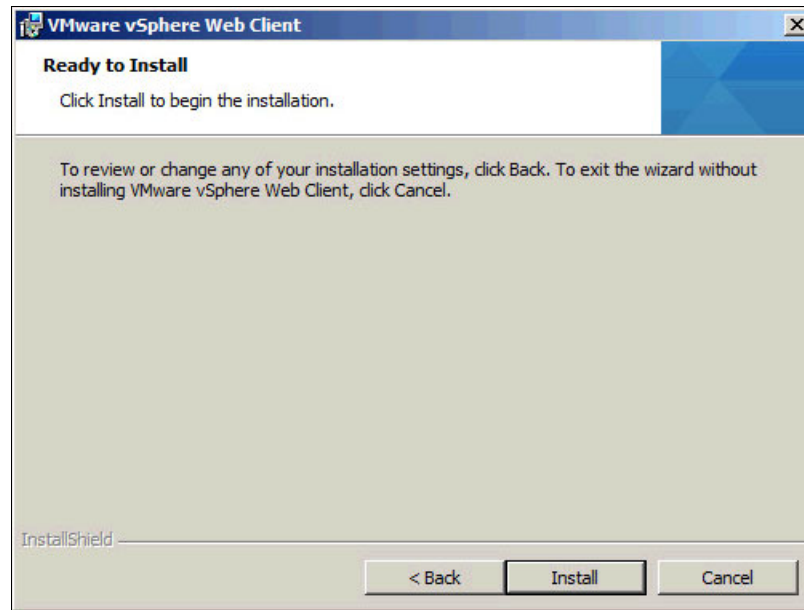


Figure 7-28 VMware vSphere Web Client Ready to Install window



6. Open a browser and browse to the following web address:  
`https://<web client IP or hostname>:9443/vsphere-client/`  
For the example in this book, the following address is used:  
`https://10.20.20.4:9443/vsphere-client/`  
The Web Client login page opens, as shown in Figure 7-29.

**Adobe Flex is required:** The vSphere Web Client uses Adobe Flex. You are required to install Adobe Flex on the system where the browser is installed.




Figure 7-29 VMware vSphere Web Client login page

7. Enter the credentials of a user who has access to vCenter.

**Client Integration plug-in required:** At the Web Client login window, you see a link to download and install the Client Integration plug-in. The Client Integration plug-in is required to use the vSphere Web Client to access a VM console. Installing the Client Integration plug-in might require a system restart.

## 7.2 Configuring vSphere

After the vCenter Server is installed, you must configure the vSphere environment. All of these steps can be performed by using the vSphere Client or the vSphere Web Client.

Complete the following steps configure vSphere:

1. Connect to vCenter by using vSphere Client or vSphere Web Client. Add licenses for vCenter and ESXi.
2. Create a datacenter object that is named DatacenterA.
3. Create the following host clusters:
  - Management
  - VDI
4. Add esxi01 and esxi02 to the Management cluster and add esxi03 and esxi04 to the VDI cluster.

Now, the configuration looks as shown in Figure 7-30.

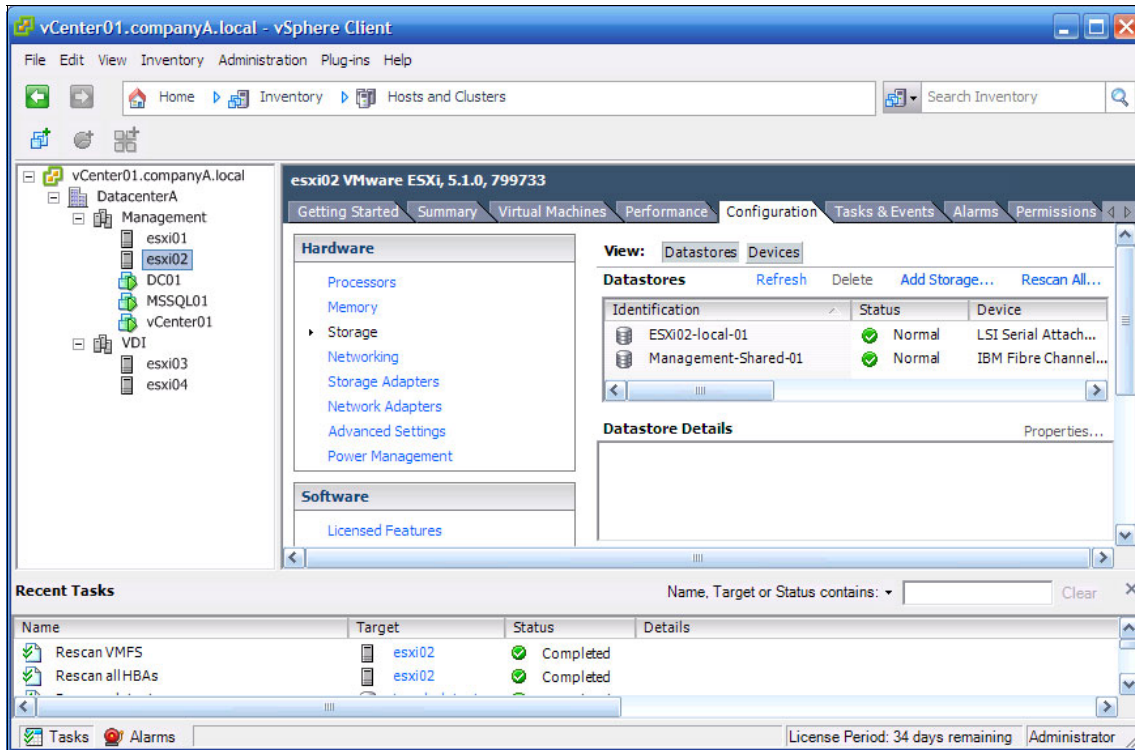


Figure 7-30 VDI cluster configuration

5. Ensure that all local and the two shared data stores are created, as shown in Figure 7-31.




Figure 7-31 Shared datastores

6. Configure virtual networking on esxi01 and esx02, as shown in Figure 7-32. The examples in this book use standard vSwitches for the Management hosts. Every switch has a single vmnic that is presented as a physical interface to ESXi but is, in fact, a vNIC that was created in Chapter 6, “Deploying IBM Flex System” on page 135. The adapter speed reflects the vNIC configuration on the physical switch.

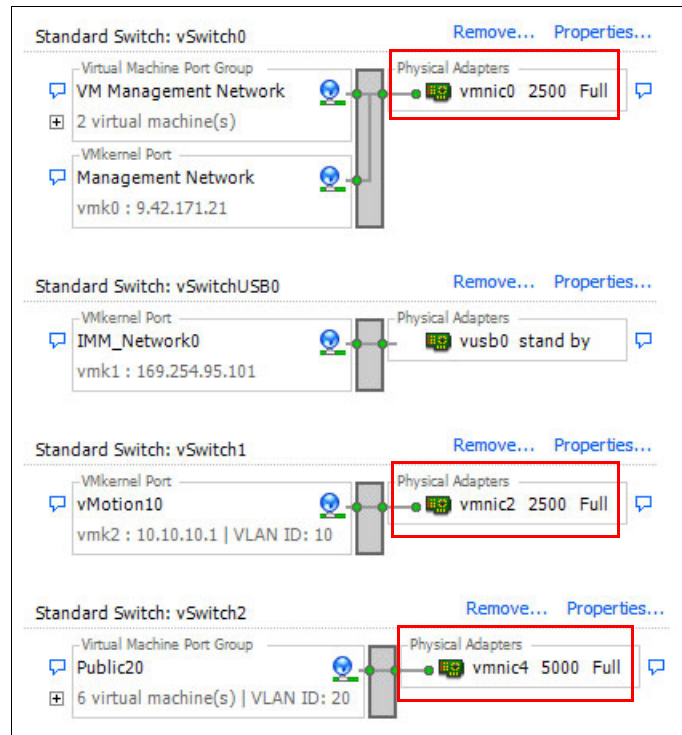


Figure 7-32 Management Cluster virtual networking configuration

**LAN over USB interface note:** IBM servers that are running ESXi 5.x with IBM customization have a standard virtual switch that is called *vSwitchUSB0*, as shown in Figure 7-32 on page 312. This vSwitch is generated automatically if the integrated management module (IMM) Ethernet over USB interface is enabled. If you remove this virtual switch, the vSwitch is re-created within 60 seconds.

The LAN over USB interface enables in-band communications to the IMM. The IMM hardware on the system board presents an internal Ethernet NIC from the IMM to the operating system. LAN over USB is also called the *USB in-band interface* in the IMM web interface.

The IMM IP address for the LAN over USB interface often is set to a static address of 169.254.95.101 with a subnet mask of 255.255.0.0. If there is an IP address collision on the network, the IMM might obtain a different IP address in the 169.254.xxx.xxx range.

This interface is required for in-band flashing of the IMM, UEFI, and dynamic system analysis (DSA) preboot firmware and for the Advance Setting Utility (ASU). This interface is enabled by default and should remain enabled.

7. Configure virtual networking for the VDI cluster for the esxi03 and esxi04 hosts, as shown in Figure 7-33. Add port groups and vmkernel ports to a single vSphere distributed switch (VDS). Remember the following points:
  - Uplink 1 is active only for the Management portgroup and is unused for the others.
  - Uplink 2 is only active for the vMotion portgroup and unused for the others.
  - Uplink 3 is only active for the Public portgroup and unused for the others.

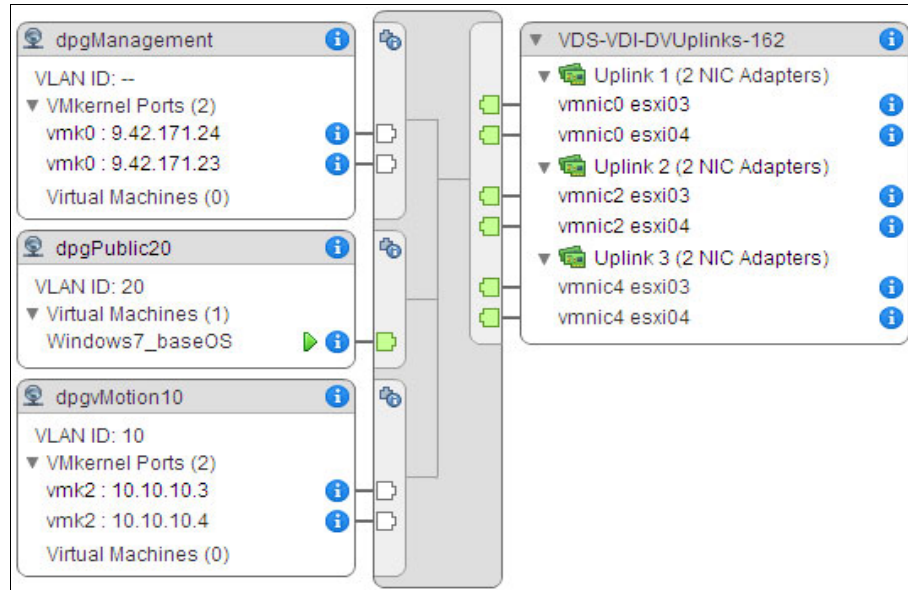


Figure 7-33 vSphere Distributed Switch configuration

8. Test vMotion and overall networking and storage configuration.

## 7.3 Installing View Composer

You install View Composer on the same server where vCenter is running. Before you start the View Composer installation package, create a DSN that points to the Composer database instance that you created in advance on the MS SQL server.

Complete the following steps to install View Composer:

1. Log in to the vCenter server OS. Browse to **Start** → **Administrative Tools** → **Data Sources (ODBC)**. Go to the System DSN tab, as shown in Figure 7-34. Click **Add**.

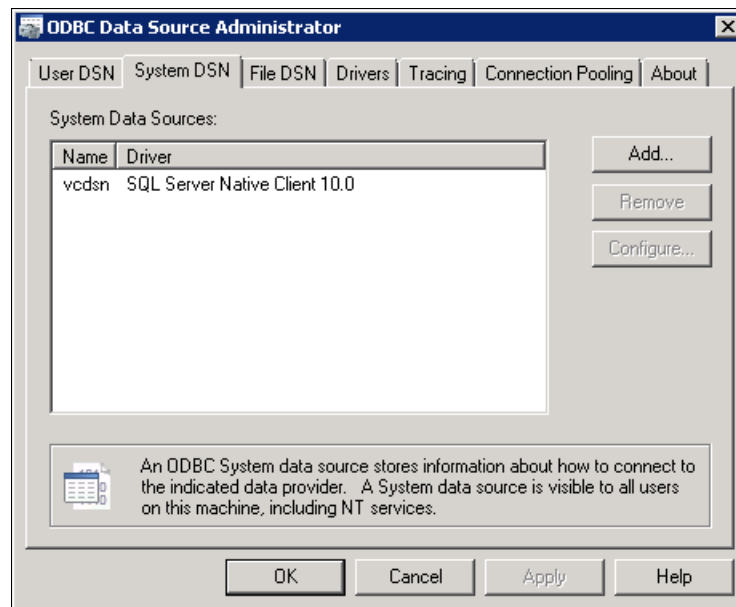


Figure 7-34 ODBC Data Source Administrator window

2. Select **SQL Server Native Client 10.0** and click **Finish**, as shown in Figure 7-35.

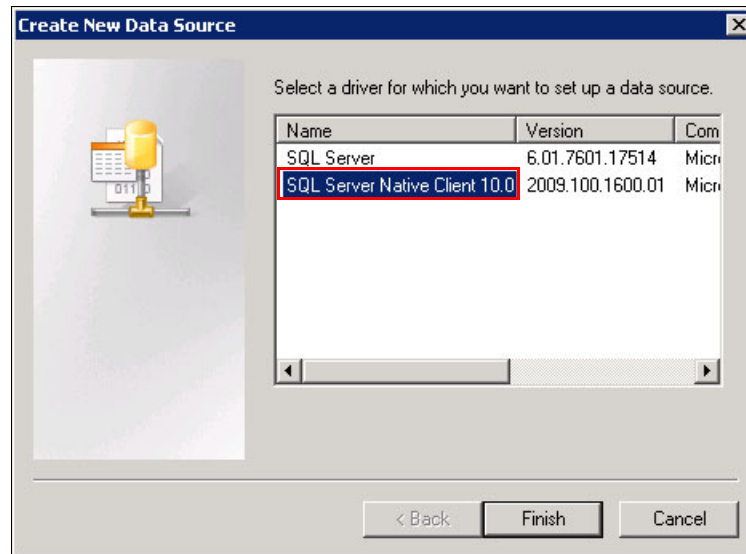


Figure 7-35 Create New Data Source driver selection window



3. Enter a name and description to which the data source and SQL Server connect. Enter `cmpdsn` as Name and `10.20.20.3` as the SQL Server address, as shown in Figure 7-36. Click **Next**.

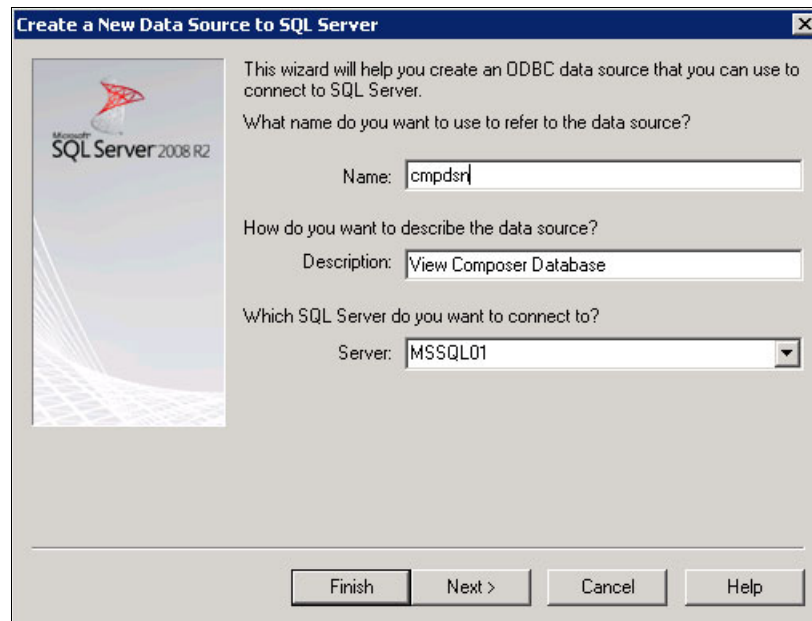


Figure 7-36 Create New Data Source wizard server information

4. Select the SQL Server authentication option and enter the Login ID and Password that has sufficient access rights to the View Composer DB instance, as shown in Figure 7-37. Click **Next**.

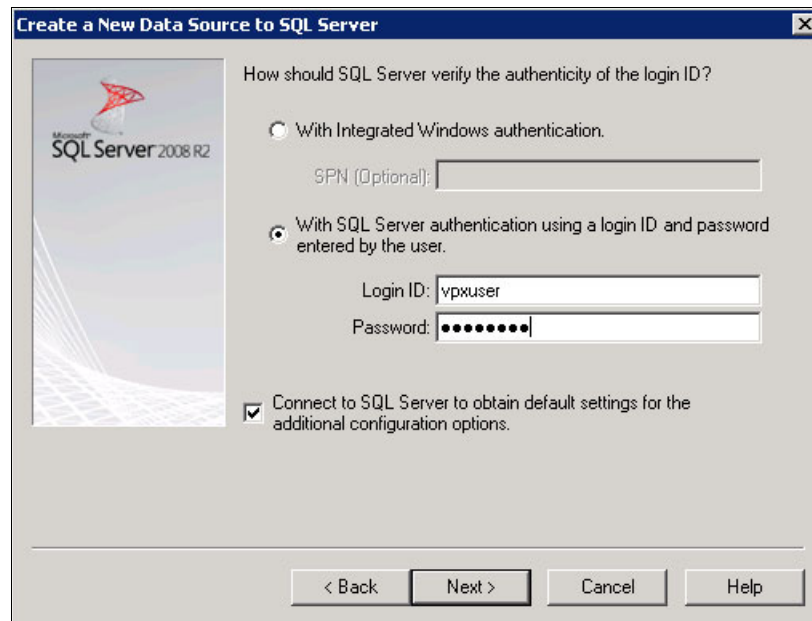


Figure 7-37 Create New Data Source wizard login information

5. Select **Change the default database to** and choose the correct View Composer DB instance from the drop-down menu. In this case, select **ViewCMPDB**, as shown in Figure 7-38. Click **Next**.

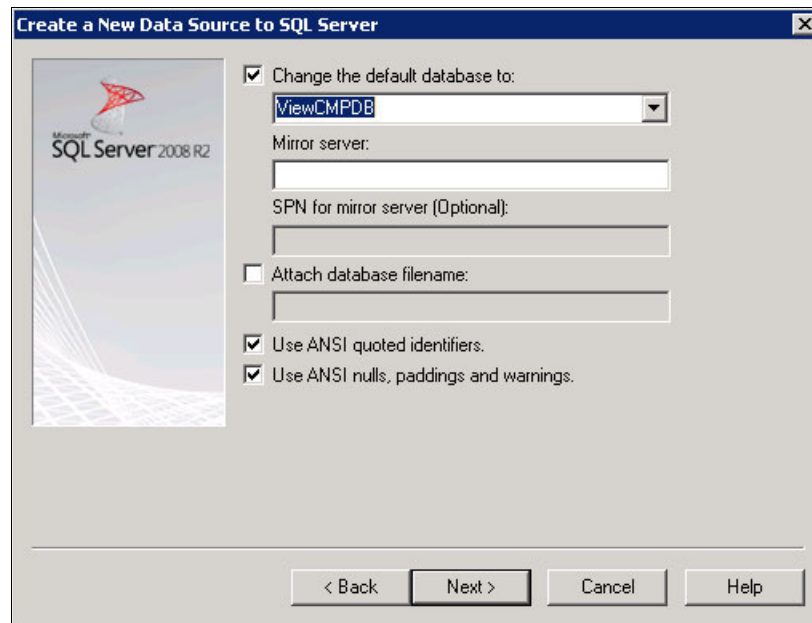
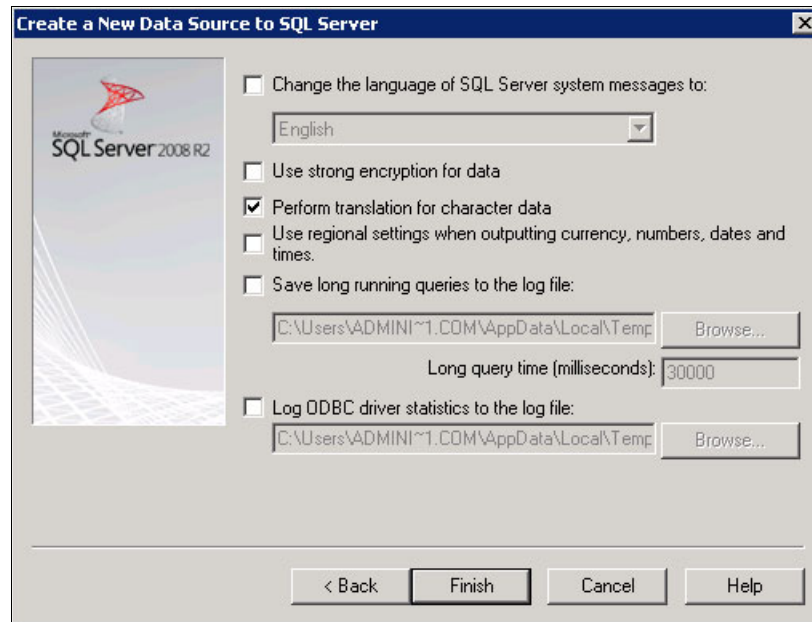


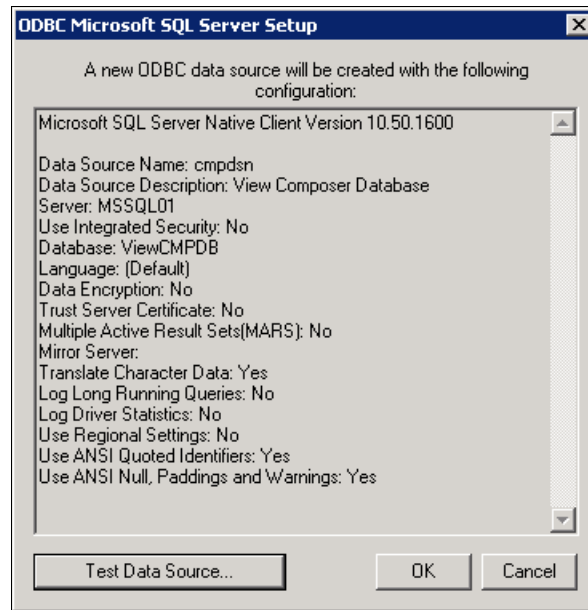
Figure 7-38 Create New Data Source wizard database instance

6. Leave the default configuration settings (as shown in Figure 7-39) and click **Next**.



*Figure 7-39 Create New Data Source wizard settings*

7. Review the data source creation summary window (as shown in Figure 7-40) and click **Test Data Source** to ensure that the configuration is correct.



*Figure 7-40 Create New Data Source wizard completion window*

8. When the DSN is created, proceed with the View Composer Installation. Copy the installation package to the vCenter server where View Composer is installed. The examples for this book used the current View Composer version, VMware-viewcomposer-5.2.0-983460.exe. Run the file. In the Welcome window, click **Next**, as shown in Figure 7-41.



Figure 7-41 View Composer Installation wizard

9. Read and accept license agreement and click **Next** to proceed. Specify the installation destination folder (as shown in Figure 7-42), and click **Next**. For the examples in this book, the View Composer destination folder is on the E: drive; however, you can install it on any partition that has free space available.


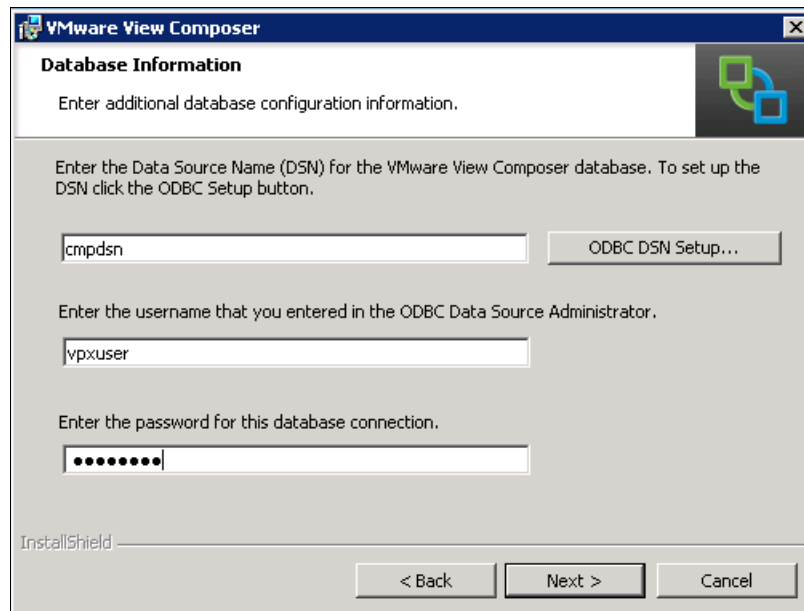


Figure 7-42 View Composer Destination Folder

10. Enter the View Composer DSN that you configured previously and specify the username and password for the connection. The configuration for the examples in this book is shown in Figure 7-43.



The screenshot shows the 'VMware View Composer' window with the 'Database Information' tab selected. The window has a title bar with the VMware logo and the text 'VMware View Composer'. Below the title bar is a section titled 'Database Information' with a sub-header 'Enter additional database configuration information.' and a VMware logo icon. The main area contains instructions: 'Enter the Data Source Name (DSN) for the VMware View Composer database. To set up the DSN click the ODBC Setup button.' Below this is a text input field containing 'cmpdsn' and a button labeled 'ODBC DSN Setup...'. Further down, it says 'Enter the username that you entered in the ODBC Data Source Administrator.' with a text input field containing 'vpxuser'. Below that, it says 'Enter the password for this database connection.' with a password input field showing eight dots. At the bottom left is the 'InstallShield' logo. At the bottom right are three buttons: '< Back', 'Next >', and 'Cancel'.

*Figure 7-43 View Composer Database Information*



11. Accept the default port and SSL configuration settings that are shown in Figure 7-44 by clicking **Next**.

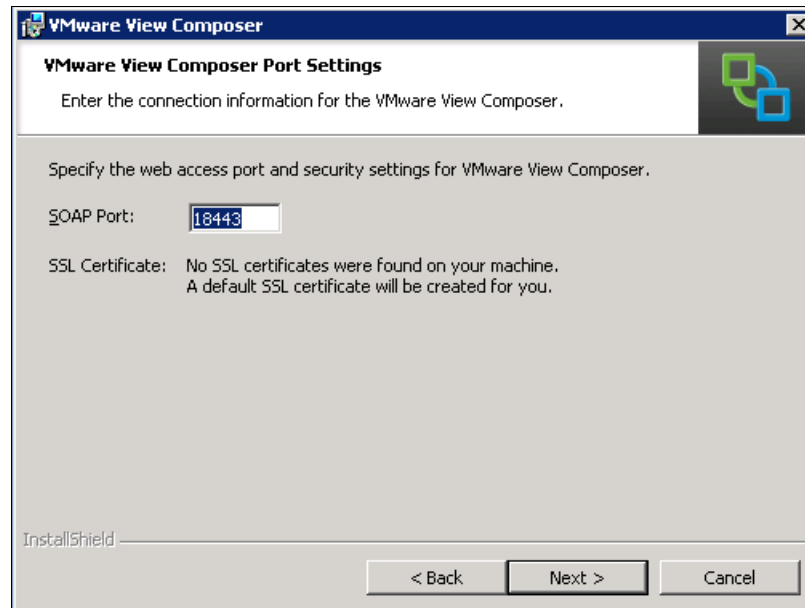
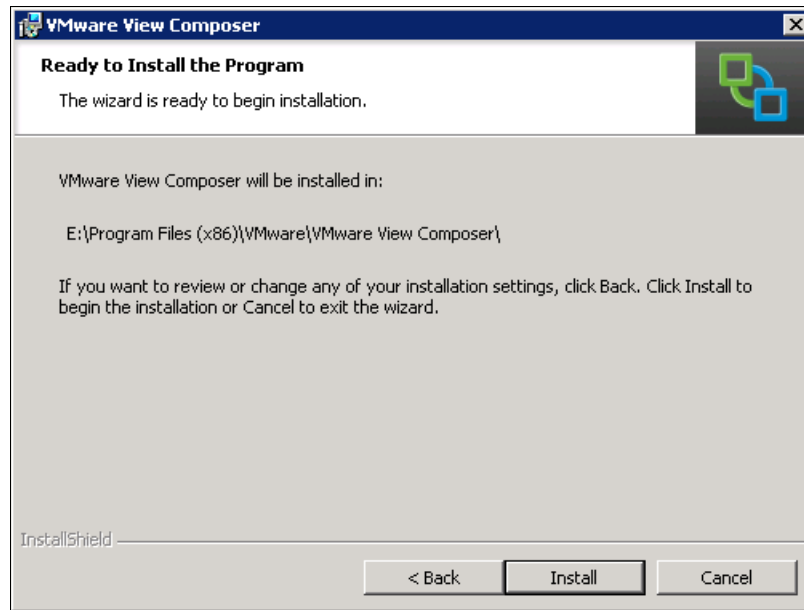


Figure 7-44 View Composer Port Settings

12. The installation configuration is complete. Click **Install** to proceed with the installation process, as shown in Figure 7-45.



*Figure 7-45 View Composer Ready to Install window*

13. When the VMware View Composer installation is complete, you are prompted to restart the system. Click **Yes** to restart now.

## 7.4 Installing View Connection Server

You provisioned a Windows operating system 2008 R2 VM that is named *CS01* where the View Connection Server is installed. Before the installation is started, ensure that the server is joined to a domain.

Complete the following steps to ensure that the server is joined to a domain:

1. Click **Start**. Right-click **Computer** and select **Properties** → **Change settings** → **Change**.
2. Complete the domain information, as shown in Figure 7-46.

**System restart:** If the server was not part of a domain, you must restart the system for the changes to take effect.

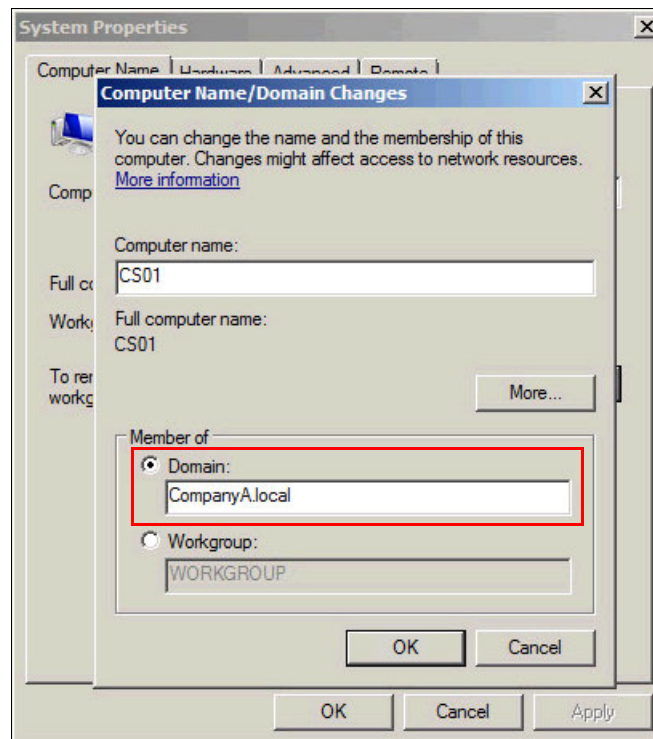


Figure 7-46 Join Computer to Domain window

Optionally, to easy administration, you can enable Remote Desktop on this VM, as shown in Figure 7-47.

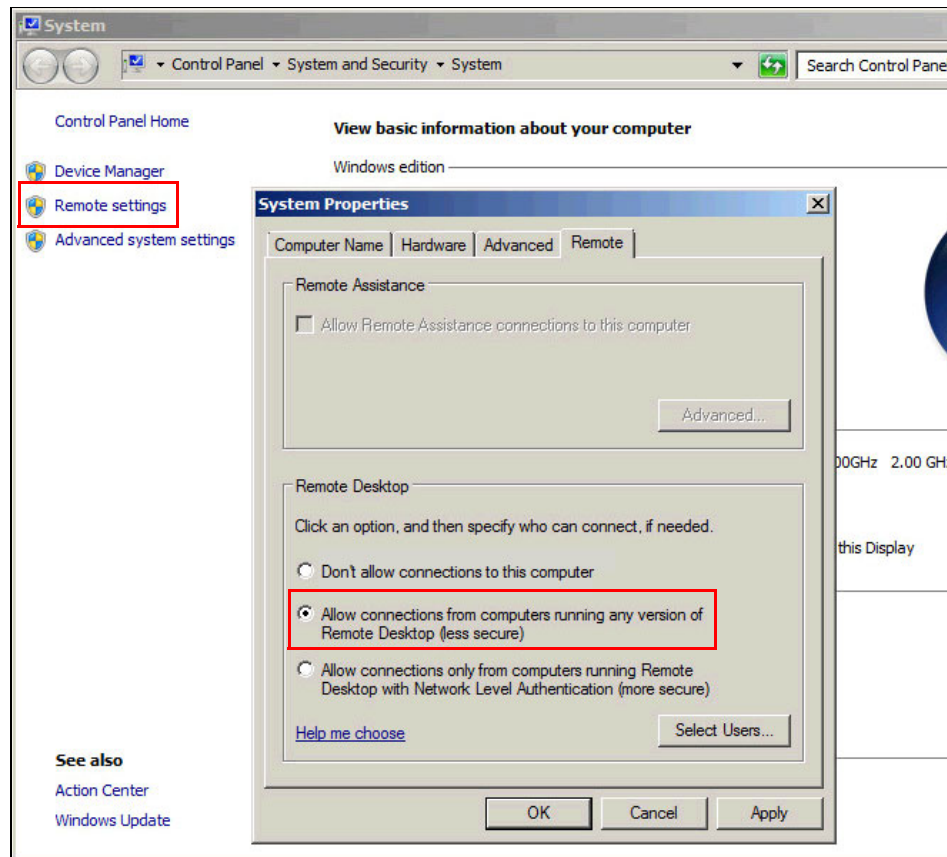


Figure 7-47 Computer Remote settings window

Complete the following steps to install View Connection Server:

1. Copy the View Connection Server installation package to CS01. Double-click the package to install it. The examples in this book used the following file:  
VMware-viewconnectionserver-x86\_64-5.2.0-987719.exe
2. In the Welcome window, click **Next**. Select the View Connection Server installation destination folder (as shown in Figure 7-48 on page 329) and click **Next**.




Figure 7-48 View Connection Server Destination Folder

3. Select the default option to install a standard stand-alone instance of View Connection Server, as shown in Figure 7-49. Click **Next**.

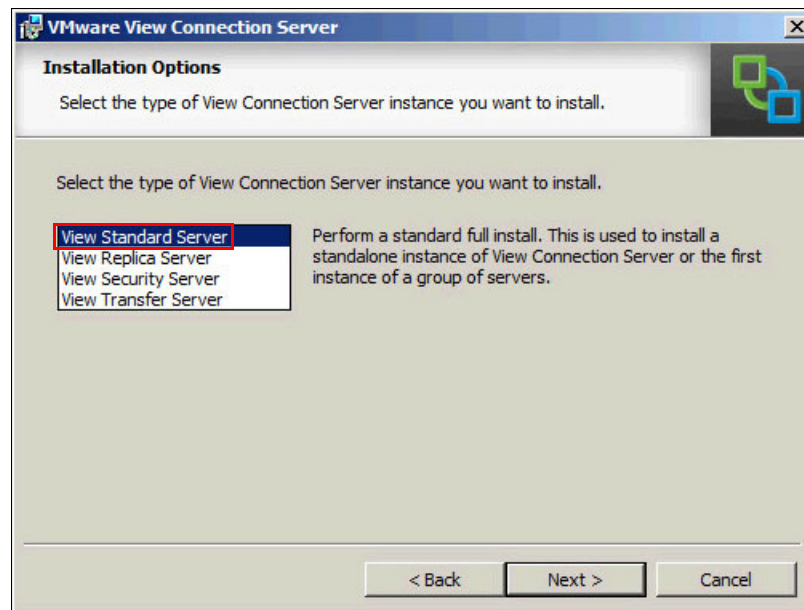


Figure 7-49 View Connection Server instance type window

4. Accept the default firewall configuration value as shown in Figure 7-50 by clicking **Next**.

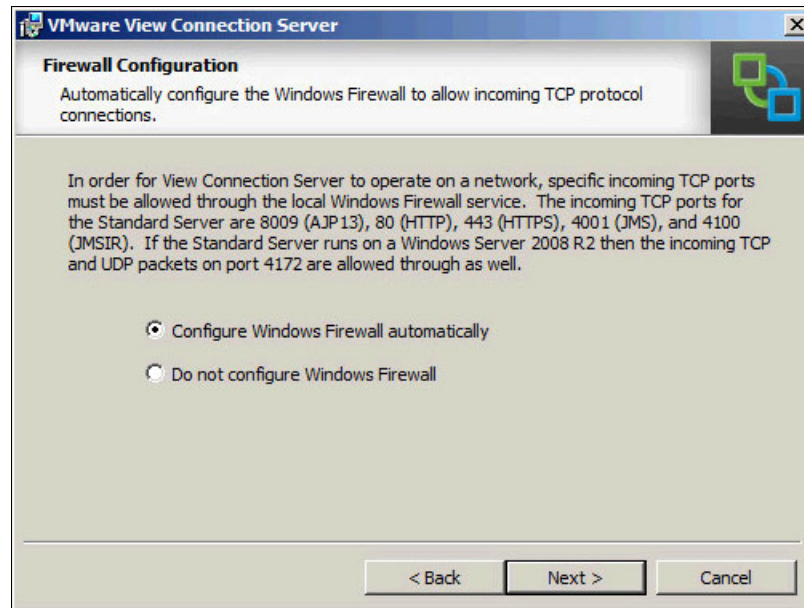


Figure 7-50 View Connection Server Firewall Configuration

5. Select **Authorize a specific domain user or domain group** and enter the View administrators group that you created in Active Directory. For this example, the administrators group is CompanyA.local\View Admins, as shown in Figure 7-51. Click **Next**.

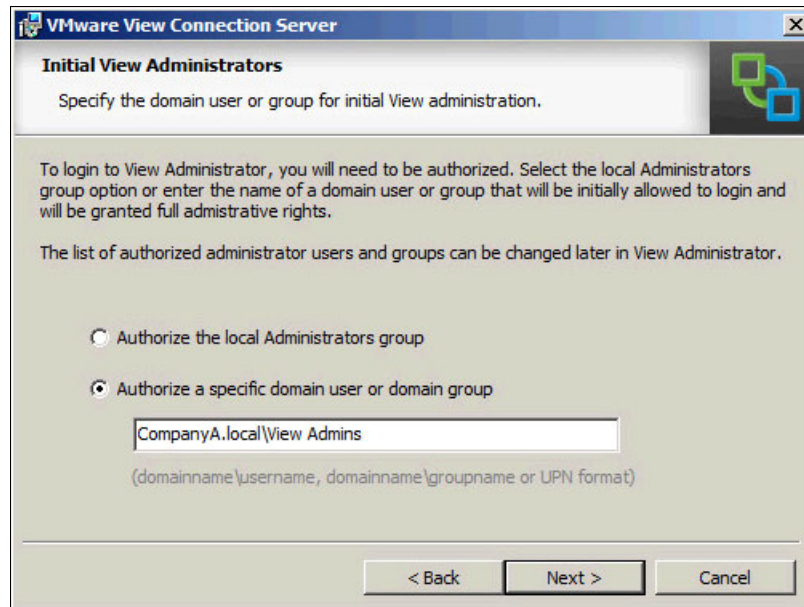


Figure 7-51 View Connection Server Administrators window



6. Specify the domain administrator details to access details of the previously created domain group, as shown in Figure 7-52. Click **Next**.

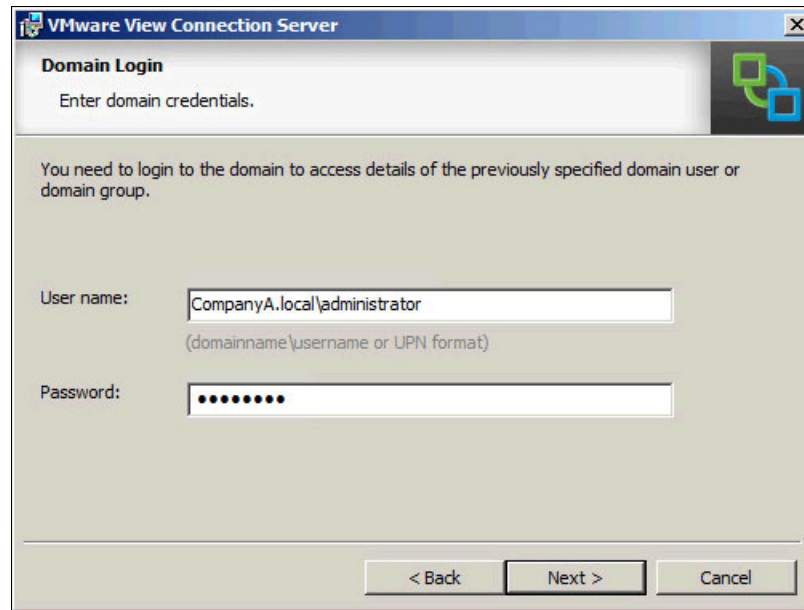


Figure 7-52 View Connection Server Domain Login window

7. Clear **Participate anonymously in the user experience improvement program**, as shown in Figure 7-53. Click **Next**.

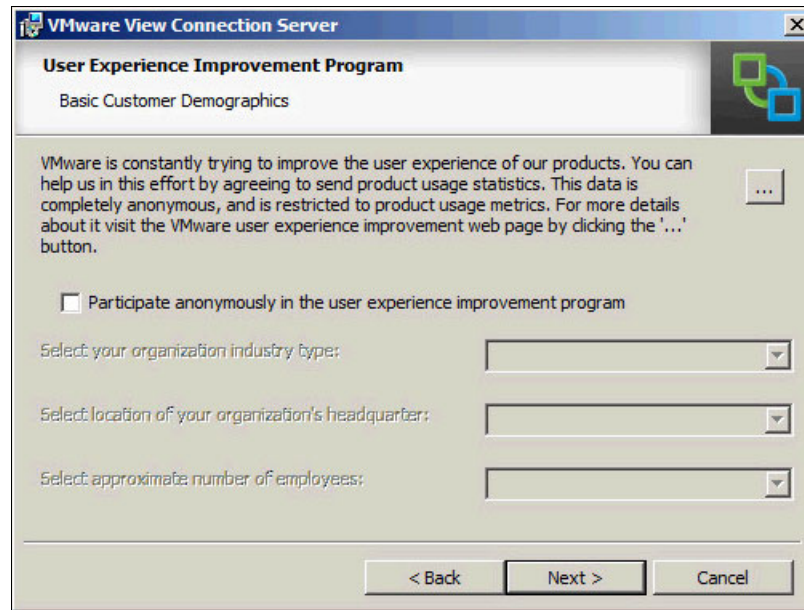


Figure 7-53 View Connection Server User Experience Improvement Program

8. The installation configuration is complete. Click **Install** to proceed with the installation process, as shown in Figure 7-54.




Figure 7-54 View Connection Server Ready to Install window

## 7.5 Configuring View Connection Server initially

VMware Horizon View Administrator is a web-based application that allows administrators to configure View Connection Server. This interface is the main interface for a View administrator to allow the deployment and management of View desktops. It is installed automatically when you install View Connection Server.

To configure View Connection Server initially:

1. Open a browser and browse to the following web address:

`https://<connection server IP or hostname>/admin/`

In this example, the address is `https://CS01:9443/admin/`. The View Administrator login page displays, as shown in Figure 7-55. Enter the user credentials of a user member of the View Admins group and click **Login**.

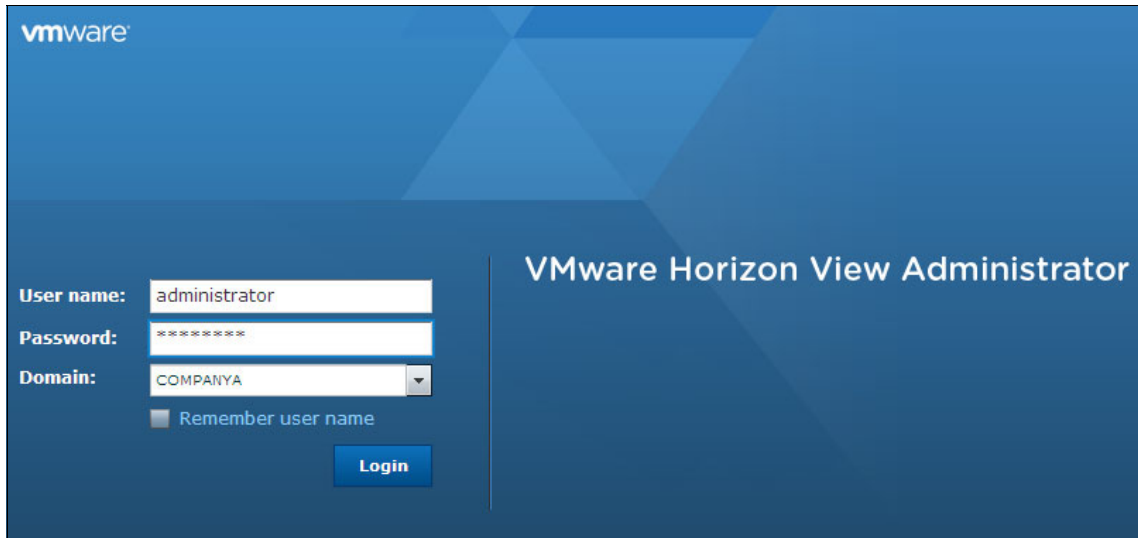


Figure 7-55 Horizon View Administrator login window

2. You must configure the View Manager license. In the Inventory section of the left pane, select **View Configuration** → **Product Licensing and Usage**. In the Licensing and Usage section of the right pane, click **Edit License**, as shown in Figure 7-56.

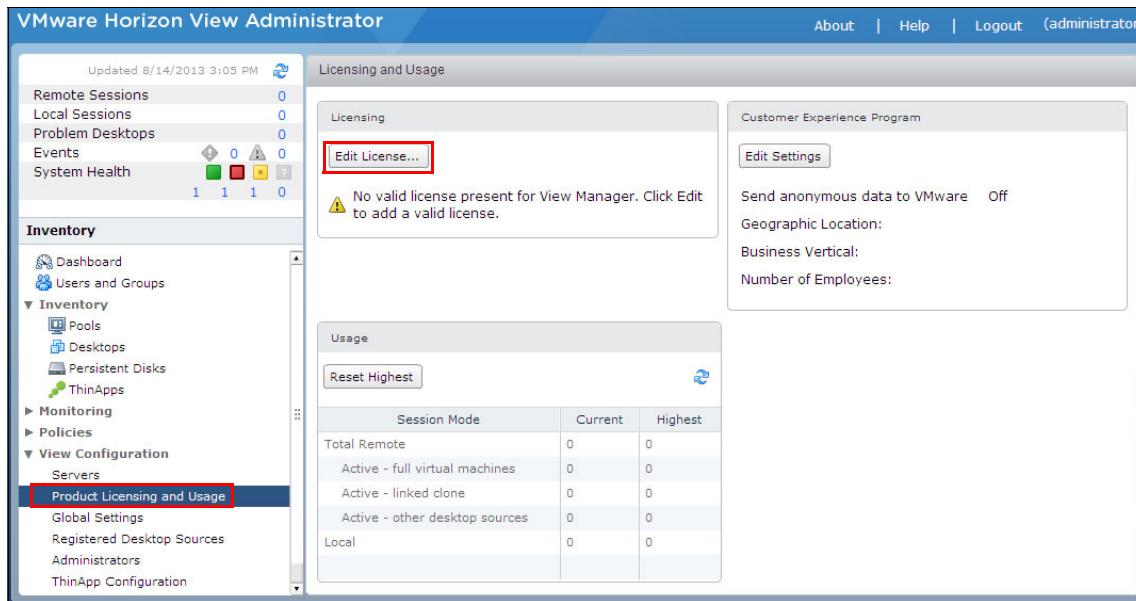


Figure 7-56 View Administrator Licensing page

3. Enter a valid license serial number, as shown in Figure 7-57. Click **OK**.



Figure 7-57 View Administrator Edit License window

4. Add a vCenter Server to the Connection Server configuration. From the left Inventory menu, select **View Configuration** → **Servers**. In the vCenter Servers tab, click **Add**, as shown in Figure 7-58.

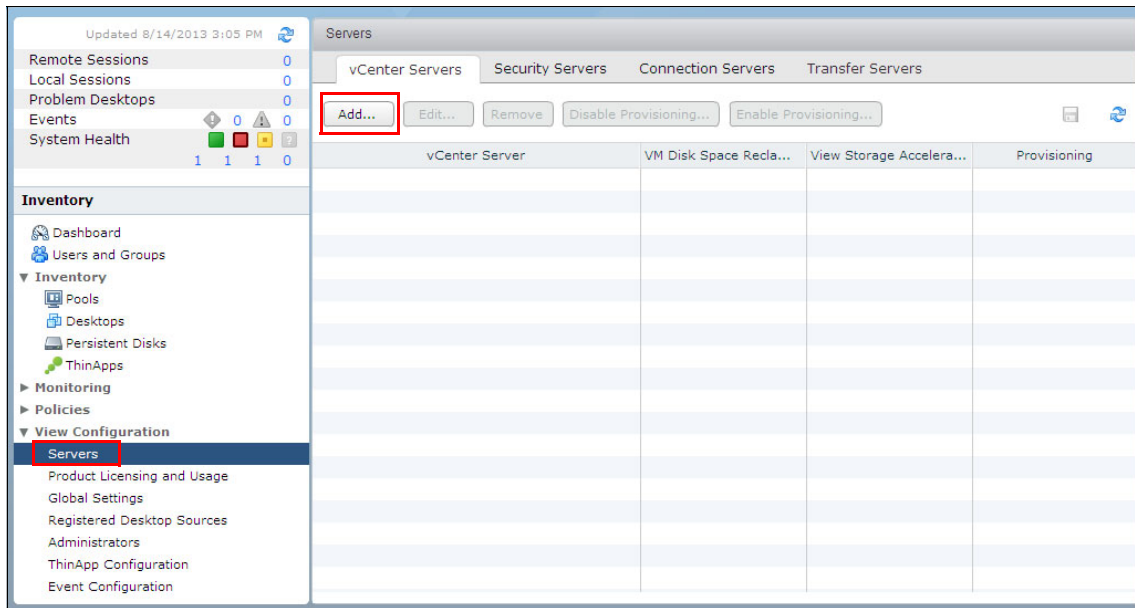


Figure 7-58 View Administrator vCenter Servers page

5. Complete the vCenter Server Settings information, as shown in Figure 7-59. Use a user account with sufficient privileges for View operations. This example uses an Administrator account that has administrator privileges for vCenter. Leave all Advanced Settings as default. Click **Next**.

**Add vCenter Server**

**VC Information**

View Composer

Storage

Ready to Complete

**vCenter Server Information**

**vCenter Server Settings**

Server address: vCenter01.CompanyA.local

User name: Administrator

Password: \*\*\*\*\*

Description:

Port: 443

**Advanced Settings**

Specify the concurrent operation limits.

Max concurrent vCenter provisioning operations: 20

Max concurrent power operations: 50

Max concurrent View Composer maintenance operations: 12

Max concurrent View Composer provisioning operations: 8

**vCenter Server Settings**

Before you add vCenter Server to View, install a valid SSL certificate signed by a trusted CA. In a test environment, you can use the default, self-signed certificate that is installed with vCenter Server, but you must accept the certificate thumbprint.

Provide the vCenter Server FQDN or IP address, user name, and password.

**Concurrent Operations Limits**

Max concurrent vCenter provisioning operations: the maximum number of concurrent VM cloning and deletion operations on this vCenter server (full clones).

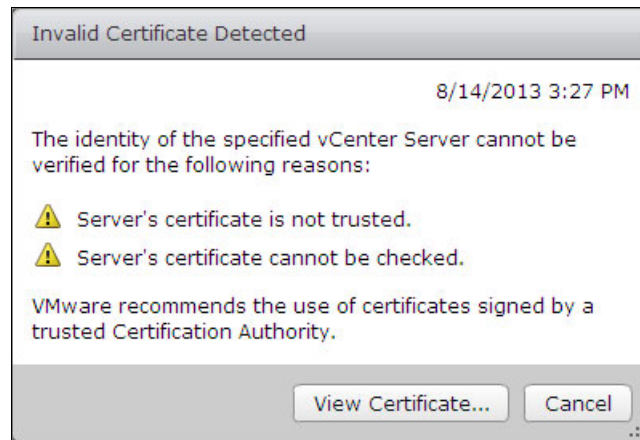
Max power operations: the maximum number of concurrent VM power-on, power-off, reset, and configuration operations (full clones and linked clones).

Max View Composer maintenance operations: the maximum number of concurrent View Composer

Next > Cancel

Figure 7-59 Add vCenter Server window

6. Unless you added a vCenter Server certificate that was signed by a trusted Certification Authority, you are prompted with the warning that is shown in Figure 7-60. Click **View Certificate**.



*Figure 7-60 vCenter Server certificate warning window*



7. This example uses the default self-signed vCenter Server certificate that is shown in Figure 7-61. Click **Accept**.

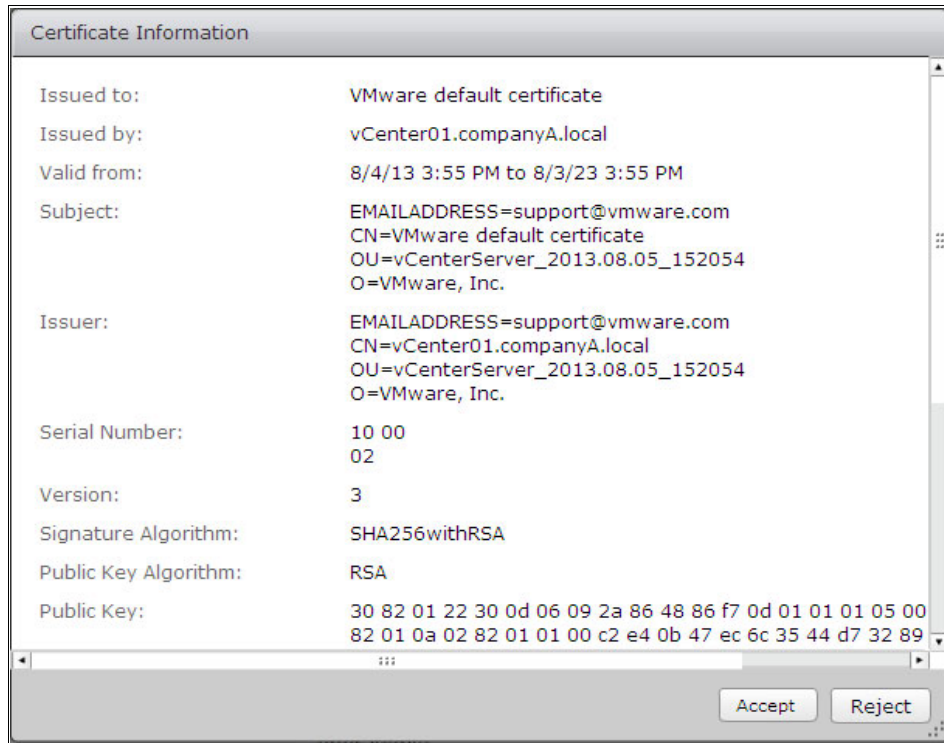


Figure 7-61 vCenter Server Certificate Information window

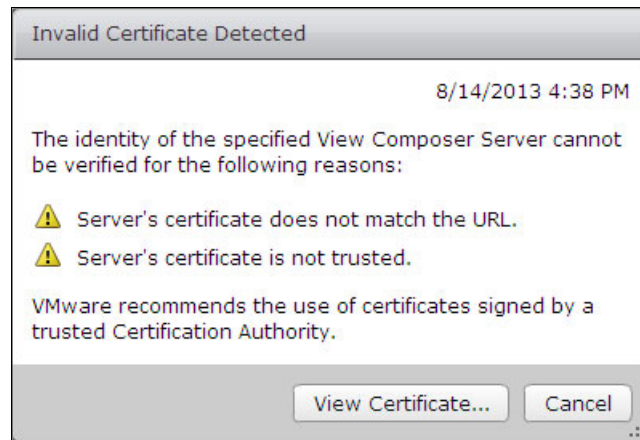
8. Configure the View Composer Settings. Select **View Composer co-installed with vCenter Server** and leave the default port, as shown in Figure 7-62. Click **Next**.

The screenshot shows the 'Add vCenter Server' wizard with the 'View Composer' tab selected. The 'View Composer Settings' section has two radio buttons: 'Do not use View Composer' and 'View Composer co-installed with vCenter Server'. The second option is selected. Below it, a text box says 'Choose this if View Composer is installed on the same server as vCenter'. The 'Port' field is set to '18443'. The 'Standalone View Composer Server' option is unselected, with a text box saying 'Choose this if View Composer is installed on a separate server from vCenter'. Below this, fields for 'Server address' (vCenter01.CompanyA.local), 'User name' (Administrator), 'Password' (masked with asterisks), and 'Port' (18443) are visible. On the right, a 'View Composer Settings' panel contains explanatory text about SSL certificates. At the bottom right are '< Back', 'Next >', and 'Cancel' buttons.

Add vCenter Server	
<b>Add vCenter Server</b> VC Information <b>View Composer</b> View Composer Domains Storage Ready to Complete	<b>View Composer</b>  <b>View Composer Settings</b>  <input type="radio"/> Do not use View Composer  <input checked="" type="radio"/> View Composer co-installed with vCenter Server Choose this if View Composer is installed on the same server as vCenter Port: <input type="text" value="18443"/>  <input type="radio"/> Standalone View Composer Server Choose this if View Composer is installed on a separate server from vCenter Server address: <input type="text" value="vCenter01.CompanyA.local"/> User name: <input type="text" value="Administrator"/> Password: <input type="password" value="*****"/> Port: <input type="text" value="18443"/>  <div>&lt; Back   Next &gt;   Cancel</div>

Figure 7-62 View Composer Settings window

9. Unless you added a View Composer Server certificate that is signed by a trusted Certification Authority, you are prompted with the warning that is shown in Figure 7-63. Click **View Certificate**.



*Figure 7-63 View Composer Server certificate warning window*

10. This example uses the default self-signed vCenter Server certificate that is shown in Figure 7-64. Click **Accept**.

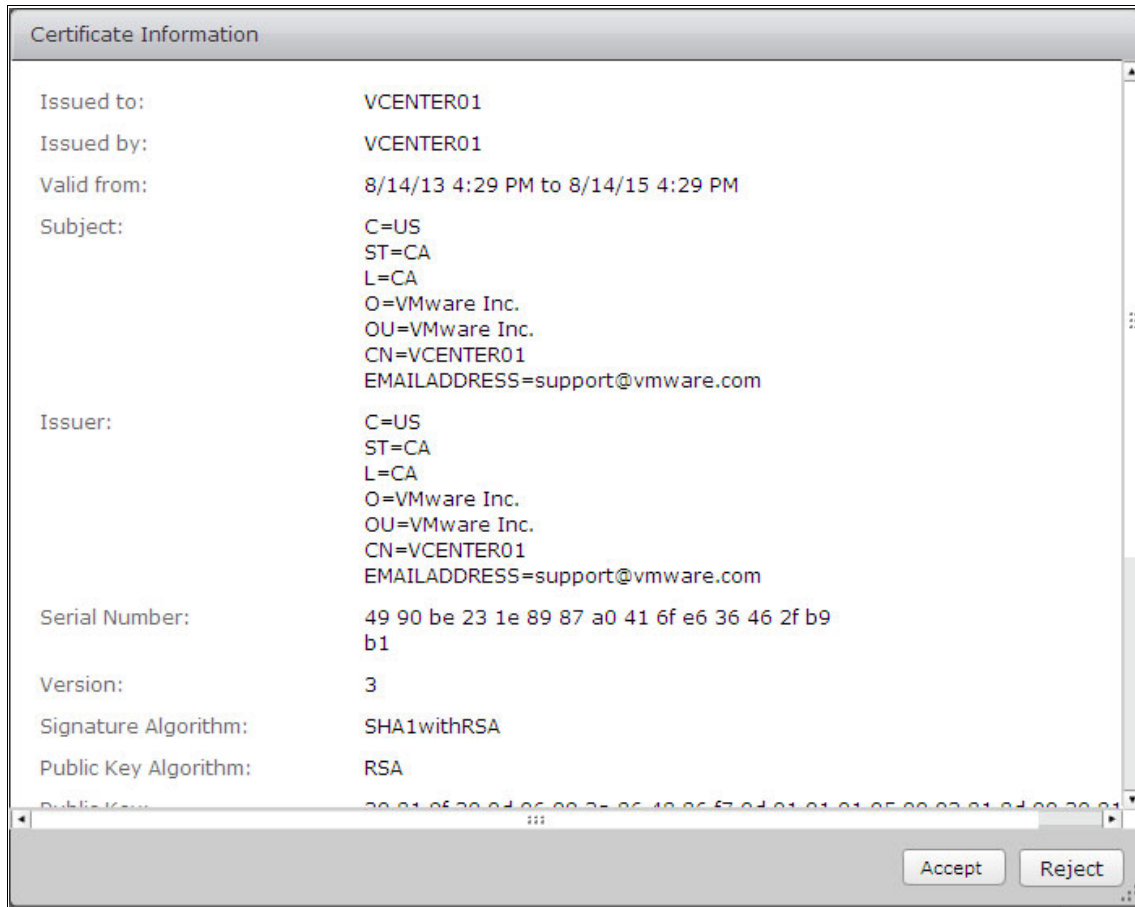


Figure 7-64 View Composer Server certificate information window

11. Click **Add** in the View Composer Domains window, as shown in Figure 7-65.


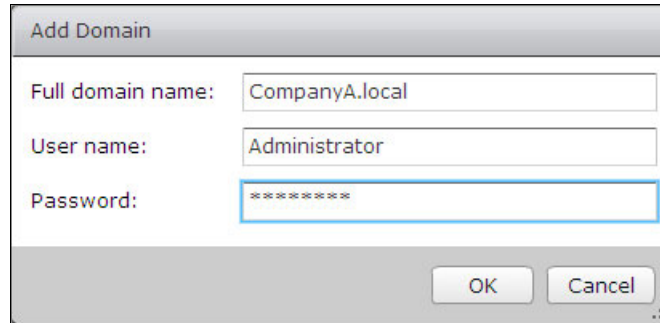


Figure 7-65 View Composer Domains window

12. Enter a domain where the virtual desktop computer accounts are created and an account with a minimum domain permission of Create Computer Objects, Delete Computer Objects, and Write All Properties. This example uses the domain administrator account, as shown in Figure 7-66. For increased security, consider creating a special domain account with minimum required privileges. Click **OK** and then click **Next**.

A screenshot of a Windows-style dialog box titled "Add Domain". It contains three input fields: "Full domain name:" with the text "CompanyA.local", "User name:" with the text "Administrator", and "Password:" with a masked password represented by ten asterisks. At the bottom right, there are "OK" and "Cancel" buttons. The dialog box has a standard Windows XP-style border and a small icon in the bottom right corner.

Add Domain

Full domain name: CompanyA.local

User name: Administrator

Password: \*

OK Cancel

*Figure 7-66 View Composer Add Domain window*

13. Under Storage Settings, accept all of the default values, as shown in Figure 7-67. Ensure that the Reclaim VM disk space and Enable View Storage Accelerator options are selected. Click **Next**.

**Add vCenter Server**

VC Information  
View Composer  
View Composer Domains  
**Storage**  
Ready to Complete

**Storage**

**Storage Settings**

☒ Reclaim VM disk space

☒ Enable View Storage Accelerator

Default host cache size:  MB

Cache must be between 100 MB and 2048 MB

**Hosts**

☐ Show all hosts

Host	Cache Size
/DatacenterA/host/Management/esxi01	Default
/DatacenterA/host/Management/esxi02	Default
/DatacenterA/host/VDI/esxi03	Default
/DatacenterA/host/VDI/esxi04	Default

**Storage Settings**

ESXi hosts can be configured to cache virtual machine disk data, which improves performance during I/O storms such as when many desktops power on and run anti-virus scans at once. Hosts read common data blocks from cache instead of reading the OS from disk.

By reducing IOPS during boot storms, View Storage Accelerator lowers the demand on the storage array and uses less storage I/O bandwidth.

< Back   Next >   Cancel

Figure 7-67 Storage Settings window

14. Review the configuration (as shown in Figure 7-68) and click **Finish**.

Add vCenter Server		
Ready to Complete		
vCenter Server		vCenter01.CompanyA.local
User name		Administrator
Password		*****
Description		
Server Port		443
Max Provision		20
Max Power		50
Max View Composer Operations		12
Max View Composer Provision		8
View Composer State		View Composer co-installed with vCenter Serv
View Composer Address		vCenter01.CompanyA.local
View Composer Password		*****
View Composer User Name		Administrator
View Composer Port		18443
Enable View Storage Accelerator		Yes
Default host cache size:		1024
VM Disk Space Reclamation		Yes

Figure 7-68 Add vCenter Server Ready To Complete window

The newly added vCenter Server now displays in the list of vCenter Servers to which the View Connection Server is connected, as shown in Figure 7-69.

Servers			
vCenter Servers			
Security Servers			
Connection Servers			
Transfer Servers			
Add... Edit... Remove Disable Provisioning... Enable Provisioning...			
vCenter Server	VM Disk Space Recla...	View Storage Accelera...	Provisioning
vCenter01.CompanyA.local(Administrator)	✓	✓	

Figure 7-69 vCenter Servers page

The base Horizon View infrastructure is now deployed. For more information, see Chapter 9, “Operating VMware Horizon View infrastructure” on page 373.





# Operating IBM Flex System

In this chapter, we describe how to operate your IBM Flex System where Horizon View is deployed and some common View-related tasks that are performed by using Flex System Manager.

This chapter includes the following topics:

- ▶ Configuring VMControl for vSphere integration
- ▶ Navigating the vSphere environment by using FSM
- ▶ Automating tasks
- ▶ Introducing IBM FSM Explorer
- ▶ Monitoring and logging in Flex System

## 8.1 Configuring VMControl for vSphere integration

In this section, we use VMControl to discover the vSphere environment. The only task that must be performed is to discover and gain access to the vCenter server. After it is connected to vCenter, VMControl pulls and synchronizes continuously all vSphere inventory and performance data.

Complete the following steps to configure VMControl for vSphere integration:

1. In the Flex System Manager (FSM) web interface navigation area, expand **Inventory** and click **System Discovery**. Enter the IP address of the vCenter Server and click **Discover Now**, as shown in Figure 8-1.

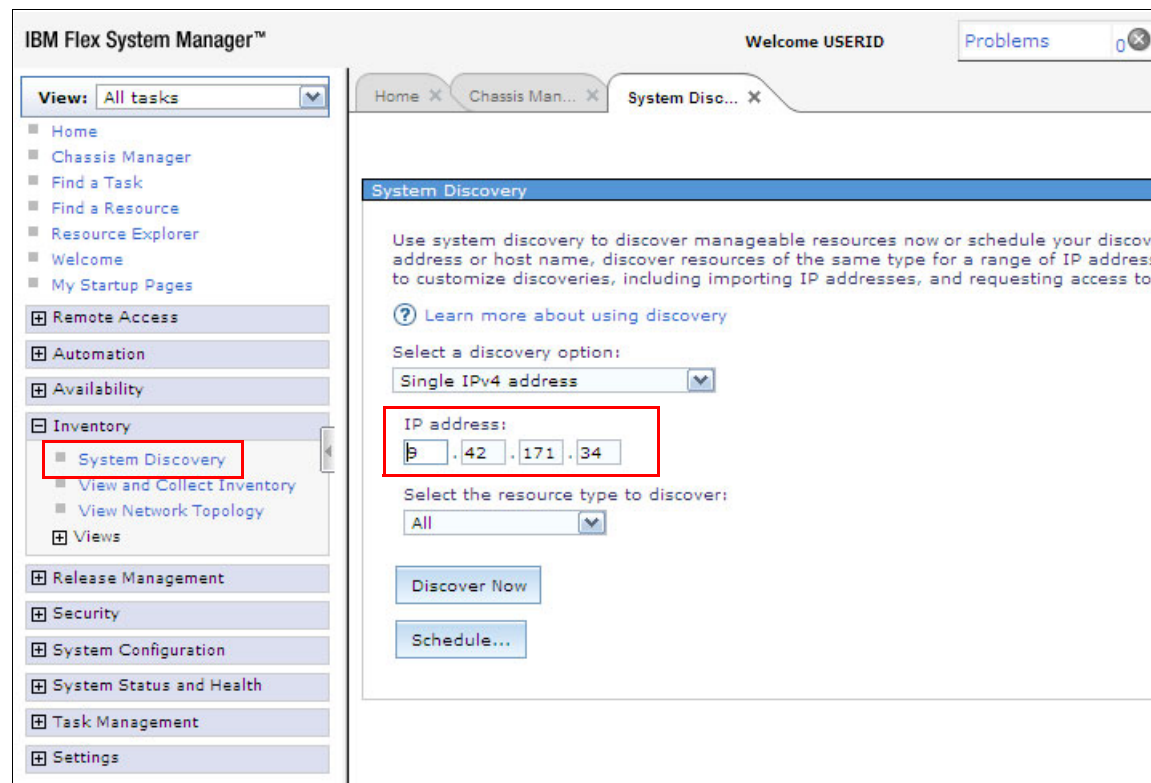


Figure 8-1 System Discovery window

2. With the system successfully discovered, click **No access** to configure credentials to access vCenter, as shown in Figure 8-2.

Discovered Manageable Systems:

Actions ▼						
Name	Discovered	Type	Access	Problems	Compliance	IP
9.42.171.34	New	Operating Sys...	No access	OK	OK	9.4

Figure 8-2 Discovered Manageable Systems window

3. Enter the user ID and password of a vCenter administrator and then click **Request Access**, as shown in Figure 8-3.

Request Access

Specify the user ID and password to authenticate Flex System Manager to one or more target systems. Then click Request Access to grant all authorized Flex System Manager users access to the target system(s).

\*User ID:  
administrator

\*Password:  
\*\*\*\*\*

Request Access Close

Selected targets:

Name	Access	Trust State
9.42.171.34	No access	Not Trusted

Figure 8-3 Request Access window

If the user ID is successfully authenticated, the vCenter Server Access column shows OK, as shown in Figure 8-4.

Request Access

Specify the user ID and password to authenticate Flex System Manager to one or more target systems. Then click Request Access to grant all authorized Flex System Manager users access to the target system(s).

User ID:  
administrator

Password:  
\*\*\*\*\*

Request Access Close

Selected targets:

Name	Access	Trust State
9.42.171.34	OK	Trusted

Figure 8-4 Access granted window

VMControl now starts pulling the vCenter inventory and performance information, which becomes visible in the FSM web interface.

## 8.2 Navigating the vSphere environment by using FSM

After VMControl is configured and vCenter is properly discovered, the vCenter objects (VMs, clusters, datacenters, networks, and so on) become visible in the FSM web interface.

To view the vCenter servers, vSphere datacenters, and clusters, click **Inventory** → **Views** in the navigation area and then click **Platform Managers and Members**, as shown in Figure 8-5.

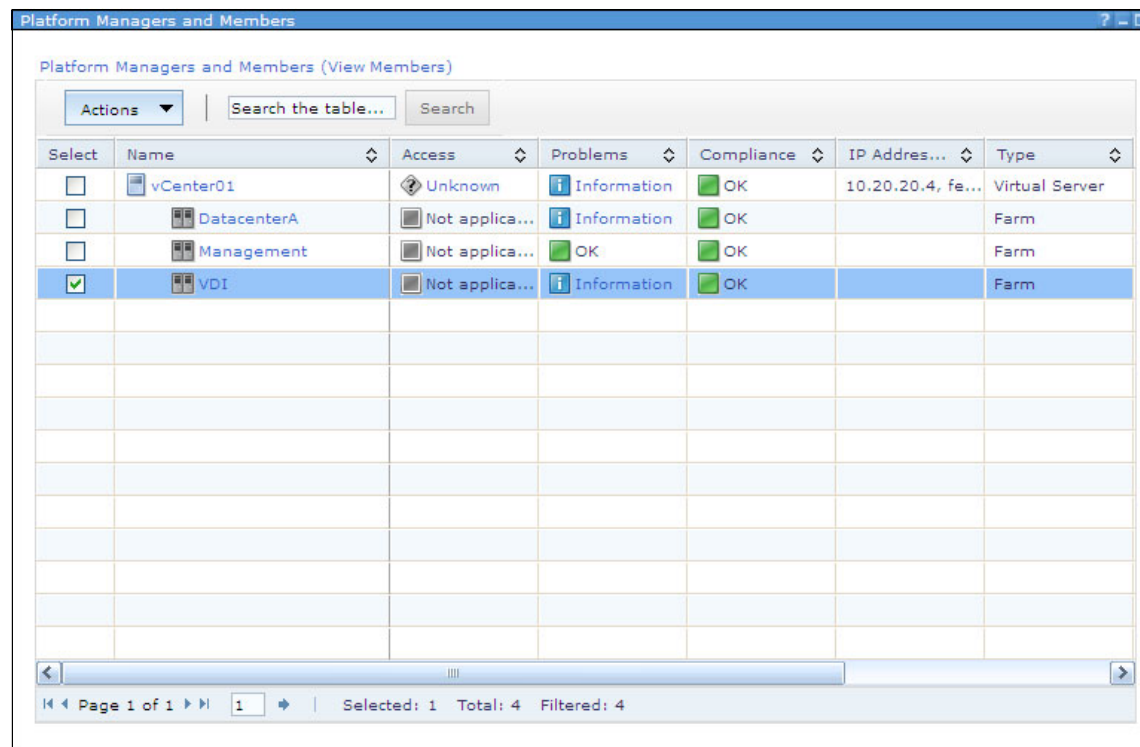


Figure 8-5 Platform Managers and Members window

You can perform the following actions on these objects:

- ▶ Enable or disable HA
- ▶ Enable or disable DRS
- ▶ Add a host to the cluster

To view ESXi hosts and virtual machines, click **Inventory** → **Views** in the navigation area and then click **Virtual Servers and Hosts**, as shown in Figure 8-6.

Select	Name	State	Access	Problems	Compliance	OS Name
<input type="checkbox"/>	Server-7954-24X-SN10778...	Started	OK	Information	OK	
<input type="checkbox"/>	itsoAIX1	Stopped	OK	OK	OK	
<input type="checkbox"/>	itsoVIO6A	Started	OK	OK	OK	
<input type="checkbox"/>	x240_Node_1	Started	OK	OK	OK	ESXi_Node_1
<input type="checkbox"/>	CMP01	Started	OK	OK	OK	
<input type="checkbox"/>	CS01	Started	OK	OK	OK	
<input type="checkbox"/>	DC01	Started	OK	OK	OK	
<input type="checkbox"/>	FS01	Started	OK	OK	OK	
<input type="checkbox"/>	MSSQL01	Started	OK	OK	OK	
<input type="checkbox"/>	vCenter01	Started	Unknown	Information	OK	9.42.171.34
<input type="checkbox"/>	x240_Node_2	Started	OK	OK	OK	ESXi_Node_2
<input type="checkbox"/>	x240_Node_3	Started	OK	OK	OK	ESXi_Node_3
<input type="checkbox"/>	FVM-1-VIP	Started	OK	OK	OK	
<input type="checkbox"/>	FVM-2-VIP	Started	OK	OK	OK	
<input type="checkbox"/>	FVM-3-VIP	Started	OK	OK	OK	

Page 1 of 2 | Selected: 0 Total: 24 Filtered: 24

Figure 8-6 Virtual Servers and Hosts window

You can perform the following actions on these objects:

- ▶ Enter host maintenance mode
- ▶ Remove a host from a cluster
- ▶ Relocate VMs
- ▶ Create VMs
- ▶ Power On/Off hosts or VMs

These actions also are available in vCenter. FSM communicates with vCenter to perform these actions. This integration between FSM and vCenter is useful for task automation, especially if the tasks are related to the hardware and software layers. The next section includes an example of how a response to a hardware alert can be automated with action on vSphere.

# 8.3 Automating tasks

In this section, we describe how to automate tasks that can prevent service outages. The process relies on *event automation plans* to kick off and control the relevant tasks in your systems management environment.

Event automation plans consist of *event filters* and *event actions* and are triggered by *events*. You create event automation plans and apply them to specific systems to trigger a specific action (such as email notification or script execution) when a certain threshold is reached or a specified event occurs. Event automation plans are a powerful feature to automate many manual tasks in your environment.

Flex System Manager can monitor all hardware elements of the Flex System chassis while it integrates with vSphere to perform certain vSphere tasks. This unique role of FSM allows you to combine hardware and software automation.

For example, you can configure automation to proactively protect the two Management cluster ESXi nodes from sudden failure. By using this approach, a host is put into maintenance mode whenever a Predictive Failure Alert event occurs on any hardware element of the specified node. To automate such preventive actions, complete the following steps:

- 1. In the FSM web interface navigation area, expand **Automation** and then click **Event Automation Plans**. Click **Create** for a new automation plan, as shown in Figure 8-7.

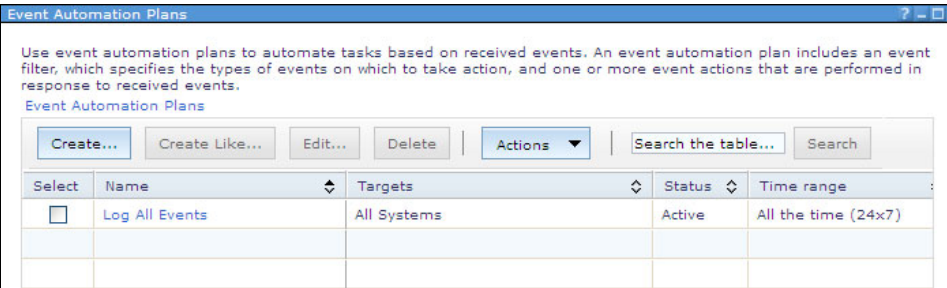


Figure 8-7 Event Automation Plans window

- 2. In the Welcome window, click **Next**.

3. Enter a Name and Description for the automation plan (see Figure 8-8) and then click **Next**.

The screenshot shows a web-based configuration interface for an automation plan. On the left is a vertical navigation pane with the following items: 'Welcome' (checked), 'Name and Description' (highlighted with a blue arrow), 'Targets', 'Events', 'Event Actions', 'Time Range', and 'Summary'. The main content area is titled 'Name and Description' and contains the instruction 'Type a name and a description for this event automation plan.' Below this, there are two input fields. The first is labeled '\*Name:' and contains the text 'PFA Maintenance'. The second is labeled 'Description:' and contains the text 'Place an ESXi host from the Management cluster in maintenance mode if PFA is detected. Entering maintenance mode will automatically evacuate all VMs off the problematic host, as long as the host is part of a DRS cluster.' At the bottom right of the main area are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 8-8 Entering a name and description for a new automation plan

4. In the Targets window, select the systems where the event automation plan is to be applied. For this example, select the two Management cluster ESXi servers (see Figure 8-9) and then click **Add**. Click **Next** to proceed.

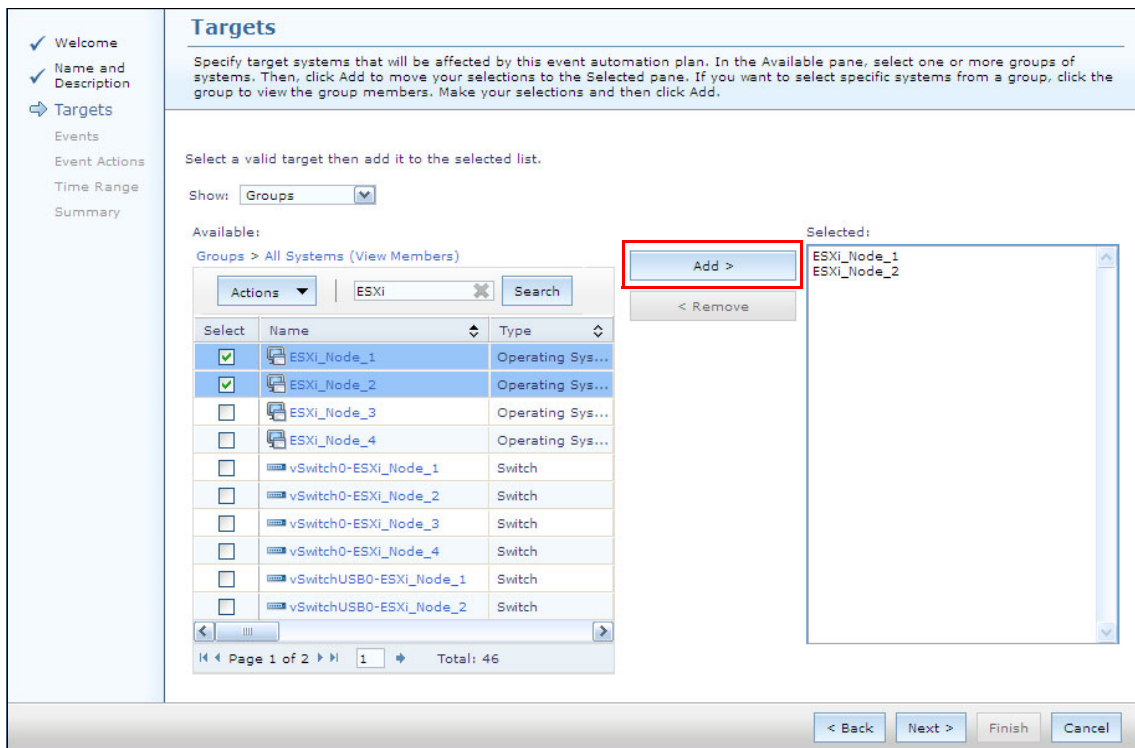


Figure 8-9 Selecting the target systems for a new automation plan



- On the Events menu, select **Advanced Event Filters**, and from the Event Filters list, select **Hardware Predictive Failure Alert events** to process all events that have a Critical severity, as shown in Figure 8-10.

**Events**

Specify one or more events from a list of commonly used events. The selected events will trigger this event automation plan. Or, select Advanced Event Filters in the Events list to use an advanced event filter.

Events:

Advanced Event Filters ▼  
Common  
Advanced Event Filters

to monitor for specific events that are not included in the common event filters or to monitor for only one event. For example, instead of monitoring for all fan event types, you can monitor for only the Fan Predictive Failure Analysis (PFA) event. Also, you can create more sophisticated event filters that are triggered when duplicates of an event are received, when a specific number of instances of an event is received over a range of time, or when a specific event is received but you want to exclude another event.

Event Filters

Create... Create Like... Edit... Delete Actions ▼ Search the table... Search

Select	Name	Description
<input type="radio"/>	All Events	Processes any events that occur on any system, except for Windows-specific an...
<input type="radio"/>	Audit Events	Processes only those events that are generated as a part of auditing.
<input type="radio"/>	Common Agent offline	Processes only those events that are generated by the Common Agent when it ...
<input type="radio"/>	Critical Events	Processes only those events that have a Critical severity
<input type="radio"/>	Disk use	Processes only those events that are generated when the currently available ha...
<input type="radio"/>	Electronic Service and Support Events	Processes only those events that are generated by Electronic Service and Suppo...
<input type="radio"/>	Electronic Service Requests	Processes only those events that are associated with detection and reporting of...
<input type="radio"/>	Environmental sensor events	Processes only those events that are associated with the condition of a system ...
<input type="radio"/>	Fatal Events	Processes only those events that have a Fatal severity
<input checked="" type="radio"/>	Hardware Predictive Failure Alert events	Processes only those events that are generated when a Predictive Failure Analy...
<input type="radio"/>	Informational Events	Processes only those events that have a Informational severity
<input type="radio"/>	Management server security events	Processes only those events that are generated by management server security...
<input type="radio"/>	Memory use	Processes only those events that are generated when the currently available m...
<input type="radio"/>	Minor Events	Processes only those events that have a Minor severity
<input type="radio"/>	Physical hardware security events	Processes only those events that are generated by physical hardware security pr...

< Page 1 of 2 > 1 Selected: 1 Total: 21 Filtered: 21

< Back Next > Finish Cancel

Figure 8-10 Designating the type of events to be processed by a new automation plan

6. For a new event action, click **Create**, as shown in Figure 8-11.

✓ Welcome

✓ Name and Description

✓ Targets

✓ Events

➔ Event Actions

Time Range

Summary

Event Actions

Specify one or more actions that will occur when this event automation plan is triggered.

Event Actions

Create...

Create Like...

Edit...

Delete

Actions ▾

Search the table...

Search

Select	Name	Type	History
<input type="checkbox"/>	Add to the event log	Add to the event log	Not saved

⏪

Page 1 of 1

1

Selected: 0 Total: 1 Filtered: 1

⏩

< Back

Next >

Finish

Cancel

Figure 8-11 Creating an event action

- Go to Page 2 of the actions list and select **Start a task on a system that generated the event**, as shown in Figure 8-12. When you are done, click **OK**.

The screenshot shows a 'Create Action' dialog box with a table of actions. The table has three columns: 'Select', 'Name', and 'Type'. The action 'Start a task on a system that generated the event' is selected, indicated by a radio button and a blue highlight. The pagination controls at the bottom show 'Page 2 of 2' and the number '2' is highlighted in a red box. The status bar at the bottom indicates 'Selected: 1 Total: 19 Filtered: 19'.

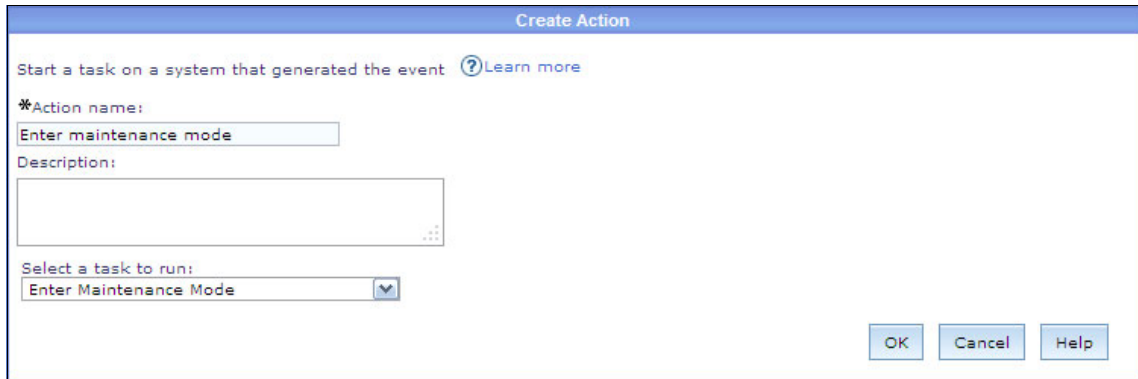
Select	Name	Type
<input type="radio"/>	Modify an event and send it	Advanced
<input type="radio"/>	Send a Tivoli Enterprise Console event	Advanced
<input checked="" type="radio"/>	Start a task on a system that generated the event	Advanced
<input type="radio"/>	Send an SNMP trap reliably to a NetView host	Advanced

Page 2 of 2 | 2 | Selected: 1 Total: 19 Filtered: 19

OK Cancel Help

Figure 8-12 Selecting an action from the actions list

8. Enter an Action name and Description for the event action and then go to the “Select a task to run” menu and choose **Enter Maintenance Mode**, as shown in Figure 8-13. The menu shows the selection of tasks that can be run as actions. After you make your selection, click **OK**.



Create Action

Start a task on a system that generated the event [? Learn more](#)

\*Action name:  
Enter maintenance mode

Description:

Select a task to run:  
Enter Maintenance Mode

OK Cancel Help

Figure 8-13 Choosing action properties

9. Select the event action that you just created, **Enter maintenance mode**, as shown in Figure 8-14, and then click **Next**.

✓ Welcome

✓ Name and Description

✓ Targets

✓ Events

➔ Event Actions

Time Range

Summary

Event Actions

Specify one or more actions that will occur when this event automation plan is triggered.

Event Actions

Create...Create Like...Edit...Delete

Actions ▾

Search the table...Search

Select	Name	Type	History
<input type="checkbox"/>	Add to the event log	Add to the event log	Not saved
<input checked="" type="checkbox"/>	Enter maintenance mode	Start a task on a system that generated the event	Not saved

⏪

Page 1 of 1

1

Selected: 1 Total: 2 Filtered: 2

⏩

< Back

Next >

Finish

Cancel

Figure 8-14 Selecting the new action event

10. In the Time Range window, click **Next**.

11. Review the information in the Summary window (see Figure 8-15) and then click **Finish** to create and apply the new event automation plan.

The screenshot shows the 'Summary' window of the 'Create Event Automation Plan' wizard. On the left is a vertical sidebar with a list of steps: 'Welcome', 'Name and Description', 'Targets', 'Events', 'Event Actions', 'Time Range', and 'Summary'. Each step is preceded by a checkmark, and 'Summary' is highlighted with a blue arrow. The main area of the window is titled 'Summary' and contains the text 'You have specified the following settings for this event automation plan:'. Below this, the following settings are listed: Name: PFA Maintenance; Description: Place an ESXi host from the Management cluster in maintenance mode if PFA is detected. Entering maintenance mode will automatically evacuate all VMs off the problematic host, as long as the host is part of a DRS cluster; Time range: All the time (24x7); Targets: ESXi\_Node\_1, ESXi\_Node\_2; Event filter: Hardware Predictive Failure Alert events; Event actions: Enter maintenance mode. At the bottom of the main area is a checkbox labeled 'Apply this event automation plan when I click Finish.' which is checked. At the bottom right of the window are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Summary

You have specified the following settings for this event automation plan:

Name: PFA Maintenance

Description: Place an ESXi host from the Management cluster in maintenance mode if PFA is detected. Entering maintenance mode will automatically evacuate all VMs off the problematic host, as long as the host is part of a DRS cluster.

Time range: All the time (24x7)

Targets: ESXi\_Node\_1  
ESXi\_Node\_2

Event filter: Hardware Predictive Failure Alert events

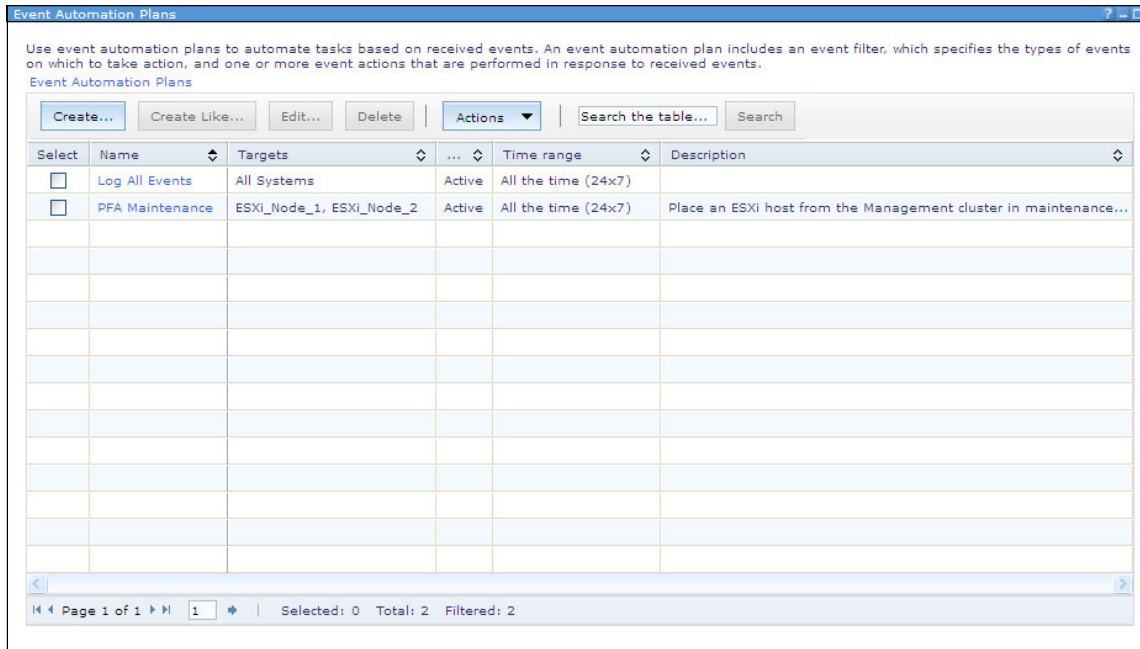
Event actions: Enter maintenance mode

☒ Apply this event automation plan when I click Finish.

< Back Next > Finish Cancel

Figure 8-15 Summary window of the Create Event Automation Plan wizard

When the process is complete, the new event automation plan is displayed in the Event Automation Plans window, as shown in Figure 8-16.



*Figure 8-16 New automation plan that is listed in the Event Automation Plans window*

## 8.4 Introducing IBM FSM Explorer

The IBM FSM Explorer console provides an alternative view of your resources and helps you manage your systems management environment.

FSM Explorer provides a resource-based view of your environment, with intuitive ways to browse through the resources. The following capabilities also are available:

- ▶ You can view basic information about your resources by hovering over their listings; you do not have to click a listing to learn about the resource.
- ▶ You can use standard browser features, such as the Back and Forward buttons, to browse through FSM Explorer pages. You can also bookmark pages that you use frequently, so you can return to them easily.
- ▶ You can work on multiple pages at once by opening the pages in separate browser tabs.

- ▶ You can paste a page URL and send it to a co-worker in an email or instant message. Your co-worker can paste the URL into their browser and view the page (after authenticating to the server).

In addition to learning about resources, you can perform the following tasks in FSM Explorer:

- ▶ Configure local storage, network adapters, boot order, and Integrated Management Module (IMM) and Unified Extensible Firmware Interface (UEFI) settings for one or more compute nodes before you deploy operating-system or virtual images to them.
- ▶ Install operating system images on IBM X-Architecture® compute nodes.
- ▶ Browse resources to view their properties and perform basic management tasks, such as powering on and off, collecting inventory, and working with LEDs.
- ▶ Use the Chassis Map to edit compute node details, view server properties, and manage compute node actions.
- ▶ Work with resource views, such as All Systems, Chassis and Members, Hosts, Virtual Servers, Network, Storage, and Favorites.
- ▶ Visually monitor statuses and events.
- ▶ View event history and active status.
- ▶ View inventory.
- ▶ Visually monitor job status.

For other tasks, you can start Flex System Manager in a separate browser window or tab and then return to the FSM Explorer tab after you finish the tasks. As more tasks are made available in FSM Explorer, you must start Flex System Manager less often.



## 8.5 Monitoring and logging in Flex System

In this section, we describe some of the features that FSM Explorer offers to monitor the Flex System chassis and all managed systems. What we describe here are not really steps to be followed, in the traditional sense. Instead, consider the following steps as one way to explore the available monitoring and logging capabilities:

1. From the FSM web interface home page, click **Launch IBM FSM Explorer**, as shown in Figure 8-17.

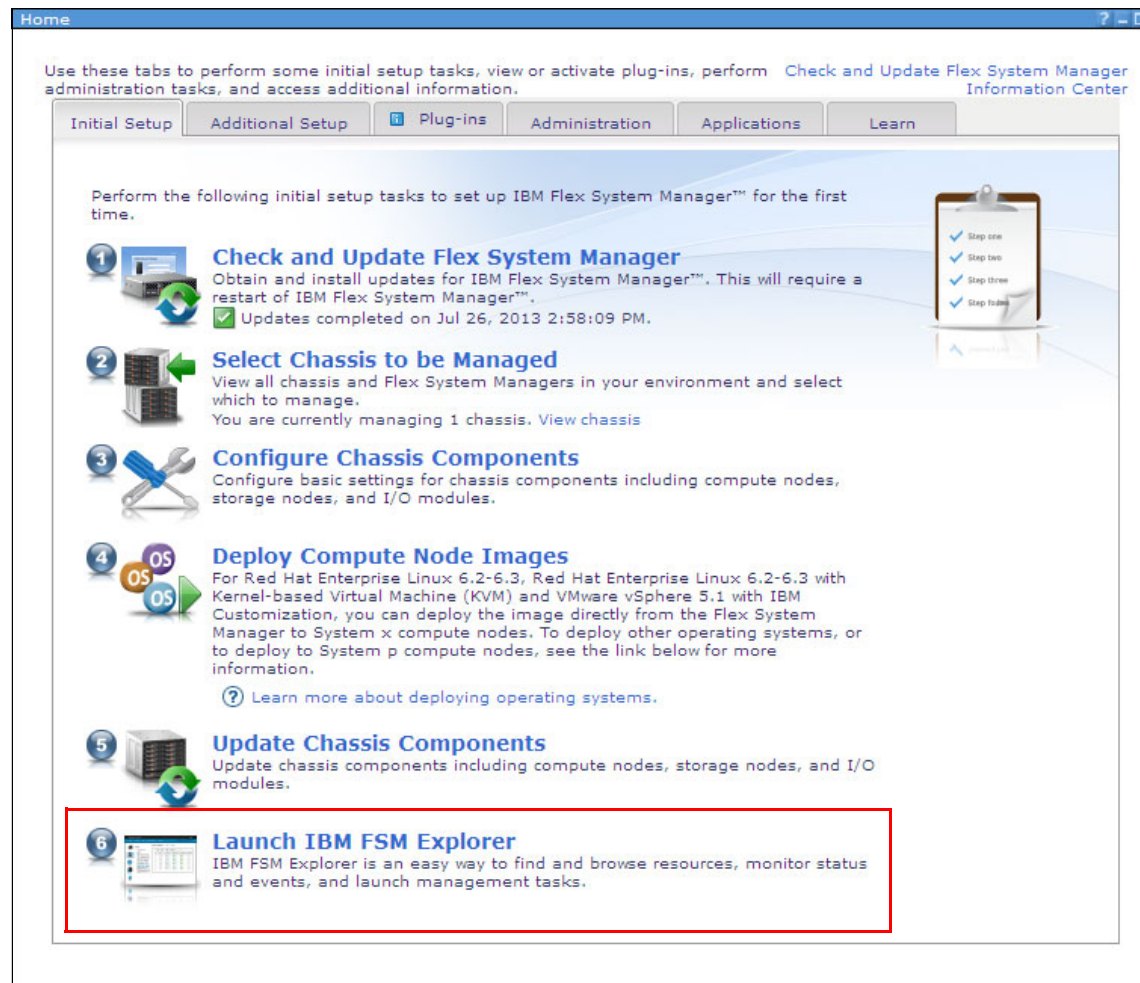


Figure 8-17 Launching FSM Explorer from the Flex System Manager home page

2. In the Chassis column of the FSM Explorer Dashboard, click the wanted chassis name, as shown in Figure 8-18. The graphical Chassis Map opens.

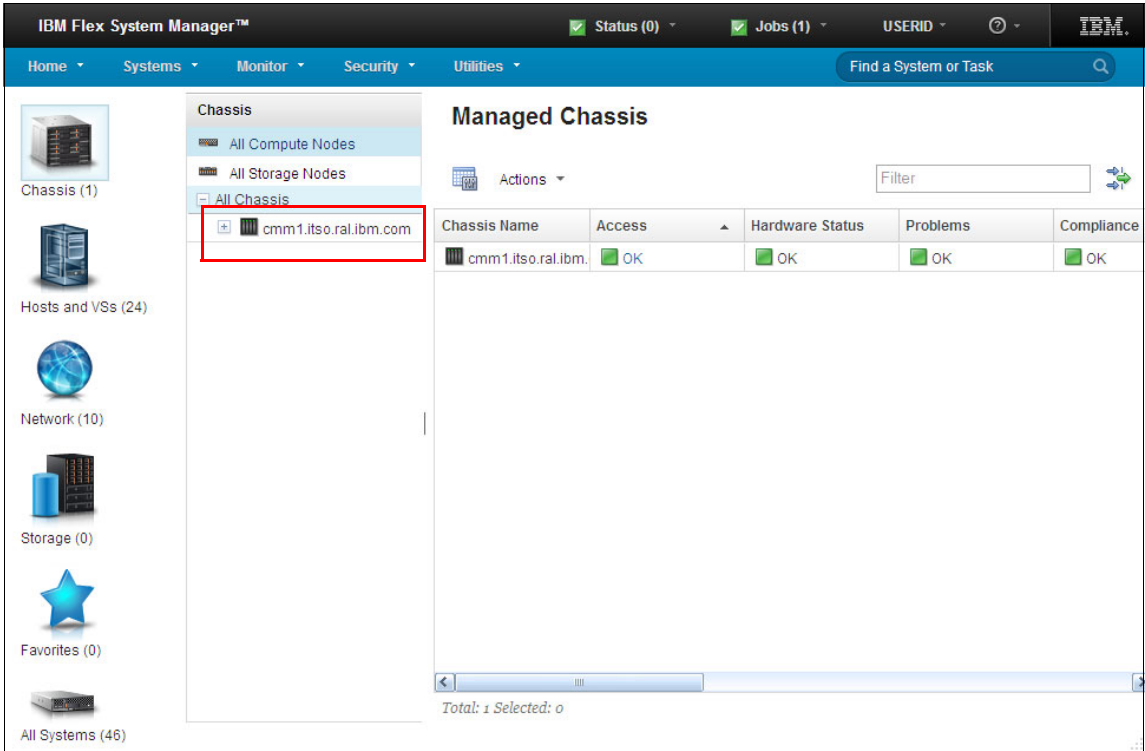


Figure 8-18 Selecting a chassis from the FSM Explorer Dashboard

The graphical Chassis Map (as show in Figure 8-19) is a visual representation of the front and back of the chassis and its components. The map shows you where your hardware components are and is a central point of management from which you can get hardware configuration and status information. You can also take actions on nodes, such as working with server-related resources, showing and installing updates, submitting service requests, and starting the remote access tools.



Figure 8-19 FSM Chassis Map

When a chassis resource shows an *active status* that requires administrative attention, the resource is highlighted on the Chassis Map, as shown in Figure 8-20.

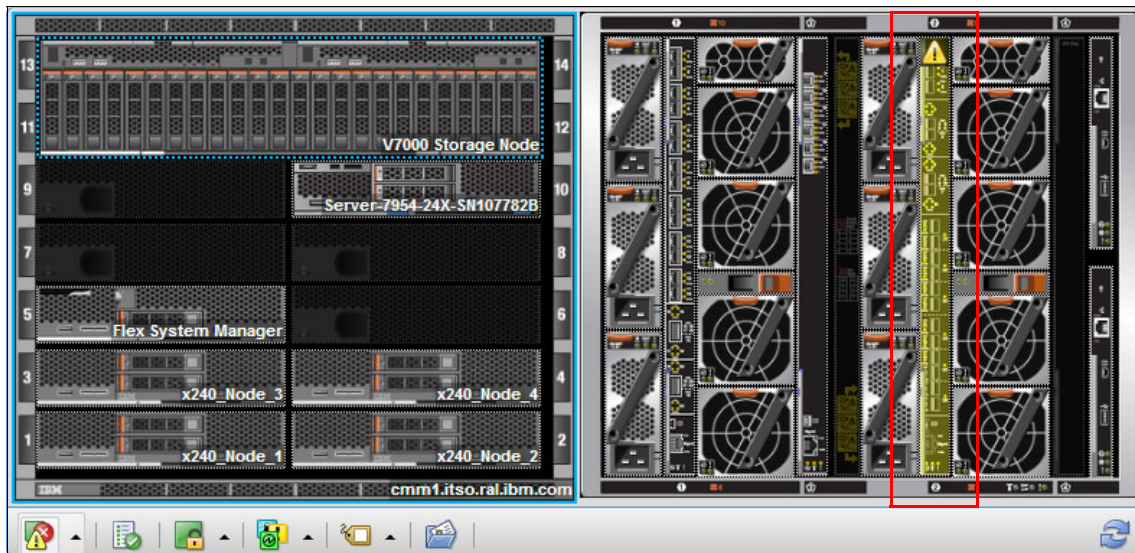


Figure 8-20 FSM Chassis Map with active status warning

**Note:** Active status contains all events that are designated as problems by the Flex System Manager management node, including problems on the management node.

The problems that are displayed in the Active Status view do not always correlate directly to specific events that are listed in the event log. Flex System Manager determines which problems to display that are based on the kinds of events that are reported.

A fly-over window provides an instant view of the chassis resource details and active status, as shown in Figure 8-21. The administrator can easily spot problems on any resource in the chassis. To see a list of all status messages for the particular resource, click a resource (in this case, the **10Gb Converged Switch**) and then click **View All Status**.

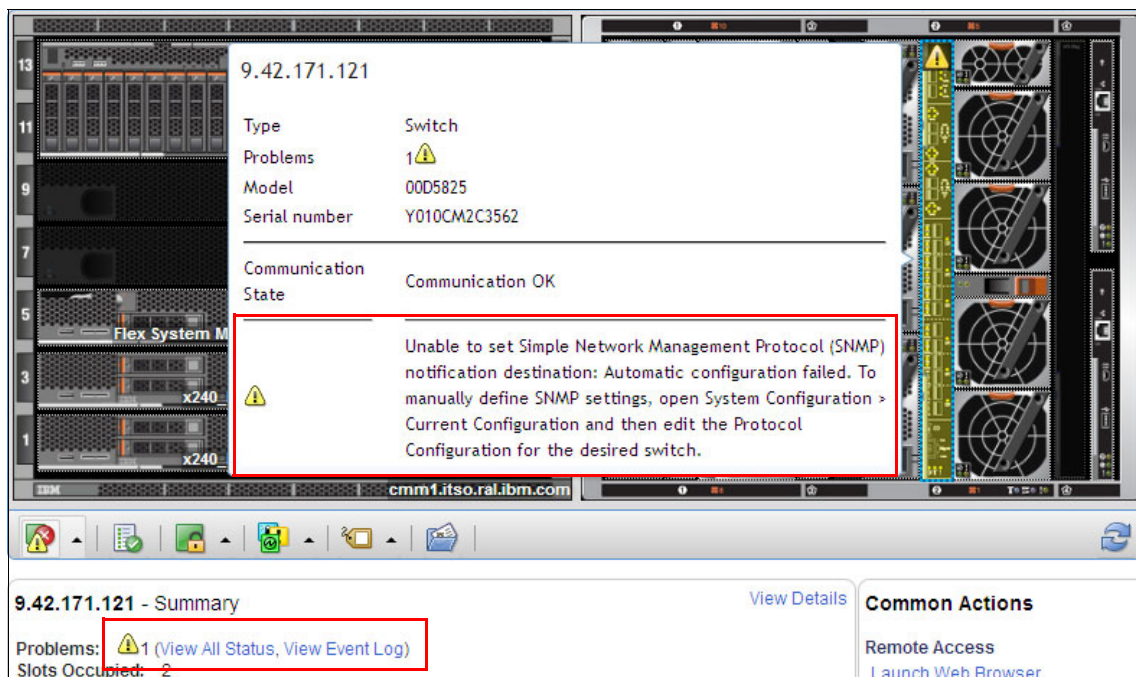


Figure 8-21 Chassis Map warning fly-over window



3. To open the Active Status view for all systems, start at the FSM Explorer menu and click **Monitor** → **Active Status**, as shown in Figure 8-22.

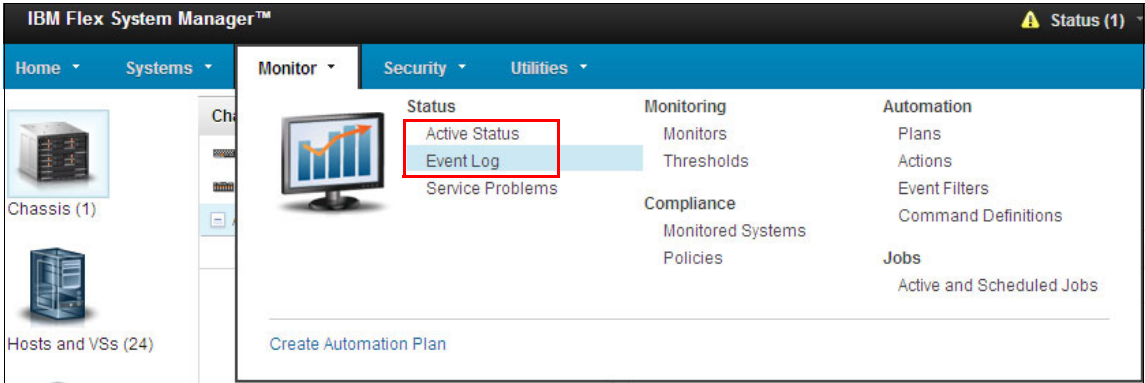


Figure 8-22 Using the FSM Explorer menu to access the Active Status view

The Active Status view (as shown in Figure 8-23) covers all chassis components and managed endpoints. The Information statuses are not shown on the Chassis Map.

**All Systems: Active Status**

Active Alerts Ignored Alerts Filter by: All Alerts

Actions

Name	Severity	System	Component	Category	Time
A virtual server has been powered on.	Information	LCVM-5-STDUSR	Server/VSM	Hardware Status	Aug 1
A virtual server has been powered on.	Information	LCVM-4-STDUSR	Server/VSM	Hardware Status	Aug 1
A virtual server has been powered on.	Information	LCVM-3-STDUSR	Server/VSM	Hardware Status	Aug 1
A virtual server has been powered on.	Information	vCenter01	Server/VSM	Hardware Status	Aug 1
A virtual server was deleted.	Information	vCenter01	Server/VSM	Hardware Status	Aug 1
A virtual server has been powered off.	Information	vCenter01	Server/VSM	Hardware Status	Aug 1
Error Log analysis has detected multiple link errors	Information	Server-7954-24X-SN10	Server-7954-24X-SN10	Service Status	Aug 8
LED Status: Informational	Information	cmm1.itso.ral.ibm.com	cmm1.itso.ral.ibm.com	LED Status	Aug 7
Unable to set Simple Network Management Protocol	Warning	9.42.171.121	Switch	Hardware Status	Jul 26

Figure 8-23 Active Status view

- To open the Event Log view for all systems, click **Monitor** → **Event Log** on the FSM Explorer menu (as shown in Figure 8-22 on page 370). The list that appears is shown in Figure 8-24.

**All Systems: Event Log**

Filter by: All Events Viewing a maximum of 500 events from the last 24 Hours. [Event Log Preferences](#)

Last Updated: 4:18:33 PM Eastern Standard Time, Tuesday Aug 20, 2013

Actions Filter

Event Text	Source	Severity	Category	Date and Time
A virtual server was deleted.	21171	Information	Alert	Aug 20, 2013 4:18:33 PM
A virtual server has been powered off.	vCenter01	Information	Alert	Aug 20, 2013 4:18:33 PM
A virtual server has been powered off.	21171	Information	Alert	Aug 20, 2013 4:18:33 PM
System 10.3.0.1 is offline	10.3.0.1	Information	Alert	Aug 20, 2013 4:18:33 PM
Management Console surveillance detected no networking available	Flex System Manager	Information	Alert	Aug 20, 2013 4:18:33 PM
System FVM is online	FVM	Information	Resolution	Aug 20, 2013 4:18:33 PM
System LCVm-4-STDUSR is online	LCVM-4-STDUSR	Information	Resolution	Aug 20, 2013 4:18:33 PM
System LCVm-5-STDUSR is online	LCVM-5-STDUSR	Information	Resolution	Aug 20, 2013 4:18:33 PM
System FVM-1-VIP is online	FVM-1-VIP	Information	Resolution	Aug 20, 2013 4:18:33 PM
System FVM-3-VIP is online	FVM-3-VIP	Information	Resolution	Aug 20, 2013 4:18:33 PM
System LCVm is online	LCVM	Information	Resolution	Aug 20, 2013 4:18:33 PM
System FVM-2-VIP is online	FVM-2-VIP	Information	Resolution	Aug 20, 2013 4:18:33 PM
System LCVm-3-STDUSR is online	LCVM-3-STDUSR	Information	Resolution	Aug 20, 2013 4:18:33 PM
System replica-f6f465c4-c78b-4844-ae2b-5b8cc609dd2a is online	replica-f6f465c4-c78b-4844-ae2b-5b8cc609dd2a	Information	Resolution	Aug 20, 2013 4:18:33 PM

← 129 Selected: 0 →

Figure 8-24 Event Log window

**Note:** An event is an occurrence of significance to a task or resource. Examples of events include the completion of an operation, the failure of a hardware component, or exceeding a processor threshold. The Event Log displays all events that FSM receives from any resource whose events the user has the authority to view.

5. To filter for events that you are most interested in, you select a specific event category. For example, choose **Warning Events** in the “Filter by” drop-down menu and then enter logon in the text filter field (as shown in Figure 8-25) to see a list of all logon-related warning events.

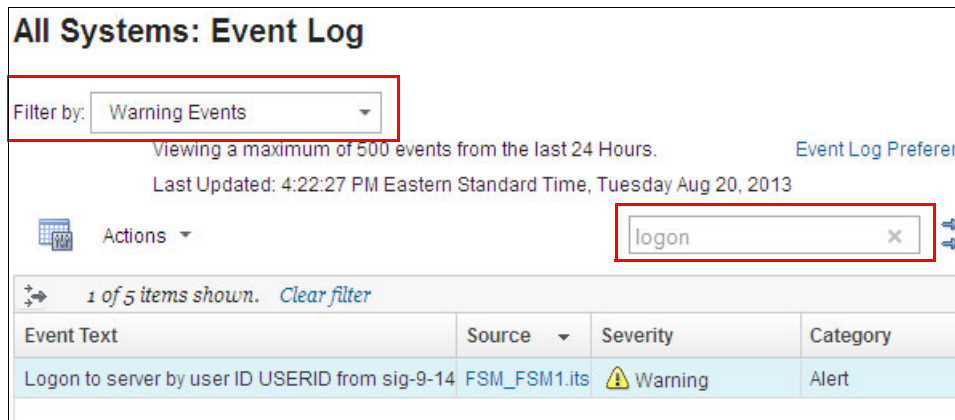


Figure 8-25 Event Log filtered window

For more information about managing Flex System, see *Implementing Systems Management of IBM PureFlex System*, SG24-8060, which is available at this website:

<http://www.redbooks.ibm.com/abstracts/sg248060.html>





# Operating VMware Horizon View infrastructure

In this chapter, we describe the process that is used to operate the VMware Horizon View infrastructure to deliver virtual desktops to users.

This chapter includes the following topics:

- ▶ Preparing the base Microsoft Windows 7 operating system x64 image to deploy
- ▶ Installing the VMware Horizon View Agent
- ▶ Installing the VMware Horizon View Agent
- ▶ Configuring active directory policies
- ▶ VMware View Manager and desktop pools
- ▶ Operating View Composer

## 9.1 Preparing the base Microsoft Windows 7 operating system x64 image to deploy

You can distribute an operating system for multiple users or multiple users group by using several methods. The method that you choose is inherent on how VMware virtual desktop infrastructure (VDI) treats virtual desktops and the purposes for that particular virtual desktop.

In this chapter, we describe the following methods to deliver a virtual desktop to a specific set of users:

- ▶ Automatically provision a full virtual machine by using a specific desktop pool
- ▶ Automatically provision a linked clone virtual machine by using a specific desktop pool

Full virtual machine desktop distribution is based on a pre-configured virtual machine (VM) template. Each time the provision server must provision a new full virtual machine, it deploys a new VM from that template to a fully working VM. This method offers a full persistent virtual desktop but uses more disk space because of the monolithic source virtual disk.

Linked clone virtual machine is based on the snapshot of a running VM, which is also called *parent VM*. This method shares the parent's operating system disk for all distributed VMs, which results in less disk space being used.

To successfully automate virtual desktops provisioning, a custom specification file is needed.

### 9.1.1 Creating a customization specification file

You create customization specification files from vCenter. A customization specification includes the required settings to join unattended VMs to the domain, assign the correct Windows license product key, set network settings, and so on.

Complete the following steps to create a customization specification file:

1. Click **View** → **Management** → **Customization Specifications Manager**, as shown in Figure 9-1.

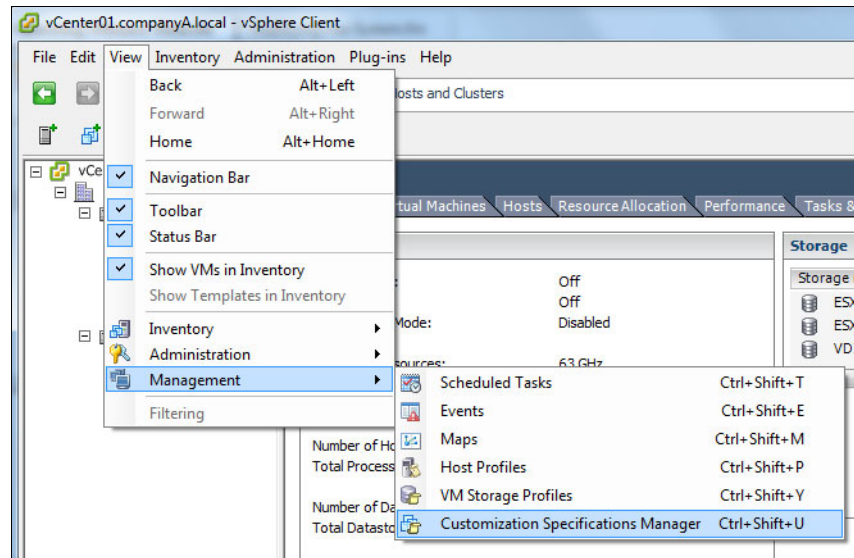


Figure 9-1 Starting the Customization Specifications Manager

2. Click **New** to start the customization wizard, as shown on Figure 9-2.

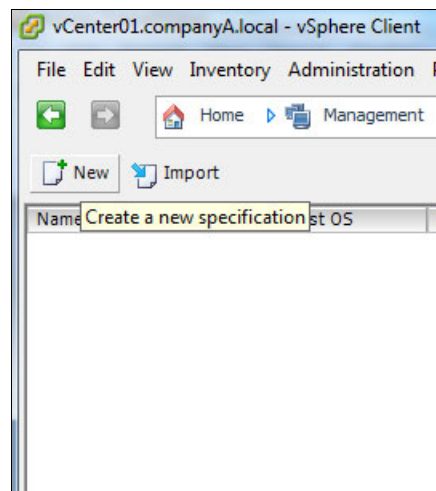


Figure 9-2 Create a customization wizard

3. Complete the New Customization Specification Properties panel in the wizard, as shown on Figure 9-3. Click **Next**.

The screenshot shows the 'vSphere Client Windows Guest Customization' window. The title bar includes the VMware logo and the text 'vSphere Client Windows Guest Customization'. The main content area is titled 'New Customization Specification' with a subtitle 'Enter a name for the new customization specification and select the OS of the target.'.

On the left, there is a 'Properties' sidebar with a list of options: Registration Information, Computer Name, Windows License, Administrator Password, Time Zone, Run Once, Network, Workgroup or Domain, Operating System Options, and Ready to Complete. 'Registration Information' is selected.

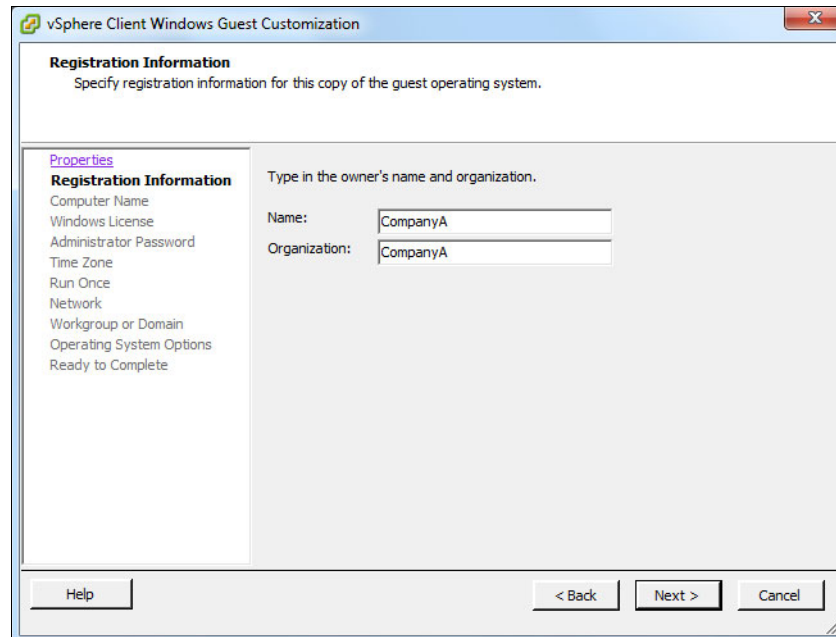
The main area contains two sections:

- Target Virtual Machine OS:** A dropdown menu showing 'Windows'. Below it is a checkbox labeled 'Use Custom Sysprep Answer File' which is unchecked.
- Customization Specification Information:** A section with a 'Name:' field containing 'Windows7\_Domain\_and\_network\_specs' and a 'Description:' text area containing 'Customized Wizard for automatic AD Domain and network specification assignment'.

At the bottom, there are three buttons: 'Help', '< Back', and 'Next >', followed by a 'Cancel' button.

Figure 9-3 New Customization Specification Properties

4. Complete the Registration Information panel, as shown in Figure 9-4. Click **Next**.



The image shows a screenshot of the 'vSphere Client Windows Guest Customization' window. The title bar reads 'vSphere Client Windows Guest Customization'. The main content area is titled 'Registration Information' with the subtitle 'Specify registration information for this copy of the guest operating system.' On the left, there is a 'Properties' sidebar with a list of options: 'Registration Information' (selected), 'Computer Name', 'Windows License', 'Administrator Password', 'Time Zone', 'Run Once', 'Network', 'Workgroup or Domain', 'Operating System Options', and 'Ready to Complete'. The main area on the right is titled 'Type in the owner's name and organization.' and contains two text input fields: 'Name:' with the value 'CompanyA' and 'Organization:' with the value 'CompanyA'. At the bottom of the window, there are three buttons: 'Help', '< Back', and 'Next >', followed by a 'Cancel' button.

Figure 9-4 Registration Information

5. In the Computer Name panel, select **Use the virtual machine name**, as shown in Figure 9-5. Click **Next**.

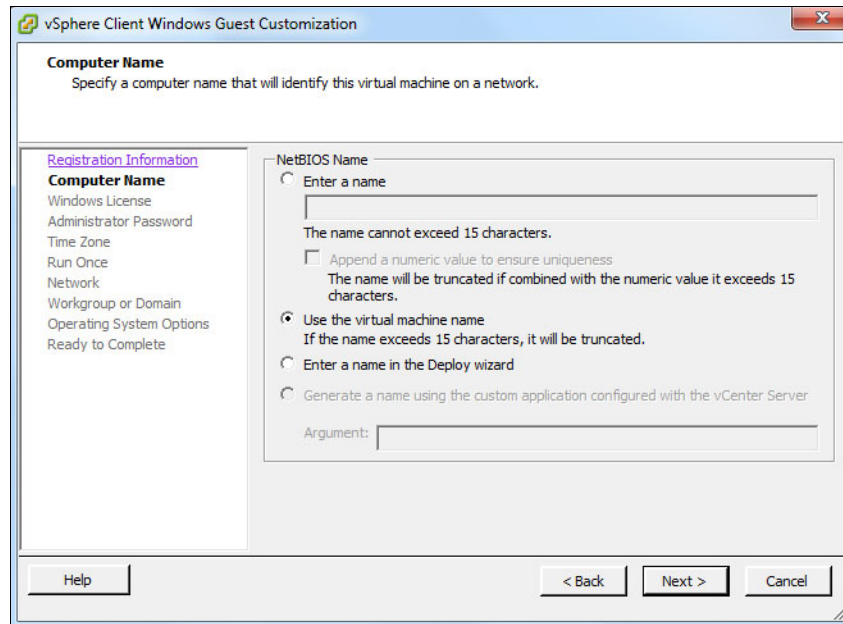
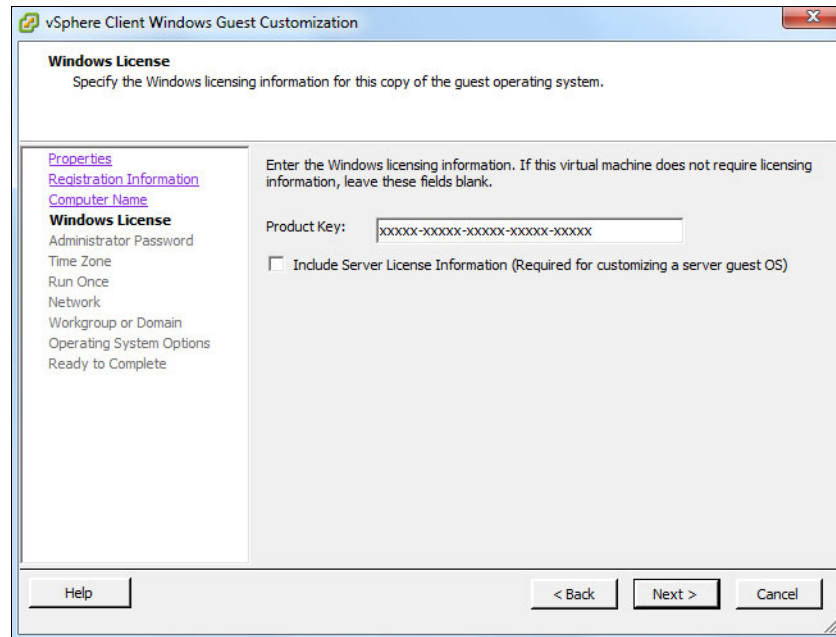


Figure 9-5 Select the Computer Name

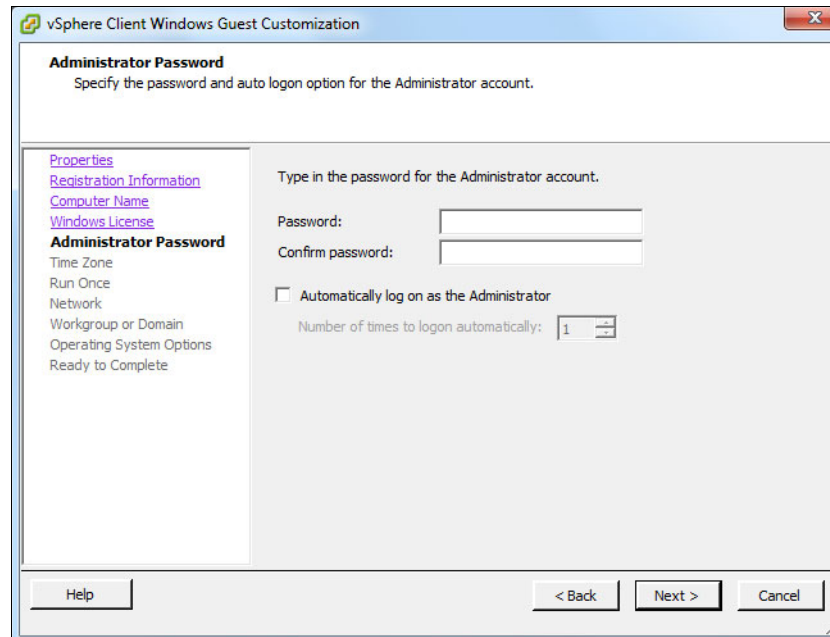
6. Insert the Windows License Product Key, as shown in Figure 9-6. Click **Next**.



The image shows a screenshot of the 'vSphere Client Windows Guest Customization' window, specifically the 'Windows License' tab. The window title bar reads 'vSphere Client Windows Guest Customization'. The main heading is 'Windows License' with a subtitle 'Specify the Windows licensing information for this copy of the guest operating system.' On the left, there is a navigation pane with links: 'Properties', 'Registration Information', 'Computer Name', 'Windows License' (which is selected), 'Administrator Password', 'Time Zone', 'Run Once', 'Network', 'Workgroup or Domain', 'Operating System Options', and 'Ready to Complete'. The main area contains the instruction 'Enter the Windows licensing information. If this virtual machine does not require licensing information, leave these fields blank.' Below this, there is a 'Product Key:' label followed by a text box containing the placeholder 'xxxxxx-xxxxxx-xxxxxx-xxxxxx-xxxxxx'. A checkbox labeled 'Include Server License Information (Required for customizing a server guest OS)' is present and unchecked. At the bottom, there are three buttons: 'Help', '< Back', and 'Next >', along with a 'Cancel' button.

Figure 9-6 Windows License Product Key information

7. Set a default local administrator account password, as shown in Figure 9-7, and click **Next**.



The image shows a screenshot of the 'vSphere Client Windows Guest Customization' window, specifically the 'Administrator Password' step. The window has a title bar with the VMware logo and the text 'vSphere Client Windows Guest Customization'. Below the title bar, the main heading is 'Administrator Password' with a subtitle 'Specify the password and auto login option for the Administrator account.' On the left side, there is a navigation pane with links: 'Properties', 'Registration Information', 'Computer Name', 'Windows License', 'Administrator Password' (which is highlighted), 'Time Zone', 'Run Once', 'Network', 'Workgroup or Domain', 'Operating System Options', and 'Ready to Complete'. The main area on the right contains the following fields and options: 'Type in the password for the Administrator account.' followed by 'Password:' and 'Confirm password:' labels, each with a text input field. Below these is a checkbox labeled 'Automatically log on as the Administrator'. If checked, there would be a 'Number of times to logon automatically:' field with a spinner box currently showing '1'. At the bottom of the window, there are three buttons: 'Help', '< Back', and 'Next >', and a 'Cancel' button on the far right.

Figure 9-7 Set local administrator password



8. Select the current time zone, as shown in Figure 9-8, and click **Next**.

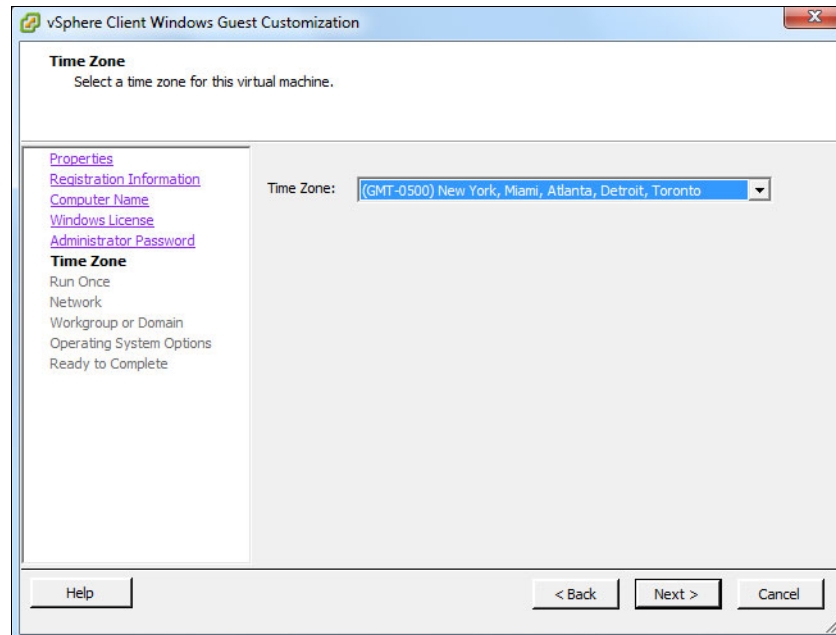


Figure 9-8 Time zone settings

9. In the Run Once panel, click **Next**, as shown on Figure 9-9.

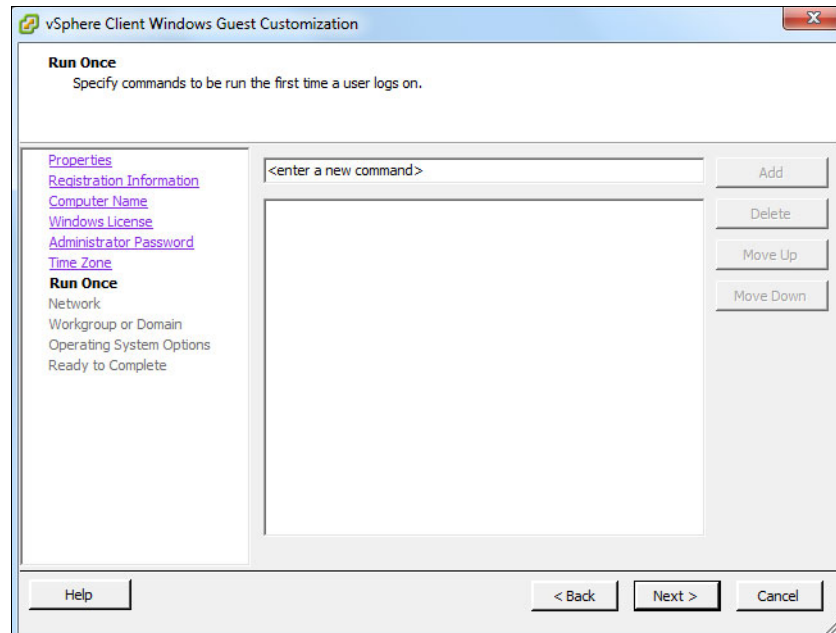


Figure 9-9 Run Once panel

10. Select **Typical settings**, as shown in Figure 9-10. Click **Next**.

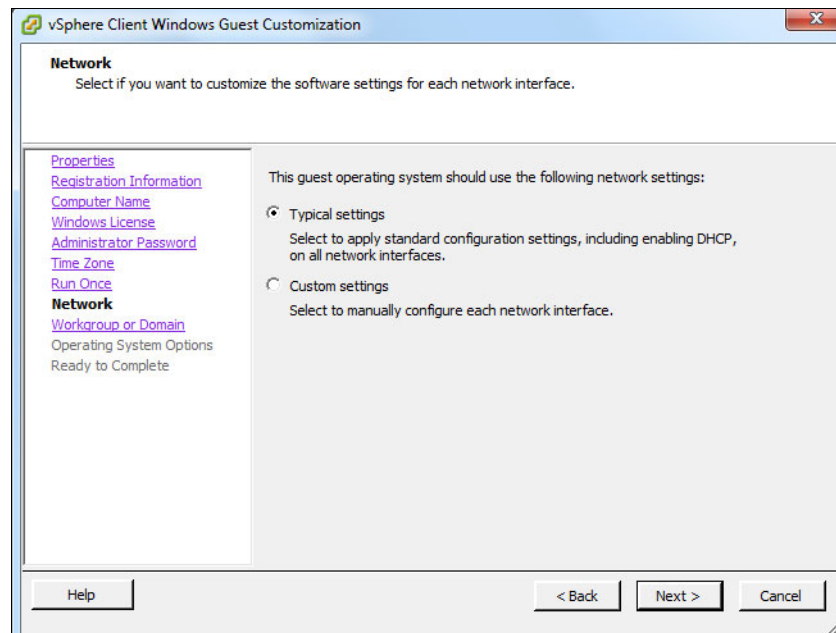
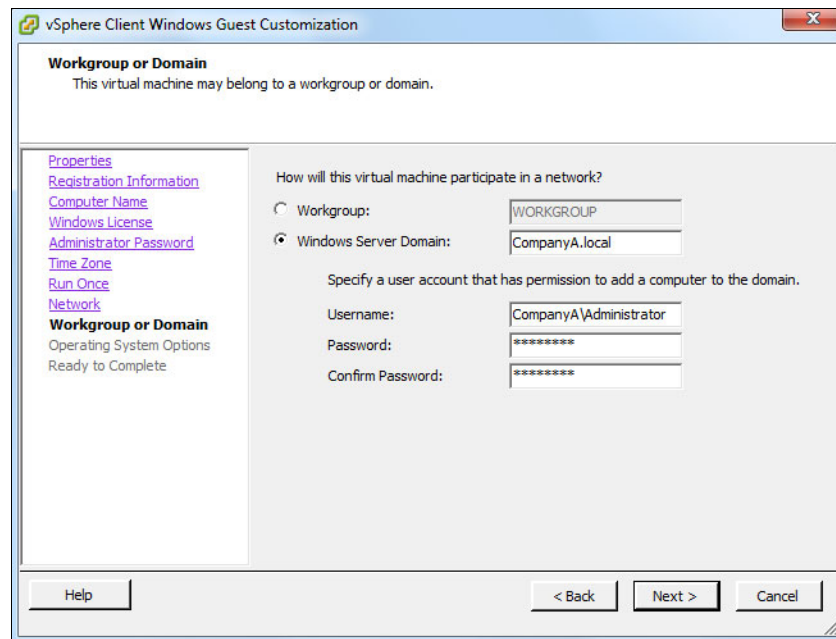


Figure 9-10 Network settings

11. Complete the information in the Workgroup or Domain panel as shown in Figure 9-11, and click **Next**.



The image shows a screenshot of the 'vSphere Client Windows Guest Customization' window, specifically the 'Workgroup or Domain' panel. The window has a title bar with the VMware logo and the text 'vSphere Client Windows Guest Customization'. Below the title bar, the panel is titled 'Workgroup or Domain' with a subtitle 'This virtual machine may belong to a workgroup or domain.' On the left side, there is a navigation pane with links: 'Properties', 'Registration Information', 'Computer Name', 'Windows License', 'Administrator Password', 'Time Zone', 'Run Once', 'Network', 'Workgroup or Domain' (which is selected), 'Operating System Options', and 'Ready to Complete'. The main area of the panel is titled 'How will this virtual machine participate in a network?'. It contains two radio buttons: 'Workgroup:' (unselected) and 'Windows Server Domain:' (selected). The 'Workgroup:' option has a text box containing 'WORKGROUP'. The 'Windows Server Domain:' option has a text box containing 'CompanyA.local'. Below these, there is a text label 'Specify a user account that has permission to add a computer to the domain.' followed by three text boxes: 'Username:' containing 'CompanyA\Administrator', 'Password:' containing '\*\*\*\*\*', and 'Confirm Password:' containing '\*\*\*\*\*'. At the bottom of the window, there are three buttons: 'Help', '< Back', and 'Next >', and a 'Cancel' button on the far right.

Figure 9-11 Domain login information

12. In the Operating System Options panel, accept the default setting by clicking **Next** to generate a new security ID window, as shown in Figure 9-12.

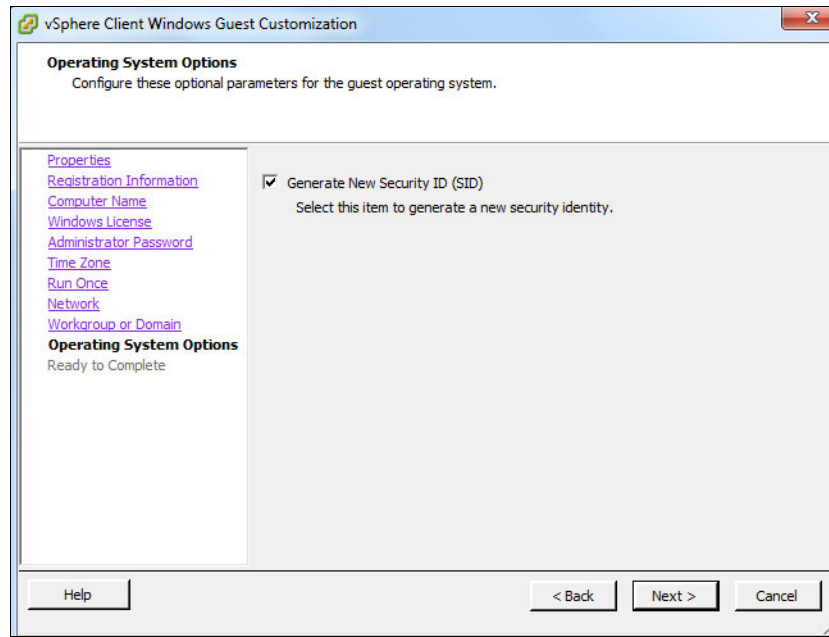


Figure 9-12 Generate a new SID

13. In the Summary window that opens, click **Finish**, as shown in Figure 9-13.

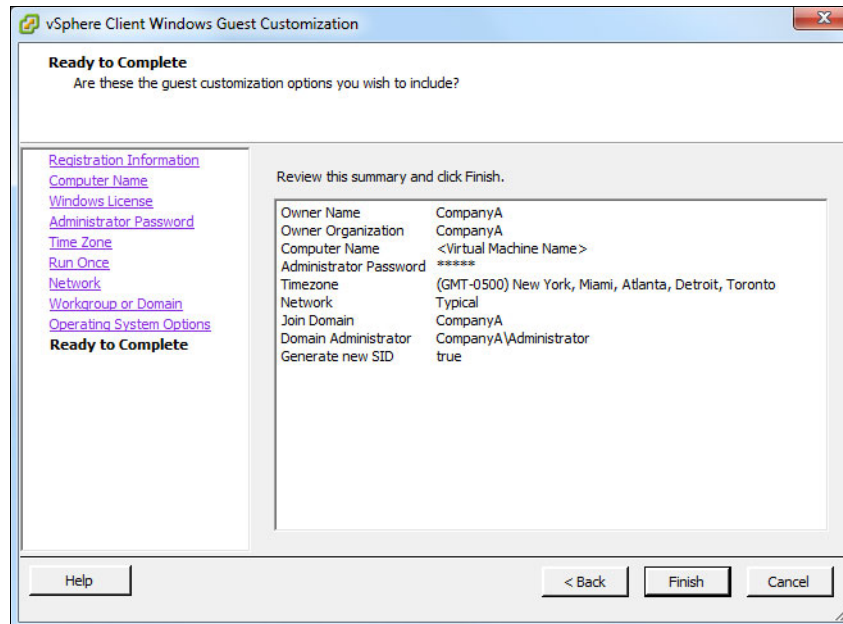


Figure 9-13 Summary window

You can now use the custom specification file for VM deployment.

### 9.1.2 Creating vCenter folders for a VDI

VMware Composer and VMware Administrator use vCenter folders to organize virtual desktops that are generated automatically from a pool. Because desktop pools are related to the type of users, you must create the following distinct vCenter folders:

- ▶ A folder for standard users
- ▶ A folder for VIP users

Complete the following steps to create these folders:

1. From a vSphere client that is connected to vCenter, change the view to turn on VM and Templates, by selecting **Inventory** → **VMs and Templates**, as shown in Figure 9-14.

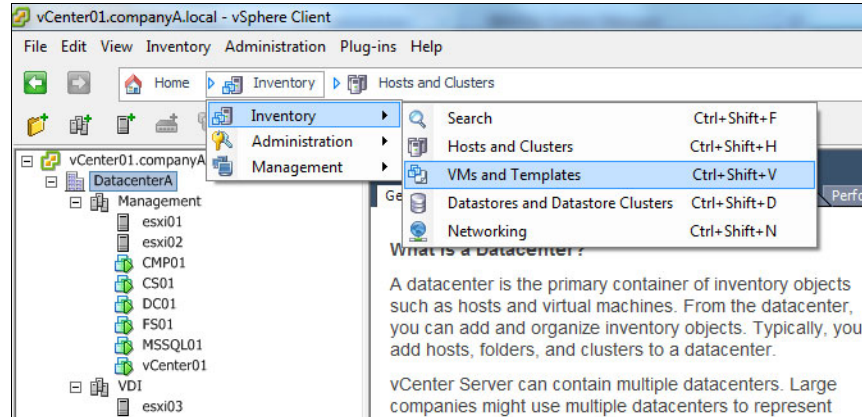


Figure 9-14 Changing the vCenter view

2. Right-click **DatacenterA** and select **New Folder**, as shown in Figure 9-15.




Figure 9-15 Select New Folder option

3. Name the folder **VIP Users**. Repeat step 2 to create another folder that is named **Standard Users**. The final result is shown in Figure 9-16.




Figure 9-16 Final folder structure

### 9.1.3 VMware View Administrator check

VMware View Administrator centralizes the management for the virtual desktops. To use local VMware ESXi disks for View Storage Accelerator as described in 5.6.2, “VDI Cluster component model” on page 132, you first must ensure that the related feature is enabled on vCenter.

Complete the following steps to verify that VMware View Administrator is configured correctly:

1. Open a web browser and enter the following URL:

[https://view\\_connection\\_server\\_IP/admin](https://view_connection_server_IP/admin)

Figure 9-17 shows the main VMware Horizon View Administrator page.

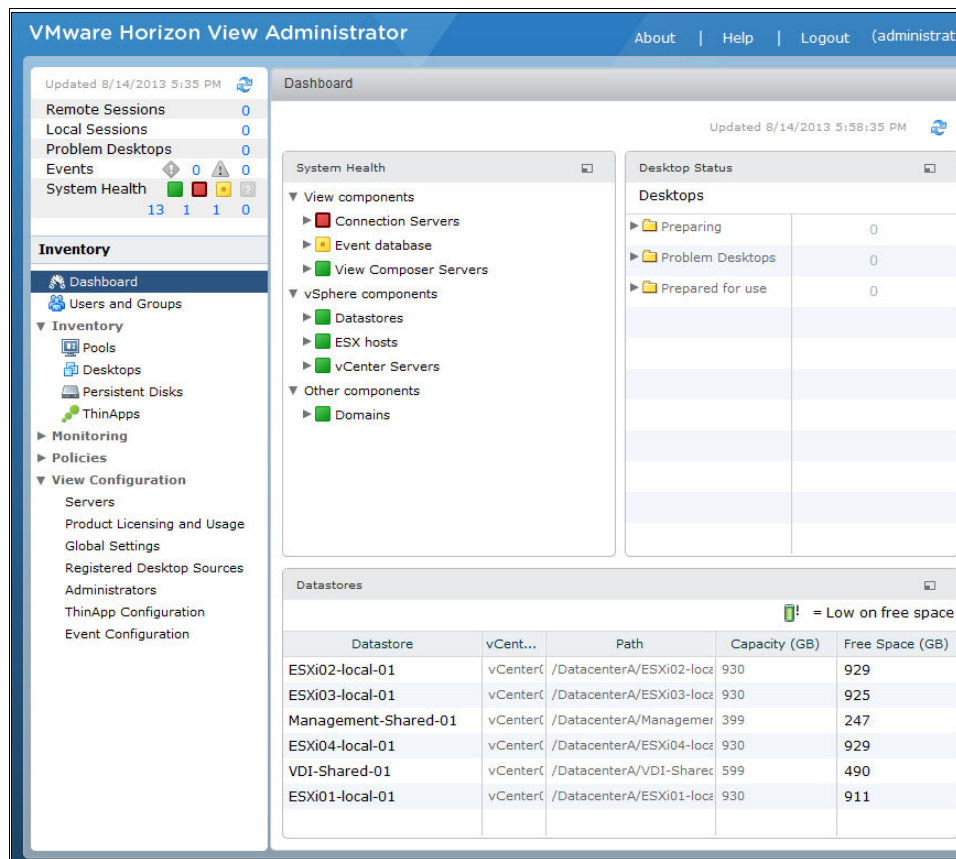


Figure 9-17 Main VMware Horizon View Administrator page



- From this page, expand **View Configuration** in the left pane, and select **Servers**. Then, go to the vCenter Servers tab and click **Edit**, as shown in Figure 9-18.

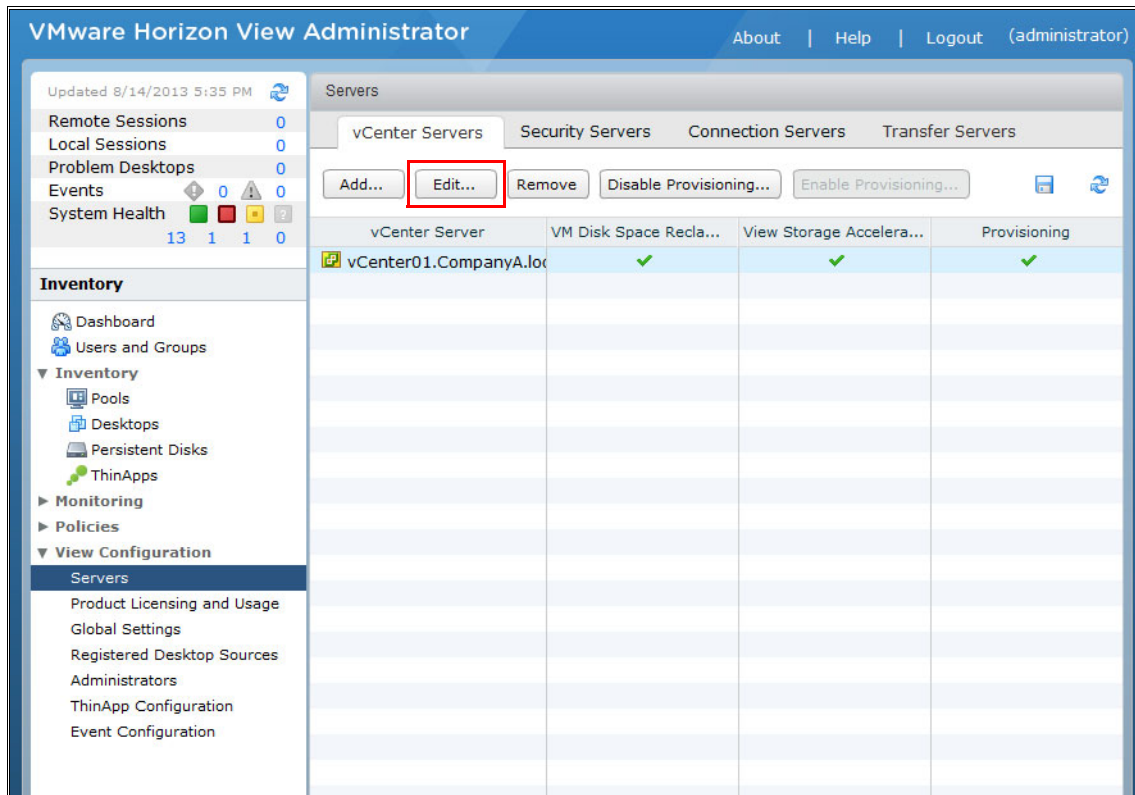


Figure 9-18 VMware View Administrator vCenter Servers tab

3. On the Edit vCenter Server page, click the **Storage** tab. Verify that the “Reclaim VM disk space” and “Enable View Storage Accelerator” options are selected, as shown in Figure 9-19.

The screenshot shows the 'Edit vCenter Server' dialog box with the 'Storage' tab selected. The 'Storage Settings' section has two checked options: 'Reclaim VM disk space' and 'Enable View Storage Accelerator'. Below these, the 'Default host cache size' is set to 1024 MB, with a note that the cache must be between 100 MB and 2048 MB. The 'Hosts' section includes a 'Show all hosts' checkbox and an 'Edit cache size...' button. A table lists four hosts, all with a 'Default' cache size.

Host	Cache Size
/DatacenterA/host/Management/esxi01	Default
/DatacenterA/host/Management/esxi02	Default
/DatacenterA/host/VDI/esxi03	Default
/DatacenterA/host/VDI/esxi04	Default

Figure 9-19 vCenter settings page

4. Click **OK** to finish.

### 9.1.4 Provisioning a full VM image

To successfully provision a full VM, a base operating system image must be installed and then the resulting VM must be transformed to a VM template. Before that template is assigned to View Manager, you must make some customizations in the base operating system to ensure that the resulting VM is faster in deployment.

**Operating system note:** The tasks that are described in this book use Microsoft Windows 7 operating system x64 Professional. Examples assume that the VM already is sized and installed.

Complete the following steps prepare a VM to be ready for full VM provisioning:

1. Install or Upgrade VMware Tools on the VM.
2. Install all Windows operating system updates and service packs.
3. Install all needed applications.
4. Set the control panel and power options to turn off display to Never, as shown in Figure 9-20. Click **Save Changes**.

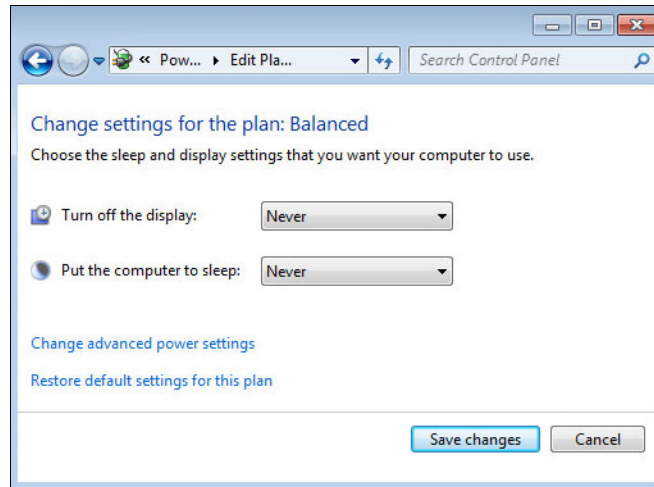


Figure 9-20 Power options

- From the control panel, click **Advanced system settings**. In the System Properties window, go to the Advanced tab and click **Settings** in the Performance section. On the Visual Effects tab, modify the VM appearance settings for “Adjust for best performance” as shown in Figure 9-21. You must click **OK** twice to return to the control panel.

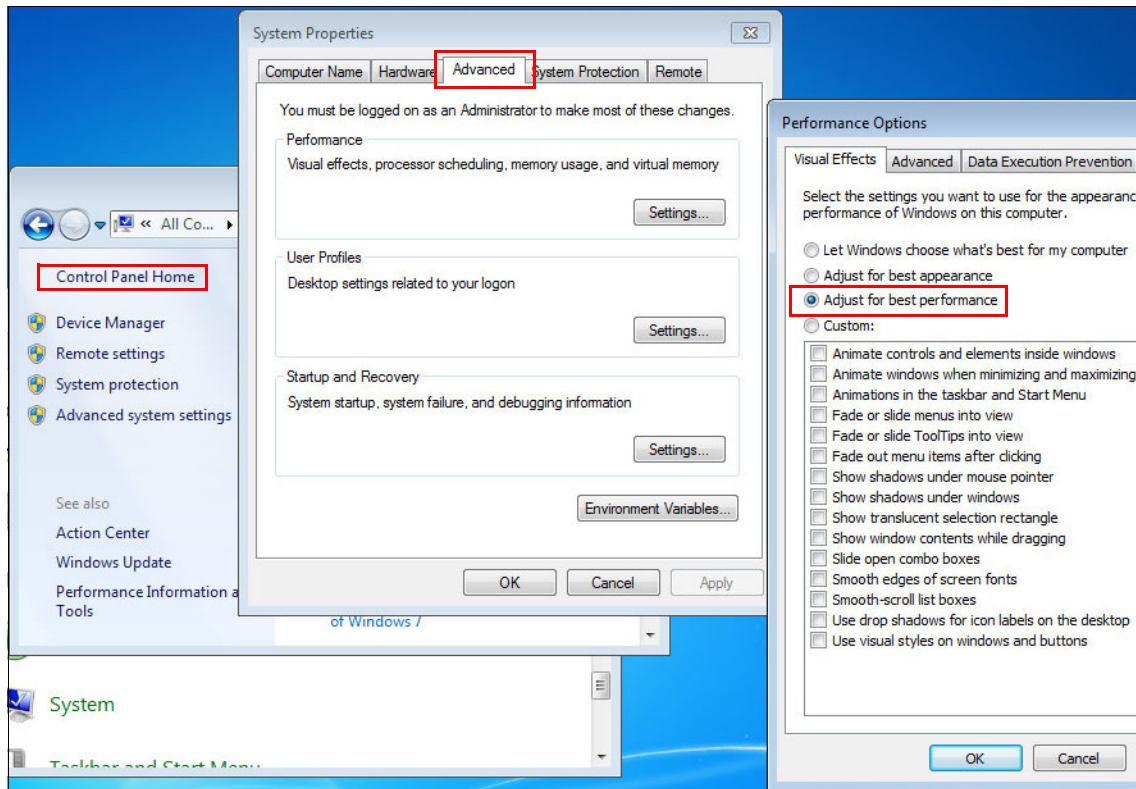


Figure 9-21 Appearance settings: Adjust for best performance

The Virtual Desktop is now ready for deployment.

You can make other Windows operating system optimizations for this base operating system image for a full VM. For more information about implementing other operating system-related optimizations, see 9.1.5, “Provisioning a linked clone virtual desktop image” on page 393.

**Optimization:** For more information about optimizing Windows 7 and Windows 8 operating systems, see *VMware Horizon View Optimization Guide for Windows 7 and Windows 8*, which is available at this website:

<http://bit.ly/1e09kxf>

## 9.1.5 Provisioning a linked clone virtual desktop image

Linked clone virtual desktops are based on a parent VM snapshot. Therefore, you must optimize the parent VM before you create the snapshot that you use to create every linked clone virtual desktop. In addition, there are adjustments to reduce the size of the operating system disk and to optimize the overall performance. These adjustments are related to the process that is described in 9.1.4, “Provisioning a full VM image” on page 390 and other specific settings, which are also applicable to the full VM images.

Consider the following settings or services to review or optimize:

- ▶ Windows hibernation
- ▶ Windows disk defragmentation
- ▶ Windows update service
- ▶ Diagnostic policy service
- ▶ Prefetch and superfetch registry key
- ▶ Windows registry backup scheduled task
- ▶ System restore
- ▶ Feed synchronization task
- ▶ Operating system components or features that are not used

**Multiple partitions are not supported:** Ensure that the parent VM has a single volume because multiple partitions on the operating system volume are not supported for linked clones.

**Optimization:** For more information about optimizing Windows 7 and Windows 8 operating systems, see *VMware Horizon View Optimization Guide for Windows 7 and Windows 8*, which is available at this website:

<http://bit.ly/1e09kxf>

## 9.2 Installing the VMware Horizon View Agent

VMware Horizon View Agent is the component that interacts between the VMware Horizon View Server infrastructure and the desktop operating system. You must install the VMware Horizon View Agent in each base image or template.

**Important:** Before VMware Horizon View Agent is installed, be sure to have VMware Tools installed correctly.

The procedure that is used to install linked clone VMs or full VMs is the same.

Complete the following steps to install and setup VMware Horizon View Agent:

1. Mount or copy the base VM the Horizon View Agent executable file, and then double-click the file to run it, as shown in Figure 9-22.

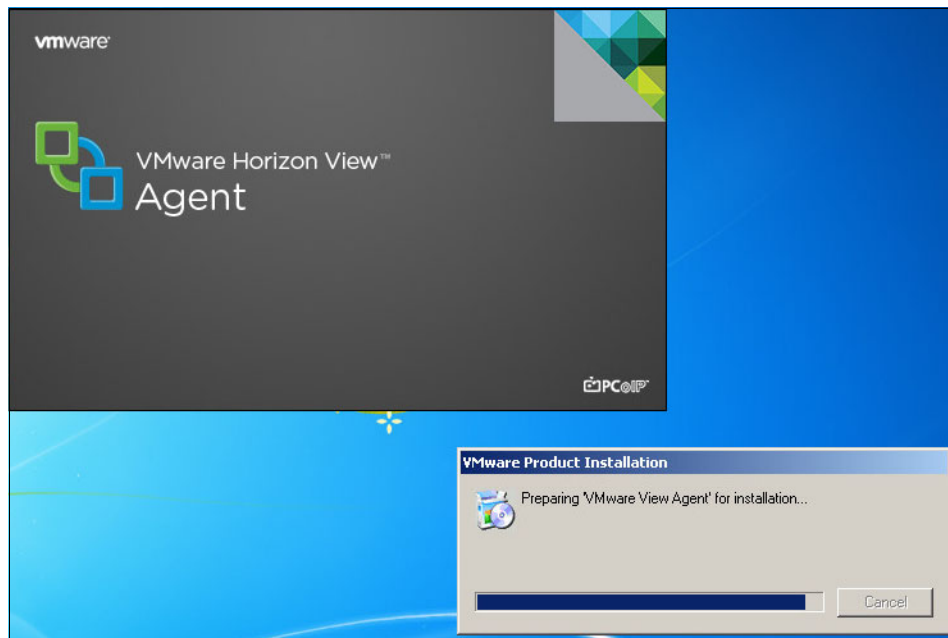


Figure 9-22 Initial VMware Horizon View installation

2. Click **Next** in the welcome window, as shown in Figure 9-23.

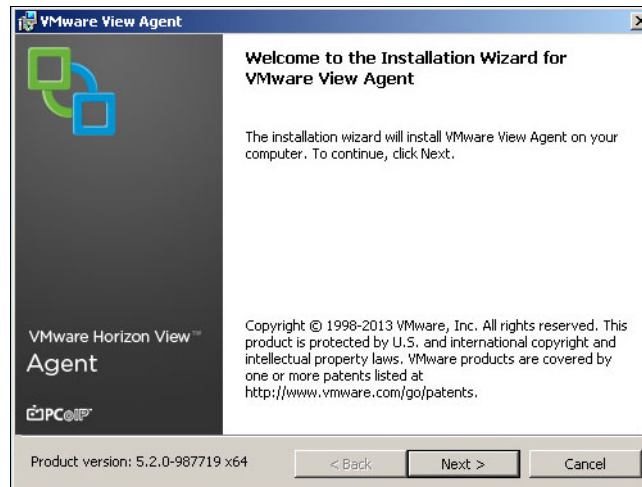


Figure 9-23 VMware Horizon View Agent welcome window

3. Read and accept the VMware View Agent License Agreement and then click **Next**, as shown in Figure 9-24.



Figure 9-24 Read and accept the VMware View Agent License Agreement

4. Keep the Custom Setup features at their default settings by clicking **Next**, as shown in Figure 9-25.

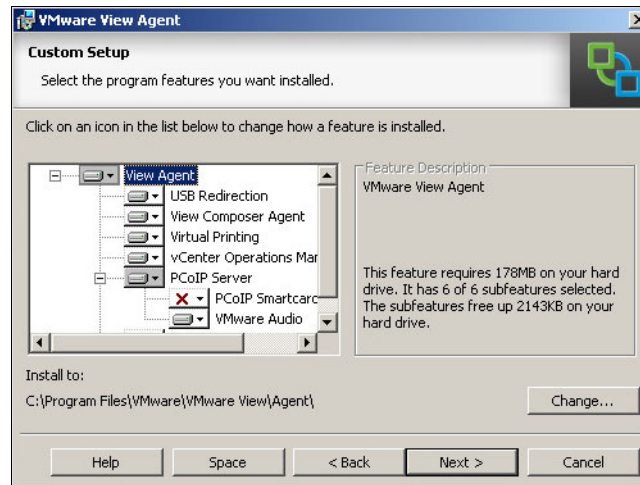


Figure 9-25 Custom Setup program features

5. Click **Install** to begin the installation process, as shown in Figure 9-26.




Figure 9-26 Begin installation



6. When the installation process ends, click **Finish**, as shown in Figure 9-27.




Figure 9-27 VMware Horizon View agent: Installation completed

7. When you are prompted for a reboot, click **Yes**, as shown in Figure 9-28.




Figure 9-28 Reboot

If users play full-screen videos for a better multimedia user experience, you must optimize the VM NIC to use all the bandwidth for multimedia contents after VM reboots.

Verify on the VM registry that the following key is present:

HKLM\System\CurrentControlSet\Services\Afd\Parameters

The key should contain a REG\_DWORD FastSendDatagramThreshold with a value of 1500, as shown in Figure 9-29.

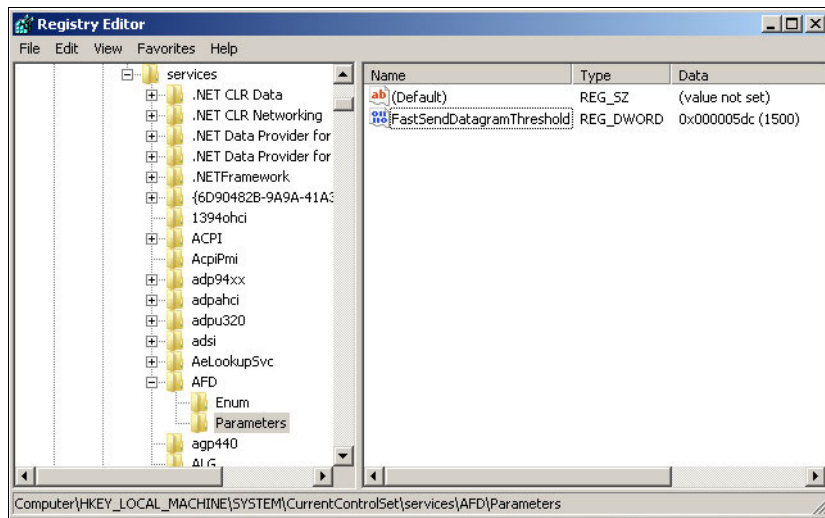


Figure 9-29 Datagram threshold

If the registry key is not present, you can find information about how to implement the necessary registry modifications on the VM at this website:

<http://support.microsoft.com/kb/235257>

Consider changing or disabling all features and services on this base operating system image that might use processor or memory resources, even if the VM is idle for a long time. For example, you might consider features, such as disk defragmentation, prefetch and superfetch, system restore points, and hibernation.

## 9.3 Configuring active directory policies

You can configure each user and each active directory computer object by using a *group policy object* (GPO). The organizational unit (OU) to which this GPO refers is the OU that was created in Chapter 7, “Deploying VMware Horizon View infrastructure” on page 277.

When you install View Connection Server, several component-specific Group Policy Administrative (ADM) template files are also installed. You can optimize and secure desktops by adding the policy settings in these ADM template files to apply a new GPO in the VDI OU in the active directory. The new GPO is applied on the desktop startup and when users log in.

The ADM template files are installed in the following directory on View Connection Server:

C:\Program Files\VMware\VMware View\Server\extras\GroupPolicyFiles

Table 9-1 lists the details about each ADM template file.

*Table 9-1 ADM template files list*

Template File	Template Name	Purpose
vdm_agent.adm	View Agent configuration	Authentication and environmental components for View Agent
vdm_client.adm	View Client Configuration	Policy settings that are related to View Client configuration
vdm_server.adm	View Server Configuration	Policy settings that are related to View Connection Server
vdm_common.adm	View common configuration	Policy settings that are common to all View components
pcoip.adm	PCoIP session variables configuration	Policy settings that are related to the PCoIP display protocol
viewPM.adm	View Persona Management configuration	Policy settings that are related to View Persona Management

Based on your current active directory environment, you might choose to add one or more ADM templates to your existing GPOs.

### 9.3.1 Configuring View Persona Management active directory policies

With View Persona Management, you can configure user profiles that are dynamically synchronized with a remote profile repository. View Persona Management expands the functionality and improves the performance of Windows roaming profiles but does not require Windows roaming profiles to operate.

With viewPM.adm administrative template, you can configure group policy settings to enable View Persona Management and control various aspects of View Persona Management.

To set up View Persona Management correctly, you need access to the following components:

- ▶ A viewPM.adm administrative template file
- ▶ Domain controller where you can configure GPOs
- ▶ File server where you can store domain user's profile files

Complete the following steps to enable and configure VMware View Persona Management to point to a network share for storing user's profiles:

1. Connect to your file server and create a network share that is named Profiles, as shown in Figure 9-30.

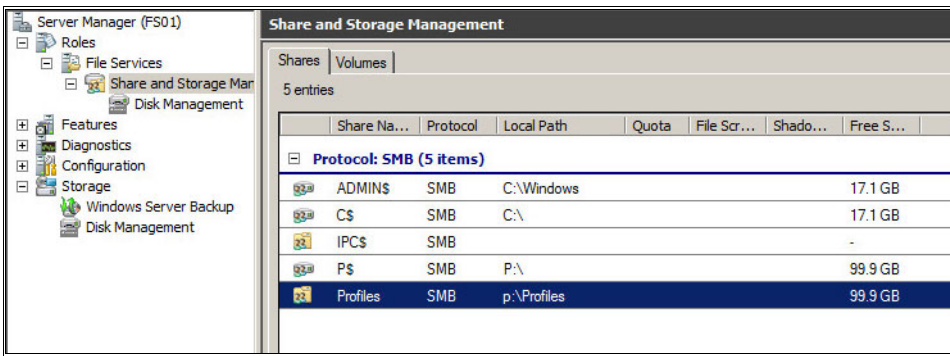



Figure 9-30 Shared folder for user Profiles

**Important:** The file server must be reachable by the desktops to apply the correct users profiles.

2. Connect to your domain controller and copy the administrative template file `viewPM.adm` from the VMware View Connection server to your domain controller, as shown in Figure 9-31.



*Figure 9-31 Local copy of View Persona Management ADM file*

3. On the Domain Controller, open Group Policy Management by clicking **Start** → **Administrative Tools** → **Group Policy Management**, as shown in Figure 9-32.

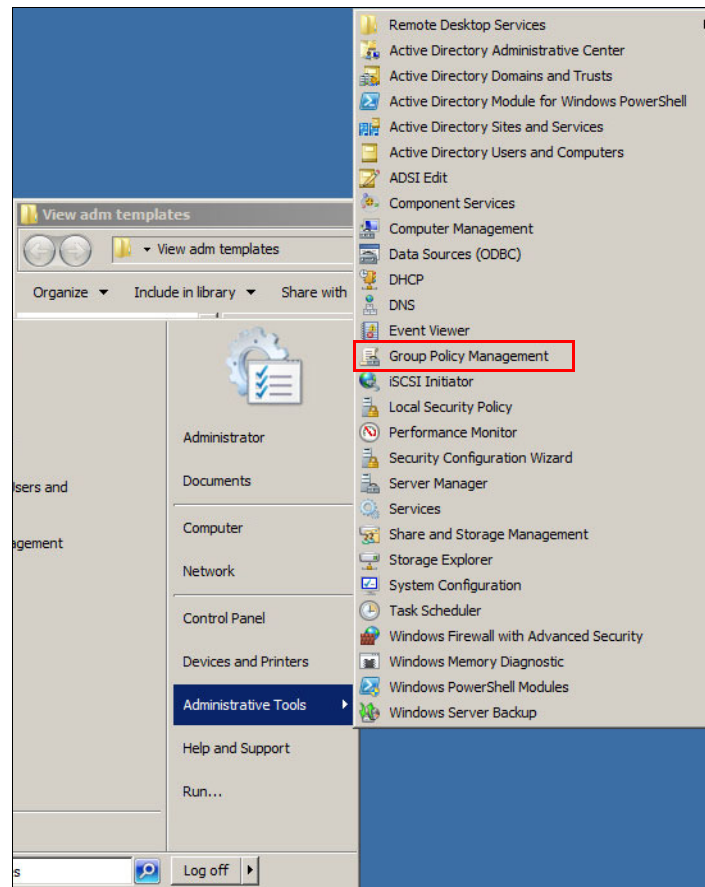


Figure 9-32 Open Group Policy Management

4. Right-click the VDI organizational unit and select **Create a GPO in this container, and link it here**, as shown in Figure 9-33.

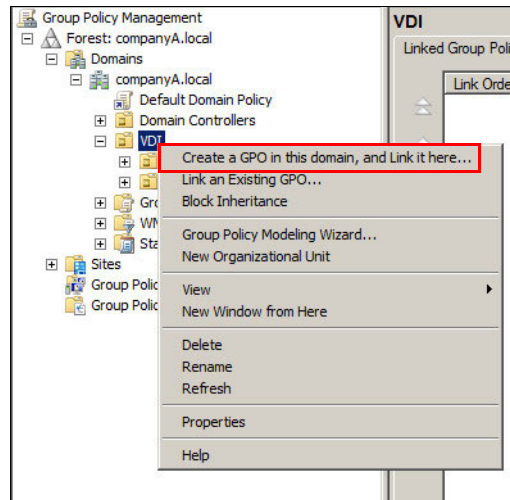


Figure 9-33 Create GPO

5. Name the GPO as shown in Figure 9-34 and then click **OK**.




Figure 9-34 GPO name

6. Right-click the newly created GPO and select **Edit**, as shown in Figure 9-35.

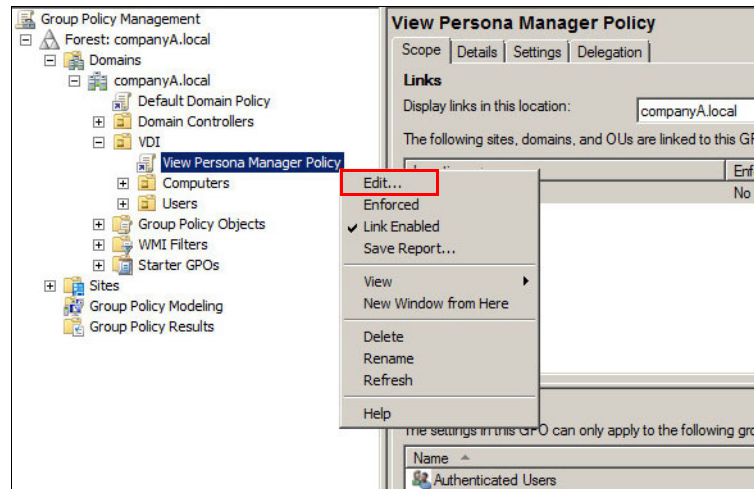


Figure 9-35 Edit GPO

7. Click **Computer Configuration** → **Policies**. Right-click **Administrative Templates** and select **Add/Remove Templates**, as shown in Figure 9-36.

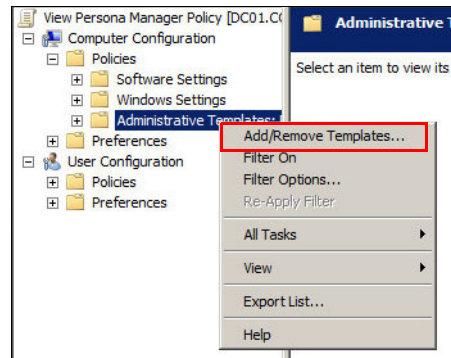


Figure 9-36 Adding a template to a GPO



8. On the Current Policy Templates window, click **Add**, as shown in Figure 9-37.

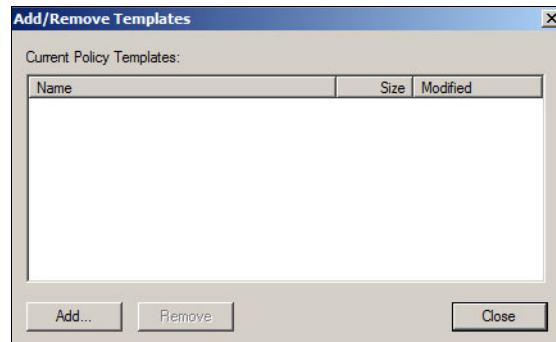


Figure 9-37 Current Policy Templates window

9. Select the folder where you previously saved the `view_PM.adm` file and click **Open**, as shown in Figure 9-38.

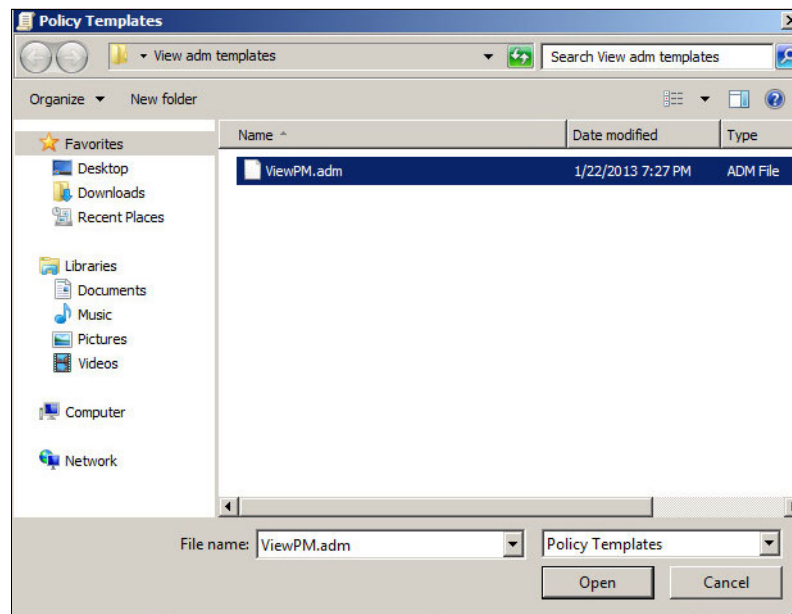


Figure 9-38 Load the ADM template

10. On the Add/Remove Templates window, click **Close**, as shown in Figure 9-39.

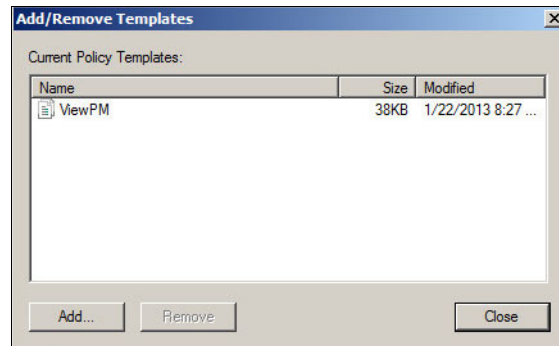


Figure 9-39 Templates added list

11. Click **Administrative Templates** → **VMware View Agent Configuration** → **Persona Management**, as shown in Figure 9-40.

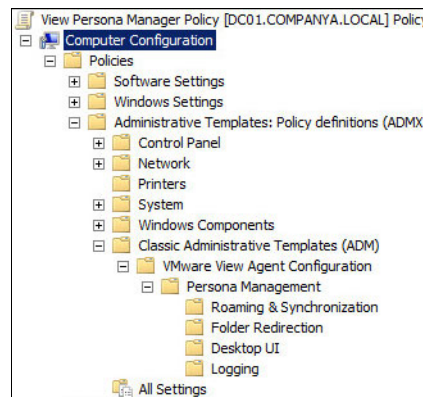


Figure 9-40 Persona Management ADM Policy

12. Select **Roaming & Synchronization** in the left pane, then double-click **Manage User Persona** in the right pane, as shown on Figure 9-41.

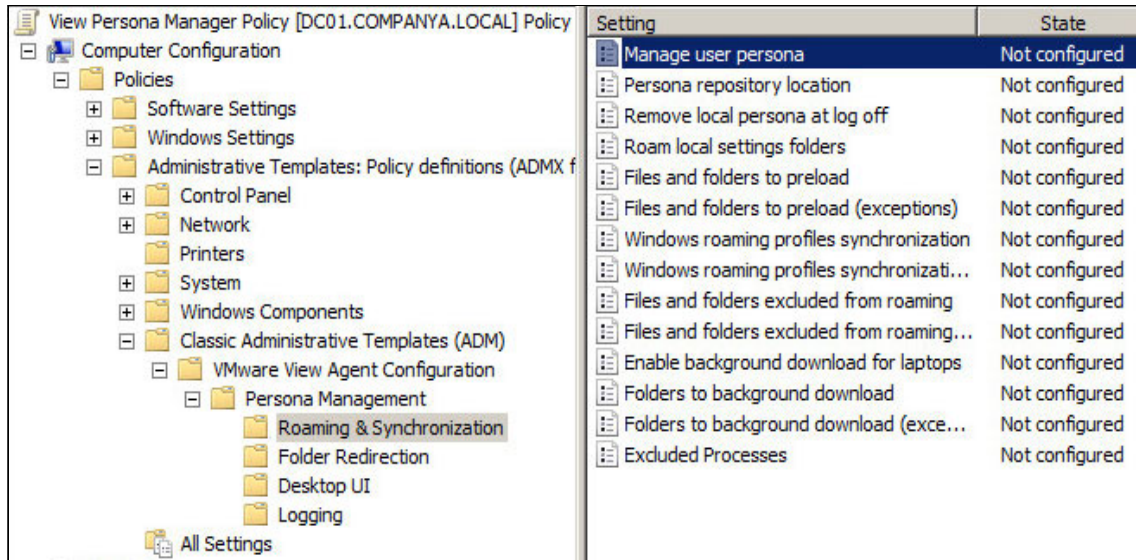


Figure 9-41 Manage user persona

13. Select the **Enabled** option and then click **OK**, as shown in Figure 9-42.

The screenshot shows a Windows-style dialog box titled "Manage user persona". At the top right are "Previous Setting" and "Next Setting" buttons. Below the title bar, there are three radio buttons: "Not Configured", "Enabled" (which is selected), and "Disabled". To the right of these is a "Comment:" text box. Below the radio buttons is a "Supported on:" text box. In the center, there are two sections: "Options:" and "Help:". The "Options:" section contains a "Profile upload interval (in minutes):" label and a spinner box set to "10". The "Help:" section contains a text box with the following text: "When enabled, the user's persona will be managed dynamically. When disabled, the user's persona will be managed by Windows. The profile upload interval is used to determine how often to upload persona changes to the network." At the bottom right are "OK", "Cancel", and "Apply" buttons.

Figure 9-42 User persona GPO

14. At the main GPO settings, double-click **Persona repository location** in the right pane, as shown in Figure 9-43.

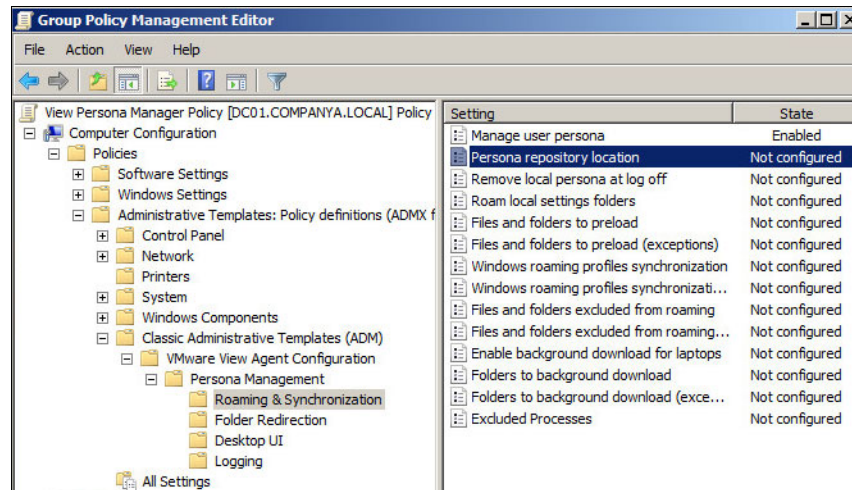


Figure 9-43 Persona repository location GPO

15. Select **Enabled** and complete the full path to the share that you created earlier, as shown in Figure 9-44. Click **OK**.

The screenshot shows a Windows-style dialog box titled "Persona repository location". At the top right are "Previous Setting" and "Next Setting" buttons. Below the title bar, there are three radio buttons: "Not Configured", "Enabled" (which is selected), and "Disabled". To the right of these is a "Comment:" text box. Below the radio buttons is a "Supported on:" text box. In the "Options:" section, there is a "Share path:" label followed by a text box containing "\\fs01\profiles". This text box is highlighted with a red rectangular border. Below the "Share path:" text box is a checkbox labeled "Override Active Directory user profile path if it is configured." To the right of the "Options:" section is a "Help:" section with two paragraphs of text: "The UNC path to the repository where user personas will be stored." and "If this path is left blank, the user profile path in Active Directory will be used." At the bottom right of the dialog are "OK", "Cancel", and "Apply" buttons.

Figure 9-44 Persona profiles path

16. Close the Group Policy Editor Console.

17. Right-click the View Persona Management Policy and select **Enforced**, as shown in Figure 9-45.

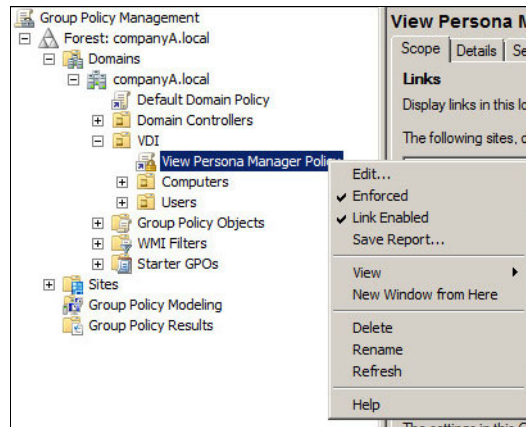


Figure 9-45 Policy enforced

### 9.3.2 Allowing a connection to a remote desktop users group policy

To allow connection to the remote desktop, you must create a domain policy to automatically add the specific group of users (standard users or VIP users) to the local virtual desktop remote desktop users.

Complete the following steps to create the GPO:

1. Connect to a domain controller and create a policy or edit an existing policy.
2. In policy editor, expand **Computer Configuration**. Then, click **Policies** → **Windows Settings** → **Security Settings**. Right-click **Restricted Groups**, as shown in Figure 9-46, and select **Add Group**.

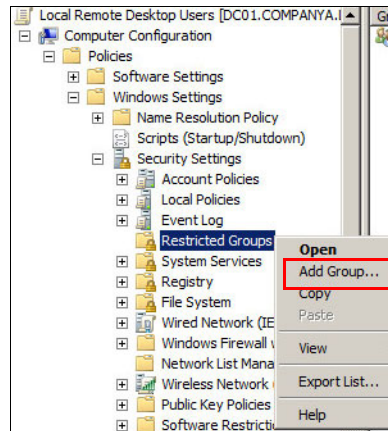


Figure 9-46 Add Group

3. On the Add Group window, click **Browse**. Then, on the Select Groups window, enter Remote Desktop Users and click **Check Names**, as shown in Figure 9-47. When the object name is underlined, click **OK**.

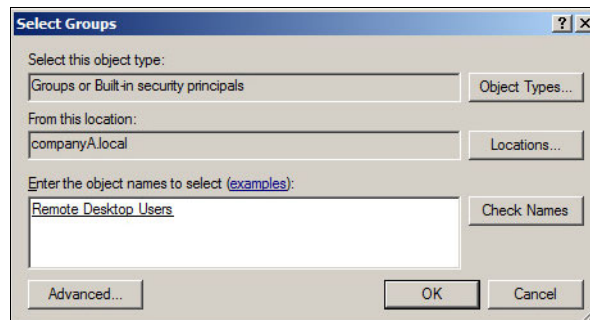


Figure 9-47 Group query

4. Click **OK** to close the Add Group window and return to Group Policy editor.



5. Double-click **Remote Desktop Users** in the group policy editor and add both Standard Users and VIP users to the group, as shown in Figure 9-48.

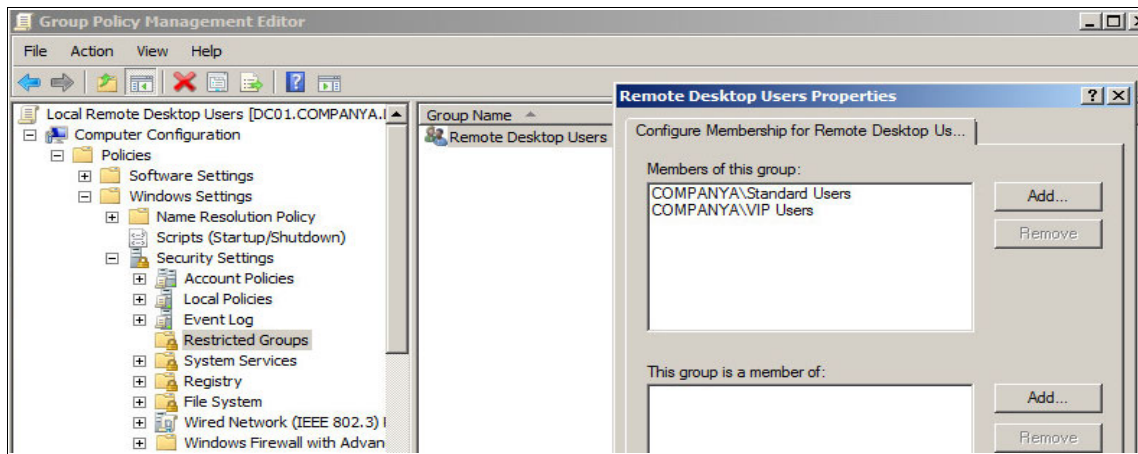


Figure 9-48 Add Remote Desktop Users

6. Click **OK** and close Group Policy Editor.
7. Reboot the virtual desktops to apply the policy.

## 9.4 VMware View Manager and desktop pools

VMware View Manager is the main web-based interface to interact with virtual desktops. It contains all of the settings and the automation scripts to interact with VMware Composer to create linked clone virtual desktops or with vCenter to create full virtual machine desktops.

### 9.4.1 Configuring Event DB

When installed, VMware View Manager does not have any Event DB configured. Therefore, there is no record at all of events, apart from those tasks that can be monitored by using the Tasks view. For example, there is no history on overnight scheduled power off and power-on tasks or if a specific virtual desktop had issues.

Complete the following steps to configure the Event DB:

1. Create an SQL DB on your existing SQL server. You can use the same SQL server that was used for the Composer database.
2. Log on to the VMware Horizon View Administrator web interface.

3. In the VMware Horizon View Administrator main window, expand **View Configuration** in the left pane and select **Event Configuration**, as shown in Figure 9-49. Click **Edit** in the right pane under Event Database.

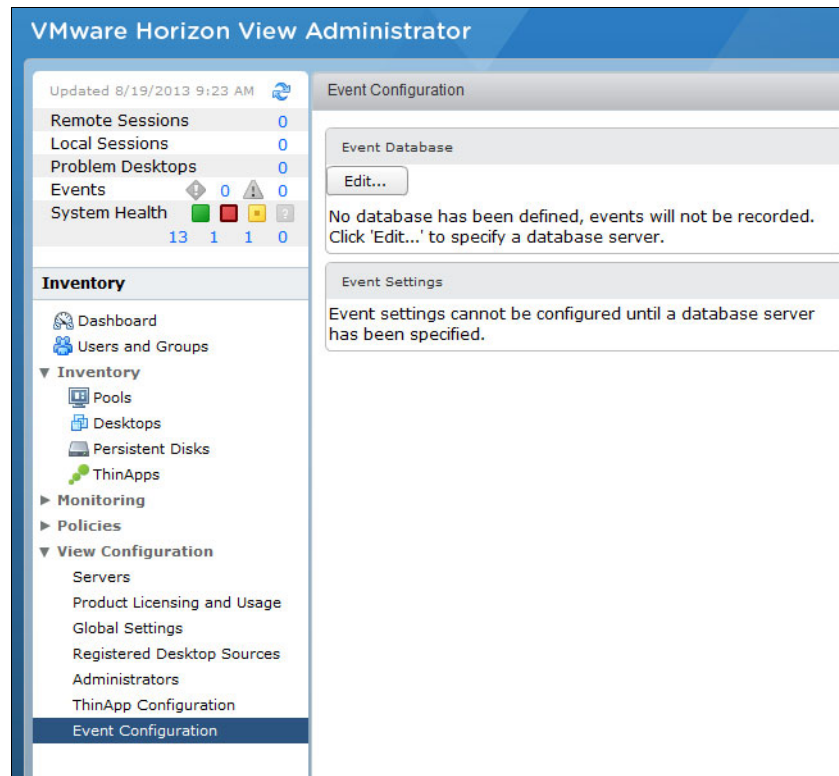
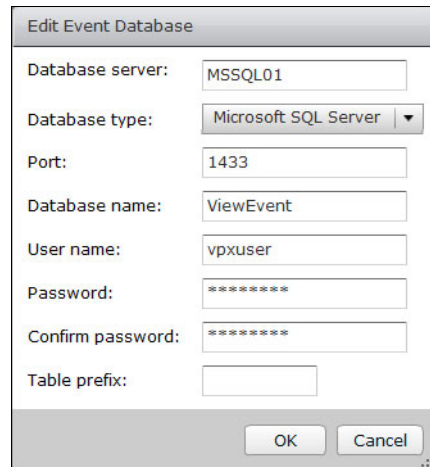


Figure 9-49 Event Configuration

4. Complete the Event DB-related information, as shown in Figure 9-50. Click **OK**.



The screenshot shows a dialog box titled "Edit Event Database". It contains the following fields and values:

- Database server: MSSQL01
- Database type: Microsoft SQL Server (dropdown menu)
- Port: 1433
- Database name: ViewEvent
- User name: vpxuser
- Password: \*\*\*\*\*
- Confirm password: \*\*\*\*\*
- Table prefix: (empty)

At the bottom right, there are "OK" and "Cancel" buttons.

*Figure 9-50 Event DB connection*

The Event DB is now configured.

## 9.4.2 Provisioning a linked clone virtual desktop

**Note:** Before you provision a linked clone virtual desktop, follow the steps that are described in 9.1.5, “Provisioning a linked clone virtual desktop image” on page 393.

To provision a linked clone VM, follow these steps:

1. Shut down the linked clone VM from the vSphere client by selecting the linked clone VM (LCVM) and then clicking **Power** → **Shut Down Guest**, as shown in Figure 9-51.

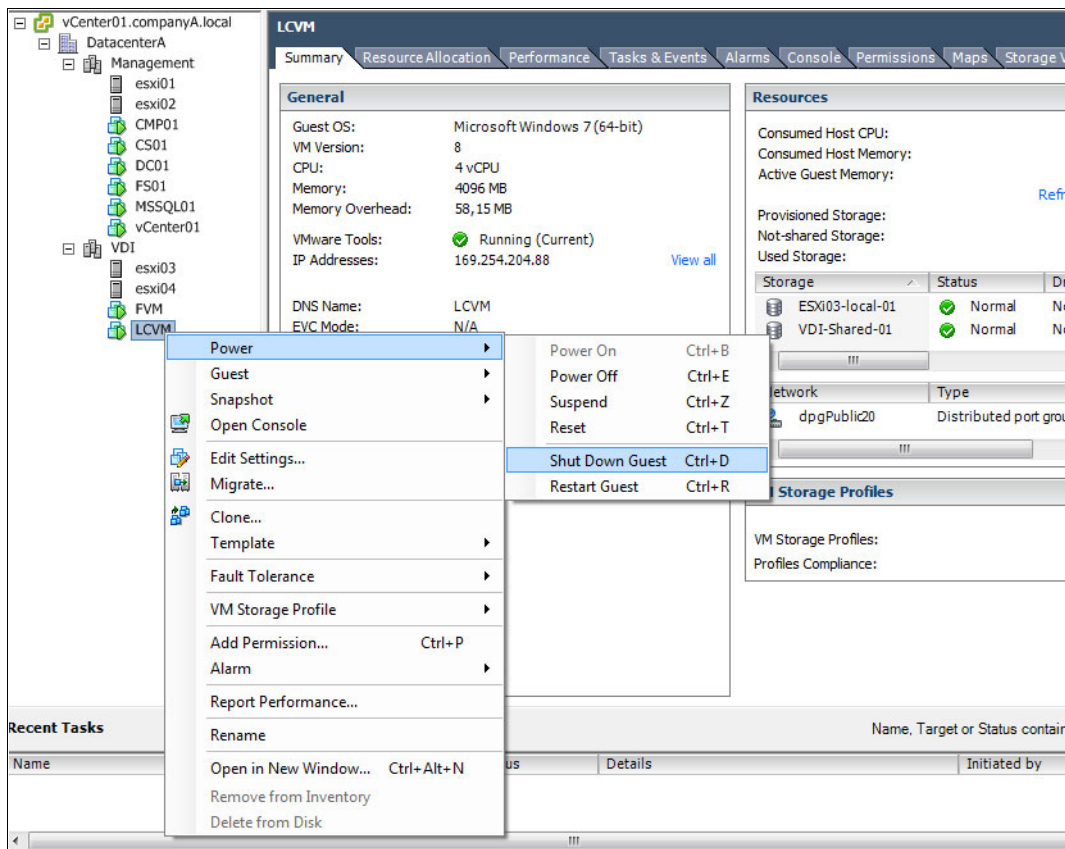


Figure 9-51 Shut down the linked clone VM

2. When you are prompted to shut down the operating system VM, click **Yes**, as shown in Figure 9-52.




Figure 9-52 Confirmation message

3. Take a snapshot of the linked clone VM (LCVM) by selecting **Snapshot** → **Take Snapshot**, as shown in Figure 9-53.

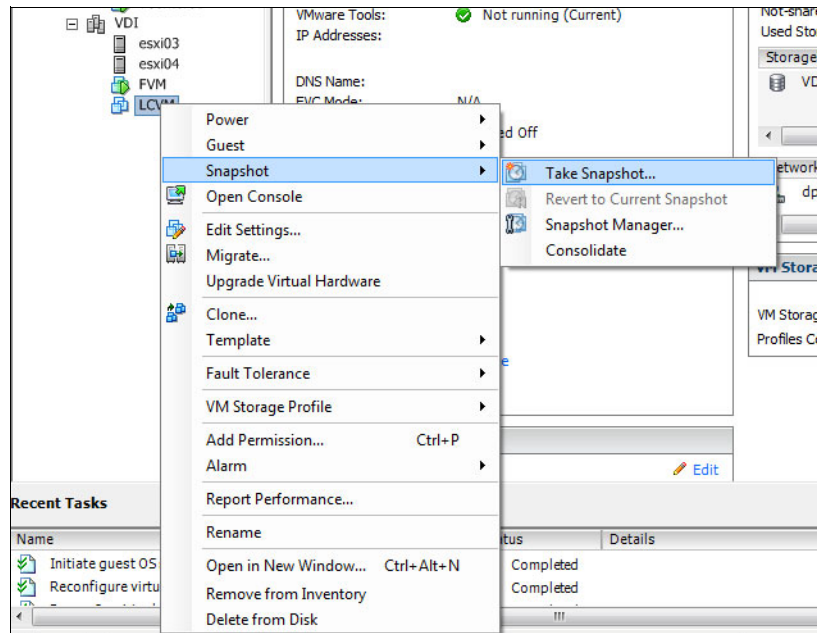


Figure 9-53 Take snapshot

4. Complete the VM snapshot information, as shown in Figure 9-54. Click **OK**.




Figure 9-54 Snapshot name and description

5. Log on to the Horizon View Administrator console. Click **Inventory** on the left pane and then click **Pools**. Click **Add** under Pools in the right pane, as shown in Figure 9-55.

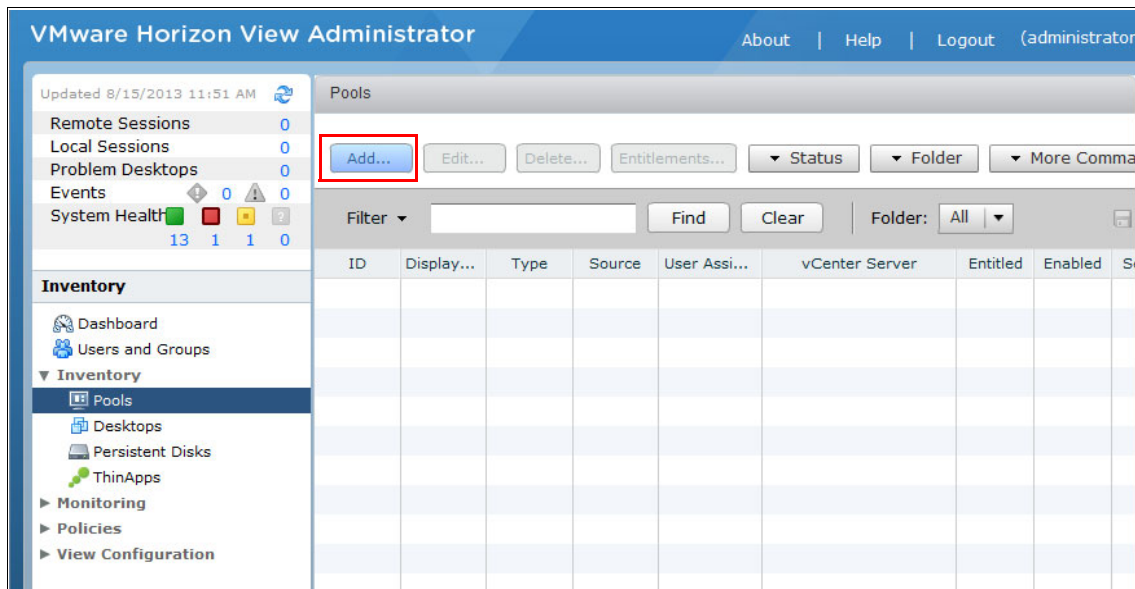


Figure 9-55 Pools main window

6. Select the **Automated Pool** option and then click **Next**, as shown in Figure 9-56.

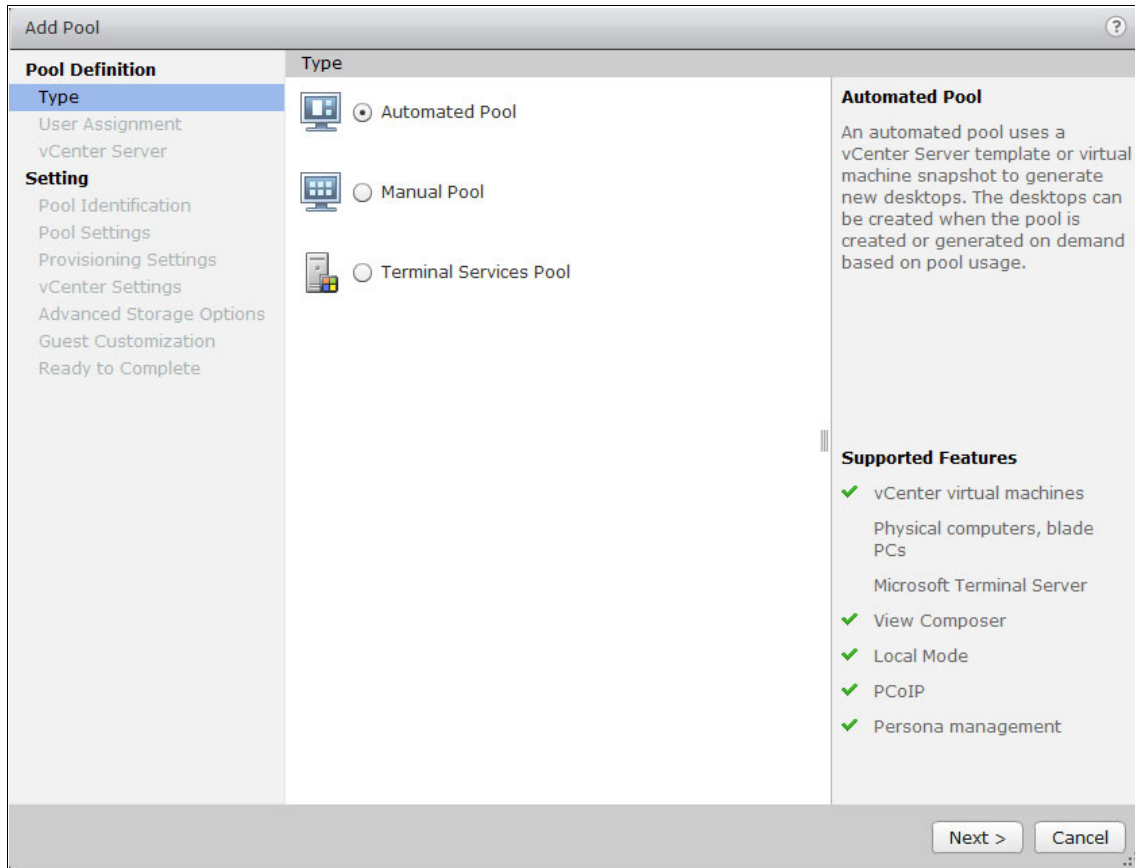


Figure 9-56 Define the pool type

7. Select **Floating**, as shown in Figure 9-57. Click **Next**.

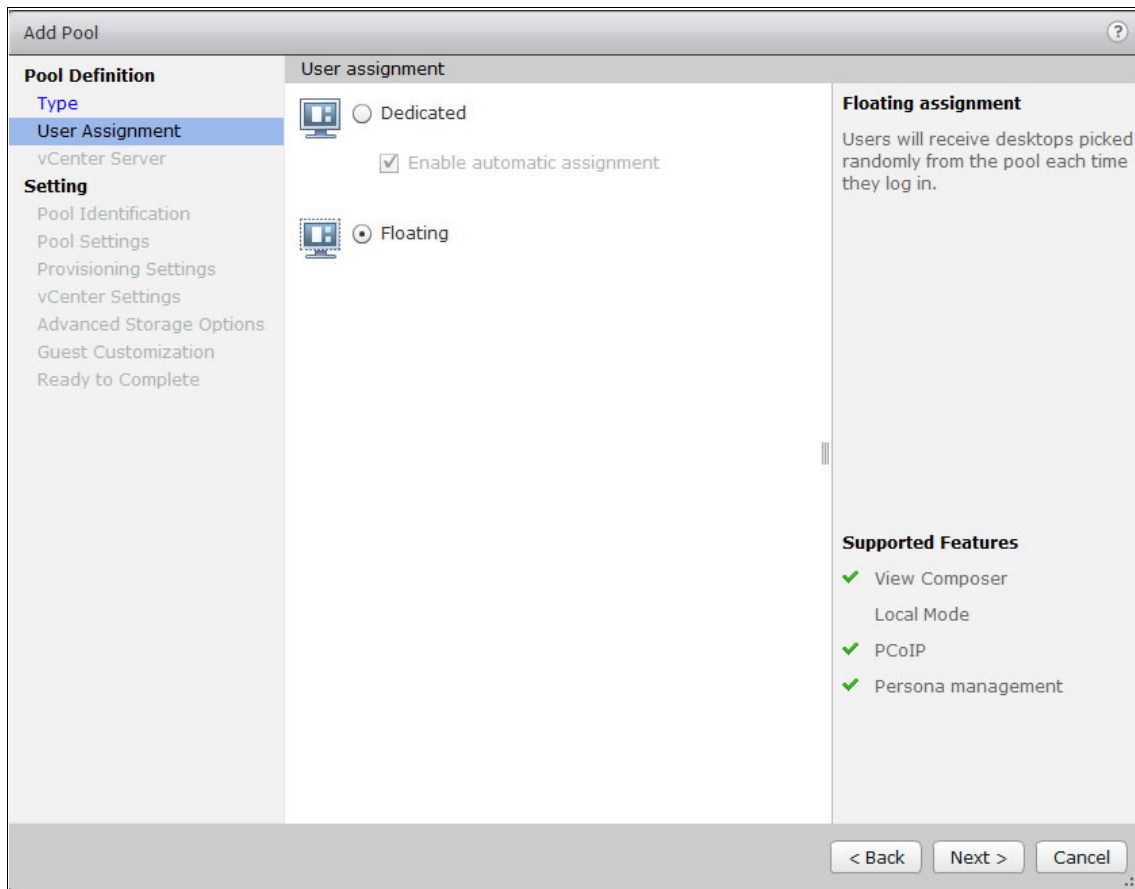


Figure 9-57 Define user assignment type



8. Select the **View Composer linked clones** option and then click **Next**, as shown in Figure 9-58.

The screenshot shows the 'Add Pool' wizard in VMware Horizon View. The 'vCenter Server' tab is selected, and the 'View Composer linked clones' radio button is chosen. The wizard is divided into three main sections: Pool Definition, Setting, and View Composer. The Pool Definition section includes links for Type, User Assignment, and vCenter Server. The Setting section lists various configuration options like Pool Identification, Pool Settings, Provisioning Settings, View Composer Disks, Storage Optimization, vCenter Settings, Advanced Storage Options, Guest Customization, and Ready to Complete. The View Composer section provides information about linked clones, including their shared base image and storage space, and the user profile redirection. A table lists the vCenter Server and View Composer details. The Supported Features section lists various features with checkmarks indicating they are supported.

**Add Pool**

**Pool Definition**

- Type
- User Assignment
- vCenter Server**

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**vCenter Server**

☐ Full virtual machines

☒ View Composer linked clones

vCenter Server	View Composer
vCenter01.CompanyA.local(Administrator)	vCenter01.CompanyA.local

Description: *None*

**View Composer**

View Composer linked clones share the same base image and use less storage space than full virtual machines.

The user profile for linked clones can be redirected to persistent disks that will be unaffected by OS updates and refreshes.

**Supported Features**

- Local Mode
- ✓ PCoIP
- ✓ Storage savings
- ✓ Recompose and refresh
- ✓ QuickPrep guest customization
- ✓ Sysprep guest customization (vSphere 4.1 or higher)
- ✓ Persona management

< Back   Next >   Cancel

Figure 9-58 Define type of virtual desktop

9. Complete the information about the pool, as shown in Figure 9-59. Click **Next**.

**Add Pool - LCVM**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification**
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**Pool Identification**

ID:

Display name:

View folder:

Description:

**ID**

The pool ID is the unique name used to identify this pool.

**Display Name**

The display name is the name that users will see when they connect to View Client. If the display name is left blank, the ID will be used.

**View Folder**

View folders can organize the pools in your organization. They can also be used for delegated administration.

**Description**

This description is only shown on the Settings tab for a pool within View Administrator.

< Back   Next >   Cancel

Figure 9-59 Pool information

10. Adjust the required pool settings, as shown in Figure 9-60. Click **Next**.

**Add Pool - LCVM**

**Pool Definition**

- Type
- User Assignment
- vCenter Server
- Setting**
- Pool Identification
- Pool Settings**
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**Pool Settings**

**General**

State:

Connection Server restrictions:  

**Remote Settings**

Remote Desktop Power Policy:  ?

Automatically logoff after disconnect:

Allow users to reset their desktops:

Allow multiple sessions per user:

Delete or refresh desktop on logoff:  ?

**Remote Display Protocol**

Default display protocol:

Allow users to choose protocol:

3D Renderer:   ?

Max number of monitors:  ?

May require power-cycle of related virtual machines ?

Max resolution of any one monitor:  ?

May require power-cycle of related virtual machines ?

HTML Access: ☐ Enabled ?

Figure 9-60 Pool settings

11. Set up a virtual desktop naming convention by following the example that is shown in Figure 9-61. Click **Next**.

**Add Pool - LCVm**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings**
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**Provisioning Settings**

**Basic**

- ☒ Enable provisioning
- ☒ Stop provisioning on error

**Virtual Machine Naming**

- ☐ Specify names manually
  - 0 names entered
  - Enter names...
- ☐ Start desktops in maintenance mode
- # Unassigned desktops kept powered on: 1
- ☒ Use a naming pattern
- Naming Pattern: LCVm-{n}-STDUSR

**Pool Sizing**

- Max number of desktops: 20
- Number of spare (powered on) desktops: 2
- Minimum number of ready (provisioned) desktops during View Composer maintenance operations: 0

**Provisioning Timing**

- ☒ Provision desktops on demand
- Min number of desktops: 1
- ☐ Provision all desktops up-front

**Naming Pattern**

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.).

The unique number can also be made a fixed length. (For example: vm-{n:fixed=3}-sales).

See the help for more naming pattern syntax options.

< Back   Next >   Cancel

Figure 9-61 Provisioning settings

12. Accept the default options for the View Composer Disks settings as shown in Figure 9-62 by clicking **Next**.

The screenshot shows the 'Add Pool - LCVM' wizard window. The left sidebar contains a tree view with the following items: 'Pool Definition' (expanded), 'Type', 'User Assignment', 'vCenter Server', 'Setting' (expanded), 'Pool Identification', 'Pool Settings', 'Provisioning Settings', 'View Composer Disks' (selected), 'Storage Optimization', 'vCenter Settings', 'Advanced Storage Options', 'Guest Customization', and 'Ready to Complete'. The main area is titled 'View Composer Disks' and contains two sections. The 'Disposable File Redirection' section has a radio button selected for 'Redirect disposable files to a non-persistent disk'. Below this, the 'Disk size' is set to '4096 MB (minimum 512 MB)' and the 'Drive letter' is set to 'Auto'. The second radio button, 'Do not redirect disposable files', is unselected. The 'Disposable File Redirection' section on the right contains the text: 'Use this option to redirect disposable files to a non-persistent disk that will be deleted automatically when a user's session ends.' At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Add Pool - LCVM

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks**
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**View Composer Disks**

**Disposable File Redirection**

☒ Redirect disposable files to a non-persistent disk

Disk size: 4096 MB (minimum 512 MB)

Drive letter: Auto

☐ Do not redirect disposable files

**Disposable File Redirection**

Use this option to redirect disposable files to a non-persistent disk that will be deleted automatically when a user's session ends.

< Back Next > Cancel

Figure 9-62 View Composer disposable disks

13. Accept the default settings for the Storage Optimization options as shown in Figure 9-63 by clicking **Next**.

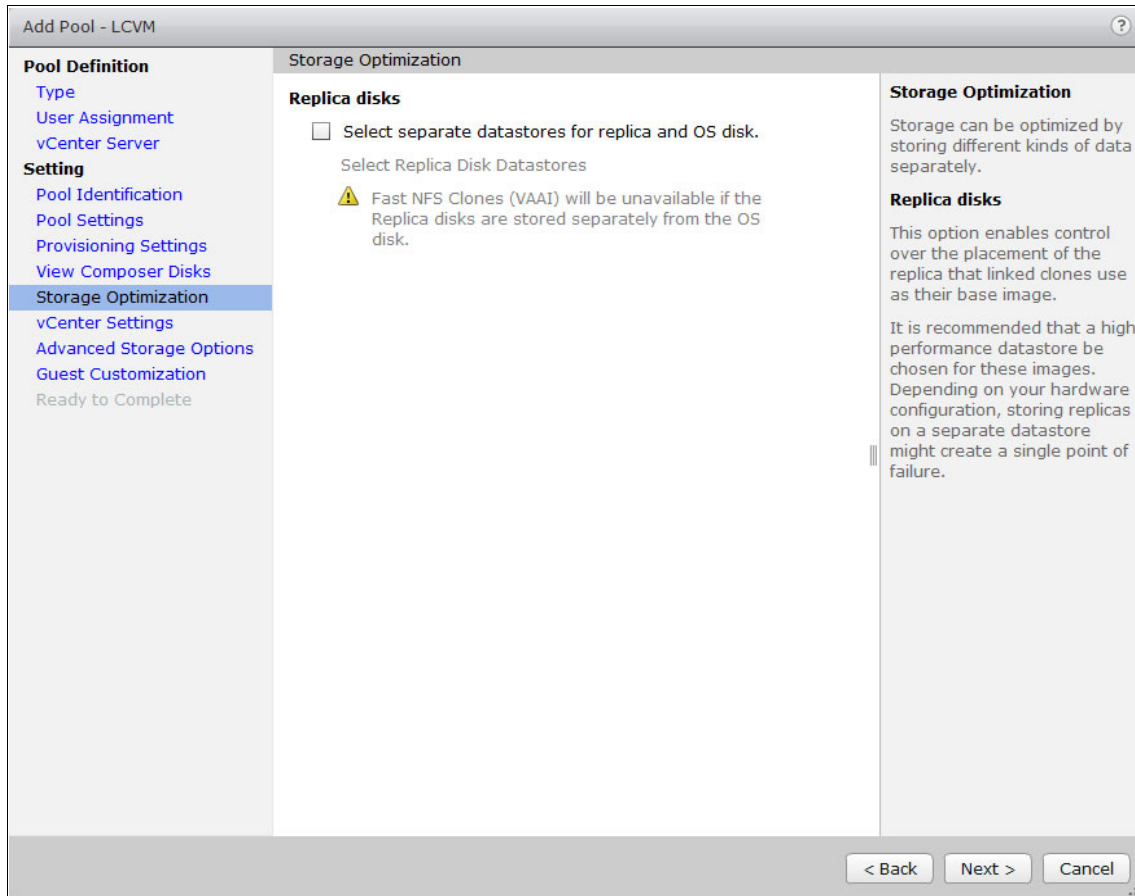


Figure 9-63 Storage optimization

14. Use the vCenter Settings page (as shown in Figure 9-64) to set several vCenter settings. To change the settings, click **Browse** to the right of the setting that you want to change.

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings**
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**vCenter Settings**

**Default Image**

1 Parent VM: <Click Browse...> Browse...

2 Snapshot: <Click Browse...> Browse...

**Virtual Machine Location**

3 VM folder location: <Click Browse...> Browse...

**Resource Settings**

4 Host or cluster: <Click Browse...> Browse...

5 Resource pool: <Click Browse...> Browse...

6 Datastores: Click Browse to select Browse...

< Back Next > Cancel

Figure 9-64 vCenter settings page

When you are changing vCenter settings, make the following changes:

- a. For the Parent VM setting, select the **LCVM** name and path, as shown in Figure 9-65, and click **OK**.




Figure 9-65 Select parent LCVM name and path

- b. For the Snapshot details, enter the Snapshot file details as shown in Figure 9-66 and click **OK**.

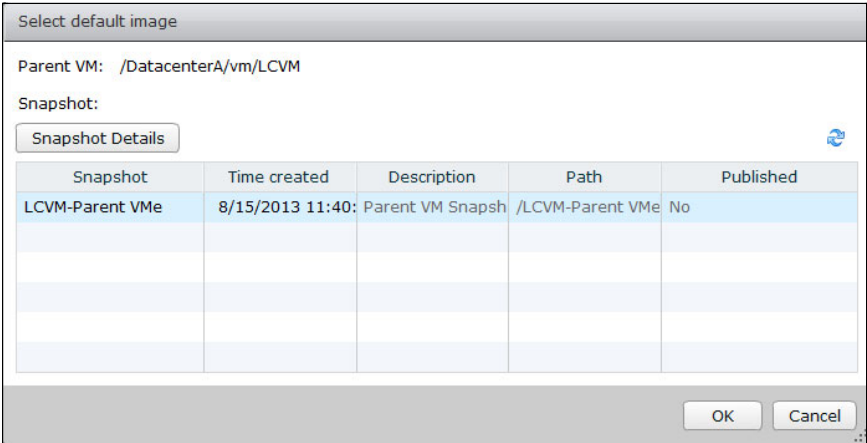



Figure 9-66 Snapshot file details

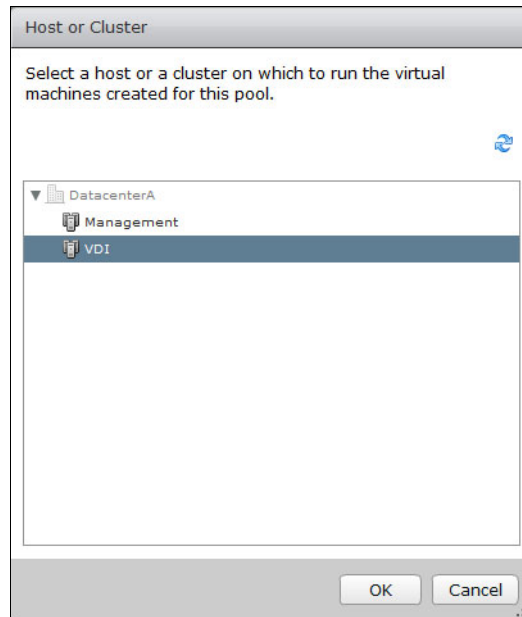


- c. Enter the VM folder location as shown in Figure 9-67 and click **OK**.



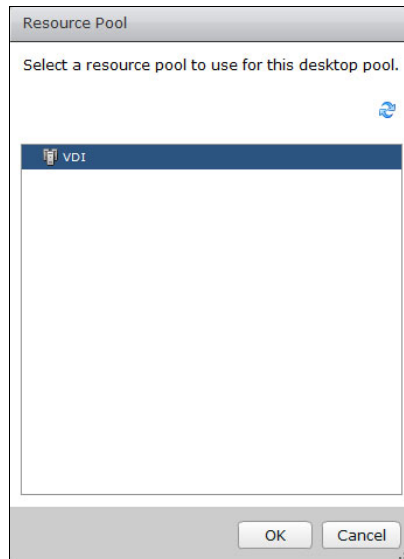
*Figure 9-67 VM folder location*

- d. Select the VDI Cluster location as shown in Figure 9-68 and click **OK**.



*Figure 9-68 VDI Cluster selection*

- e. Select the VDI pool as shown in Figure 9-69 and then click **OK**.



*Figure 9-69 Select a resource pool*

- f. Select the two ESXi SSD local data stores, as shown in Figure 9-70. Click **OK**.

Select Linked Clone Datastores

Select the linked clone datastores to use for this pool. Only datastores that can be used by the selected host or cluster can be selected.

☒ Show all datastores (including local datastores)
 

Local datastore

Shared datastore

	Datastore	Capacity (GB)	Free (GB)	Type	Storage Overcommit
<input checked="" type="checkbox"/>	ESXi03-local-01	930.25	880.16	VMFS 5	Conservative
<input checked="" type="checkbox"/>	ESXi04-local-01	930.25	883.44	VMFS 5	Conservative
<input type="checkbox"/>	VDI-Shared-01	599.75	498.79	VMFS 5	

Data Type	Selected Free Space (GB)	Min Recommended (GB)	50% utilization (GB)	Max Recommended (GB)
Linked clones	1,763.60	108.00	129.00	154.00

OK

Cancel

Figure 9-70 Select linked clone data stores

15. On the main wizard page, click **Next**. When you are prompted with a warning message about the pool configuration (as shown in Figure 9-71), click **OK**.

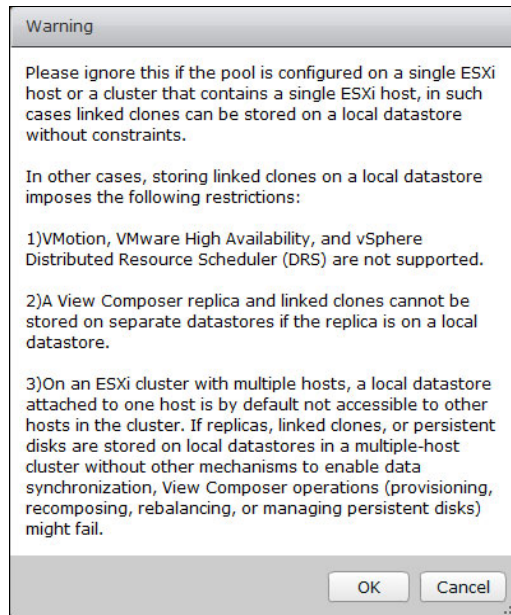


Figure 9-71 Warning message

16. On the Advanced Storage Options page, verify the settings as shown in Figure 9-72. Click **Next**.

**Add Pool - LCVN**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options**
- Guest Customization

Ready to Complete

**Advanced Storage Options**

Based on your resource selection, the following features are recommended. Options that are not supported by the selected hardware are disabled.

☒ Use View Storage Accelerator

Disk Types: OS disks

Regenerate storage accelerator after: 7 Days

☐ Other Options

- ☐ Use native NFS snapshots (VAAI) Tech Preview ?
- ☐ Reclaim VM disk space ?

Initiate reclamation when unused space on VM exceeds: 1 GB

**Blackout Times**

Storage accelerator regeneration and VM disk space reclamation do not occur during blackout times. The same blackout policy applies to both operations.

Add... Edit... Remove

Day	Time

**View Storage Accelerator**

vSphere 5.x hosts can be configured to improve performance by caching certain pool data. Enable this option to use View Storage Accelerator for this pool. View Storage Accelerator is most useful for shared disks that are read frequently, such as View Composer OS disks.

**Native NFS Snapshots (VAAI)**

VAAI (vStorage API for Array Integration) is a hardware feature of certain storage arrays. It uses native snapshotting technology to provide linked clone functionality. Choose this option only if you have appropriate hardware devices.

**Disk Space Reclamation**

With vSphere 5.x, virtual machines can be...

< Back Next > Cancel

Figure 9-72 Advanced Storage Options page

17. In the Guest Customization page, click **Browse** to the right of the AD container field, and then select the linked clone VM OU (in this example, OU=LCVM,OU=Computers,OU=VI), as shown on Figure 9-73.

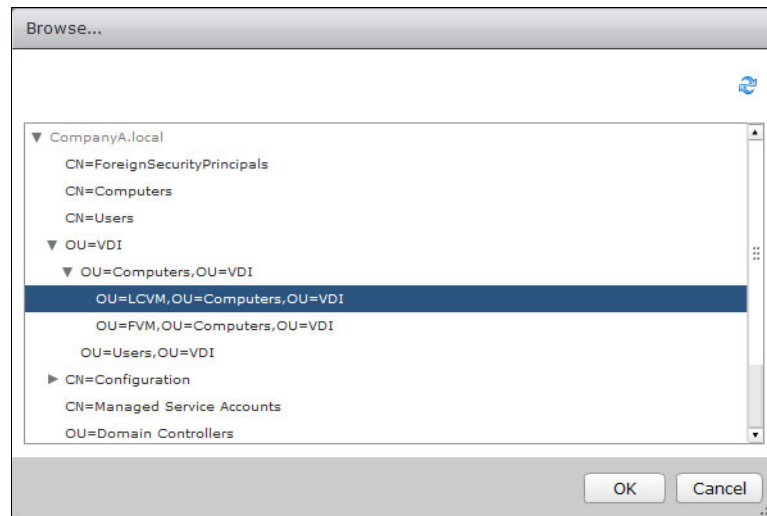


Figure 9-73 AD container selection

18. Complete the Guest Customization window by selecting the custom specification file that you created in 9.1.1, “Creating a customization specification file” on page 374, as shown on Figure 9-74. Click **Next**.

**Add Pool - LCVN**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization**
- Ready to Complete

**Guest Customization**

Domain:

AD container:

☐ Allow reuse of pre-existing computer accounts ?

☐ Use QuickPrep

Power-off script name:  ?

Power-off script parameters:  Example: p1 p2 p3

Post-synchronization script name:  ?

Post-synchronization script parameters:  Example: p1 p2 p3

☒ Use a customization specification (Sysprep)

Name	Guest OS	Description
CompanyA	Windows	
Windows7_Domain_and_network_s	Windows	Customized Wizard for automatic

< Back   Next >   Cancel

Figure 9-74 Customization specification file select window

19. In the Ready to Complete window, click **Entitle users after this wizard finishes**, as shown in Figure 9-75. Click **Finish**.

**Add Pool - LCVM**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- View Composer Disks
- Storage Optimization
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete**

**Ready to Complete**

☒ Entitle users after this wizard finishes

Type:	Automated
User assignment:	Floating assignment
vCenter Server:	vCenter01.CompanyA.local(Administrator)
Use View Composer:	Yes
Unique ID:	LCVM
Display name:	Standard User Pool
View Folder:	/
Desktop pool state:	Enabled
Remote Desktop Power Policy:	Take no power action
Automatic logoff after disconnect:	Immediately
Connection Server restrictions:	None
Allow users to reset their desktop:	No
Allow multiple sessions per user:	No
Delete or refresh desktop on logoff:	Never
Default display protocol:	PCoIP
Allow users to choose protocol:	No
3D Renderer:	Disabled
Max number of monitors:	2
Max resolution of any one monitor:	1920x1200
HTML Access:	Disabled

< Back Finish Cancel

Figure 9-75 Ready to Complete page



20. Click **Add** to query Active Directory for the users group that is entitled to use the LCV virtual desktop pool, as shown in Figure 9-76.

Entitlements

Entitled users and groups can use this pool

Add... Remove

Name	Domain	Email

OK Cancel

Figure 9-76 Pool entitlement page

21. Select the Standard Users group, as shown in Figure 9-77. Click **OK**.

Find User or Group

Type: ☐ Users ☒ Groups

Domain: Entire Directory

Name/User name: Contains

Description: Contains

Find

Name	User Name	Email	Description	In Folder
Server Operators	Server Operators/		Members can adm	companyA.local/Bu
Standard Users	Standard Users/cc			companyA.local/VC
Terminal Server Li	Terminal Server Li		Members of this gr	companyA.local/Bu
Users	Users/companyA.l		Users are prevent	companyA.local/Bu
vCenter Admins	vCenter Admins/cc			companyA.local/Us
View Admins	View Admins/comp			companyA.local/Us

OK Cancel

Figure 9-77 Select the AD Group

22. The final result is shown in Figure 9-78. Click **OK**.

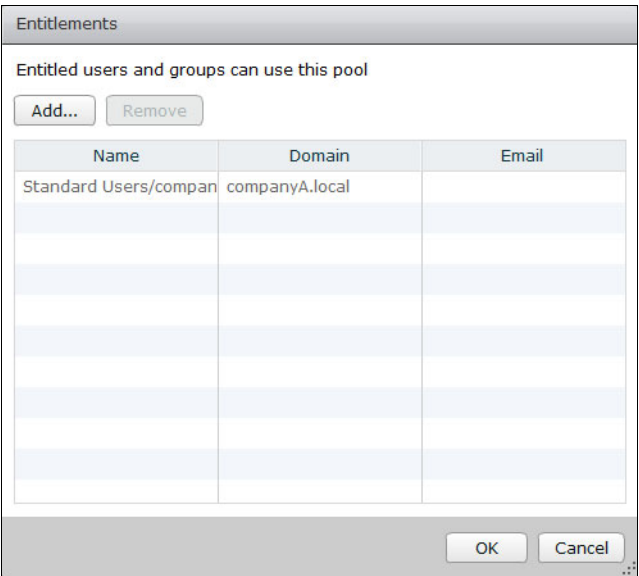


Figure 9-78 Entitlement complete window

The LCVN Desktop Pool is now created, as shown in Figure 9-79.

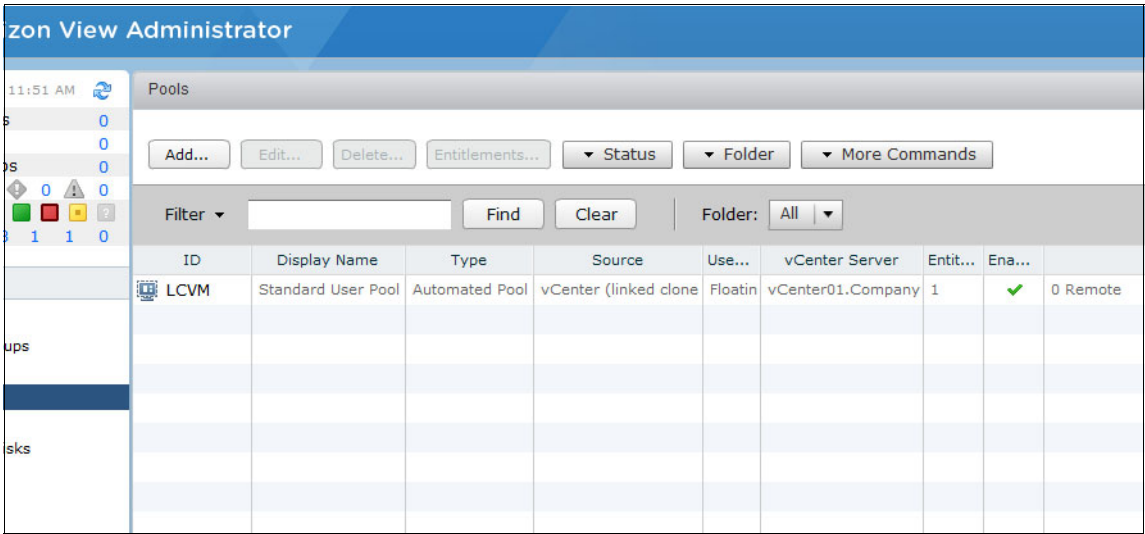


Figure 9-79 Pool view

Clicking **Desktops** in the left pane shows virtual desktops creation and customization options, as shown in Figure 9-80.

Inventory

Dashboard

Users and Groups

▼ Inventory

Pools

Desktops

Persistent Disks

ThinApps

► Monitoring

► Policies

► View Configuration

Filter

Find

Clear

Folder: All

Desktop	1 ▲	Pool	DNS Name	User	H.	Agent ...	Datastore	Mode
LCVM-10-STDUSR		LCVM			es	Unknown	ESXi04-local-01	Remote
LCVM-11-STDUSR		LCVM			es	Unknown	ESXi04-local-01	Remote
LCVM-12-STDUSR		LCVM			es	Unknown	ESXi03-local-01	Remote
LCVM-13-STDUSR		LCVM				Unknown		Remote
LCVM-14-STDUSR		LCVM				Unknown		Remote
LCVM-15-STDUSR		LCVM				Unknown		Remote
LCVM-16-STDUSR		LCVM				Unknown		Remote
LCVM-17-STDUSR		LCVM				Unknown		Remote
LCVM-18-STDUSR		LCVM				Unknown		Remote
LCVM-1-STDUSR		LCVM			es	Unknown	ESXi04-local-01	Remote
LCVM-2-STDUSR		LCVM			es	Unknown	ESXi03-local-01	Remote
LCVM-3-STDUSR		LCVM			es	Unknown	ESXi03-local-01	Remote
LCVM-4-STDUSR		LCVM			es	Unknown	ESXi03-local-01	Remote
LCVM-5-STDUSR		LCVM			es	Unknown	ESXi04-local-01	Remote
LCVM-6-STDUSR		LCVM			es	Unknown	ESXi04-local-01	Remote
LCVM-7-STDUSR		LCVM			es	Unknown	ESXi04-local-01	Remote
LCVM-8-STDUSR		LCVM			es	Unknown	ESXi03-local-01	Remote
LCVM-9-STDUSR		LCVM			es	Unknown	ESXi03-local-01	Remote

Figure 9-80 Desktop provisioning

**Firewall rules:** For more information about firewall rules, see this website:

<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1027217>

### 9.4.3 Provisioning a full virtual machine virtual desktop

**Important:** Before you provision a full virtual machine virtual desktop, follow the steps that are described in 9.1.4, “Provisioning a full VM image” on page 390.

Complete the following steps to provision a full virtual machine virtual desktop:

1. Connect with VSphere client to vCenter and shutdown the reference full virtual machine VM.
2. Right-click **FVM** and select **Template** → **Clone to Template**, as shown in Figure 9-81.

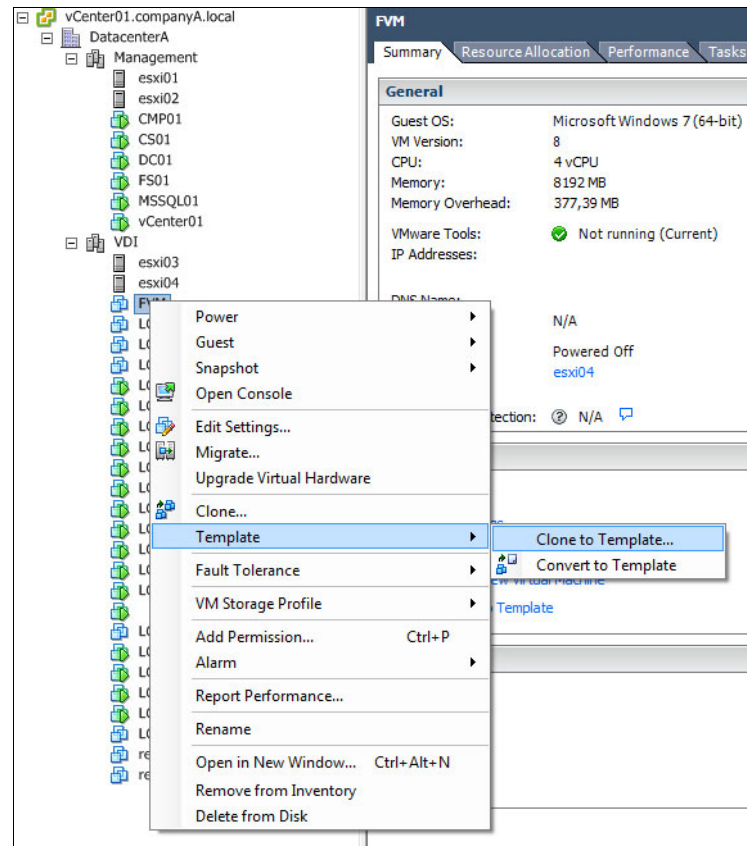


Figure 9-81 Clone to template task

3. The Clone Virtual Machine wizard opens. Enter the information in the Name and Location page, as shown in Figure 9-82. Click **Next**.

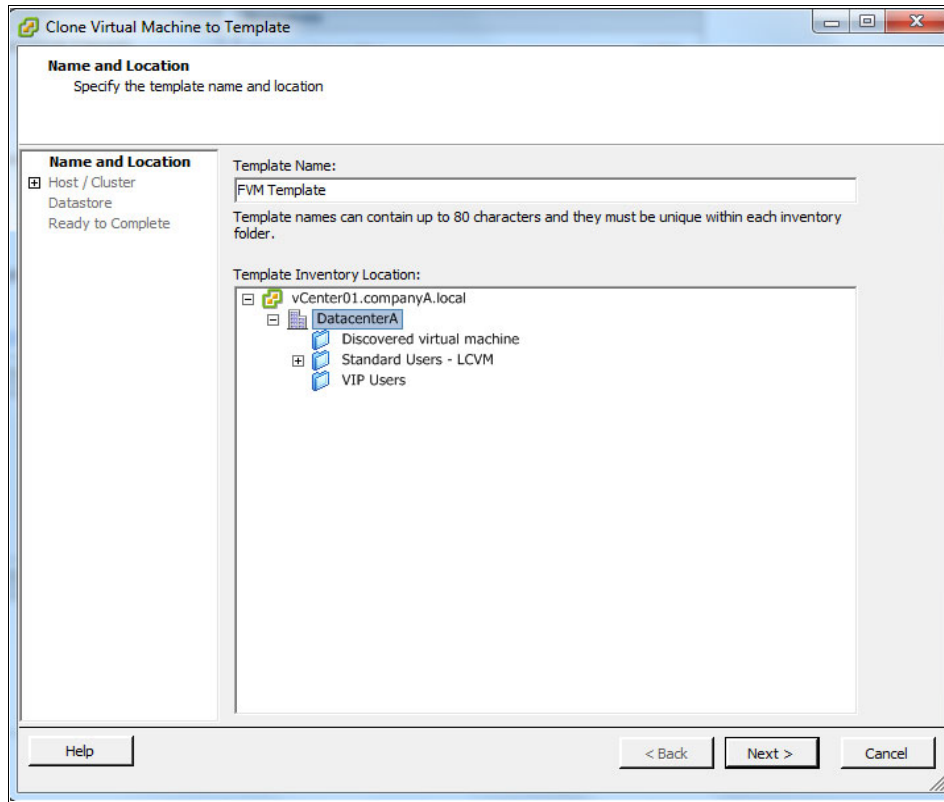


Figure 9-82 Enter a template name

4. Select the VDI cluster, as shown in Figure 9-83. Click **Next**.

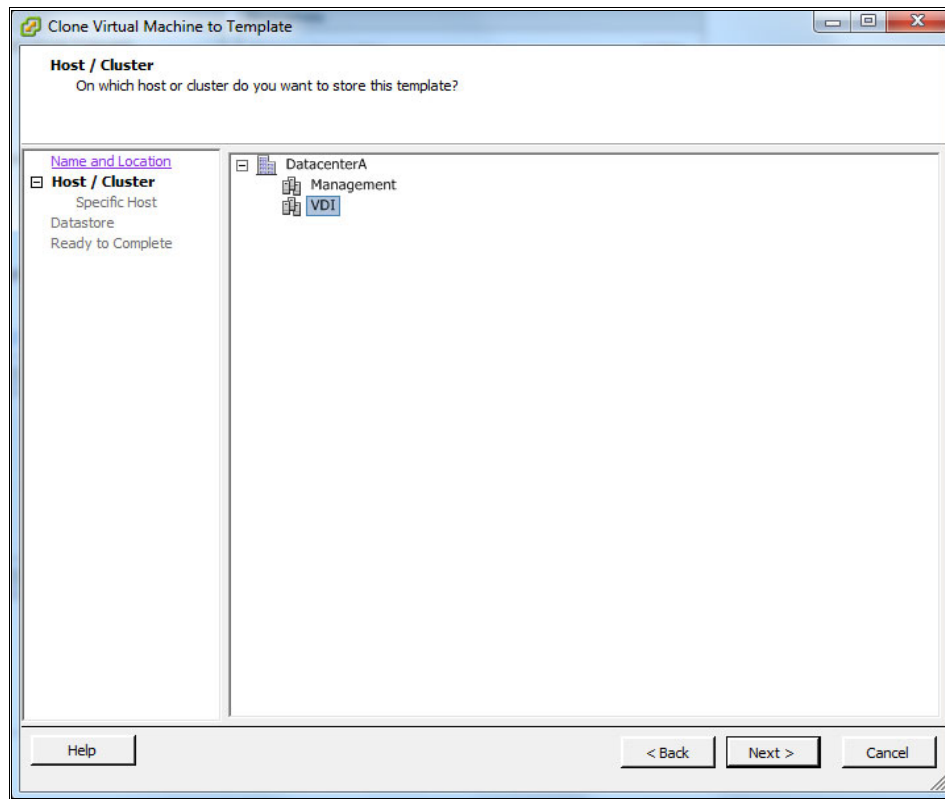


Figure 9-83 Select the VDI cluster

5. Select one of the two nodes for the VDI cluster, as shown in Figure 9-84. Click **Next**.

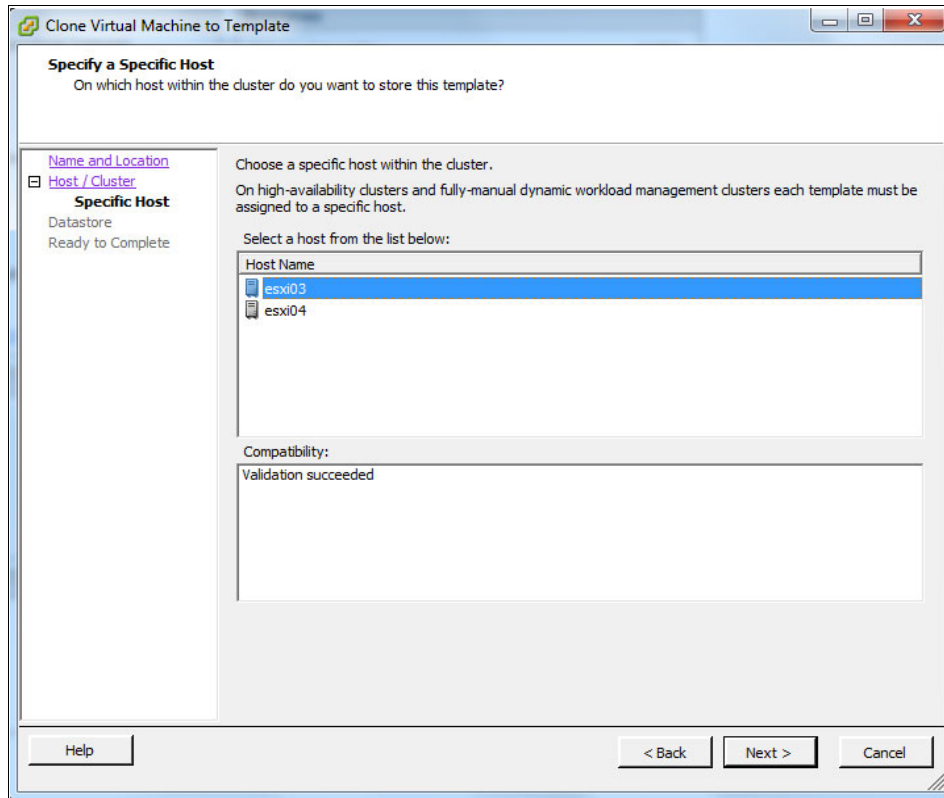


Figure 9-84 Select the node

6. Select the shared data store, as shown in Figure 9-85. Click **Next**.

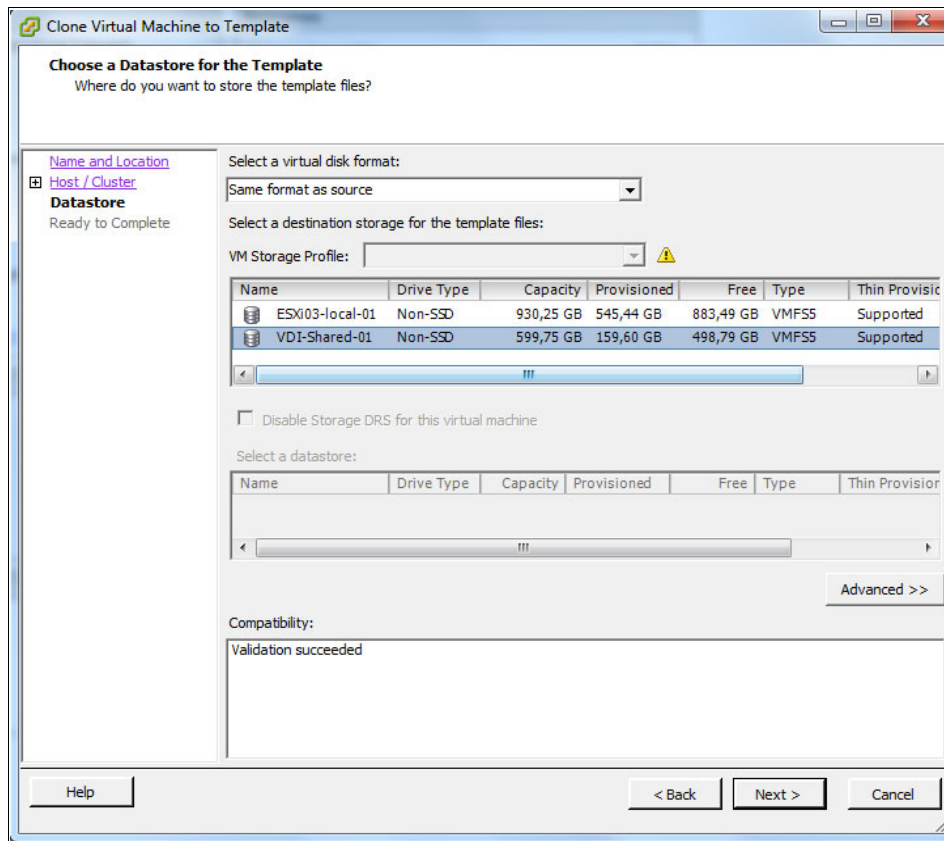


Figure 9-85 Select the shared data store



7. In the Ready to Complete window, click **Finish**, as shown in Figure 9-86.

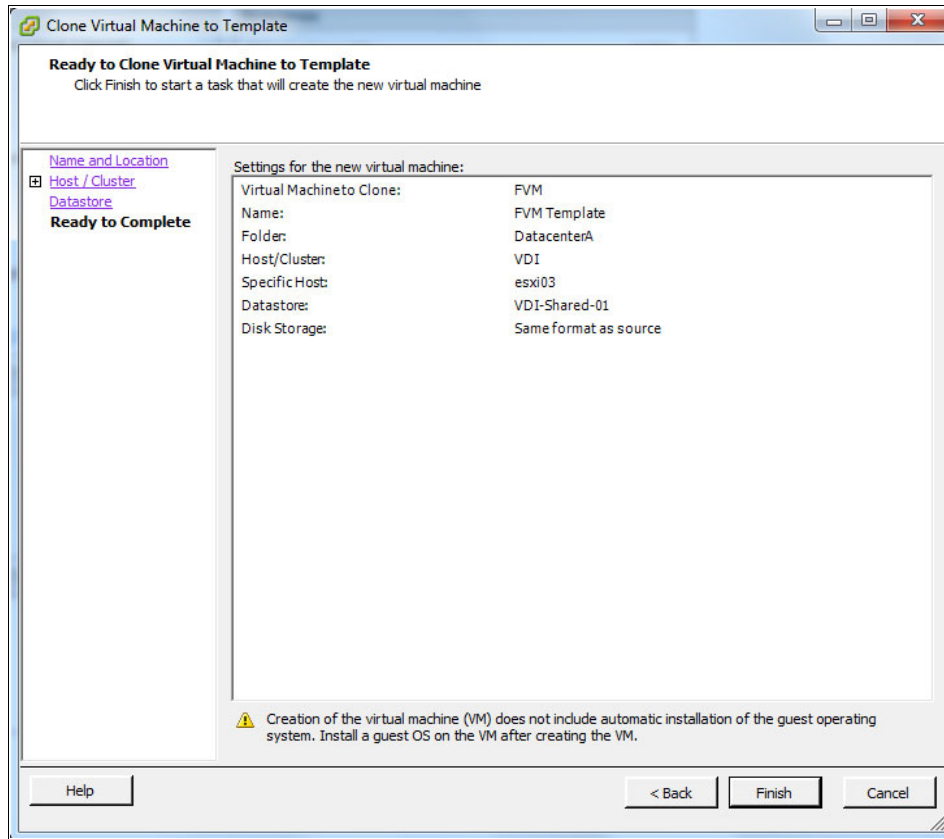


Figure 9-86 Summary page

8. Log on to the VMware Horizon View Administrator web console. Then, in the left pane, click **Inventory** → **Pools**, then click **Add**, as shown in Figure 9-87.

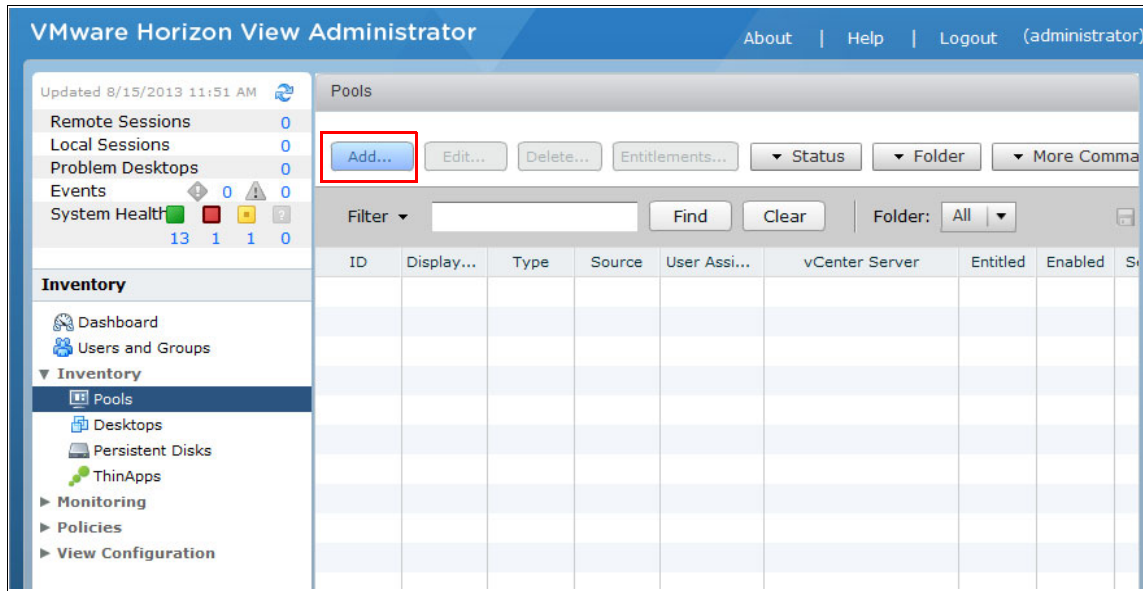


Figure 9-87 Pools main page

9. Select **Automated Pool** type, as shown in Figure 9-88. Click **Next**.

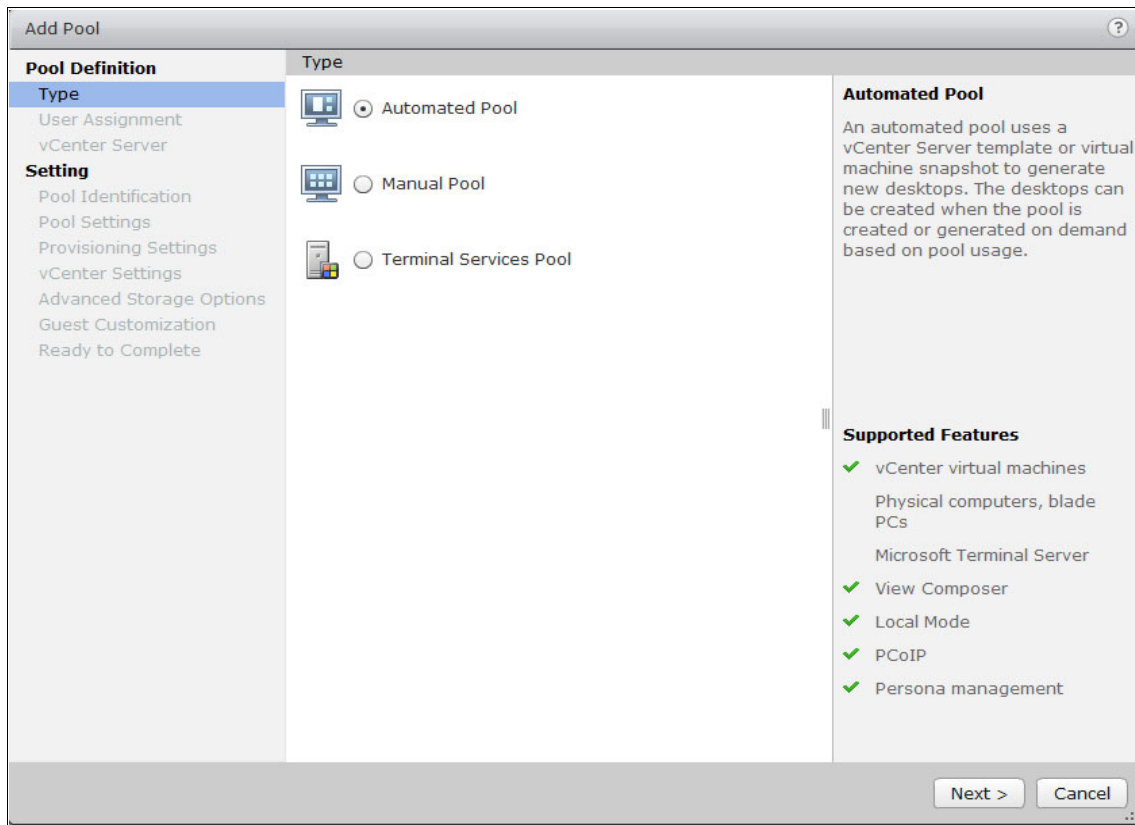


Figure 9-88 Select pool type

10. Select **Dedicated** as the user assignment, as shown in Figure 9-89. Click **Next**.

The screenshot shows the 'Add Pool' wizard with the 'User assignment' step selected. The left pane shows the 'Pool Definition' section with 'User Assignment' highlighted. The main area shows two options: 'Dedicated' (selected) and 'Floating'. The 'Dedicated' option has a checked box for 'Enable automatic assignment'. The right pane provides details for the 'Dedicated assignment' method, including a description, instructions on enabling automatic assignment, and a list of supported features.

**Add Pool**

**Pool Definition**

- Type
- User Assignment**
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**User assignment**

☒ **Dedicated**

☒ Enable automatic assignment

☐ Floating

**Dedicated assignment**

Users receive the same desktops each time they log into the pool.

**Enable automatic assignment**

If a user connects to a pool to which the user is entitled, but does not have a desktop, View automatically assigns a spare desktop to the user. In an automated pool, a new desktop may be created if no spare desktops exist.

If automatic assignment is not enabled, users must be assigned to desktops manually in View Administrator. Manual assignment can still be done even if automatic assignment is enabled.

**Supported Features**

- ✓ View Composer
- ✓ Local Mode
- ✓ PCoIP
- ✓ Persona management

< Back   Next >   Cancel

Figure 9-89 Select the user assignment

11. Select **Full virtual machines**, as shown in Figure 9-90. Click **Next**.

**Add Pool**

**Pool Definition**

- Type
- User Assignment
- vCenter Server**

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**vCenter Server**

☒ Full virtual machines  
☐ View Composer linked clones

vCenter Server	View Composer
vCenter01.CompanyA.local(Administrator)	vCenter01.CompanyA.local

Description:

**Full Virtual Machine**

Desktops sources will be full virtual machines that are created from a vCenter Server template.

**Supported Features**

- ✓ Local Mode
- ✓ PCoIP
- Storage savings
- Recompose and refresh
- QuickPrep guest customization
- ✓ Sysprep guest customization
- ✓ Persona management

< Back   Next >   Cancel

Figure 9-90 Select the virtual machine type

12. Complete the Pool Identification information as shown in Figure 9-91. Click **Next**.

Add Pool - FVM

Pool Definition

Type

User Assignment

vCenter Server

Setting

Pool Identification

Pool Settings

Provisioning Settings

vCenter Settings

Advanced Storage Options

Guest Customization

Ready to Complete

Pool Identification

ID:

FVM

Display name:

VIP Users Pool

View folder:

/

Description:

ID

The pool ID is the unique name used to identify this pool.

Display Name

The display name is the name that users will see when they connect to View Client. If the display name is left blank, the ID will be used.

View Folder

View folders can organize the pools in your organization. They can also be used for delegated administration.

Description

This description is only shown on the Settings tab for a pool within View Administrator.

< Back

Next >

Cancel

*Figure 9-91 Pool Identification information*

13. Complete the Pool Settings information as shown in Figure 9-92. Click **Next**.

**Add Pool - FVM**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings**
- Provisioning Settings
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**Pool Settings**

**General**

State:

Connection Server restrictions:  

**Remote Settings**

Remote Desktop Power Policy:  ?

Automatically logoff after disconnect:

Allow users to reset their desktops:

**Remote Display Protocol**

Default display protocol:

Allow users to choose protocol:

3D Renderer:   ?

Max number of monitors:  ?

May require power-cycle of related virtual machines ?

Max resolution of any one monitor:  ?

May require power-cycle of related virtual machines ?

HTML Access: ☐ Enabled ?

Requires installation of the HTML Desktop Access feature pack.

**Adobe Flash Settings for Remote Sessions**

Adobe Flash quality:  ?

Adobe Flash throttling:  ?

< Back   Next >   Cancel

Figure 9-92 Pool Settings information

14. Complete the Provisioning Settings information as shown in Figure 9-93.  
Click **Next**.

The screenshot shows the 'Add Pool - FVM' wizard with the 'Provisioning Settings' tab selected. The left sidebar contains a tree view with the following items: 'Pool Definition' (expanded), 'Type', 'User Assignment', 'vCenter Server', 'Setting' (expanded), 'Pool Identification', 'Pool Settings', 'Provisioning Settings' (selected), 'vCenter Settings', 'Advanced Storage Options', 'Guest Customization', and 'Ready to Complete'. The main content area is divided into three sections: 'Basic', 'Virtual Machine Naming', and 'Pool Sizing'. The 'Basic' section has two checked options: 'Enable provisioning' and 'Stop provisioning on error'. The 'Virtual Machine Naming' section has two radio buttons: 'Specify names manually' (unselected) and 'Use a naming pattern' (selected). The 'Specify names manually' option has a text box showing '0 names entered' and an 'Enter names...' button. The 'Use a naming pattern' option has a text box showing 'FVM-{n}-VIP'. Below this is a section for 'Pool Sizing' with two text boxes: 'Max number of desktops:' (20) and 'Number of spare (powered on) desktops:' (3). At the bottom of the main content area is a section for 'Provisioning Timing' with two radio buttons: 'Provision desktops on demand' (unselected) and 'Provision all desktops up-front' (selected). The 'Provision all desktops up-front' option has a text box showing '1'. On the right side of the main content area is a 'Naming Pattern' section with explanatory text. At the bottom of the wizard are three buttons: '< Back', 'Next >', and 'Cancel'.

**Add Pool - FVM**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings**
- vCenter Settings
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**Provisioning Settings**

**Basic**

- ☒ Enable provisioning
- ☒ Stop provisioning on error

**Virtual Machine Naming**

- ☐ Specify names manually
  - 0 names entered
  - Enter names...
- ☐ Start desktops in maintenance mode
  - # Unassigned desktops kept powered on: 1
- ☒ Use a naming pattern
  - Naming Pattern: FVM-{n}-VIP

**Pool Sizing**

- Max number of desktops: 20
- Number of spare (powered on) desktops: 3

**Provisioning Timing**

- ☐ Provision desktops on demand
  - Min number of desktops: 1
- ☒ Provision all desktops up-front

**Naming Pattern**

Virtual machines will be named according to the specified naming pattern. By default, View Manager appends a unique number to the specified pattern to provide a unique name for each virtual machine.

To place this unique number elsewhere in the pattern, use '{n}'. (For example: vm-{n}-sales.).

The unique number can also be made a fixed length. (For example: vm-{n:fixed=3}-sales.).

See the help for more naming pattern syntax options.

< Back   Next >   Cancel

Figure 9-93 Provisioning Settings information



15. Use the vCenter Settings page (as shown in Figure 9-94) to set several vCenter settings. To change the settings, click **Browse** to the right of the setting that you want to change.

The screenshot shows the 'Add Pool - FVM' wizard with the 'vCenter Settings' page selected in the left-hand navigation pane. The main content area is divided into three sections: 'Virtual Machine Template', 'Virtual Machine Location', and 'Resource Settings'. Each section contains a numbered step (1-5) and a 'Browse...' button to the right of a text field. The text fields contain the placeholder '<Click Browse...>'. The 'Resource Settings' section has a text label 'Click Browse to select' instead of a text field for the 'Datastores' step. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

**Add Pool - FVM**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- vCenter Settings**
- Advanced Storage Options
- Guest Customization
- Ready to Complete

**vCenter Settings**

**Virtual Machine Template**

1 Template: <Click Browse...> **Browse...**

**Virtual Machine Location**

2 VM folder location: <Click Browse...> **Browse...**

**Resource Settings**

3 Host or cluster: <Click Browse...> **Browse...**

4 Resource pool: <Click Browse...> **Browse...**

5 Datastores: Click Browse to select **Browse...**

< Back Next > Cancel

Figure 9-94 vCenter Settings main page

When you are changing vCenter settings, make the following changes:

- a. For the Template setting, select the template, as shown in Figure 9-95.  
Click **OK**.

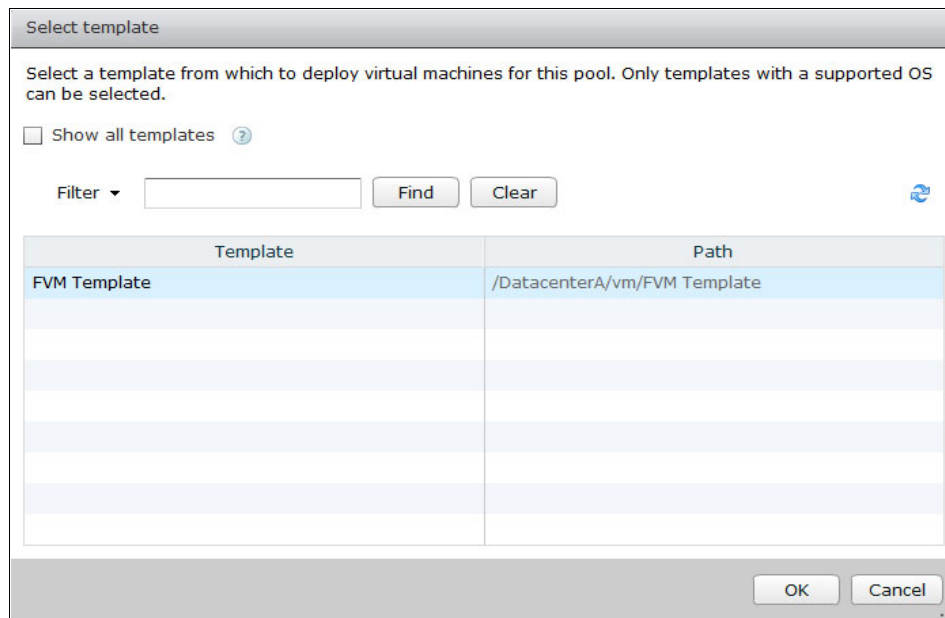


Figure 9-95 Select the template

- b. For the VM folder location, complete the information as shown in Figure 9-96. Click **OK**.


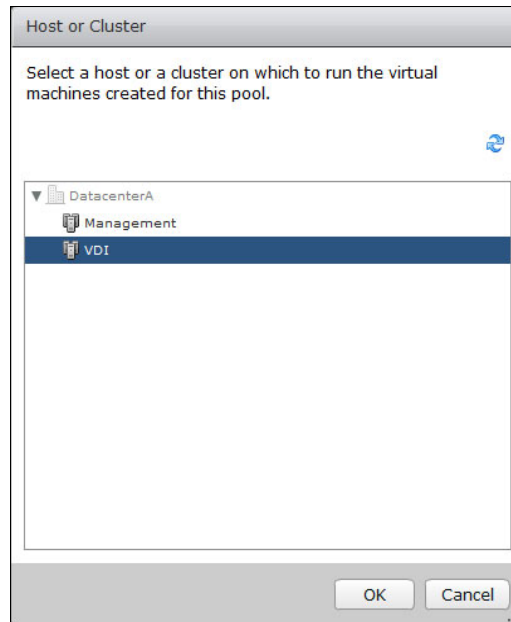


Figure 9-96 VM folder location settings

- c. For the Host or cluster settings, complete the information as shown in Figure 9-97. Click **OK**.



*Figure 9-97 Select the cluster*

- d. For the Resource pool settings, complete the information as shown in Figure 9-98. Click **OK**.

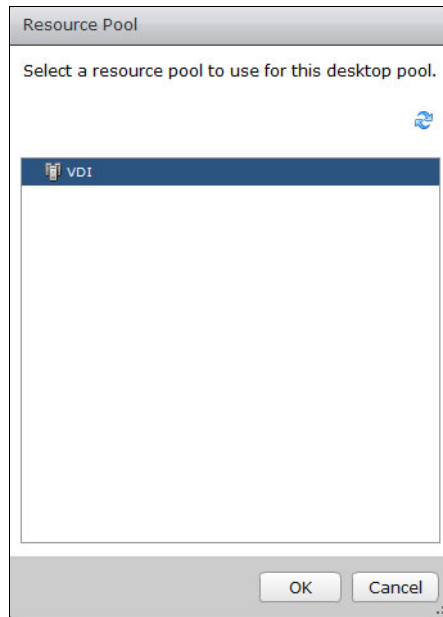


Figure 9-98 Select the resource pool

- e. For the Datastores settings, complete the information that is shown in Figure 9-99. Click **OK**.

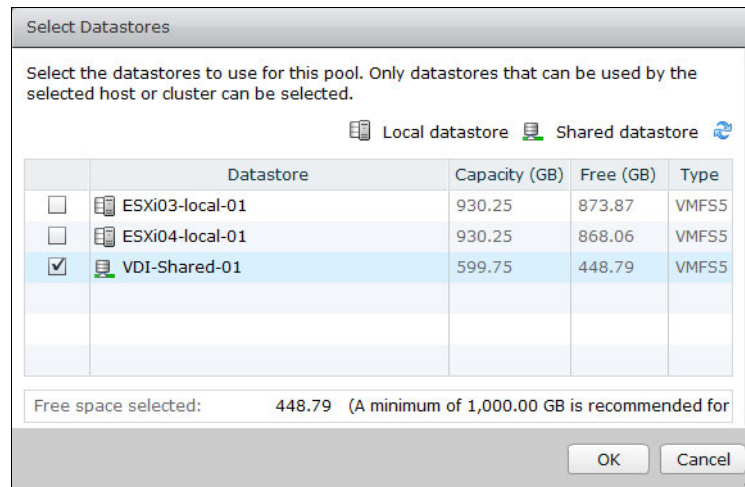


Figure 9-99 Selecting the data stores

16. Return to the vCenter settings main page and then click **Next** to continue with the wizard. In the Advanced Storage Options window (as shown in Figure 9-100), accept the default options by clicking **Next**.

The screenshot shows the 'Add Pool - FVM' wizard window. The left sidebar contains a tree view with the following items: 'Pool Definition' (expanded), 'Type', 'User Assignment', 'vCenter Server', 'Setting' (expanded), 'Pool Identification', 'Pool Settings', 'Provisioning Settings', 'vCenter Settings', 'Advanced Storage Options' (selected), 'Guest Customization', and 'Ready to Complete'. The main area is titled 'Advanced Storage Options' and contains the following content:

Based on your resource selection, the following features are recommended. Options that are not supported by the selected hardware are disabled.

☒ Use View Storage Accelerator

Regenerate storage accelerator after:  Days

**Blackout Times**

Storage accelerator regeneration and VM disk space reclamation do not occur during blackout times. The same blackout policy applies to both operations.

Buttons: Add..., Edit..., Remove

Day	Time

On the right side, there is a section titled 'View Storage Accelerator' with the following text: 'vSphere 5.x hosts can be configured to improve performance by caching certain pool data. Enable this option to use View Storage Accelerator for this pool. View Storage Accelerator is most useful for shared disks that are read frequently, such as View Composer OS disks.'

At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Figure 9-100 Advanced storage options

17. For the Guest Customization, select **Use this Customization Specification** and then select the customization that you created in 9.1.1, “Creating a customization specification file” on page 374, as shown in Figure 9-101. Click **Next**.

**Add Pool - FVM**

**Pool Definition**

- Type
- User Assignment
- vCenter Server

**Setting**

- Pool Identification
- Pool Settings
- Provisioning Settings
- vCenter Settings
- Advanced Storage Options
- Guest Customization**

Ready to Complete

**Guest Customization**

☐ None - Customization will be done manually

☐ Do not power on virtual machines after creation

☒ Use this customization specification:

Name	Guest OS	Description
CompanyA	Windows	
Windows7_Dom	Windows	Customized Wizard for automatic AD Domain and network spe

< Back   Next >   Cancel

Figure 9-101 Select the customization specification

18. In the Ready to Complete window, click **Entitle users after this wizard finishes**, as shown in Figure 9-102. Click **Finish**.

**Add Pool - FVM**

**Ready to Complete**

☒ Entitle users after this wizard finishes

Type:	Automated
User assignment:	Dedicated assignment
Assign on first login:	Yes
vCenter Server:	vCenter01.CompanyA.local(Administrator)
Use View Composer:	No
Unique ID:	FVM
Display name:	VIP Users Pool
View Folder:	/
Desktop pool state:	Enabled
Remote Desktop Power Policy:	Take no power action
Automatic logoff after disconnect:	Never
Connection Server restrictions:	None
Allow users to reset their desktop:	No
Default display protocol:	PCoIP
Allow users to choose protocol:	No
3D Renderer:	Automatic
VRAM Size:	512 MB
Max number of monitors:	2
Max resolution of any one monitor:	1920x1200
HTML Access:	Disabled
Adobe Flash quality:	Disabled
Enable provisioning:	Yes
Stop provisioning on error:	Yes
Virtual Machine Naming:	Use a naming pattern
VM naming pattern:	FVM-{n}-VIP
Provision all desktops up-	Yes

< Back   Finish   Cancel

Figure 9-102 Summary page



19. In the entitlement window (see Figure 9-103), click **Add**.

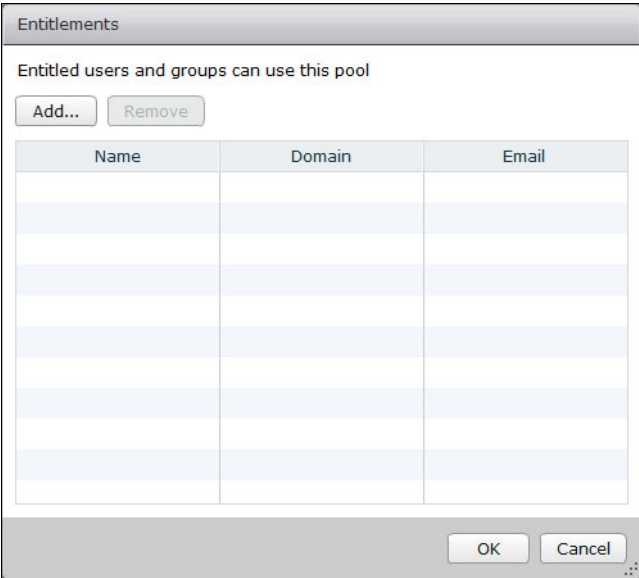


Figure 9-103 Entitlement page

20. Select the VIP Users group, as shown in Figure 9-104. Click **OK**.

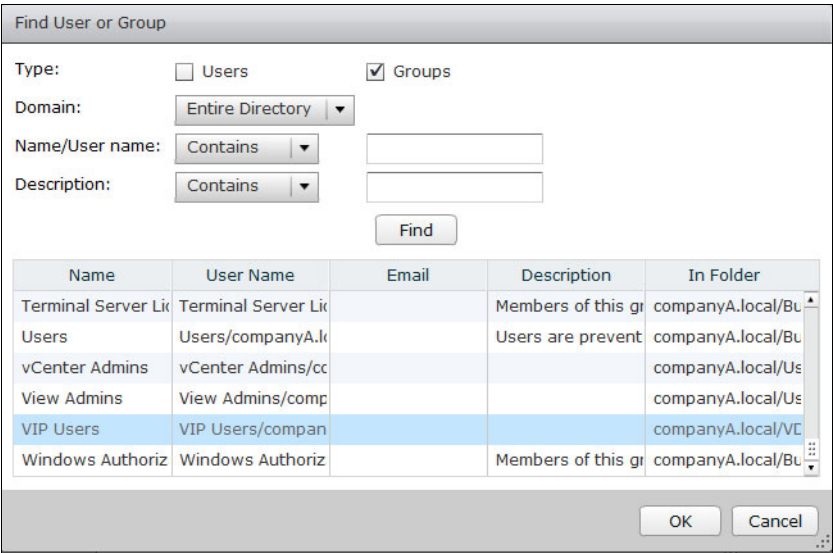


Figure 9-104 Desktop pool group entitlements

21. Figure 9-105 shows the final group configuration. Click **OK** to return to VMware Horizon View Administrator.

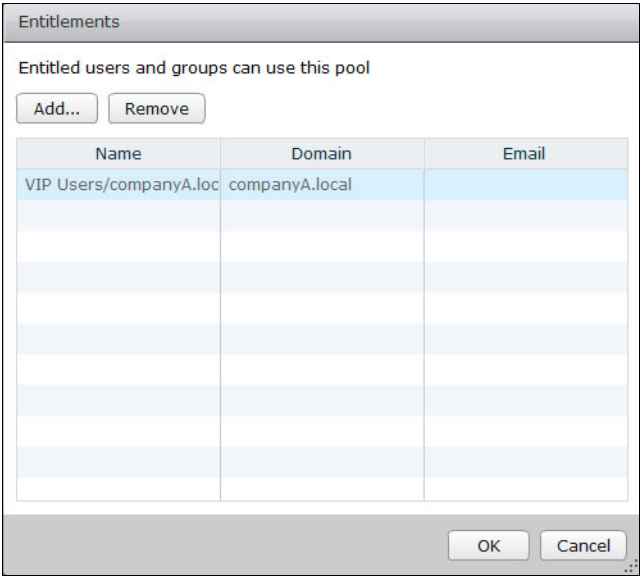


Figure 9-105 Entitlement complete

To connect to the entitled desktop, you must use VMware Horizon View client.

**Firewall rules:** For more information about firewall rules, see this website:  
<http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1027217>

## 9.5 Operating View Composer

To maintain service quality and compliance in virtual desktops environment, an administrator must complete some basic tasks. As users use their virtual desktops, the virtual desktop size tends to increase. In addition, as times goes by, new operating system updates or applications must be installed or replaced on the base operating system image.

To efficiently maintain the linked clones virtual desktop environment, an administrator uses the following tasks most frequently:

- ▶ Desktop refresh
- ▶ Desktop recompose
- ▶ Desktop rebalance

**Important:** Do not use vCenter to migrate virtual desktops. Use the rebalance feature instead.

Full virtual machines, as with normal desktops, offer less flexibility than linked clones.

### 9.5.1 Performing a desktop refresh operation

A desktop *refresh* operation resets all of the virtual desktops in a pool to their original state by reapplying the snapshot image of the parent VM. This option completely resets the desktops to their factory settings and restores the operating system disk of each linked clone to its original state and size. It reduces storage costs and can be used regularly to improve system responsiveness.

**Important:** Performing a desktop refresh task disconnects all connected users, but you can schedule this operation to occur overnight.

Complete the following steps to perform a desktop refresh:

1. Connect to VMware Horizon View Administrator web interface. Then, click **Inventory** in the left pane and select **Pools**. Double-click the **LCVM** desktop pool to display the details of this desktop pool, as shown in Figure 9-106.

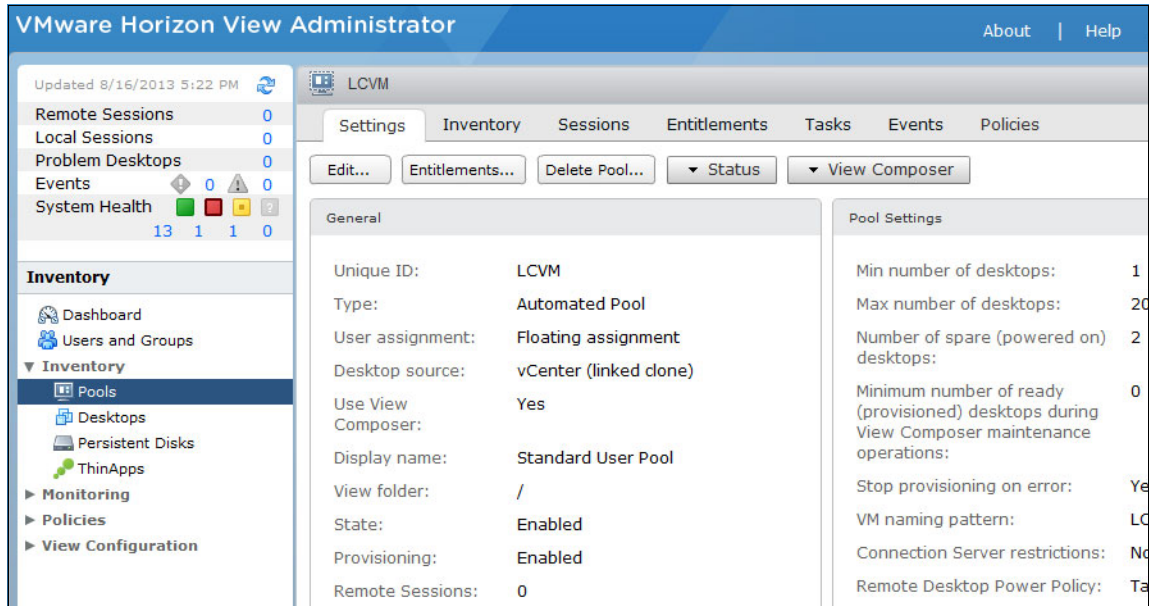


Figure 9-106 Desktop pool details

- To refresh the entire pool, on the Settings tab, click **View Composer** → **Refresh**, as shown in Figure 9-107.

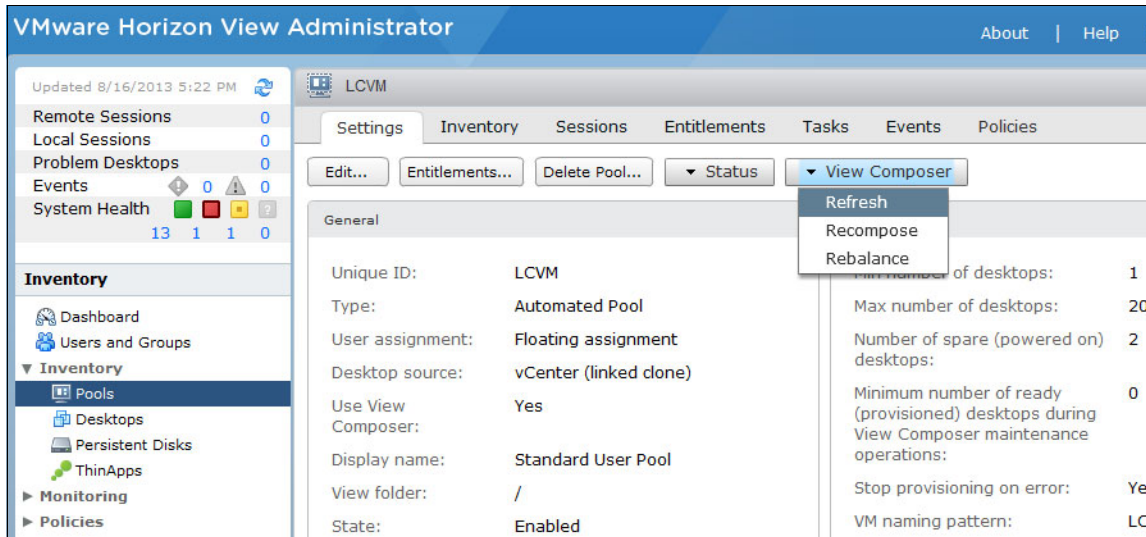


Figure 9-107 Refresh operation for the entire pool

- To refresh a single virtual desktop, go to the Inventory tab, click the virtual desktop to refresh, and then click **View Composer** → **Refresh**, as shown in Figure 9-108.

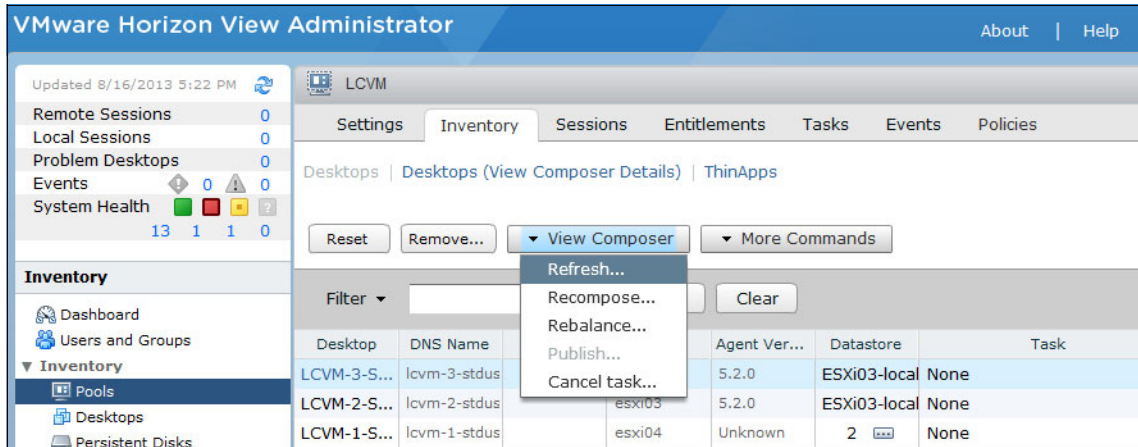
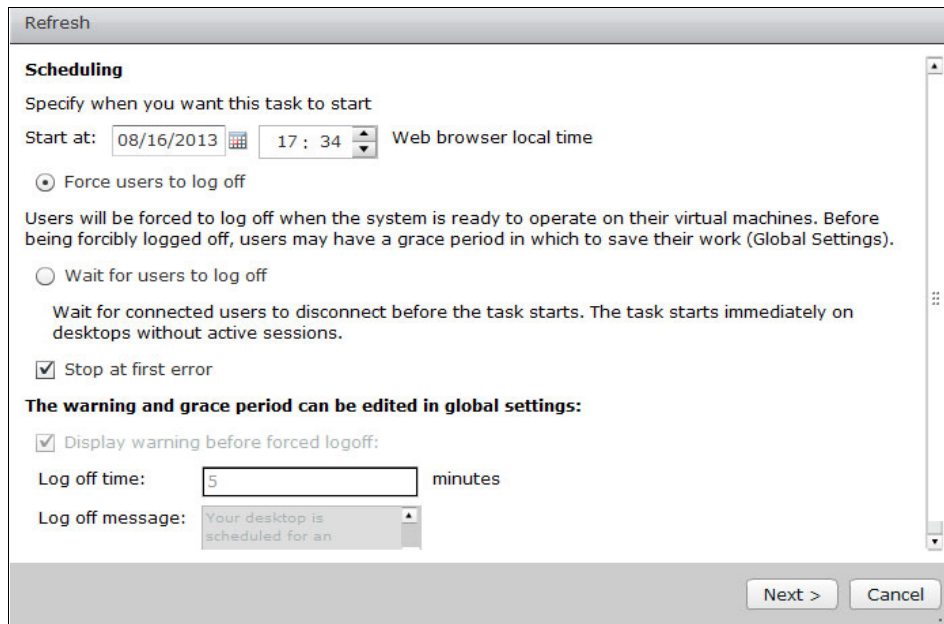


Figure 9-108 Refresh a single desktop

4. A refresh wizard opens, as shown in Figure 9-109. Select the option to force users to log off or to wait for users to log off. Click **Next**.



The screenshot shows a 'Refresh' wizard window with a 'Scheduling' section. It includes a 'Start at' field with a date of 08/16/2013 and a time of 17:34, with a 'Web browser local time' label. There are two radio button options: 'Force users to log off' (selected) and 'Wait for users to log off'. Below these, there is a checkbox for 'Stop at first error' which is checked. A note states 'The warning and grace period can be edited in global settings:'. Another checkbox 'Display warning before forced logoff:' is checked. The 'Log off time' is set to 5 minutes. The 'Log off message' field contains the text 'Your desktop is scheduled for an'. At the bottom right, there are 'Next >' and 'Cancel' buttons.

Refresh

**Scheduling**

Specify when you want this task to start

Start at: 08/16/2013 17:34 Web browser local time

☒ Force users to log off

Users will be forced to log off when the system is ready to operate on their virtual machines. Before being forcibly logged off, users may have a grace period in which to save their work (Global Settings).

☐ Wait for users to log off

Wait for connected users to disconnect before the task starts. The task starts immediately on desktops without active sessions.

☒ Stop at first error

**The warning and grace period can be edited in global settings:**

☒ Display warning before forced logoff:

Log off time: 5 minutes

Log off message: Your desktop is scheduled for an

Next > Cancel

Figure 9-109 Select log off options

5. At the Ready to Complete window (see Figure 9-110), click **Finish**.

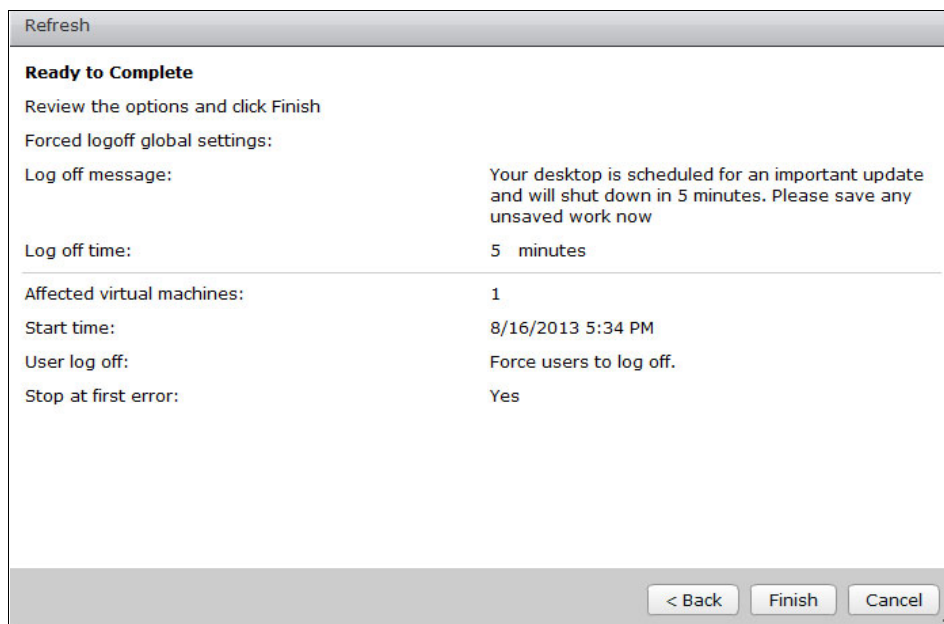
A screenshot of the 'Ready to Complete' window in VMware Horizon View. The window has a title bar with a 'Refresh' button. Below the title bar, the text 'Ready to Complete' is followed by 'Review the options and click Finish'. The 'Forced logoff global settings:' section contains the following details: 'Log off message:' with the text 'Your desktop is scheduled for an important update and will shut down in 5 minutes. Please save any unsaved work now'; 'Log off time:' set to '5 minutes'; 'Affected virtual machines:' set to '1'; 'Start time:' set to '8/16/2013 5:34 PM'; 'User log off:' set to 'Force users to log off.'; and 'Stop at first error:' set to 'Yes'. At the bottom right, there are three buttons: '< Back', 'Finish', and 'Cancel'.

Figure 9-110 Ready to Complete window

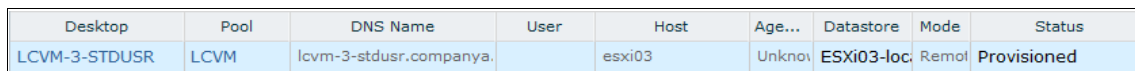
From the Desktops view, the status changes to In progress, as shown in Figure 9-111.

A screenshot of the 'Desktops' view in VMware Horizon View. It shows a table with two rows. The first row is for 'LCVM-3-STDUSR' in the 'LCVM' pool, with status 'In progress'. The second row is for 'FVM-2-VIP' in the 'FVM' pool, with status 'Available'.

Desktop	Pool	DNS Name	User	Host	Age...	Datastore	Mode	Status
LCVM-3-STDUSR	LCVM	lcvm-3-stdusr.companya.		esxi03	5.2.0	ESXi03-loc	Remot	In progress
FVM-2-VIP	FVM	fvm-2-vip		esxi04	5.2.0	VDI-Share	Remot	Available

Figure 9-111 In progress status

6. When the operation completes, the status changes to Provisioned, as shown in Figure 9-112.

A screenshot of the 'Desktops' view in VMware Horizon View, showing the status after the operation is complete. The table now shows the 'LCVM-3-STDUSR' desktop with a status of 'Provisioned'.

Desktop	Pool	DNS Name	User	Host	Age...	Datastore	Mode	Status
LCVM-3-STDUSR	LCVM	lcvm-3-stdusr.companya.		esxi03	Unknov	ESXi03-loc	Remot	Provisioned

Figure 9-112 Desktop provisioned

## 9.5.2 Performing a desktop recompose operation

After the parent VM snapshot is updated with new applications or an operating system patch, the *recompose* task allows you to re-create all of the linked clone desktops in a desktop pool from the updated snapshot. Its use depends on specific needs and from the dynamic nature of the infrastructure. Before you recompose a linked clone desktop pool, you must update the parent virtual machine that is used as a base image for the linked clones by installing operating system patch, new applications, and so on. After all of changes are complete, switch off the virtual machine and take a new snapshot, as described in step 3 on page 417.

Complete the following steps to perform a desktop recompose:

1. Connect to VMware Horizon View Administrator web interface. Then, click **Inventory** in the left pane and select **Pools**. Double-click the **LCVM** desktop to display the details of that pool, as shown in Figure 9-113.

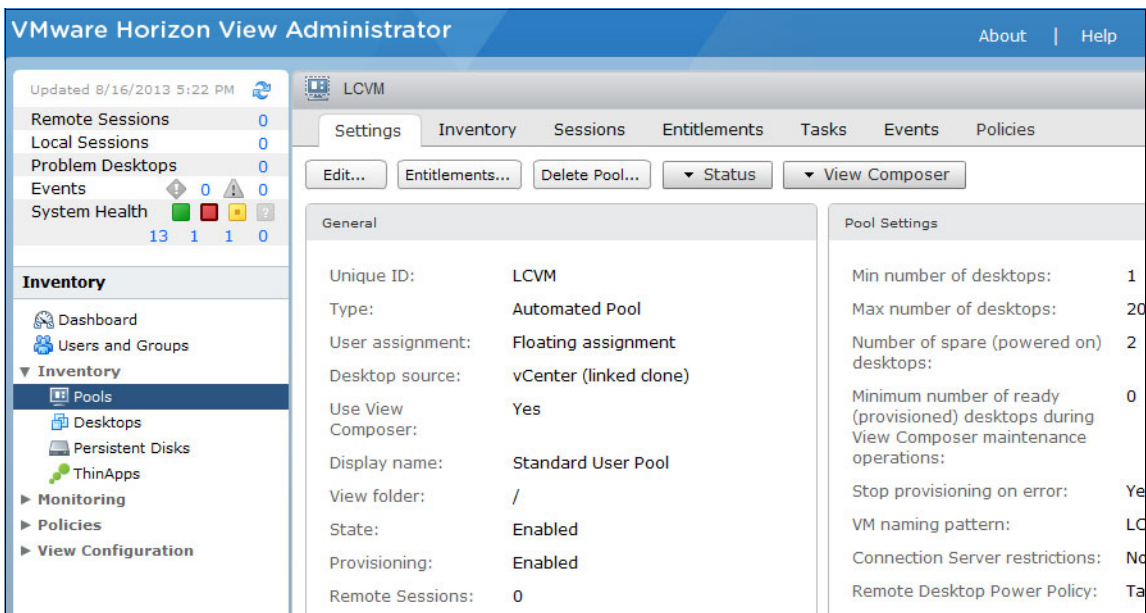


Figure 9-113 Desktop pool details



- To recompose the entire pool, on the Settings tab, click **View Composer** → **Recompose**, as shown in Figure 9-114.

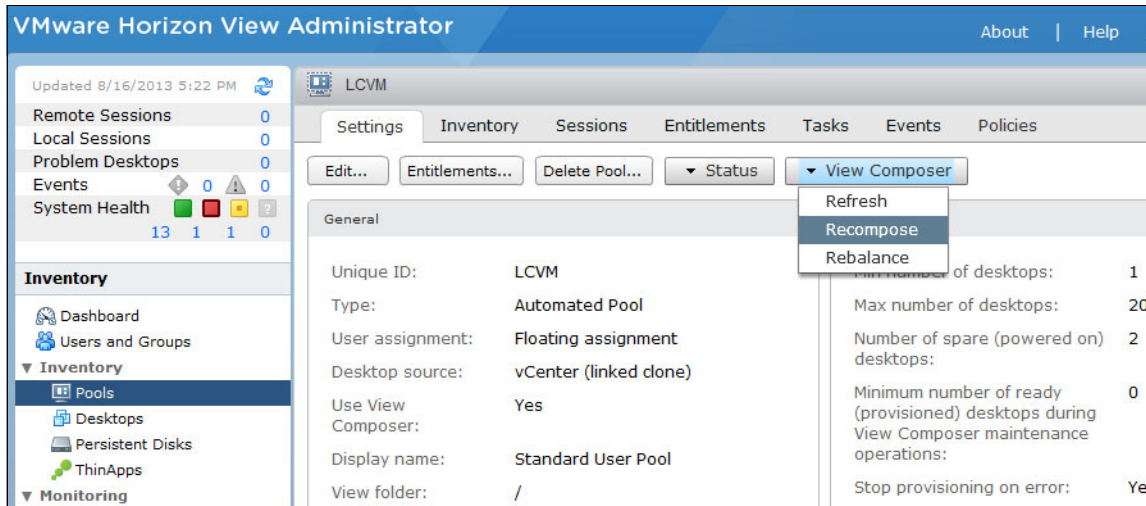


Figure 9-114 Recompose operation for the entire pool

- To recompose a single virtual desktop, go to the Inventory tab, hold down the Ctrl key and click each virtual desktop that you want to recompose, and then click **View Composer** → **Recompose**, as shown in Figure 9-115.

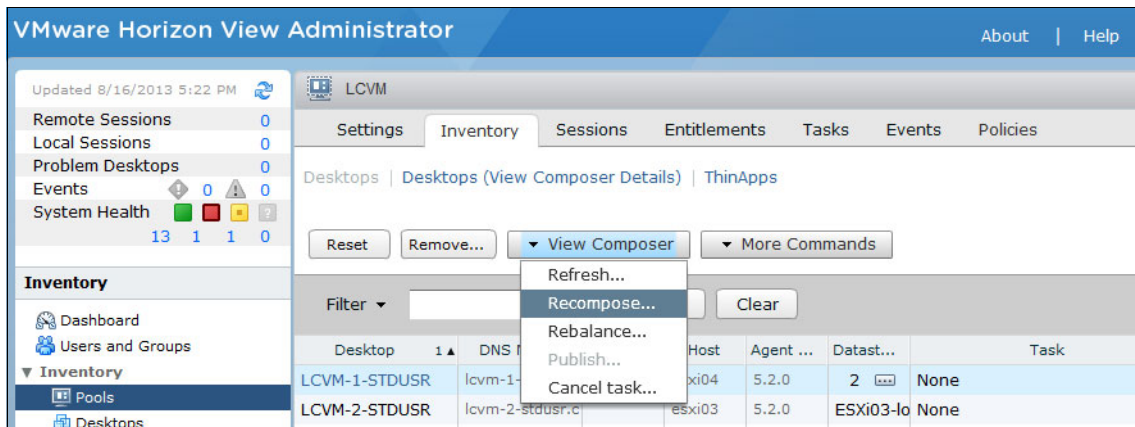


Figure 9-115 Recompose operation on selected virtual desktop

4. A recompose wizard opens. Select the new snapshot, as shown in Figure 9-116. Click **Next**.

Recompose

Image

Performing a recompose operation will cause sysprep to be re-run on the selected desktops, potentially resulting in changes to aspects of each machine's identity (e.g. SID or third party GUID).

Select the snapshot that will be used as the image. This snapshot can be on the current parent VM or a different one.

The desktops created in this pool will use the information in the image as their baseline system configuration.

Parent VM: 

/DatacenterA/vm/LCVM

Change...

Snapshot:

Snapshot Details

Snapshot	Time created	Description	Path	Published
LCVMv2	8/15/2013 9:08:00		/LCVMv2	No
LCVMv3	8/16/2013 9:08:00		/LCVMv2/LCVMv3	No

Next >

Cancel

Figure 9-116 Select a new snapshot

5. Select the option to force users to log off or to wait for users to log off, as shown in Figure 9-117. Click **Next**.

**Recompose**

**Scheduling**

Specify when you want this task to start

Start at: 08/16/2013 18 : 10 Web browser local time

☒ Force users to log off

Users will be forced to log off when the system is ready to operate on their virtual machines. Before being forcibly logged off, users may have a grace period in which to save their work (Global Settings).

☐ Wait for users to log off

Wait for connected users to disconnect before the task starts. The task starts immediately on desktops without active sessions.

☒ Stop at first error

**The warning and grace period can be edited in global settings:**

☒ Display warning before forced logoff:

Log off time: 5 minutes

Log off message: Your desktop is scheduled for an

< Back Next > Cancel

Figure 9-117 Scheduling options for recompose

6. At the Ready to Complete window (see Figure 9-118), click **Finish**.

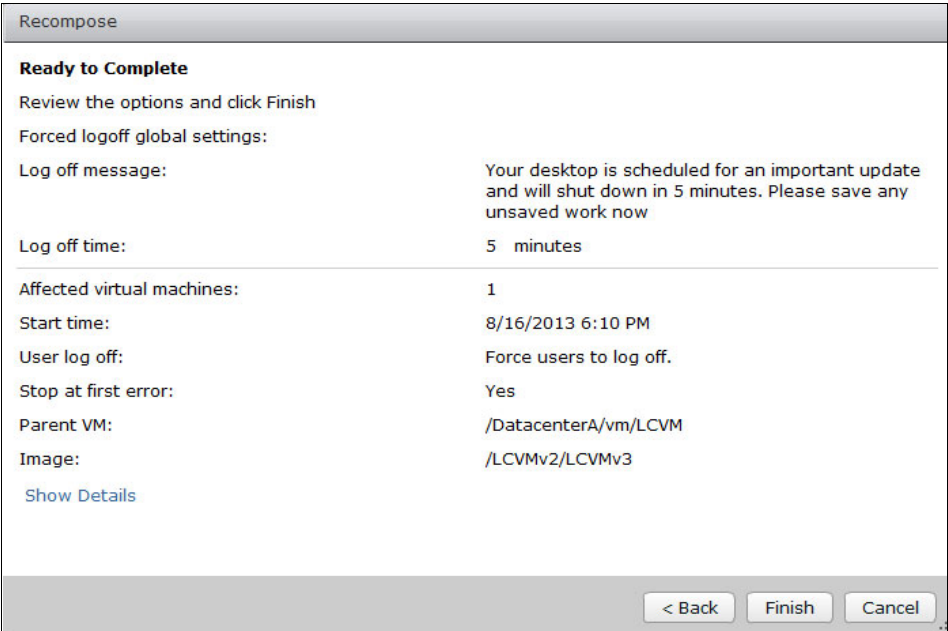


Figure 9-118 Ready to Complete

7. You can monitor the operation by clicking **Pool**, selecting the LCVM pool, and then going to the Tasks tab, as shown in Figure 9-119.

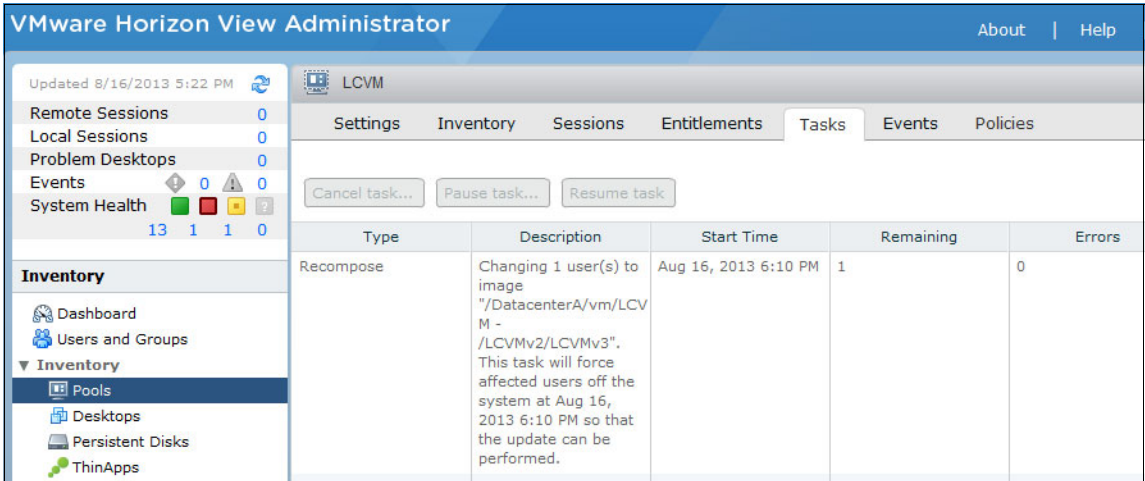


Figure 9-119 Recompose operation monitoring

### 9.5.3 Performing a desktop rebalance operation

A rebalance operation redistributes linked clone desktops among configured data stores and migrates virtual desktops to another data store. When you create large linked clone desktop pools and use multiple logical unit numbers (LUNs), the space might not be used efficiently. If an aggressive storage overcommit level is set, the linked clones can grow quickly and use all the free space on the data stores.

Desktop rebalance also refreshes the linked clones, which reduces the size of their OS disks. You can use the rebalance to migrate linked clone desktops to another data store.

Complete the following steps to rebalance virtual desktops:

1. Connect to VMware Horizon View Administrator web interface. Then, click **Inventory** in the left pane and select **Pools**. Double-click the **LCVM** desktop pool to display the details of this desktop, as shown in Figure 9-120.

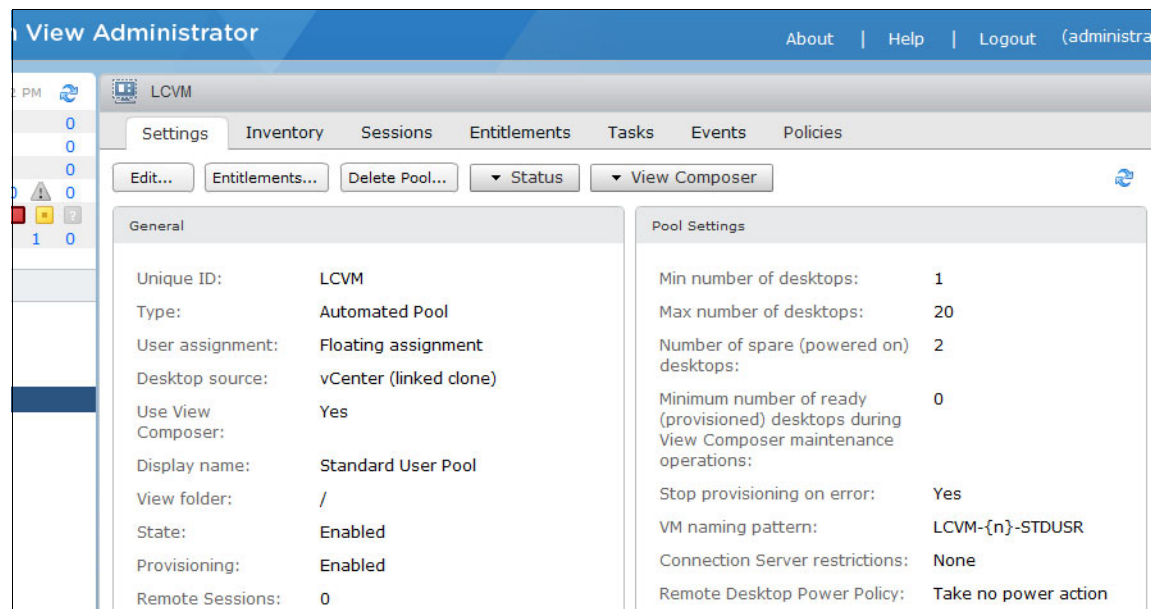


Figure 9-120 Desktop pool details

- To rebalance the entire pool, on the Settings tab, click **View Composer** → **Rebalance**, as shown in Figure 9-121.

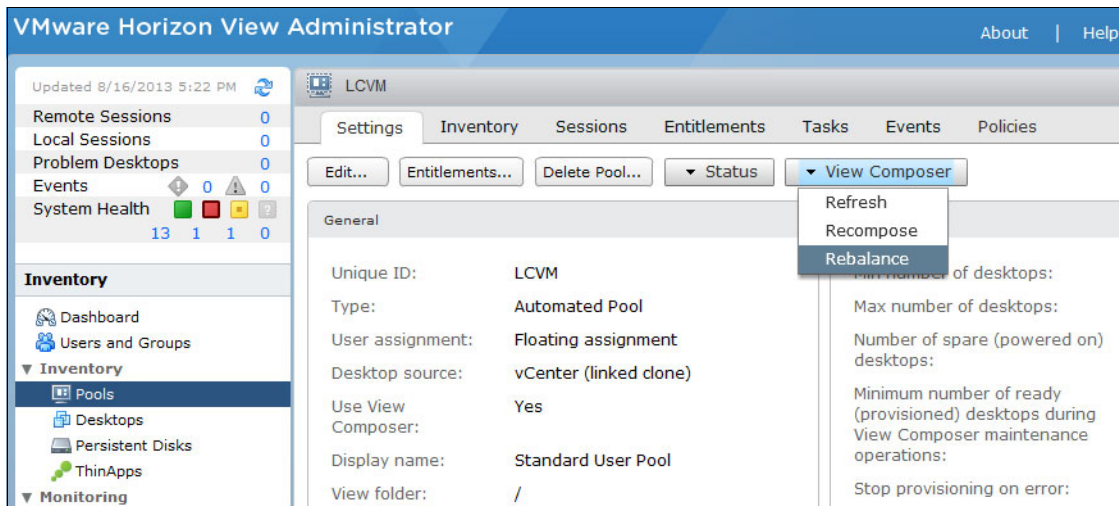


Figure 9-121 Rebalance operation for the entire pool

- To rebalance one or more virtual desktops, go to the Inventory tab, hold down the Ctrl key and click each virtual desktop that you want to rebalance, and click **View Composer** → **Rebalance**, as shown in Figure 9-122.

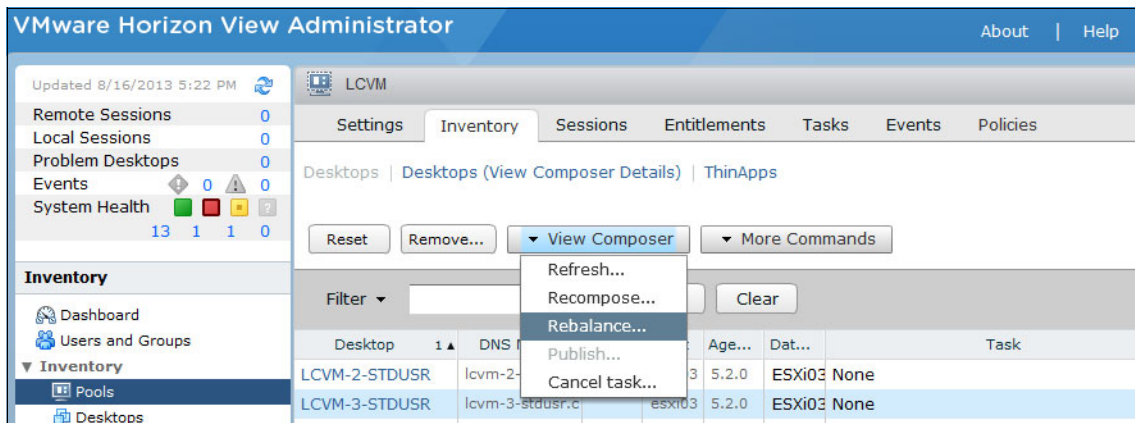


Figure 9-122 Rebalance multiple desktops

4. A rebalance wizard opens. Click **Next** to continue. Then, select the option to force users to log off or to wait for users to log off, as shown in Figure 9-123. Click **Next**.

The screenshot shows a window titled "Rebalance" with a "Scheduling" section. It prompts the user to "Specify when you want this task to start" with a date field set to "08/16/2013" and a time field set to "18 : 33", with a "Web browser local time" label. There are two radio button options: "Force users to log off" (selected) and "Wait for users to log off". Below the second option is a descriptive text: "Wait for connected users to disconnect before the task starts. The task starts immediately on desktops without active sessions." There is a checked checkbox for "Stop at first error". A bolded note states "The warning and grace period can be edited in global settings:". Below this is another checked checkbox for "Display warning before forced logoff:". The "Log off time:" is set to "5 minutes" in a text field. The "Log off message:" is shown in a text area with the text "Your desktop is scheduled for an". At the bottom right are three buttons: "< Back", "Next >", and "Cancel".

Rebalance

**Scheduling**

Specify when you want this task to start

Start at: 08/16/2013 18 : 33 Web browser local time

☒ Force users to log off

Users will be forced to log off when the system is ready to operate on their virtual machines. Before being forcibly logged off, users may have a grace period in which to save their work (Global Settings).

☐ Wait for users to log off

Wait for connected users to disconnect before the task starts. The task starts immediately on desktops without active sessions.

☒ Stop at first error

**The warning and grace period can be edited in global settings:**

☒ Display warning before forced logoff:

Log off time: 5 minutes

Log off message: Your desktop is scheduled for an

< Back Next > Cancel

Figure 9-123 Select the rebalance options

5. At the Ready to Complete window, click **Finish**, as shown in Figure 9-124.

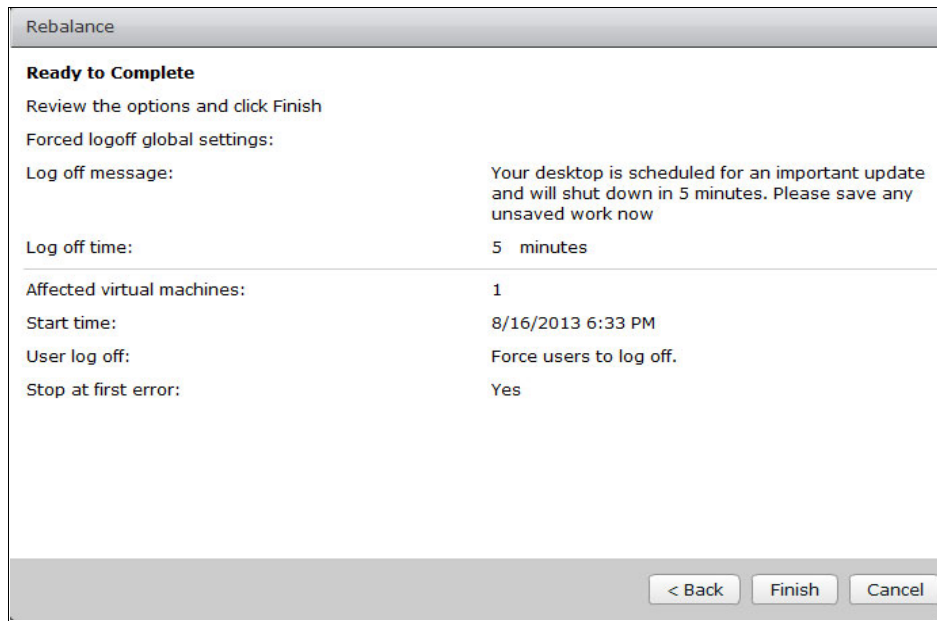


Figure 9-124 Rebalance summary

## 9.5.4 Migrating virtual desktops to another data store

To migrate virtual desktops to another data store, do not use vCenter.

Instead, use the rebalance feature that is described in 9.5.3, “Performing a desktop rebalance operation” on page 473 with editing the pool settings from VMware Horizon View Administration.



Complete the following steps:

1. Connect to VMware Horizon View Administrator web interface. Then, click **Inventory** in the left pane, and select **Pools**. Double-click the **LCVM** desktop pool to show the details for this desktop pool, as shown on Figure 9-125. Click **Edit**.

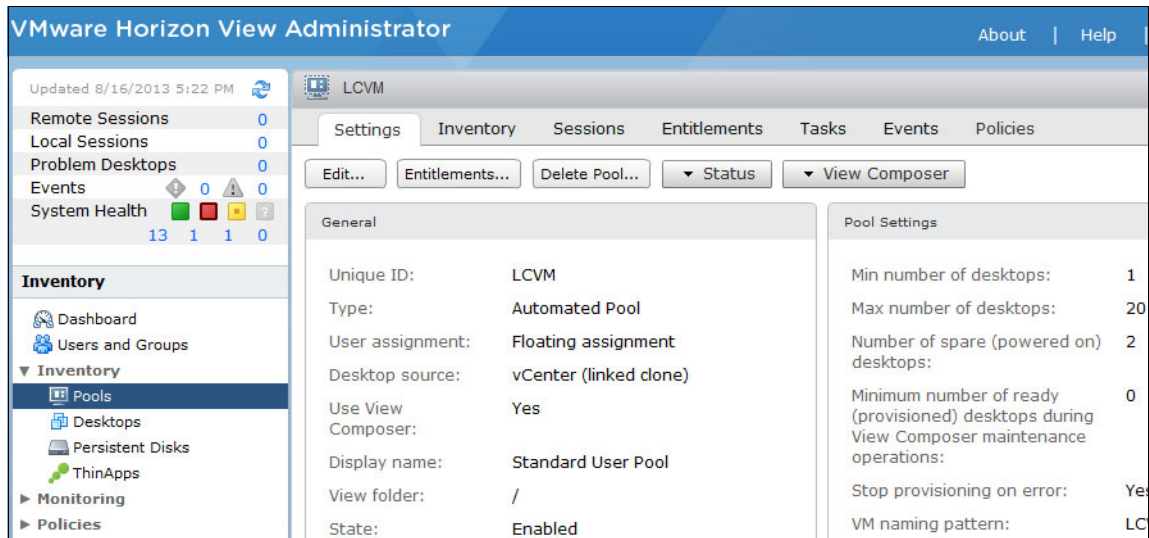


Figure 9-125 Desktop pool details

2. From the vCenter Settings page, you set only the Datastores setting. To change the Datastores settings, click **Browse**. Then, clear the data store selections and select the new data store, as shown in Figure 9-126. Click **OK**.

Select Linked Clone Datastores

Select the linked clone datastores to use for this pool. Only datastores that can be used by the selected host or cluster can be selected.

The table of minimum, maximum and 50% values only reflects the amount of storage needed for new virtual machines. It does not factor in the amount of storage space required for the disk growth of current virtual machines

Local datastore

Shared datastore

	Datastore	Capacity (GB)	Free (GB)	Type	Desktop	Storage Overcommit ?
<input type="checkbox"/>	<div>ESXi03-local-01</div>	930.25	910.51	VMFS 5	1	
<input type="checkbox"/>	<div>ESXi04-local-01</div>	930.25	914.00	VMFS 5	1	
<input checked="" type="checkbox"/>	<div>VDI-Shared-01</div>	599.75	272.86	VMFS 5	0	Conservative

Data Type	Selected Free Space (GB)	Min Recommended (GB)	50% utilization (GB)	Max Recommended (GB)
Linked clones	272.86	144.00	522.00	972.00

OK

Cancel

Figure 9-126 Select the data store

3. At the vCenter settings page click **OK**, as shown in Figure 9-127.

The screenshot shows the 'Edit LCMV' dialog box with the 'vCenter Settings' tab selected. The dialog has a title bar with a question mark icon. Below the title bar are tabs: 'General', 'Pool Settings', 'Provisioning ...', 'vCenter Setti...', 'Guest Custo...', and 'Advanced Sto...'. A warning message states: 'Changing the vCenter settings affects newly created virtual machines only. The new settings do not affect existing virtual machines.' The 'Default Image' section contains two fields: 'Parent VM:' with the value '/DatacenterA/vm/LCVM' and a 'Browse...' button, and 'Snapshot:' with the value '/LCVMv2/LCVMv3' and a 'Browse...' button. The 'Virtual Machine Location' section contains one field: 'VM folder:' with the value '/DatacenterA/vm/Standard Users - LC\'. The 'Resource Settings' section contains three fields: 'Host or cluster:' with the value '/DatacenterA/host/VDI' and a 'Browse...' button, 'Resource pool:' with the value '/DatacenterA/host/VDI/Resources' and a 'Browse...' button, and 'Datastores:' with the value '1 selected' and a 'Browse...' button. At the bottom right are 'OK' and 'Cancel' buttons.

1 Parent VM: /DatacenterA/vm/LCVM Browse...

2 Snapshot: /LCVMv2/LCVMv3 Browse...

3 VM folder: /DatacenterA/vm/Standard Users - LC\

4 Host or cluster: /DatacenterA/host/VDI Browse...

5 Resource pool: /DatacenterA/host/VDI/Resources Browse...

6 Datastores: 1 selected Browse...

OK Cancel

Figure 9-127 vCenter settings page

- Returning to VMware Horizon View Administrator page, click **View Composer** → **Rebalance**, as shown in Figure 9-128.

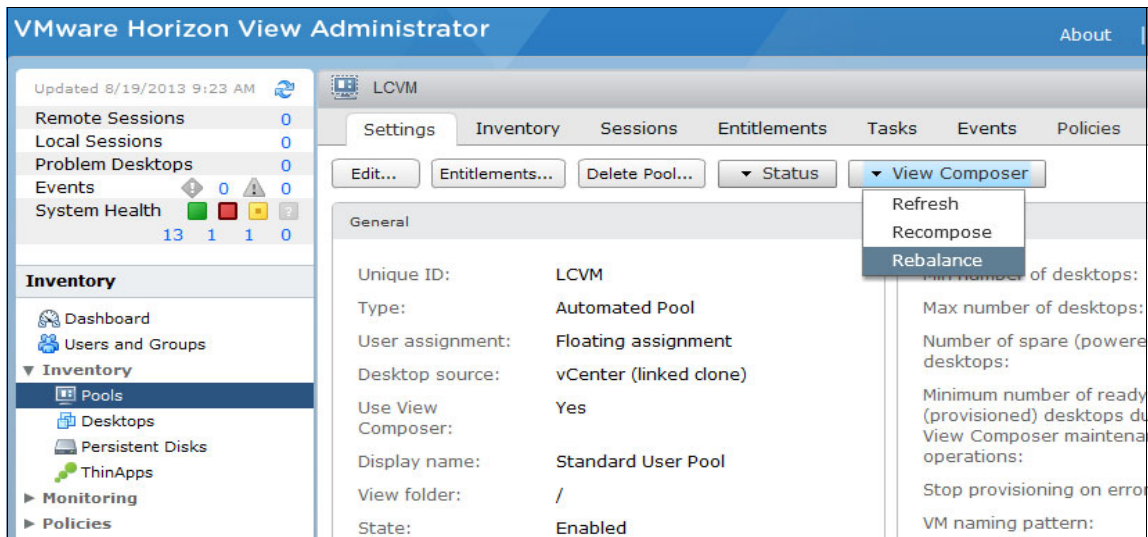


Figure 9-128 Rebalance the entire desktop pool

- When the wizard starts, click **Next**.

6. Select the option to force users to log off or to wait for users to log off, as shown in Figure 9-129. Click **Next**.

**Rebalance**

**Scheduling**

Specify when you want this task to start

Start at: 08/16/2013 18 : 33 Web browser local time

☒ Force users to log off

Users will be forced to log off when the system is ready to operate on their virtual machines. Before being forcibly logged off, users may have a grace period in which to save their work (Global Settings).

☐ Wait for users to log off

Wait for connected users to disconnect before the task starts. The task starts immediately on desktops without active sessions.

☒ Stop at first error

**The warning and grace period can be edited in global settings:**

☒ Display warning before forced logoff:

Log off time: 5 minutes

Log off message: Your desktop is scheduled for an

< Back Next > Cancel

Figure 9-129 Select the rebalance options

7. Click **Finish** at the Ready to Complete window, as shown in Figure 9-130.

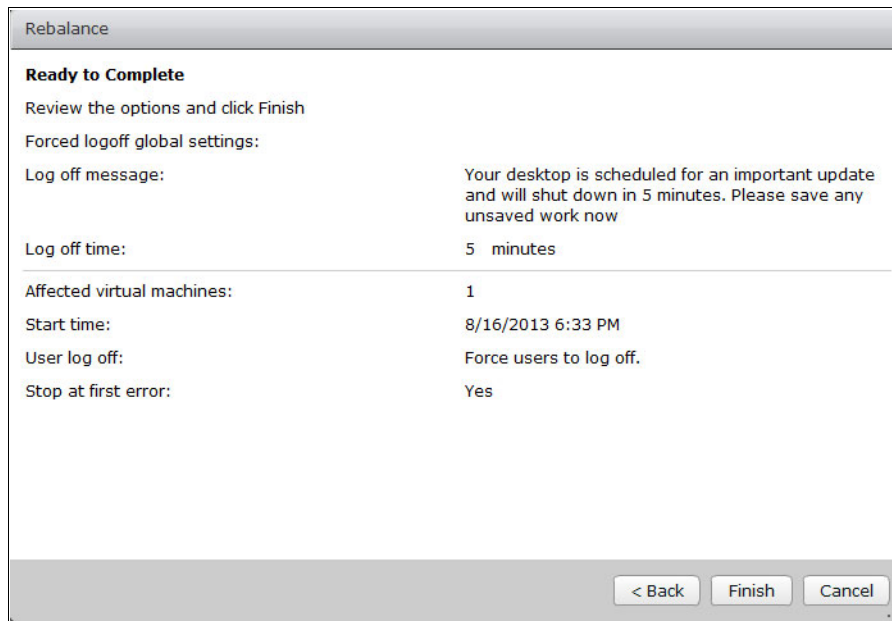


Figure 9-130 Rebalance summary

You can monitor for errors in the rebalancing operation from the LCVN Tasks tab, as shown in Figure 9-131.

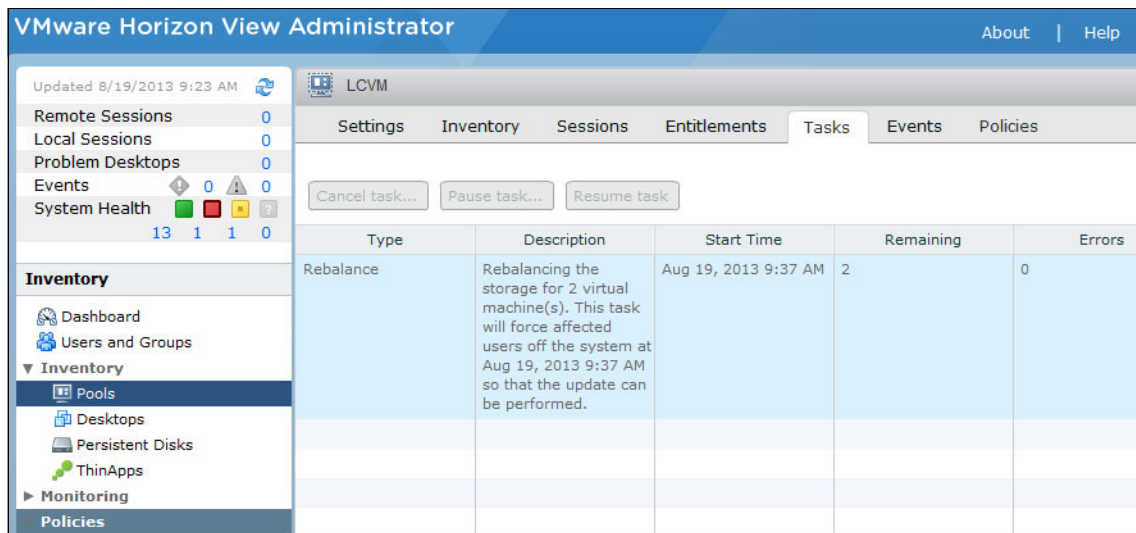


Figure 9-131 LCVN Tasks Tab Monitoring VMware View with Flex System

# Abbreviations and acronyms

<b>AD</b>	Active Directory	<b>HBAs</b>	host bus adapter
<b>ASU</b>	Advance Setting Utility	<b>HDD</b>	hard disk drive
<b>ATM</b>	automated teller machine	<b>HVD</b>	hosted virtual desktop
<b>BE3</b>	BladeEngine 3	<b>IBM</b>	International Business Machines Corporation
<b>BYOD</b>	bring-your-own-device	<b>IMM</b>	Integrated Management Module
<b>CAD</b>	computer-aided design	<b>IMM2</b>	Integrated Management Module II
<b>CBRC</b>	Content Based Read Cache	<b>IOPS</b>	input/output operations per second
<b>CDP</b>	Cisco Discovery Protocol	<b>IPC</b>	interprocess communication
<b>CIFS</b>	Common Internet File System	<b>ITSO</b>	International Technical Support Organization
<b>CIM</b>	Common Information Model	<b>JF</b>	jumbo frames
<b>CLI</b>	command line interface	<b>LCVM</b>	linked clone virtual machine
<b>CMM</b>	Chassis Management Module	<b>LLDP</b>	Link Layer Discovery Protocol
<b>CMMs</b>	Chassis Management Modules	<b>LRO</b>	Large Receive Offload
<b>CNA</b>	converged network adapter	<b>LUN</b>	logical unit number
<b>COM</b>	Component Object Model	<b>LUNs</b>	logical unit numbers
<b>DCOM</b>	distributed component object model	<b>MSRP</b>	Microsoft Roaming Profiles
<b>DPM</b>	Distributed Power Management	<b>NAS</b>	network attached storage
<b>DRS</b>	Distributed Resource Scheduler	<b>NFS</b>	Network File System
<b>DSA</b>	dynamic system analysis	<b>NIC</b>	network interface card
<b>FC</b>	Fibre Channel	<b>NPIV</b>	N_Port ID Virtualization
<b>FCoE</b>	Fibre Channel over Ethernet	<b>OU</b>	organizational unit
<b>FSM</b>	Flex System Manager	<b>OVA</b>	Open Virtualization Format Archive
<b>FT</b>	Fault Tolerance	<b>RA</b>	Reference Architecture
<b>FVM</b>	Full Virtual Machine	<b>RDC</b>	Remote Desktop Connection
<b>GPO</b>	group policy object	<b>RDP</b>	Remote Desktop Protocol
<b>GPO</b>	Group Policy Object	<b>SAN</b>	storage area network
<b>GPU</b>	graphics processing units	<b>SAS</b>	serial-attached SCSI
<b>GUI</b>	graphical user interface		
<b>HA</b>	High Availability		

<b>SEN</b>	System Storage Expansion Node
<b>SFP</b>	small form-factor pluggable
<b>SID</b>	security identifier
<b>SLP</b>	Service Location Protocol
<b>SNIA</b>	Storage Networking Industry Association
<b>SNMP</b>	Simple Network Management Protocol
<b>SSD</b>	solid-state drive
<b>SSDs</b>	solid-state drives
<b>SSH</b>	Secure Shell
<b>SSO</b>	single sign-on
<b>TCO</b>	total cost of ownership
<b>TCP</b>	Transmission Control Protocol
<b>TOE</b>	TCP offload engine
<b>TOR</b>	terminal-owning region
<b>TSO</b>	TCP Segmentation Offload
<b>UEFI</b>	Unified Extensible Firmware Interface
<b>UIM</b>	Upward Integration Module
<b>VDI</b>	virtual desktop infrastructure
<b>VDS</b>	virtual distributed switch
<b>VLANs</b>	Virtual Local Area Networks
<b>VM</b>	virtual machine
<b>VMDK</b>	VMware Virtual Machine Disk
<b>VMDq</b>	Virtual Machine Device Queues
<b>VMs</b>	virtual machines
<b>VSS</b>	virtual standard switch
<b>eMLC</b>	enterprise multilevel cell
<b>iSCSI</b>	Internet Small Computer System Interface
<b>pNIC</b>	Physical NIC mode
<b>vNIC</b>	virtual network interface card
<b>vNICs</b>	virtual network interface cards



# 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.

## IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only:

- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- ▶ *Implementing Systems Management of IBM PureFlex System*, SG24-8060
- ▶ *IBM Flex System Interoperability Guide*, REDP-FSIG

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials, at the following website:

<http://www.ibm.com/redbooks>

## Online resources

The following websites are also relevant as further information sources:

- ▶ IBM SmartCloud Desktop Infrastructure:  
<http://www.ibm.com/systems/virtualization/desktop-virtualization/>
- ▶ IBM Reference Architecture: SmartCloud Desktop Infrastructure:  
<http://ibm.co/186BJt7/>
- ▶ IBM Reference Architecture for VMware View:  
<http://ibm.co/17c0yaN/>
- ▶ VMware Horizon View 5.2 Architecture Planning Guide:  
<http://bit.ly/1hINKJk/>

- ▶ Storage Considerations for VMware Horizon View 5.2:  
[http://www.vmware.com/files/pdf/view\\_storage\\_considerations.pdf](http://www.vmware.com/files/pdf/view_storage_considerations.pdf)
- ▶ VMware Horizon View 5.2 Documentation Center:  
<http://pubs.vmware.com/view-52/index.jsp/>
- ▶ IBM Flex System Information Center:  
<http://pic.dhe.ibm.com/infocenter/flexsys/information/index.jsp/>

## Help from IBM

IBM Support and downloads:

<http://www.ibm.com/support>

IBM Global Services:

<http://www.ibm.com/services>

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.



# Implementing VMware Horizon View on IBM Flex System

(1.0" spine)  
0.875"<->1.498"  
460 <-> 788 pages

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.frm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.





# Implementing VMware Horizon View on IBM Flex System

**Introduces IBM Flex System and VMware Horizon View offerings**

**Describes design, planning, and deployment considerations**

**Provides a step-by-step configuration guide**

The IBM SmartCloud Desktop Infrastructure offers robust, cost-effective, and manageable virtual desktop solutions for a wide range of clients, user types, and industry segments. These solutions to help 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.

IBM SmartCloud Desktop Infrastructure with VMware Horizon View simplifies desktop and application management and increases security and control. Horizon View delivers a personalized high fidelity experience for users across sessions and devices. It also enables higher availability and agility of desktop services that are unmatched by traditional PCs, while reducing the total cost of desktop ownership.

This IBM Redbooks publication provides an overview of the SmartCloud Desktop Infrastructure solution that is based on VMware Horizon View running on IBM Flex System. It highlights key components, architecture, and benefits of this solution. It also provides planning and deployment considerations, and step-by-step instructions on how to perform specific tasks.

## INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

### BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
[ibm.com/redbooks](http://ibm.com/redbooks)