

The Lenovo logo, consisting of the word "Lenovo" in white sans-serif font on a dark grey rectangular background.

Migrating your x86 Servers from BladeCenter to Flex System

**Advice on upgrading from an x86
environment to Lenovo Flex
System**

**Hardware migration guides with
BladeCenter source examples**

**Guidance on migrating Microsoft,
VMware, and KVM system images**

**Recommended for all BladeCenter
customers**

Bin Qi Zhang

Ye Xu

Jason Brunson

Jun Zeng

Jian Guo Ma



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

Last update on December 2013

This edition applies to these products:

- ▶ Flex System Manager 1.3
- ▶ VMware vSphere 4.1
- ▶ VMware vCenter Server/Client 4.1
- ▶ Microsoft Windows 2008 x64 R2
- ▶ Microsoft Hyper-V Manager 6.1
- ▶ Red Hat Enterprise Linux 6.2
- ▶ Kernel Virtual Machine KVM-22

© Copyright Lenovo 2016. All rights reserved.

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

Contents

Notices	v
Trademarks	vi
Preface	vii
The team who wrote this paper	vii
Comments welcome	viii
Do you have the latest version?	viii
Chapter 1. Introduction	1
1.1 Comparing Flex System with BladeCenter	3
1.1.1 Chassis	3
1.1.2 Compute nodes	5
1.1.3 Management module	6
1.1.4 Systems management	7
1.1.5 Power supplies and fan modules	8
1.1.6 I/O modules	9
1.1.7 Expansion nodes	11
1.2 Introduction to the migration process	13
Chapter 2. Migrating hardware settings	15
2.1 Migrating chassis and compute node settings	16
2.1.1 Migrating AMM and CMM settings	16
2.1.2 Migrating compute node settings	26
2.1.3 Migrating UEFI settings	30
2.2 Migrating network settings	32
2.3 Migrating storage settings	41
Chapter 3. Migrating operating system images	55
3.1 Source server and Flex System are disconnected	56
3.1.1 Disconnected Physical-to-Physical	56
3.1.2 Disconnected Physical-to-Virtual	56
3.1.3 Disconnected Virtual-to-Virtual	75
3.2 Source servers and Flex System are connected	88
3.2.1 Connected Physical-to-Physical	88
3.2.2 Connected Physical-to-Virtual	88
3.2.3 Connected Virtual-to-Virtual	104
3.3 Conclusion	137
Abbreviations and acronyms	139
Related publications	141
Lenovo Press publications	141
Other publications and online resources	141

Notices

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

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

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

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

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

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

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

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

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

Trademarks

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

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

BladeCenter®	Omni Ports™	ServerProven®
Flex System™	RackSwitch™	System x®
Lenovo®	Lenovo (logo)®	TruDDR4™
Netfinity®	ServeRAID™	vNIC™

The following terms are trademarks of other companies:

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

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.

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

Preface

Lenovo Flex System is a next-generation platform ideally suited to data center environments that require flexible, cost-effective, secure, and energy-efficient server infrastructure.

This Lenovo Press paper highlights the major advances of Flex System™. Lenovo BladeCenter® is used as a comparative example on aspects such as chassis, compute node, management module, power supply, cooling, network, and I/O. Key factors users need to consider when planning workload migration also are reviewed. Some instrumentation and reference tools are offered here, with detailed steps that take the user through the entire planning process.

This paper is intended for those IT professionals who want to migrate a workload from existing hardware to the Flex System.

The team who wrote this paper

This document is produced by the following subject matter experts working from around the world.

Bin Qi Zhang is an Advisory Software Engineer and IBM System x® ToolsCenter Architect in IBM China Systems & Technology Labs (CSTL). He writes extensively on migration processes, including hardware migration and operating system image migration. Before joining CSTL in 2006, he had three years experience as a software engineer in the healthcare IT and microelectronics industries.

Ye Xu is a Staff Software Engineer in China. He has five years of experience in the IBM System & Technology Lab. He holds a Master's degree in Electronic Engineering from East China Normal University. His areas of expertise include System Management Software, System x virtualization technology, and IBM Cloud. He writes extensively on Common Information Model (CIM) technology and System x virtualization technology.

Jason Brunson is a Senior Technical Specialist with the US Advanced Technical Support organization. He has 15 years of experience in the field of systems management and two years in the virtualization field. He also has 16 years of experience at IBM. His areas of expertise include systems management, virtualization, and PureFlex System. He writes extensively about systems management and virtualization technologies.

Jun Zeng is a Staff Software Engineer in China. He has six years of experience in the software development field. He holds a Master's degree from Wuhan University. His areas of expertise include the System x software solution and virtualization technology.

Jian Guo Ma is a Staff Software Engineer in China. He has six years of experience in the software development field. He holds a Master's degree in Control Theory and Control Engineering from Tongji University. His areas of expertise include the development of the System x software solution, and in Computer-aided Design/Computer-aided Manufacturing/Computer-aided Engineering (CAD/CAM/CAE).

Comments welcome

Your comments are important to us!

We want our documents to be as helpful as possible. Send us your comments about this paper in one of the following ways:

- ▶ Use the online feedback form found at the web page for this document:
<http://lenovopress.com/redp4887>
- ▶ Send your comments in an email to:
comments@lenovopress.com

Do you have the latest version?

We update our books and papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Introduction

Flex System is a next-generation platform that is ideally suited to data center environments that require flexible, cost-effective, secure, and energy-efficient server infrastructure.

The innovative design features of the Flex System products make it possible to configure integrated, customized, highly secure solutions. These solutions meet data center needs and provide for flexible expansion capabilities. The scalable hardware features and the unprecedented power and cooling capabilities of the Flex System products help optimize hardware and power utilization, minimize operational costs, and simplify the overall management of the data center.

Figure 1-1 shows Flex System.



Figure 1-1 Flex System

Flex System is also an infrastructure that is custom-built to your requirements with more performance and bandwidth, true integrated enterprise SAN storage, and far more capability to consolidate and virtualize than previous systems.

The Flex System chassis is designed for new levels of simplicity, reliability, and upgradeability. With high performance compute nodes, enhanced networking capabilities, and sophisticated system management capabilities, you can upgrade an existing blade server system to this converged infrastructure to make your IT simpler, more flexible, more open, and more efficient.

This paper provides an overview of the process to migrate workloads from an x86 source server to the Flex System platform. All the necessary key factors that must be considered during the migration process are documented with detailed descriptions and explanations. You can take these general guidelines and apply them to your environment and make changes where appropriate.

Workload migration is not only about moving the operating system (OS) images that contain the application and data. Preparing the hardware infrastructure and ensuring it is optimally fitted to the workload that runs on that hardware also is important. Such preparation also prevents problems, such as different hardware platforms that have different software or OS support. Problem-solving techniques are presented in this paper.

This paper contains the following chapters:

- ▶ In the Introduction, the hardware components of the Flex System are introduced by comparing the components to their well-known BladeCenter counterparts.
- ▶ In the remaining chapters, the migration solutions are described in terms of the hardware settings migration process and OS images migration process, and the considerations behind those processes.

For users familiar with the OS, hypervisor, virtual machine, and application migration, you can find more information in Chapter 2, “Migrating hardware settings” on page 15 about setting up the hardware as similar to the source environment as possible.

For users who are planning for Physical-to-Virtual or Virtual-to-Virtual migration, you can find more information in Chapter 3, “Migrating operating system images” on page 55 about the different scenarios to consider.

For users unfamiliar with the workload migration process, review this entire paper and then focus on the sections that are most relevant.

This paper does not address all potential issues that can occur during a migration process, although many detailed operations are shown. Users must consider the actual environment, constraints, topology, and policies of their system, and then design a comprehensive customized transition plan to implement. Data migration is not described in this paper. Therefore, product interfaces that are shown in this paper can change because of software updates.

Important: Some migrations might share commonalities with the cases and scenarios that are introduced in this paper. However, a successful migration is not guaranteed. Lenovo is not liable for any damage, loss of data, or any other unpredictable consequences to your equipment that might result from the migration process.

1.1 Comparing Flex System with BladeCenter

Before starting any migration, it is useful to first understand the hardware infrastructure, the value proposition of the Flex System, and the key hardware components. A simple approach is to start with a familiar concept and extend the thinking to the new offerings. In this paper, the BladeCenter system is used to introduce the Flex System.

Capacity planning is not a simple mapping of one component to another. Instead, also considered in the new workload planning are the scalability, reliability, availability, serviceability, and manageability of the offerings of the new Flex System as compared to BladeCenter. This hardware component comparison is only an example that is used to begin the planning process.

1.1.1 Chassis

The Flex System chassis, compute node, storage, networking, and management components are engineered to integrate and deliver optimized, highly reliable systems.

The Flex System features blade server technology. This type of server is a hot-swap, independent server with its own processors, memory, storage, network device controllers, operation system, and applications.

The chassis is an enclosure for blade servers that supplies shared power modules, cooling components (fans), and management modules for system management. The chassis also includes I/O modules, such as Ethernet, Fibre Channel, InfiniBand, and Serial Attached Small Computer System Interface (SAS) switch modules. In some configurations, network pass-through modules also are used for different network topologies.

Shared modules are included with Flex System Enterprise chassis.

The hardware form factor differs between a BladeCenter chassis and Flex System Enterprise chassis in the following aspects:

- ▶ The BladeCenter H Chassis has 14 slots (also known as *bays*) and is 9U high. The orientation of the blade servers is vertical in most BladeCenter chassis.
- ▶ Flex System Enterprise chassis also has 14 bays, but the bays are horizontally arranged in seven rows, each with two standard-width bays. The Flex System Enterprise chassis is 10U high.

The different chassis designs are shown in Figure 1-2.

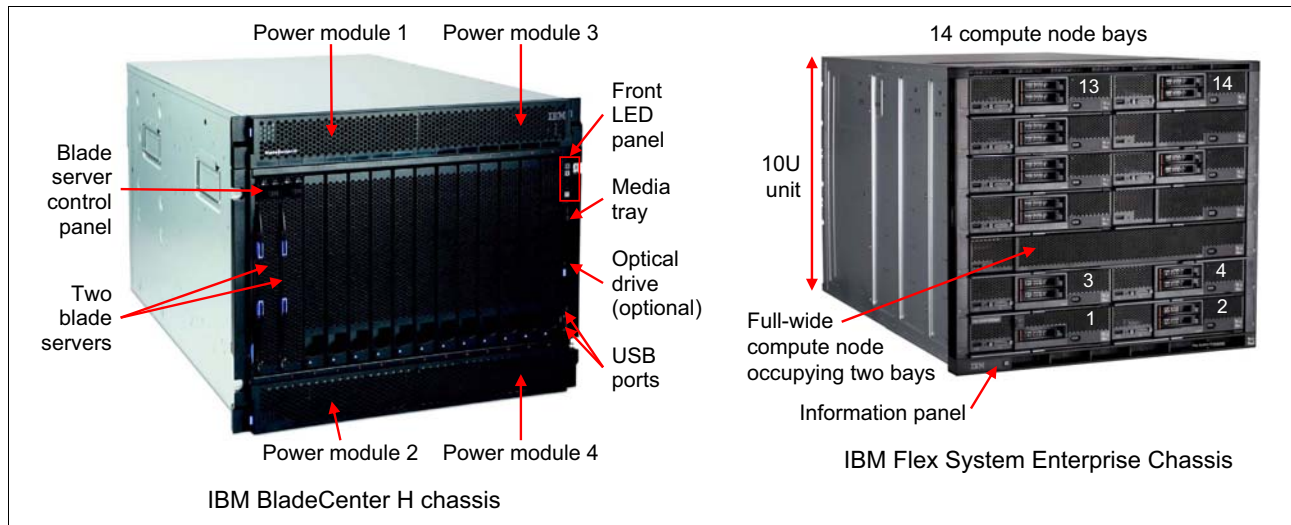


Figure 1-2 Front view of the BladeCenter H chassis and the Flex System Enterprise Chassis

The rear view of the chassis shows that, although the components are arranged differently, the components are similar. These components include I/O modules, fan modules, power modules, and so on, which are shared by all the servers within the chassis.

Figure 1-3 shows the rear view of the chassis.




Figure 1-3 Rear view of the BladeCenter H and Flex System chassis

1.1.2 Compute nodes

A BladeCenter chassis includes bays for blade servers or other BladeCenter devices, as shown in Figure 1-4. The blade server shares power, fans, switches, and ports with other blade servers.




Figure 1-4 Blade servers that are installed in a BladeCenter chassis

In the Flex System, a *compute node*, such as a blade server, contains components such as microprocessors, memory, Ethernet controllers, and hard disk drives. Power and network connections are provided by the mid-plane of the Flex System Enterprise chassis.

Flex System Enterprise Chassis supports up:

- ▶ 28 servers in 14 Flex System x222 Compute Nodes
- ▶ 14 servers in 14 Flex System x220 or x240 Compute Nodes
- ▶ 7 servers in 7 Flex System x440 Compute Nodes

Figure 1-5 shows a one-bay compute node (or *standard* compute node) and a two-bay compute node (or *double-wide* compute node).




Figure 1-5 Flex System compute nodes that are installed in the Enterprise Chassis

1.1.3 Management module

The management module is a hot-swap module that is used to configure and manage BladeCenter or Flex System components. The BladeCenter and Flex System chassis have similar modules, but the names differ:

- ▶ BladeCenter uses the *Advanced Management Module (AMM)*
- ▶ Flex System Enterprise Chassis uses the *Chassis Management Module (CMM)*

Figure 1-6 shows the form factor and installation position for each of these modules in the BladeCenter chassis and Flex System Enterprise chassis.




Figure 1-6 Comparing the management modules

The login windows are similar to the windows that are shown in Figure 1-7 on page 7.

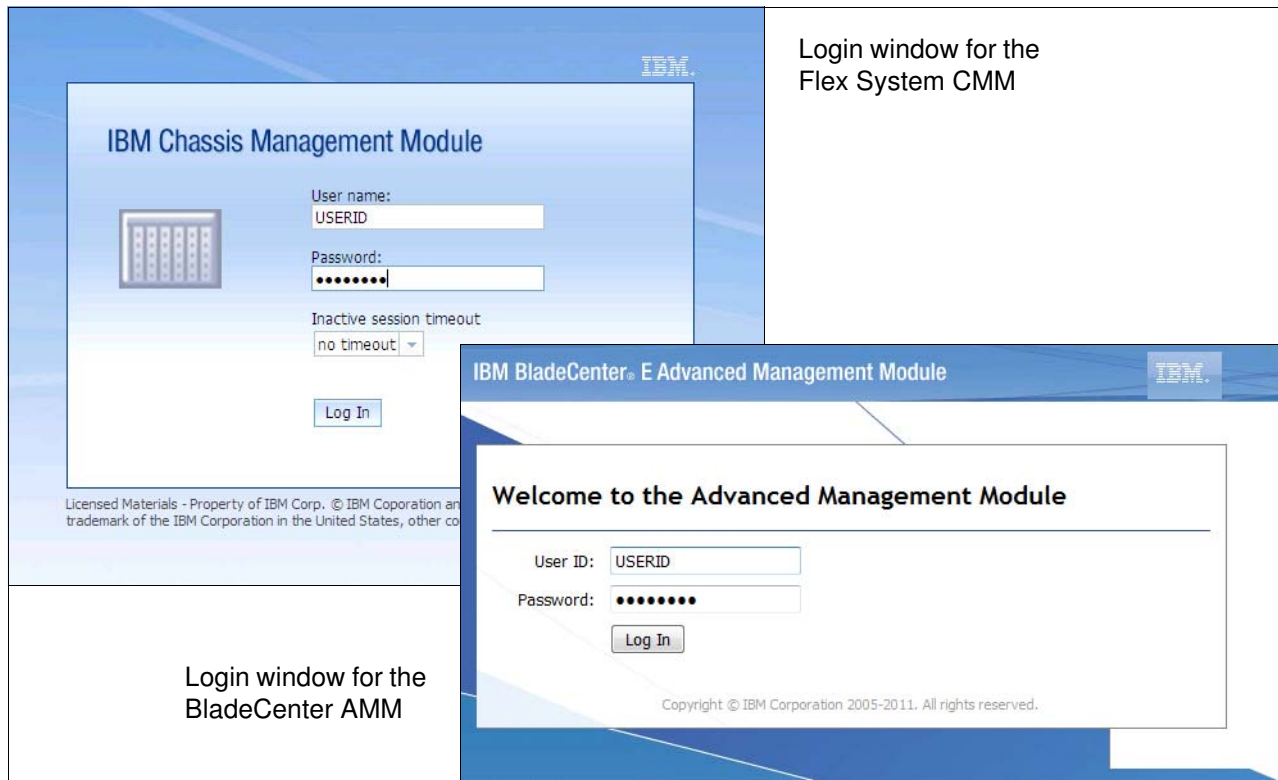


Figure 1-7 Login windows

During the setting migration process from AMM to CMM, document the settings of the power policies, user accounts, and the network so that the source and target systems are consistent after migration.

The CMM provides more security and powerful systems management functions. Some new settings can be applied to the CMM after you migrate the old settings. For more information, see the *Chassis Management Module User's Guide* at this website:

<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>

1.1.4 Systems management

Flex System Manager™ is a management node that configures and manages multiple chassis platforms remotely. Flex System Manager provides systems management functions for all compute nodes in a multiple-chassis configuration (up to four chassis and 56 nodes).

Flex System Manager is an essential component for PureFlex System, but still can be an optional component for Flex System. From a hardware perspective, it is a standard compute node that is preinstalled with an OS and management software stack.

Figure 1-8 shows the Flex System Manager node.



Figure 1-8 Flex System Manager

Because a similar component is not available for BladeCenter, there are no configuration or data settings to be migrated to the Flex System.

If your configuration does not include Flex System Manager, Lenovo ToolsCenter is another option that can provide you with basic hardware management capability, including server diagnostic tests, an UEFI configuration, a firmware update, and operating system deployment. ToolsCenter supports some basic hardware management for Flex System and for BladeCenter. For more information and to download it, go to the following website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=tool-center>

1.1.5 Power supplies and fan modules

The BladeCenter H chassis includes two or four hot-swap power modules and two hot-swap blowers for cooling, as shown in Figure 1-9.




Figure 1-9 BladeCenter H power and blower modules

The design of the Flex System Enterprise Chassis is much more integrated. It supports up to six auto-ranging AC power supplies and up to 10 fan modules (two 40 mm fan modules and eight 80 mm fan modules). It includes a minimum of six hot-swap fan modules (four 80 mm fan modules and two 40 mm fan modules).

The two smaller 40 mm fan modules at the top of the chassis provide cooling to the I/O modules and the CMMs. The larger 80 mm fan modules provide cooling to the compute nodes and the Flex System Manager.

Figure 1-10 shows the Flex System Enterprise Chassis power and fan modules.




Figure 1-10 Flex System power and blower modules

Although the power and fan modules do not directly affect the workload migration, you must ensure that the modules are sufficient for your availability, redundancy, and energy consumption needs. For more information, see the *Flex System Enterprise Chassis & PureFlex Power Requirements Guide*, found at:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102111>

1.1.6 I/O modules

The BladeCenter H chassis contains 10 hot-swap I/O module bays that can be used in the following configurations:

- ▶ The modules in bays 1 and 2, such as Ethernet switches or pass-through modules, provide a communication link for the first and second network interface controllers (NICs) on each blade server.
- ▶ The modules in bay 3 and 4 can be switch modules or bridge modules. If a switch module is installed, the module provides a communication link for the third and fourth NIC in each blade server, similar to bays 1 and 2. If bridge modules are installed, they provide links to bays 7 - 10 that can be used as additional outputs for I/O modules in those bays.

- ▶ The modules in bays 5 and 6, such as bridge modules, provide links to bays 7 - 10 that can be used as additional outputs for I/O modules in those bays. (This configuration is the same as a configuration that uses bays 3 and 4 in bridge mode.)
- ▶ Modules in bays 7 - 10, such as InfiniBand switches, provide high-speed communication links to NICs 5 - 8 in each blade server.

The Flex System I/O design is simpler and more flexible than the BladeCenter design. Four high-speed I/O modules are supported in the Flex System Enterprise Chassis, including Ethernet, Fibre Channel, and InfiniBand switches, and pass-through modules (optical and copper). These I/O module bays can be used in the following configurations:

- ▶ I/O module bays 1 and 2 support any standard Ethernet switches or pass-through modules that connect to the two integrated Ethernet controllers or an Ethernet I/O adapter in each of the compute nodes. When you install an adapter in the first slot on the compute node, the I/O module bays support any switch module with the same type of network interface that is used in the corresponding compute node adapter slot.
- ▶ I/O module bays 3 and 4 support Ethernet, Fibre Channel, InfiniBand switch modules, and pass-through modules. If you install an additional I/O module in bay 3 or bay 4, a corresponding adapter must be installed in each compute node to access the I/O bay.

With a range of available adapters and switches to support key network protocols, you can configure Flex System to fit in your infrastructure and still anticipate future needs. The networking resources in Flex System are standards-based, flexible, and fully integrated into the system, which results in a no-compromise networking solution. Network resources are virtualized and managed by workload. These capabilities are automated and optimized to make your network more reliable and simpler to manage.

Flex System features the following capabilities:

- ▶ Supports your networking infrastructure, including Ethernet, Fibre Channel, and InfiniBand
- ▶ Offers industry-leading performance with 1 Gb, 10 Gb, and 40 Gb Ethernet; 8 Gb and 16 Gb Fibre Channel and FDR InfiniBand
- ▶ Provides pay-as-you-grow scalability so you can add ports and bandwidth as needed

The BladeCenter and Flex System I/O modules are shown in Figure 1-11.




Figure 1-11 Comparison of I/O modules

During the migration process, the I/O module is important because it contains the network topology and other essential settings from your source system.

If you plan to migrate the old physical network topology to an Flex System, you must adhere to the following guidelines:

- ▶ Before starting the migration, you must consider the necessary I/O settings to migrate from the source system to the Flex System. It might not be necessary to migrate these configurations.

For example, if a Physical-to-Virtual (P2V) migration of the source server OS to the Flex Compute Node is needed, the original physical I/O settings must not be migrated. These settings must not be migrated because the original physical settings become virtual during the migration process. If you want to keep the old network topology, configure the similar settings in the virtual switch by using the hypervisor tool.

- ▶ There are two types of modules: *switch modules* and *pass-through modules*. Use caution when the source system contains modules that differ from the Flex System because the network topology might need to be changed.

The following solutions are available for several common configuration scenarios:

- Scenario 1: A switch module is installed in the source system and in the Flex System Enterprise Chassis.

Solution: Document all the network information from the source system and apply that information to the Flex System.

- Scenario 2: A switch module is installed in the source system and a pass-through module is installed in the Flex System Enterprise Chassis.

Solution: Although no additional work is required, you must ensure that the compute node and switch I/O module connectivity is consistent with what was expected.

- Scenario 3: A pass-through module is installed in the source system and a pass-through module is installed in the Flex System Enterprise Chassis.

Solution: No additional work is required, and no changes are made in the Flex System Enterprise Chassis.

- Scenario 4: A pass-through module is installed in the source system and a switch module is installed in the Flex System Enterprise Chassis.

Solution: You do not need to retrieve information from the source switch. However, you must create a large VLAN group in the Flex System configuration of which all the compute nodes must be a member.

- ▶ Ensure that the bandwidth of your Ethernet adapter, switch module, and Top of the Rack (TOR) switches match. If these switches do not match, the network operates at the lowest bandwidth that matches one of these components. Upgrade to a 10 Gb Ethernet network to make full use of the Flex System capacity.

1.1.7 Expansion nodes

Flex System offers directly attached expansion nodes, which are similar to BladeCenter expansion blades. These units provide a cost-effective way to add various industry-standard PCI Express (PCIe) adapters or storage to the attached server. They allow you to customize and balance your specific workload requirements by increasing the capabilities and usage of your existing compute node resources.

Two expansion nodes are available:

- ▶ PCIe Expansion Node

The Flex System PCIe Expansion Node supports additional PCIe adapters and I/O expansion adapters for compute nodes in the Flex System Enterprise Chassis.

The PCIe Expansion Node is physically connected to a supported compute node. The connection is made through the interposer cable and connectors on the PCIe Expansion Node and compute node system boards. The expansion adapters that are installed in the PCIe Expansion Node are available only to the attached compute node.

The PCIe Expansion Node that is connected to the x240 Compute Node is shown in Figure 1-12 on page 12 along side the PCIe Expansion Blade for BladeCenter.




Figure 1-12 Comparison of PCIe expansion units

► Storage expansion node

The Flex System Storage Expansion Node is a direct-attach storage enclosure that is dedicated to a single compute node in the Flex System Enterprise Chassis. It is similar to the BladeCenter Storage and I/O Expansion Blade, although the Storage Expansion Node provides more local storage than the BladeCenter unit.

The Storage Expansion Node is different from V7000 Storage Node in that the V7000 Storage Node is shared by all the compute node in chassis, and the Storage Expansion Node is accessible only by the compute node to which it is attached.

The storage expansion node is PCIe 3.0 and SAS 2.1 compliant. It supports up to 12 SAS or SATA 2.5-inch hot-swap hard disk drives and solid-state drives, and it supports JBOD configurations and RAID 0, 1, 5, 6, 10, 50, and 60. Some RAID configurations require a Features on Demand key, an optional 512 MB or 1 GB flash/RAID adapter, or both.

Figure 1-13 shows a comparison of the storage expansion units.



Figure 1-13 Comparison of storage expansion units

1.2 Introduction to the migration process

The migration process consists of two phases: hardware settings migration and OS and workload image migration. For the purposes of this paper, the following assumptions are made before initiating the migration:

- ▶ The Flex System chassis, servers, and I/O modules are installed, the power is on, and all cables are connected.
- ▶ Initial setup is complete for the Flex Chassis Management Module and Flex System Manager (if present).

For more information, see the Flex System Information Center website:

<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>

This paper covers the two migration phases:

- ▶ Migrating hardware settings
- ▶ Migrating operating system images

Migrating hardware settings

In this chapter, you learn how to migrate hardware settings so that the source hardware and target hardware can be set as similar to each other as possible. These settings include the compute node, networking, and storage components.

The following topics are described in these sections:

- ▶ “Migrating chassis and compute node settings” on page 16
- ▶ “Migrating network settings” on page 32
- ▶ “Migrating storage settings” on page 41

2.1 Migrating chassis and compute node settings

Several groups of settings must be considered during the migration process: Advanced Management Module (AMM), Chassis Management Module (CMM), and Compute Node/Unified Extensible Firmware Interface (UEFI) settings.

This section covers the following topics:

- ▶ Migrating AMM and CMM settings
- ▶ Migrating compute node settings
- ▶ Migrating UEFI settings

2.1.1 Migrating AMM and CMM settings

As shown in 1.1.3, “Management module” on page 6, the management module is used to configure and manage BladeCenter and Flex System components. You configure and manage the components by logging in to the AMM web page. After you are logged in, you can capture the source chassis configuration settings and deploy the settings to the target CMM.

The AMM and CMM are assigned static IP addresses. Depending upon the configuration that is used, you can choose to acquire an IP address from a Dynamic Host Configuration Protocol (DHCP) address instead of using the assigned address. In either case, enter the IP address into a web browser and log in to the AMM with your AMM credentials, as shown in Figure 2-1.

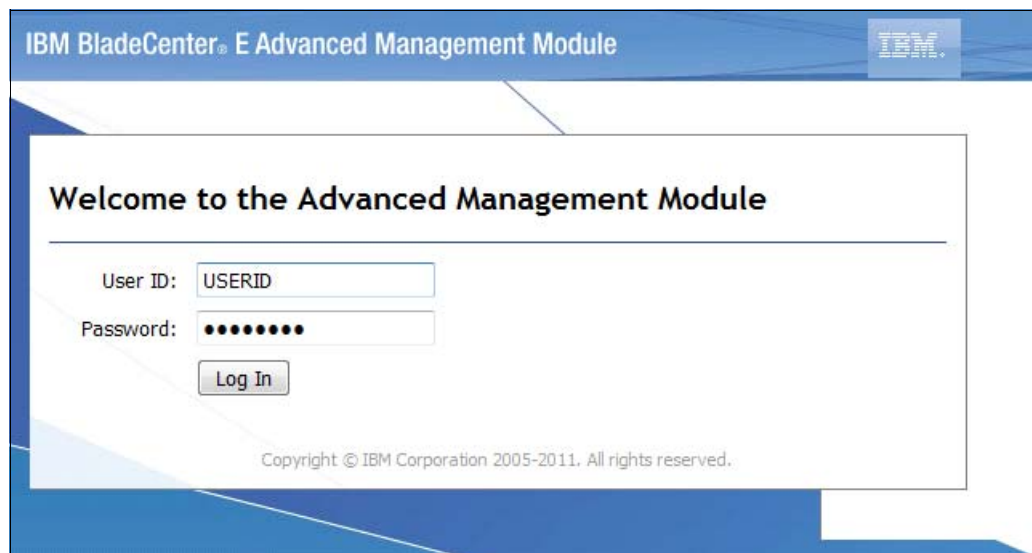


Figure 2-1 AMM login window

Many management functions are available on the AMM web page. The following sections describe the settings that must be considered for migration from the AMM to CMM.

Determining power management settings

The power management policy settings determine how the BladeCenter chassis reacts in each power domain to a failure of a power source or power module. The combination of the BladeCenter configuration, power-management policy settings, and available power might cause blade servers to reduce their power level (throttle) or fail to turn on.

Follow these steps to determine the current power management settings:

1. In the AMM navigation bar, click **Monitor** → **Power Management**. The BladeCenter Power Domain Summary page opens, as shown in Figure 2-2. This page features the power supply status and related information, including the Power Management Policy of the power domains.

BladeCenter Power Domain Summary ?		
	Power Domain 1	Power Domain 2
Status	Power domain status is good.	Power domain status is good.
Power Modules	Bay 1: 2000W Bay 2: 2000W	Bay 3: 2000W Bay 4: 2000W
Power Management Policy	Basic Power Management Total allowed power is higher than other policies and is limited only by the total power capacity of all the Power Modules up to the maximum of chassis power rating. This is the least conservative approach, since it does not provide any protection for AC power source or Power Module failure. If any single power supply fails, blade and/or chassis operation may be affected.	Basic Power Management Total allowed power is higher than other policies and is limited only by the total power capacity of all the Power Modules up to the maximum of chassis power rating. This is the least conservative approach, since it does not provide any protection for AC power source or Power Module failure. If any single power supply fails, blade and/or chassis operation may be affected.
Maximum Power Limit [†]	2650W	2650W
Power in Use ^{††}	833W	656W

BladeCenter Power Domain Planning ?		
	Power Domain 1	Power Domain 2
Maximum Power Limit [†]	2650W	2650W
- Allocated Power (Max) ^{†††}	1631W	1309W
= Remaining Power	1019W	1341W

BladeCenter Chassis Power Summary ?	
Total DC Power Available	5300W
Total AC Power In Use ^{††}	1755W
Total Thermal Output	5,988.1 BTU/Hour

Refresh

Figure 2-2 Power Domain Summary page

2. Click **Basic Power Management** under each power domain in turn to change the power policy for each power domain. (Depending on the power policy you chose, the link name might differ from the link that is shown in Figure 2-2.) A table opens that lists the power management policies, as shown in Figure 2-3.

This table lists the power management policies ordered from most conservative to least conservative.

Select	Option Name	Power Supply Failure Limit [†]	Maximum Power Limit (Watts)
<input type="radio"/>	Power Module Redundancy Intended for a single AC power source into the chassis where each Power Module is on its own dedicated circuit. Total allowed power draw is limited to one less than the number of Power Modules when more than one Power Module is present. One Power Module can fail without affecting blade operation. Multiple Power Module failures can cause the chassis to power off. Note that some blades may not be allowed to power on if doing so would exceed the policy power limit. More...	1	2000
<input type="radio"/>	Power Module Redundancy with Blade Throttling Allowed Very similar to Power Module Redundancy. This policy allows you to draw more total power; however, capable blades may be allowed to throttle down if one Power Module fails. More...	1	2650
<input checked="" type="radio"/>	Basic Power Management Total allowed power is higher than other policies and is limited only by the total power capacity of all the Power Modules up to the maximum of chassis power rating. This is the least conservative approach, since it does not provide any protection for AC power source or Power Module failure. If any single power supply fails, blade and/or chassis operation may be affected. More...	0	2650

Figure 2-3 Power management policy in effect

3. Review and record the AMM power policy settings as shown. In our example, the policy in effect is “Basic Power Management”.

Current power policy setting: It is important to understand the current power policy setting and the effects of that setting. If the policy provides the protection and performance that is needed, migrate this policy to the CMM.

Determining user account settings

Follow these steps to determine the current user account settings:

1. In the navigation bar, click **MM Control** → **Login Profiles**, as shown in Figure 2-4.

IBM BladeCenter® E Advanced Management Module

Welcome **USERID**

Bay 1: SN#YK14807741JR

- Monitors
- Blade Tasks
- I/O Module Tasks
- MM Control
 - General Settings
 - Login Profiles**
 - Alerts
 - Serial Port
 - Port Assignments
 - Network Interfaces
 - Network Protocols
 - Chassis Int Network
 - Security
 - File Management
 - Firmware Update
 - Configuration Mgmt
 - Restart MM
 - License Manager
- Service Tools

Management Module Login Configuration ?

Use the following links to jump down to different sections on this page.

- [Login Profiles](#)
- [Group Profiles](#)
- [Account Security Management](#)

Login Profiles ?

To configure a login profile, click a link in the "Login ID" column.

	Login ID	Role	Active Sessions	Last Login	Password Compliant	Days Until Password Expires	Dormant	State
1	USERID	S	1	02/21/12 05:04:29	Yes	n/a		Active
2	xhj	S	0	09/24/09 15:21:24	Yes	n/a		Active
3	tester	S	0	09/01/11 08:35:30	Yes	n/a		Active
4	~ not used ~							
5	~ not used ~							
6	~ not used ~							
7	~ not used ~							
8	~ not used ~							
9	~ not used ~							
10	~ not used ~							
11	~ not used ~							
12	~ not used ~							

Group Profiles for Active Directory Users ?

Use this section to configure group authorization profiles.

Figure 2-4 AMM Login Profiles page

2. Click each user account. Document the information of each login profile that is defined with the login ID and the role or access level that is assigned to each profile.

3. Click **Account Security Management**. As shown in Figure 2-5, the advanced settings for Account Security Management are displayed. You document items such as authentication method, session timeout, and security level here.

Account Security Management ?

User authentication method

Local only

Web inactivity session timeout

User picks timeout

CLI inactivity session timeout (seconds)

0

Number of simultaneous active sessions for LDAP users

0

Do not log new authentication events for the same user for

5 minutes

Ignore client IP address when tracking user authentication events

☐

Account security level:

Security Level	Details
<input checked="" type="radio"/> Legacy security settings	No password required No password expiration No password re-use restrictions No password change frequency restrictions Account is locked for 2 minutes after 5 login failures Simple password rules No account inactivity monitoring
<input type="radio"/> High security settings	Password required Factory default 'USERID' account password must be changed on next login Force user to change password on first login Passwords expire in 90 days Password re-use checking enabled (last 5 passwords kept in history) Minimum 24 hour interval between password changes Account is locked for 60 minutes after 5 login failures Complex password rules with 2 degrees of difference from previous password Alert on account inactivity after 120 days Accounts disabled after 180 days of inactivity
<input type="radio"/> Custom security settings	<div>Edit Security Settings</div>

Figure 2-5 Account Security Management settings page

Determining network settings

The Network Interfaces view displays the network and IP setting for AMM. Follow these steps to determine the current network settings:

1. In the navigation bar, click **MM Control** → **Network Interfaces**. The External Network Interface page opens, as shown in Figure 2-6.

IBM BladeCenter® E Advanced Management Module

Welcome USERID

Bay 1: SN#YK14807741JR

- Monitors
 - System Status
 - Event Log
 - LEDs
 - Power Management
 - Hardware VPD
 - Firmware VPD
 - Remote Chassis
- Blade Tasks
- I/O Module Tasks
- MM Control
 - General Settings
 - Login Profiles
 - Alerts
 - Serial Port
 - Port Assignments
 - Network Interfaces
 - Network Protocols
 - Chassis Int Network
 - Security
 - File Management
 - Firmware Update
 - Configuration Mgmt
 - Restart MM
 - License Manager
- Service Tools
 - AMM Service Data
 - Blade Service Data
 - AMM Status
 - Service Advisor

External Network Interface (eth0) ?

Interface: Enabled

☒ IPv6 Enabled

☐ Hide all IPv6 configuration fields when IPv6 is disabled. [more...](#)

Primary Management Module ?

This management module is in **Bay 1** of the chassis

Hostname: MM00145EE1A0FE

Domain name:

Register this interface with DNS: ☐

[Advanced Ethernet Setup](#)

IPv4

DHCP: Disabled - Use static IP configuration

*** Currently the static IP configuration is active for this interface. *** This static configuration is active.

IPv4 Static IP Configuration

IP address	9.125.90.213
Subnet mask	255.255.255.0
Gateway address	9.125.90.1

IPv6

Link local address: fe80::214:5eff:fedf:8068

IPv6 static IP configuration: Enabled

IP address	2002:325b:1000::97d:5ad5
Address prefix length (1-128)	64
Default route	0::0

DHCPv6: Enabled

Figure 2-6 AMM Module page

2. From this page, you can find the IPv4 and IPv6 settings for the AMM, including the DHCP setting, IP address, mask, and gateway configuration information.

For more information about these and other settings in the AMM, see the *Advanced Management Module User's Guide* at this website:

http://download.boulder.ibm.com/ibmdl/pub/systems/support/system_x_pdf/00d3237.pdf

You now have most (if not all) of the relevant customizations that are made in the AMM. You can access the CMM in the Flex System Enterprise Chassis to apply the new settings. Log in with your CMM credentials, as shown in Figure 2-7 on page 21.

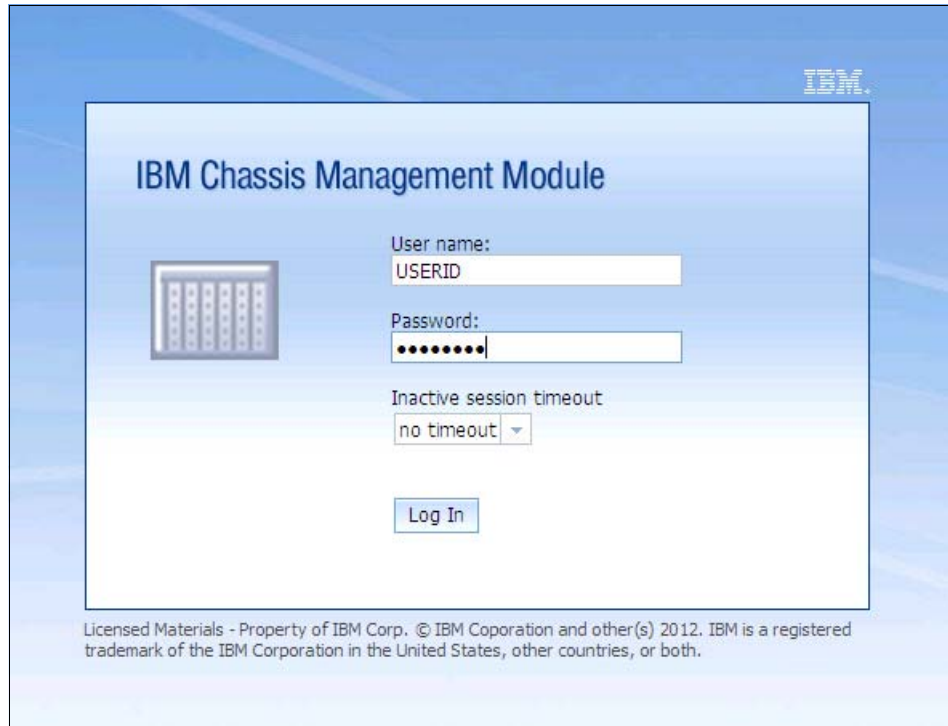


Figure 2-7 CMM login window

Applying power management settings

Follow these steps to apply the current settings to the CMM:

1. From the menu bar at the top of the main CMM window, click **Chassis Management** → **Power Modules and Management**, as shown in Figure 2-8. The Power Modules and Management window opens, as shown in Figure 2-9 on page 22.

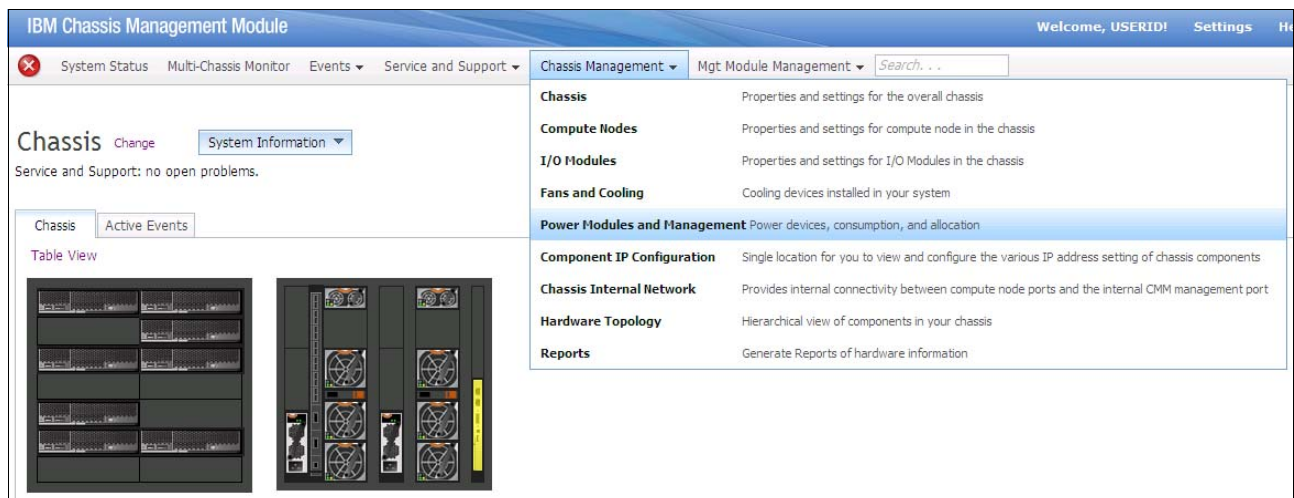


Figure 2-8 Chassis Management Module menu

2. Click **Change** under the No Power Policy section (as shown in Figure 2-9) to set or change the power policy.

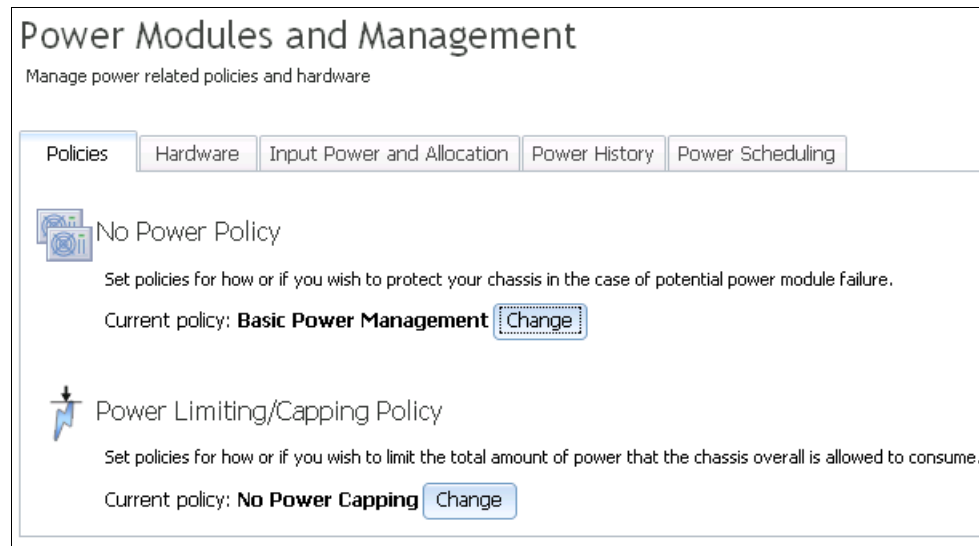


Figure 2-9 Power Modules and Management window

3. The Change Power Policy window opens, as shown in Figure 2-10 on page 23. In this window, there are five options available (three are identical to the policies on the AMM). The Flex System Enterprise Chassis includes more power options than BladeCenter. You must understand the meaning of the additional options if you want to use them.

	Power Supply Failure Limit	Maximum Input Power (Watts)	Estimated Actual Input Power for your chassis (percent)
<input type="radio"/> AC Power Source Redundancy Intended for dual AC power sources into the chassis. Maximum input power is limited to the capacity of two power modules. This is the most conservative approach and is recommended when all four power modules are installed. When the chassis is correctly wired with dual AC power sources, one AC power source can fail without affecting compute node server operation. Note that some compute nodes may not be allowed to power on if doing so would exceed the policy power limit.	1	2745	24
<input type="radio"/> AC Power Source Redundancy with Compute Node Throttling Allowed Very similar to the AC Power Source Redundancy. This policy allows higher input power, however capable compute nodes may be allowed to throttle down if one AC power source fails.	1	3538	19
<input type="radio"/> Power Module Redundancy Intended for a single AC power source into the chassis where each Power Module is on its own dedicated circuit. Maximum input power is limited to one less than the number of Power Modules when more than one Power Module is present. One Power Module can fail without affecting compute node operation. Multiple Power Module failures can cause the chassis to power off. Note that some compute nodes may not be allowed to power on if doing so would exceed the policy power limit.	1	2745	24
<input type="radio"/> Power Module Redundancy with Compute Nodes Throttling Allowed Very similar to Power Module Redundancy. This policy allows higher input power; however, capable compute nodes may be allowed to throttle down if one Power Module fails.	1	5490	12
<input checked="" type="radio"/> Basic Power Management Maximum input power is higher than other policies and is limited only by the nameplate power of all the Power Modules combined. This is the least conservative approach, since it does not provide any protection for AC power source or Power Module failure. If any single power supply fails, compute node and/or chassis operation may be affected.	0	5490	12

OK Cancel

Figure 2-10 CMM Change Power Policy window

- Click **Change** in the Power Limiting/Capping Policy section of the Power Modules and Management window, as shown in Figure 2-9 on page 22. The Change Power Policy window opens, as shown in Figure 2-11. From this window, you can set or change the power capping policy. By setting a specific power static capping limit value, you can control the chassis power consumption that is suitable for your environment.

Change Power Capping Policy	
<input type="radio"/> No Power Capping Maximum input power will be determined by the active Power Redundancy policy.	
<input checked="" type="radio"/> Static Capping Sets an overall chassis limit on the maximum input power. In a situation where powering on a component would cause the limit to be exceeded, the component would not be permitted to power on.	
0% 20% 40% 60% 80% 100% 	4057 Watts (Range 1979 - 5490) 74 % of max allocation
OK Cancel	

Figure 2-11 Change Power Capping Policy window

Applying user account settings

Follow these steps to apply the current settings to the CMM. You need to decide which user accounts and groups to migrate:

5. From the CMM, click **Mgt Module Management** → **User Accounts**. The User Accounts window opens, as shown in Figure 2-12.

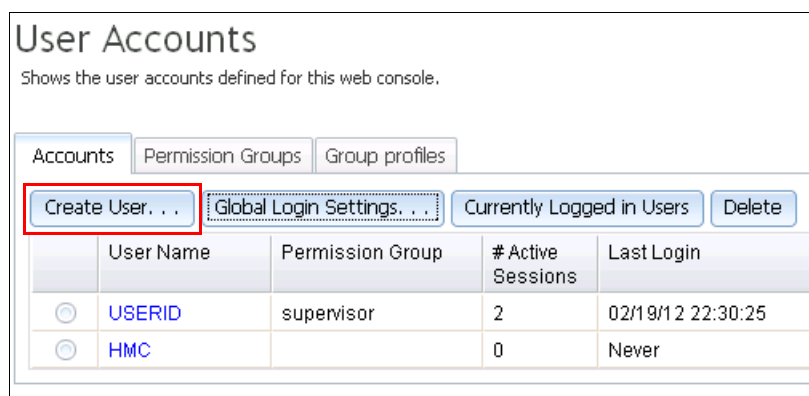


Figure 2-12 User Accounts window

6. Click **Create User** (see Figure 2-12) and re-create the same user IDs and passwords that are used with the source AMM.
7. In the Group profiles tab, click **Add a Group** (as shown in Figure 2-13) to create the same group profile as the source AMM.

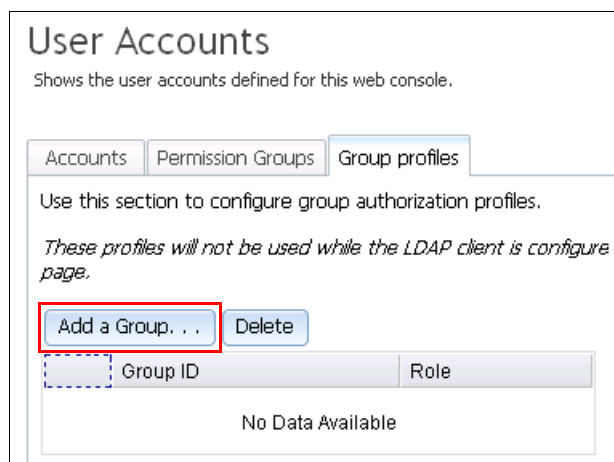


Figure 2-13 Group profiles tab

8. In the Accounts tab, click **Global Login Settings**. The Global Login Settings window opens, as shown in Figure 2-14 on page 25.

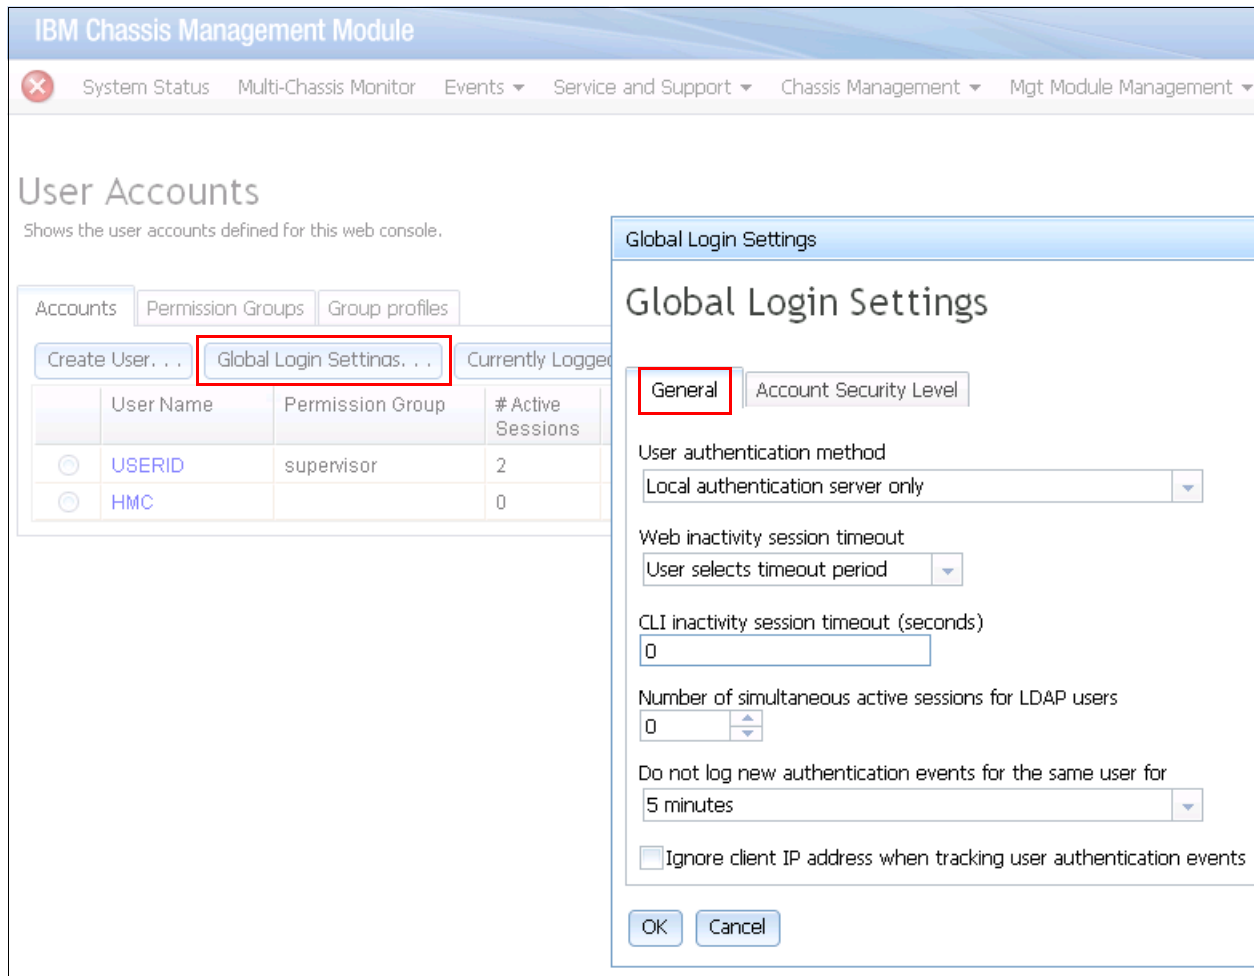


Figure 2-14 Global Login Settings window

9. Use the same settings here that you used in the source AMM.
10. Click the **Account Security Level** tab and choose the needed security settings, as shown in Figure 2-15.



Figure 2-15 Global Login Settings window

Applying network settings

Apply the network settings to the CMM, IPv4, and IPv6 by selecting **Mgt Module Management** → **Network**, as shown in Figure 2-16.

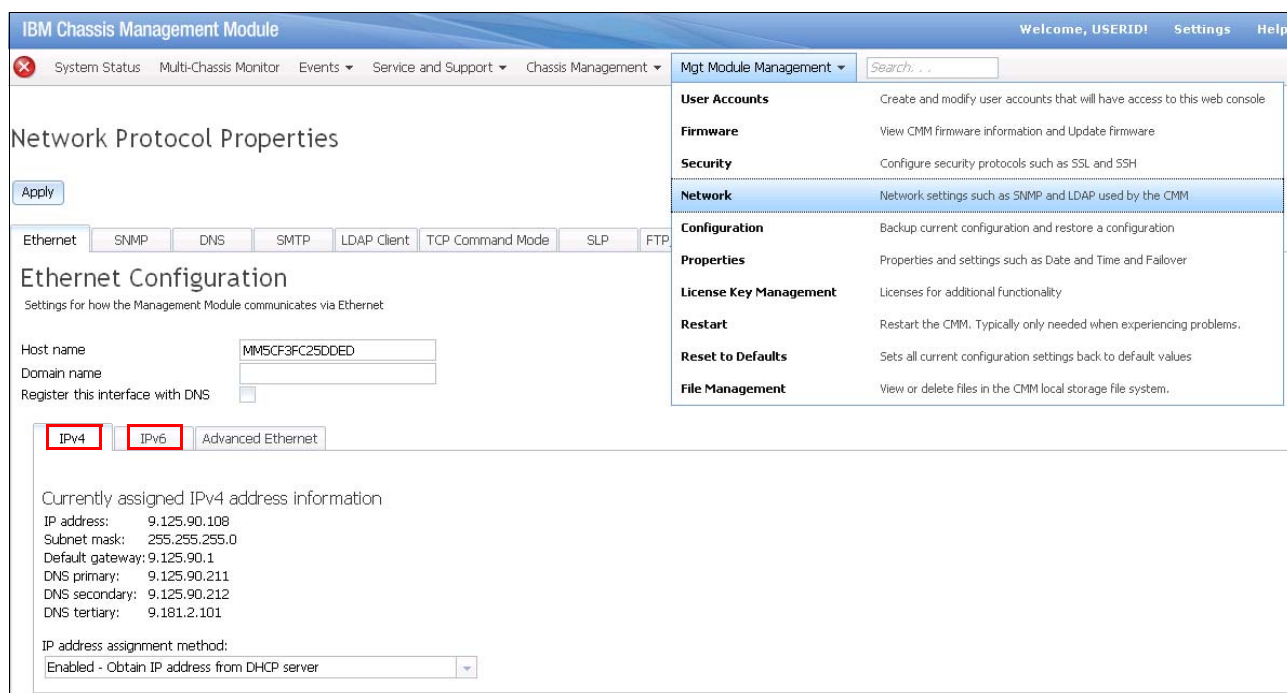


Figure 2-16 Network Protocol Properties window

After completing these migration steps, you might consider migrating more customized settings to their specific environment. For more information, see the *Chassis Management Module User's Guide* at this website:

<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>

2.1.2 Migrating compute node settings

In this section, we describe how the configuration information of a compute node (server node) is obtained. The configuration information includes the OS type, host network settings, virtual machine utilization, and host resource allocation.

You must ascertain the blade setting through the AMM interface. Knowing the blade setting saves time later in the migration process.

Follow these steps to determine the current blade settings:

1. Log in to the AMM with your AMM credentials.
2. In the navigation bar, click **System Status**. In the Blades section (as shown in Figure 2-17 on page 27), the status of all chassis bays is shown, including information about the installed blade servers.

IBM BladeCenter® H Advanced Management Module

Welcome USERID About | Help | Logout

Bay 1: SN#YK17808BD1PY

Monitors

- System Status
- Event Log
- LEDs
- Power Management
- Hardware VPD
- Firmware VPD
- Remote Chassis

Blade Tasks

- I/O Module Tasks
- MM Control
- Service Tools
- Scalable Complex

Blades ?

Click the icon in the Status column to view detailed information about each blade.

Bay	Status	Name	Pwr	Owner**		cKVM*	I/O Compatibility	WOL*	Local Cont	
				KVM	MT*				Pwr	KVM
1	✓	Mako_TCD_B034	On	✓			OK	On	✓	✓
2	✓	HX5TCTB023	On				OK	On	✓	✓
3	✗	Comm Error	---				Unknown	---	✓	✓
4	✓	Crichton5	On				OK	On	✓	✓
5	✗	HS22V_Mako_SIT	On				OK	On	✓	✓
6	✓	Crichto_TCTB007	On				OK	On	✓	✓
7	✓	SN#Y010UF13F050	Off				OK	On	✓	✓
8	✓	None	On				OK	On	✓	✓
9	✓	Mongoos_TCTB014	On				OK	On	✓	✓
10	✓	TCD8010_TBU	On				OK	On	✓	✓
11	✗	Comm Error	---		✓		Unknown	---	✓	✓
12	✓	Crichto_TCD8023	On			✓	OK	On	✓	✓
13	✓	TCD8021_TBU	On				OK	On	✓	✓
14	✓	Mongoos_TCTB013	On				OK	On	✓	✓

* MT = Media Tray (CD/ USB) , WOL = Wake on LAN , BEM = Blade Expansion Module
 BSE1 (BSE2,BSE3) = Blade Storage Expansion 1st Generation (2nd Generation, 3rd Generation)
 PEU1 = PCI Expansion Unit 1st Generation PEU2 = PCI Expansion Unit II BPE3/BPE4 = PCI Express Expansion Unit

Figure 2-17 System Status window

3. Create a mapping list that includes the host name and its host IP addresses. Use the remote presence feature in the AMM to obtain the IP address. The OS must be running before completing this step.
4. In the navigation bar, click **Blade Tasks** → **Remote Control**. The Remote Control Status window opens, as shown in Figure 2-18.

Remote Control Status ?

Firmware status: Active

KVM owner (since 12/22/2011 09:29:17): Blade3 - SN#Y030UN19R00R

Media tray owner (since 12/22/2011 09:00:21): Blade8 - HS22VTCTB031

Console redirect: No session in progress.

Refresh

Start Remote Control ?

Click "Start Remote Control" to control a blade remotely. A new window will appear that provides access to the Remote Console and Remote Disk functionality. On this window, you will have full keyboard and mouse control of the blade which currently owns the KVM. You will also be able to change KVM and media tray ownership.

Note: An Internet connection is required to download the Java Runtime Environment (JRE) if the Java Plug-in is not already installed. Remote Control is supported for Sun JRE 6.0 update 10 or later versions.

Start Remote Control

Figure 2-18 Remote Control Status window

Tip: For more information about the Remote Control features of the AMM, see the *Advanced Management Module User's Guide* at this website:

<http://ibm.com/support/entry/portal/docdisplay?lnidocid=MIGR-5073887>

5. In the Remote Control Status window, select each server in turn, and log in and use the OS windows to determine the host IP addresses of each server. For example, for Windows, you can use Network Connections by way of the Windows Control Panel to determine the host IP address, as shown in Figure 2-19.

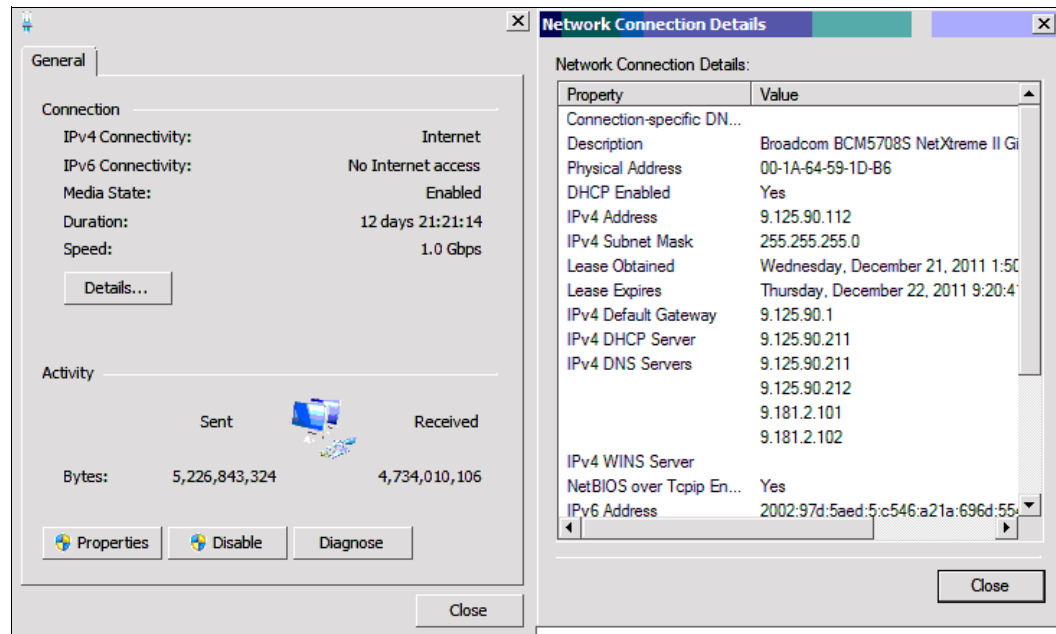


Figure 2-19 Windows networking window that shows the host IP address

For Linux, the IP address is displayed by using **ifconfig**, as shown in Figure 2-20.

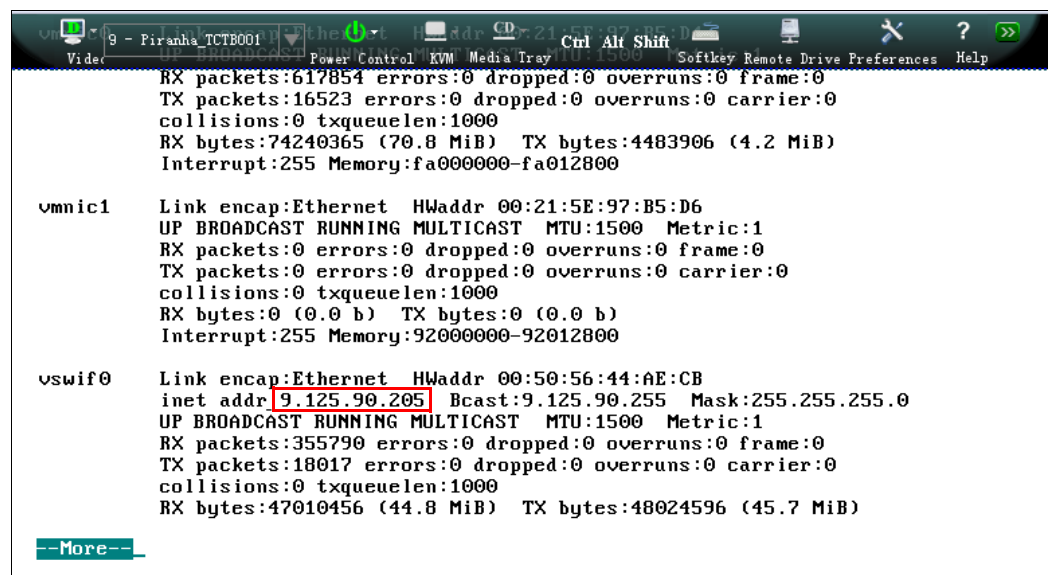


Figure 2-20 Determining the IP address of a Linux server

For servers that are installed with bare metal hypervisors, such as VMware vSphere (formerly ESX or ESXi), the method to obtain the IP address is different. In most cases, you can obtain the IP address directly from the panel that is shown in Figure 2-21.

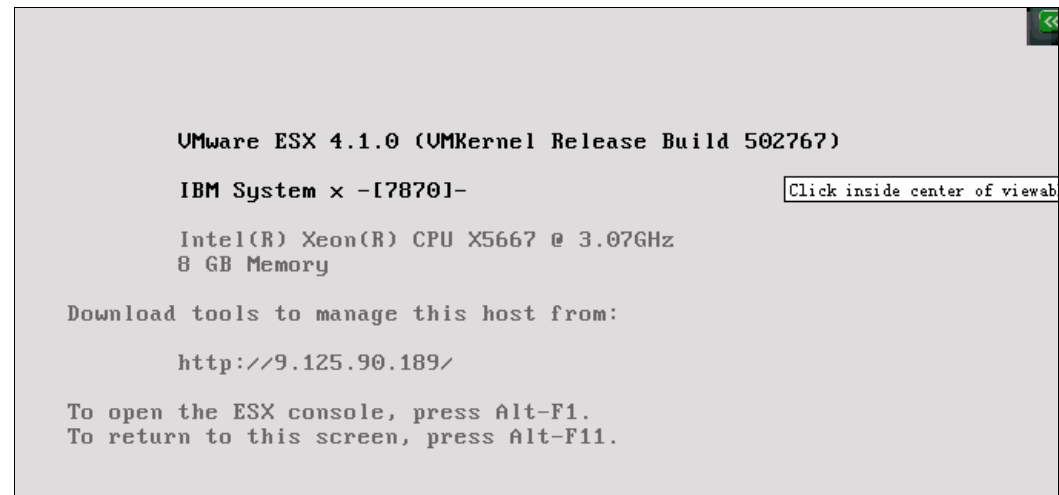


Figure 2-21 Determining the IP address of a VMware ESX server

You can build a comprehensive spreadsheet that includes the chassis identifier, blade server bay number, host name, operating system (OS) type, and host IP address. This information is critical to the migration process and is regarded as the starting point to create the workload migration plan.

You must also consider the mapping relationship of the target compute nodes, including the virtualized and physical computing nodes. For more information, see Chapter 3, “Migrating operating system images” on page 55. Example 2-1 shows an example of Physical-to-Virtual (P2V) system information.

Example 2-1 Sample collection of system information

Flex System Enterprise Chassis (IP of CMM, related setting from AMM)	
Compute Node1	
OS (VMware vSphere/ESXi 5.0)	
Virtual Machine: P2V from source blade 1 (Windows 2008: 192.168.1.1)	
Virtual Machine: P2V from source blade 2 (Windows 2003: 192.168.1.2)	
Virtual Machine: P2V from source blade 3 (Windows 2003: 192.168.1.3)	
Compute Node2	
OS (Red Hat 5.4)	
Virtual Machine: P2V from source blade 4 (SLES 11: 192.168.1.4)	
Virtual Machine: P2V from source blade 5 (SLES 10.3: 192.168.1.5)	

2.1.3 Migrating UEFI settings

An optional migration step is to copy BIOS or UEFI settings that might be preserved from the source server to the target server. (The source server can be BIOS-based or UEFI-based.) For this task, the first step is to determine the settings that need to be migrated.

Follow these steps to determine the current UEFI settings:

1. When prompted, press **F1** during the system boot process. A System Configuration and Boot Management menu opens that is similar to the menu shown in Figure 2-22.

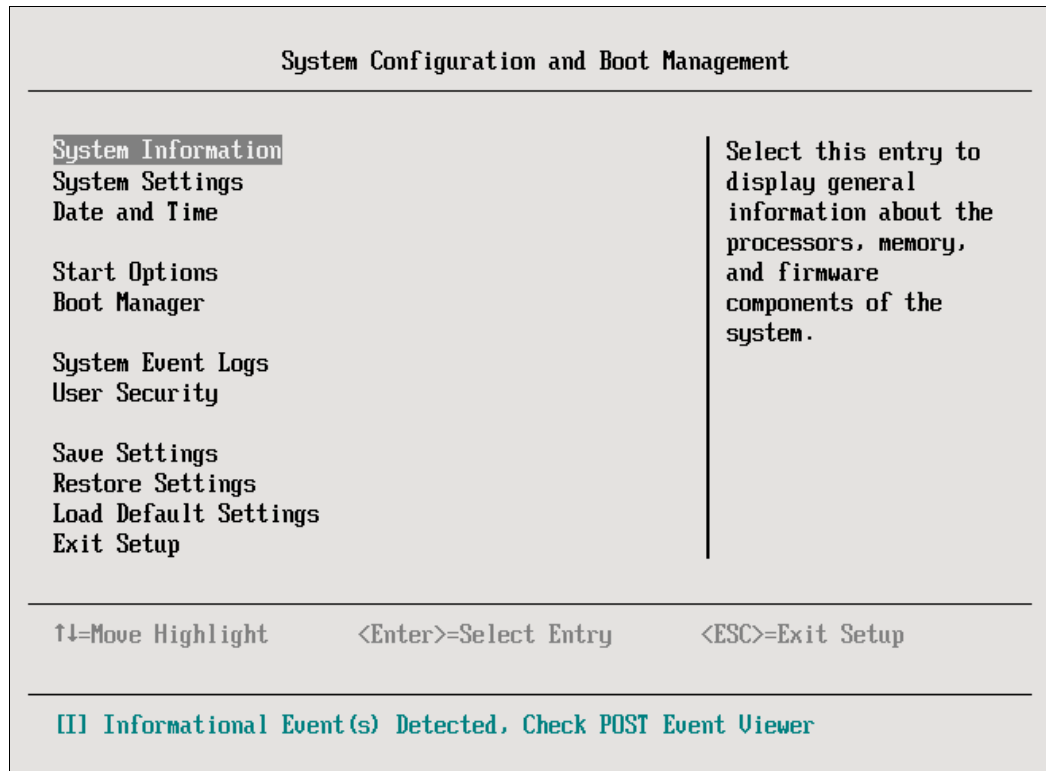


Figure 2-22 System Configuration and Boot Management window

This menu includes the settings for the compute node. A similar configuration menu is available for the compute nodes in a Flex System configuration.

2. Choose the settings that you deem important and use the same or similar settings between the source and target server. Using these same settings helps to ensure that both servers have the same performance, power, capability, and reliability features.

Pay attention to the customized server boot order and security settings, as shown in Figure 2-23.

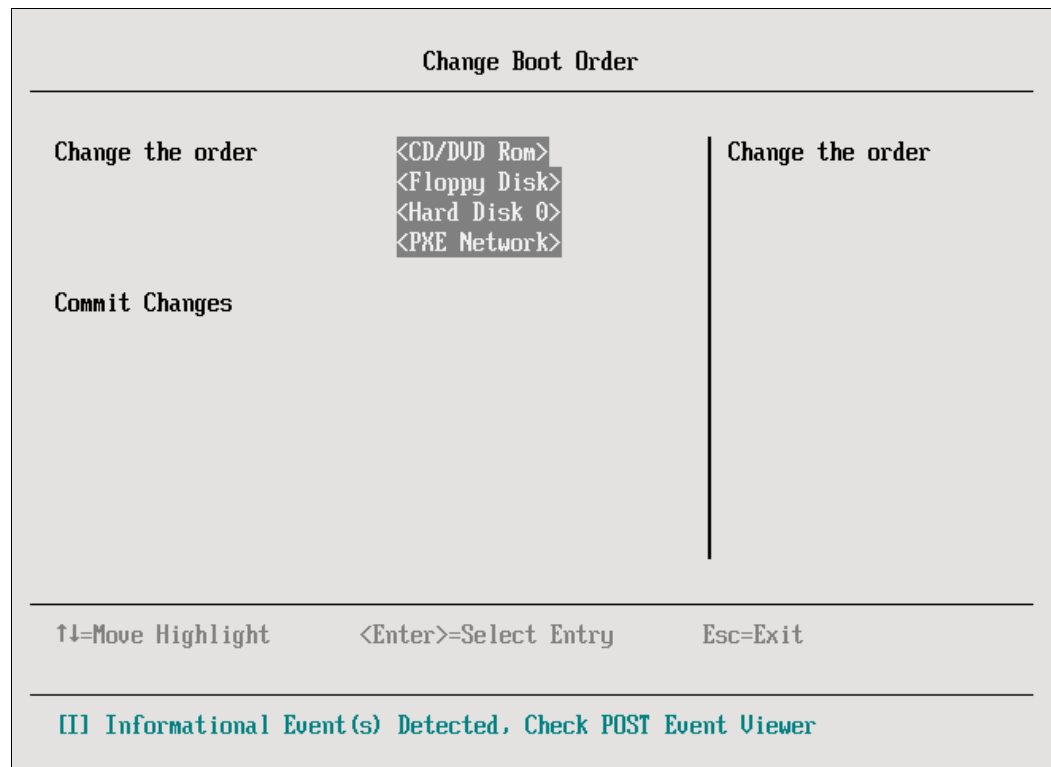


Figure 2-23 Change Boot Order menu

3. Review the Intel Virtualization Technology (Intel VT) setting in the Processors setting page, as shown in Figure 2-24. Most Flex System clients run a virtualization environment, so you must ensure that the Intel VT settings are enabled.




Figure 2-24 Virtualization settings in UEFI

2.2 Migrating network settings

When planning for a server network setting migration, compare the network settings of the server nodes and network switch modules.

Assuming that you have the networking settings from the source servers, a mapping relationship might be needed in the case of static allocation or some required setting changes (for example, DHCP). These new settings are applied to target compute nodes within the utilities by OS, hypervisor, or other methods. You might want to migrate the NIC Adapter Teaming setting to the target system.

The following examples show how to configure Adapter Teaming on different platforms:

- Linux: Adapter Teaming is displayed in the command shell output of `ifconfig`, as shown in Figure 2-25 on page 33. You also can see that two physical adapters, `eth0` and `eth1`, are bonded as `bond0`.

```

bond0      Link encap:Ethernet  HWaddr 00:0C:29:C6:BE:59
          inet addr:192.168.1.20  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::200:ff:fe00:0/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
          RX packets:2804 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1879 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:250825 (244.9 KiB)  TX bytes:244683 (238.9 KiB)
eth0      Link encap:Ethernet  HWaddr 00:0C:29:C6:BE:59
          inet addr:192.168.1.20  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fec6:be59/64 Scope:Link
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:2809 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1390 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:251161 (245.2 KiB)  TX bytes:180289 (176.0 KiB)
          Interrupt:11 Base address:0x1400
eth1      Link encap:Ethernet  HWaddr 00:0C:29:C6:BE:59
          inet addr:192.168.1.20  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::20c:29ff:fec6:be59/64 Scope:Link
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:4 errors:0 dropped:0 overruns:0 frame:0
          TX packets:502 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:258 (258.0 b)  TX bytes:66516 (64.9 KiB)
          Interrupt:10 Base address:0x1480

```

Figure 2-25 *ifconfig* output

- Microsoft Windows Server: The NICs from different vendors have their specified network teaming tools. Use these vendor-specific tools to manage these settings.

Using a Broadcom adapter as an example, the Broadcom Advanced Control Suite can be used to create and manage the network teaming setting. Figure 2-26 shows that “Team1” is bonded by the Ethernet Adapter [0001] and Ethernet Adapter [0010]. Ethernet Adapter [0014] is still a Standby Member (stand-alone NIC adapter).

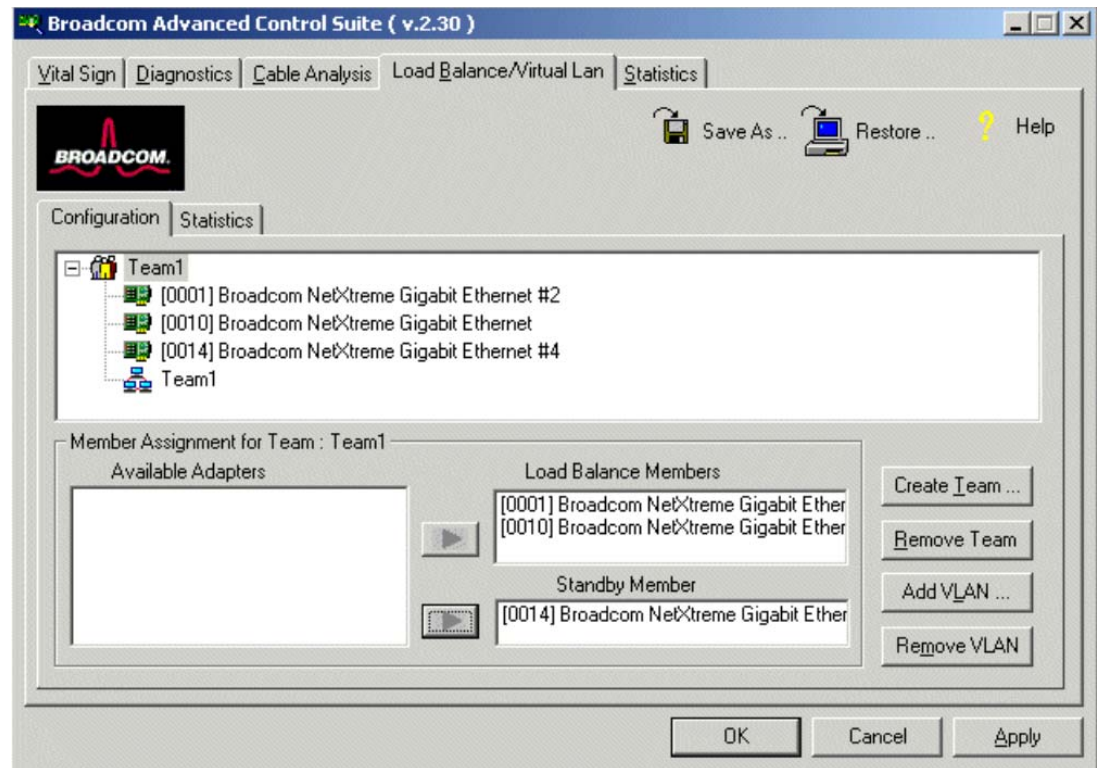


Figure 2-26 Broadcom Advanced Control Suite window

After you complete the settings migration for the network adapter, you must manage the settings for the network switch module.

For the switch modules, the virtual local area network (VLAN) settings must be carefully documented. These windows display a group of hosts with a common set of requirements that communicates as though they were attached to the same broadcast domain, regardless of their physical locations. VLANs are the basic logic network topology of the source environment.

Follow these steps to retrieve the VLAN information from the source server (for this example, BladeCenter):

1. At the source side, log in to the AMM.

2. On the left side of the AMM window, click **I/O Module Tasks** → **Configuration**. The I/O modules that are installed in the source system are displayed, as shown in Figure 2-27.

IBM BladeCenter® E Advanced Management Module

Bay 1: SN#YK14807741JR

- Monitors
 - System Status
 - Event Log
 - LEDs
 - Power Management
 - Hardware VPD
 - Firmware VPD
 - Remote Chassis
- Blade Tasks
- I/O Module Tasks
 - Admin/Power/Restart
 - Configuration**
 - Firmware Update
- MM Control
- Service Tools

I/O Module Configuration

IPv6 Support Slot 1

IPv6 Support and Status

Click the checkboxes in the first column to select one or more I/O modules; then, click one of the actions in the action list below the table and click Perform Action to perform the desired action.

<input type="checkbox"/>	Bay	Name	IPv6 state
<input type="checkbox"/>	1	Ethernet SM	<i>not supported</i>
<input type="checkbox"/>	2	<i>Not installed</i>	
<input type="checkbox"/>	3	<i>Not installed</i>	
<input type="checkbox"/>	4	<i>Not installed</i>	

Available actions

Enable IPv6

* SM = Switch Module, CM = Concentrator Module, PM = Pass-thru Module

Fri, 06 Jan 2012 07:18:45

Figure 2-27 I/O Module Configuration window

- Click the tab that is associated with the I/O module bay. In our example, this tab is **Slot 1**, as shown in Figure 2-28.

The screenshot shows the 'I/O Module Configuration' page for Slot 1. The left sidebar lists navigation options: Monitors (System Status, Event Log, LEDs, Power Management, Hardware VPD, Firmware VPD, Remote Chassis), Blade Tasks, I/O Module Tasks (Admin/Power/Restart, Configuration, Firmware Update), MM Control, and Service Tools. The 'Configuration' option under I/O Module Tasks is highlighted. The main content area is titled 'I/O Module Configuration' and has tabs for 'IPv6 Support', 'Slot 1', and 'Slot 2'. Under 'Current IP Configuration', it shows: Configuration method: Static, IP address: 9.123.198.81, Subnet mask: 255.255.255.0, and Gateway address: 9.123.198.1. A note states: 'To change the IP configuration for this I/O module, fill in the following fields and click "Save". This will save and enable the new IP configuration.' Below this is the 'New Static IP Configuration' section with fields for Configuration status (Enabled), IP address (9.123.198.81), Subnet mask (255.255.255.0), and Gateway address (9.123.198.1). A link for 'Advanced Options' is at the bottom. A footnote indicates: '* SM = Switch Module, CM = Concentrator Module, PM = Pass-thru Module'. The timestamp 'Sat, 14 Jun 2012 12:15:28' is at the bottom.

Figure 2-28 Information for the I/O module in Slot 1

- Find the switch module type. (The type of device in Figure 2-28 is an Ethernet switch module as indicated in Figure 2-27 on page 35.) For each module type, you must determine whether you want the same configuration on the target side.
- To find information such as the settings for the switch module VLAN, click **Advanced Options**. Scroll down to the Start Web Session section, as shown in Figure 2-29.

The screenshot shows the 'Start Web Session' section. It includes a heading 'Start Web Session' with a help icon. Below is a instruction: 'Choose your session parameters below, and then click Start Session. All available options for this module will be shown.' There are two input fields: 'IP Address' with the value '9.125.90.242' and 'Security' with the value 'Unsecure'. A 'Start Session' button is located at the bottom right.

Figure 2-29 Start Web Session section

- Click **Start Session**. You see the web interface of the switch you selected. In our example, the switch is a Lenovo RackSwitch switch on which BladeOS is running, as shown in Figure 2-30. If you have different switches that are installed, the interface might be different.

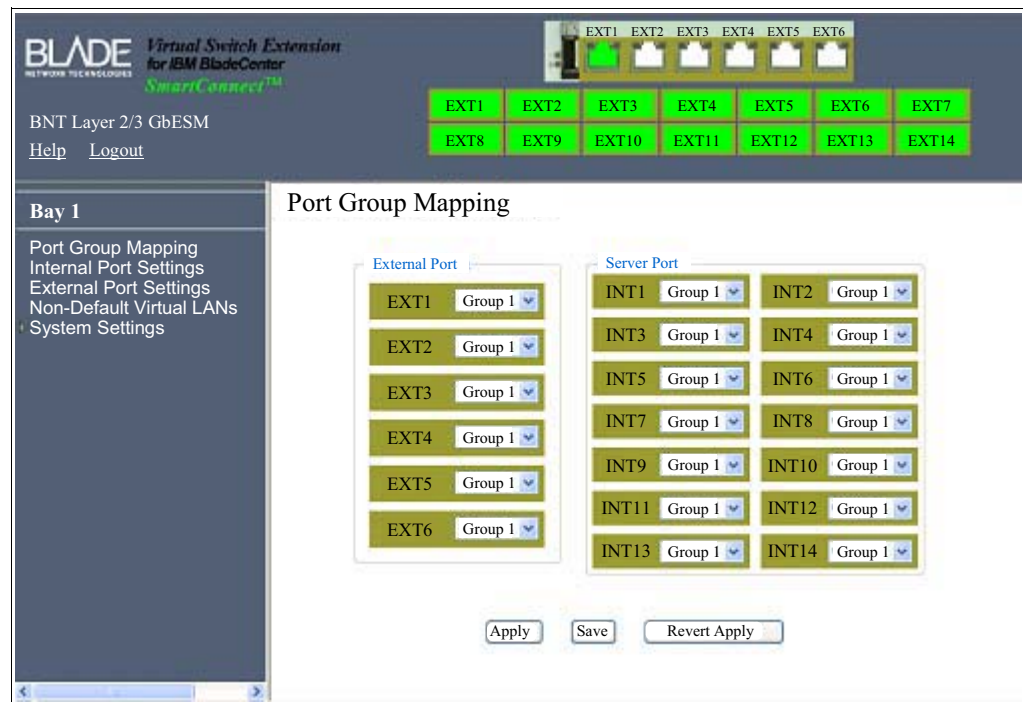


Figure 2-30 Switch web interface

Figure 2-30 shows that all the blade servers connected to the specific switch are in one group (Group 1) and use static VLAN (port-based) information. For other switch devices, you might see the dynamic VLAN setting with the Media Access Control (MAC) address of the server that is connected to the port. However, you can ascertain the logic diagram by examining the types of VLAN settings.

Using the information that is gathered from the figures in the previous steps, follow these steps to configure the same network topology on the Flex System chassis:

- Log in to the CMM.

- From the menu bar, click **Chassis Management** → **Component IP configuration**. The Component IP Configuration window opens, as shown in Figure 2-31.

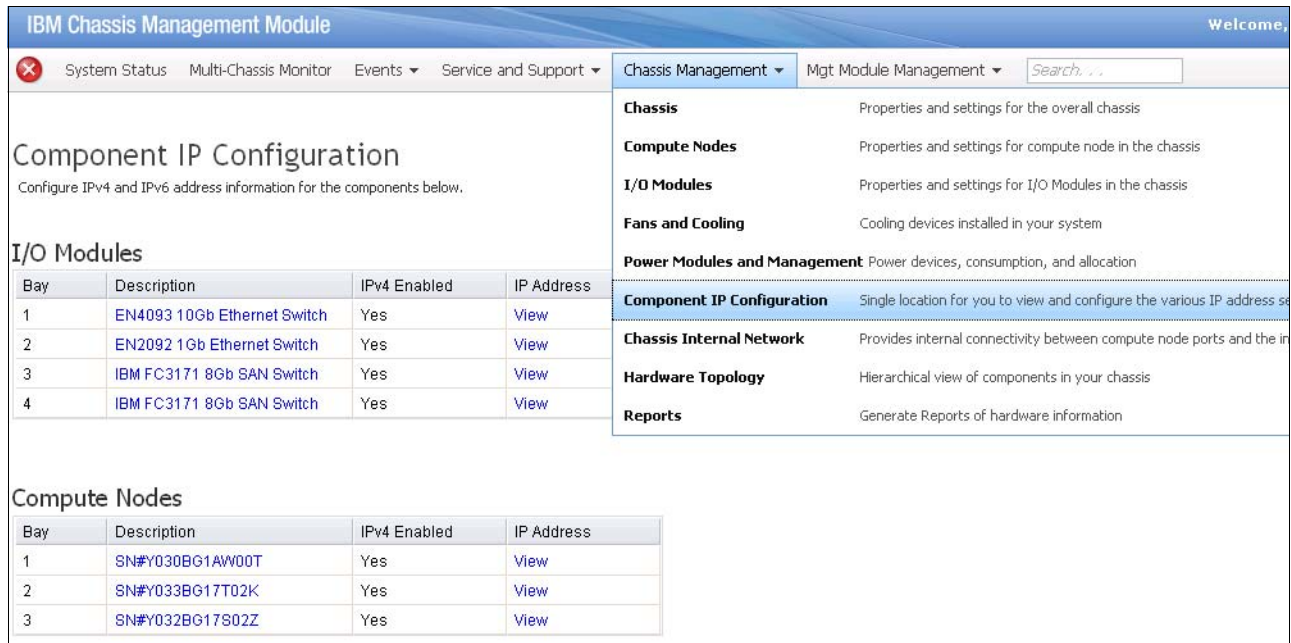


Figure 2-31 Component IP Configuration window

- Click an Ethernet switch. An Ethernet Switch Configuration wizard opens, as shown in Figure 2-32.

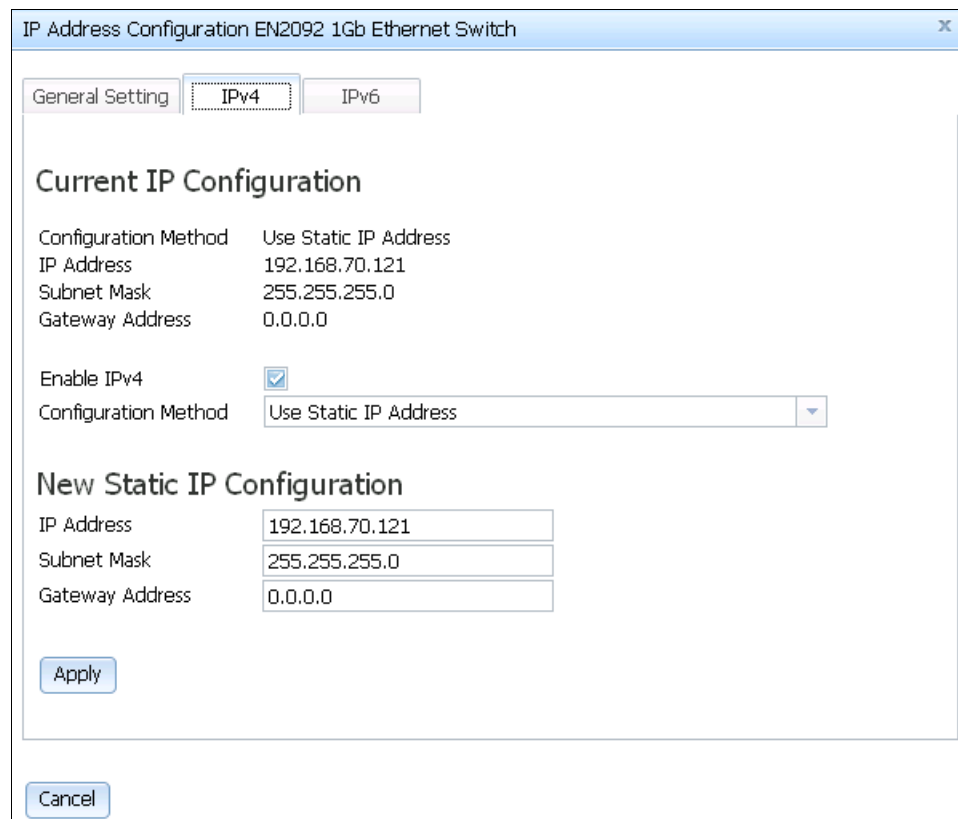


Figure 2-32 Ethernet Switch Configuration wizard

- Copy into a web browser the IP address that is listed in the Current IP Configuration section of the wizard. A login window opens, as shown in Figure 2-33.

Figure 2-33 EN2092 Ethernet Switch login window

- Log in with the credentials you customized or use the following default login information:
 - User name: USERID
 - Password: PASSW0RD (where “0” is zero)
- Click **Configure** at the top of the window. In the navigation bar, click **Flex System EN2092 1Gb ScSE** → **Switch Ports**, as shown in Figure 2-34.

IBM

Configure

Statistics

Dashboard

Networking OS

Apply

Save

Revert

Diff

Dump

Show Log

Help

Logout

67. Feb 13 17:44:21 9.125.90.113 WARNING ntp: cannot contact primary NTP server fe80::200:ff:fe00:0

IBM Flex System EN2092 1Gb ScSE

System

Switch Ports

Port-Based Port Mirroring

Layer 2

RMON Menu

Layer 3

QoS

Access Control

Virtualization

Switch Ports Configuration

Switch Port	State	VLAN Tagging	Default PVID	PVID tagging	Multicast Threshold	Broadcast Threshold	Destination Fail Threshold	Lookup Priority	802.1p Priority	FDB Learning
INT1A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT2A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT3A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT4A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT5A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT6A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT7A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT8A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT9A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT10A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT11A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT12A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT13A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
INT14A	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
MGT1	enabled	enabled	4095	disabled	disabled	disabled	disabled	0	enabled	
EXT1	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT2	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT3	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT4	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT5	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT6	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT7	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT8	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT9	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	
EXT10	enabled	disabled	1	disabled	disabled	disabled	disabled	0	enabled	

Figure 2-34 EN2092 Ethernet Switch switch ports configuration

Except for the management port, all of the switch ports belong to PVID 1. This setting is the default setting. All of the compute nodes that are connected to this switch are in one VLAN. This configuration is similar to the configuration that was captured from the source switch.

- If you require additional customization, click **Layer 2** → **Virtual LANs** → **Add vLAN**, as shown in Figure 2-35. Select the ports from the Ports Available list and click **Add** to add them to the VLAN.

Figure 2-35 Add VLAN window

- After submitting the change, you can see that an additional VLAN group is available, as shown in Figure 2-36.

VLAN ID	VLAN Name	State
1	Default VLAN	enabled
2	VLAN 2	enabled
4095	Mgmt VLAN	enabled

Figure 2-36 VLANs Configuration window

- Click **Switch Ports** in the navigation window again. As shown in Figure 2-37, ports 1 and 2 belong to the VLAN group 2 that you created.

The screenshot shows the IBM Networking OS configuration interface. The left sidebar contains a navigation tree with categories like System, Switch Ports, Port-Based Port Mirroring, Layer 2, Layer 3, QoS, Access Control, and Virtualization. The main panel displays the 'Switch Ports Configuration' table.

Switch Port	State	VLAN Tagging	Default PVID	PVID tagging	Multicast Threshold	Broadcast Threshold
INT1A	enabled	disabled	2	disabled	disabled	disabled
INT2A	enabled	disabled	2	disabled	disabled	disabled
INT3A	enabled	disabled	1	disabled	disabled	disabled
INT4A	enabled	disabled	1	disabled	disabled	disabled
INT5A	enabled	disabled	1	disabled	disabled	disabled
INT6A	enabled	disabled	1	disabled	disabled	disabled
INT7A	enabled	disabled	1	disabled	disabled	disabled
INT8A	enabled	disabled	1	disabled	disabled	disabled
INT9A	enabled	disabled	1	disabled	disabled	disabled
INT10A	enabled	disabled	1	disabled	disabled	disabled
INT11A	enabled	disabled	1	disabled	disabled	disabled
INT12A	enabled	disabled	1	disabled	disabled	disabled
INT13A	enabled	disabled	1	disabled	disabled	disabled
INT14A	enabled	disabled	1	disabled	disabled	disabled
MGT1	enabled	enabled	4095	disabled	disabled	disabled
EXT1	enabled	disabled	1	disabled	disabled	disabled
EXT2	enabled	disabled	1	disabled	disabled	disabled
EXT3	enabled	disabled	1	disabled	disabled	disabled
EXT4	enabled	disabled	1	disabled	disabled	disabled
EXT5	enabled	disabled	1	disabled	disabled	disabled
EXT6	enabled	disabled	1	disabled	disabled	disabled
EXT7	enabled	disabled	1	disabled	disabled	disabled
EXT8	enabled	disabled	1	disabled	disabled	disabled
EXT9	enabled	disabled	1	disabled	disabled	disabled
EXT10	enabled	disabled	1	disabled	disabled	disabled

Figure 2-37 The result of creating a second VLAN

2.3 Migrating storage settings

In addition to a compute node and network switch configuration, external storage is an item that you must plan for during the workload migration.

In many cases, critical data and images are on this external storage, but the data source is the same. If data is stored in external storage devices, maintain the data that is in those devices. The best approach is to map the new compute node to the same target logical unit number (LUN), where the data is stored.

In this section, the process to ensure that the images that contain the OS and applications are migrated smoothly is reviewed.

The storage settings migration processes (P2P, P2V, and V2V) for the images also are reviewed. (For more information, see Chapter 3, “Migrating operating system images” on page 55.)

Virtual-to-Physical migration: Although Virtual-to-Physical (V2P) migration is an option, it is rarely used and therefore is not covered in this document.

In the V2V and P2V migration scenarios, the underlying infrastructure must be as similar as possible, unless there is a valid reason against it. The configuration must be copied. If the image on the source server is on a local RAID array (for example, RAID-5), the image on the target server also is based on local RAID-5. The V2V scenario assumes that the same RAID level and additional configuration options are available on the new RAID controller. If no counterpart is found in the new hardware, you must decide which RAID level to use. However, using a local RAID array is still the primary consideration. But, if the image on the source server is on external storage, the target server must be configured in the same manner. Reuse the data segment by mapping the Fibre Channel host bus adapter (HBA) to the same LUN target.

The P2V process is different. When you want to move a workload from a physical system to a virtual system, most of the hardware configuration on the source server becomes meaningless. For instance, an RAID-5 configuration on the source server secures the storage with high availability. However, this configuration is no longer achievable on a virtualized disk for the target server because the hypervisor controls the virtual disk. The hypervisor might have a different policy from a local RAID-5 array and use storage area network (SAN) storage for implementing live migration.

There are a few ways to configure local RAID. If you want to configure the LSI Corporation RAID adapter, use the LSI Corporation Configuration Utility and WebBIOS.

Follow these steps to obtain the information about the RAID volumes in the source servers:

1. During system POST, press Ctrl+C when prompted to start the configuration utility, as shown in Figure 2-38. This utility features a menu style in which the RAID configuration can be defined.

LSI Corp Config Utility v6.30.00.00 (2009.11.12)							
Adapter List Global Properties							
Adapter	PCI Bus	PCI Dev	PCI Fnc	PCI Slot	FW Revision	Status	Boot Order
SAS1064	0B	00	00	00	1.30.10.00-IR	Enabled	0
Esc = Exit Menu F1/Shift+1 = Help Alt+N = Global Properties -/+ = Alter Boot Order Ins/Del = Alter Boot List							

Figure 2-38 LSI Configuration Utility

2. Highlight the selected RAID controller (if it is not selected) and press Enter. In our example, our blade server has an onboard SAS1064 controller. The Adapter Properties window opens, as shown in Figure 2-39 on page 43.

```

LSI Corp Config Utility    v6.30.00.00 (2009.11.12)
Adapter Properties - SAS1064E

```

Adapter	SAS1064
PCI Slot	00
PCI Address (Bus/Dev)	0B:00
MPT Firmware Revision	1.30.10.00-IR
SAS Address	5005076B:08A1EDAC
NVDATA Version	2D.22
Status	Enabled
Boot Order	0
Boot Support	[Enabled BIOS & OS]

```

Esc = Exit Menu      F1/Shift+1 = Help
Enter = Select Item  -/+ /Enter = Change Item

```

Figure 2-39 Adapter properties

- Highlight **RAID Properties** and press Enter. The View Array window opens, as shown in Figure 2-40.

```

LSI Corp Config Utility    v6.30.00.00 (2009.11.12)
View Array - SAS1064E

```

Array	1 of 1
Identifier	LSILOGICLogical Volume 3000
Type	IM
Scan Order	5
Size (MB)	68664
Status	Optimal

Slot Num	Device Identifier	RAID Disk	Hot Spr	Drive Status	Pred Fail	Size (MB)
0	IBM-ESXSCBRBA073C3ETSO NC49C	Yes	No	Primary	No	68664
1	IBM-ESXSCBRBA073C3ETSO NC49C	Yes	No	Secondary	No	68664

```

Esc = Exit Menu      F1/Shift+1 = Help
Enter=Select Item    Alt+N=Next Array  C=Create an array  R=Refresh Display

```

Figure 2-40 View Array - RAID properties

In our example, the RAID level is RAID-1, as indicated by Type: IM (Integrated Mirroring). The logical size of the array is 68664 MB. Two physical drives also are available for the specific machine, as shown in Figure 2-40.

As an alternative interface, you can start the LSI MegaRAID WebBIOS utility by pressing Ctrl+H during POST, as shown in Figure 2-41.

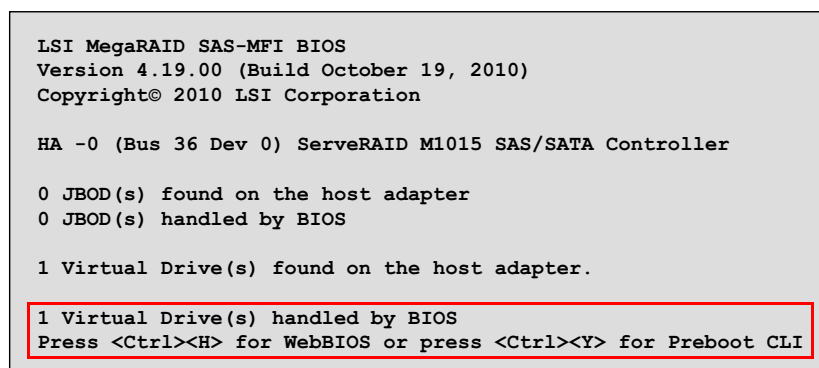


Figure 2-41 Starting WebBIOS

The WebBIOS utility shows the physical and virtual devices and provides you a graphical display to view and configure RAID. The WebBIOS interface is shown in Figure 2-42.

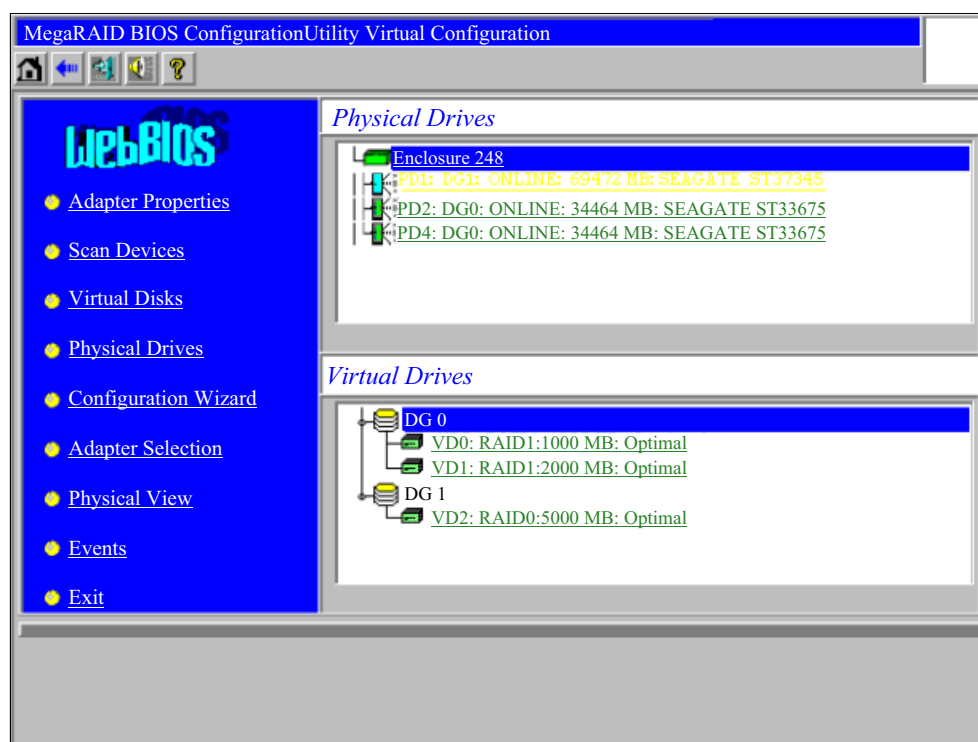


Figure 2-42 MegaRAID BIOS Configuration Utility Virtual Configuration interface

Two virtual drives are available with RAID-1 and RAID-0 enabled separately, as shown in Figure 2-42. By clicking the icons on the left side of the window, you can obtain more information about these drives.

You also can obtain the RAID and disk information of the source server. You can use the same utility to apply the same RAID settings on the compute nodes in the Flex System chassis.

Follow these steps to create new RAID arrays on the Flex System compute nodes by using WebBIOS:

1. Click **Configuration Wizard** in the left navigation window of the MegaRAID BIOS Configuration Utility Virtual Configuration interface, as shown in Figure 2-42 on page 44. The Configuration wizard opens, as shown in Figure 2-43.

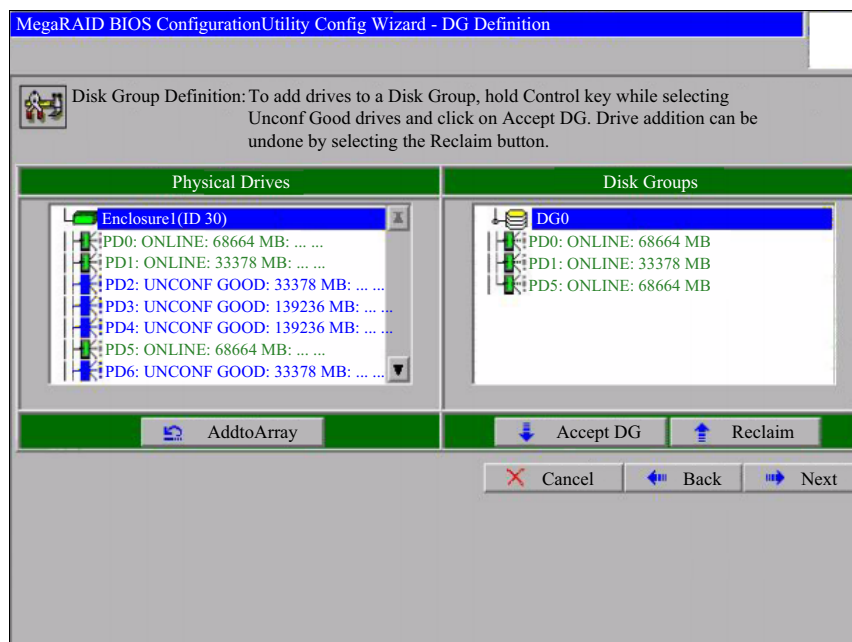


Figure 2-43 MegaRAID WebBIOS Configuration wizard

2. Choose the physical drives that you want to include in a virtual group and click **Next**. The VD Definition wizard opens, as shown in Figure 2-44.

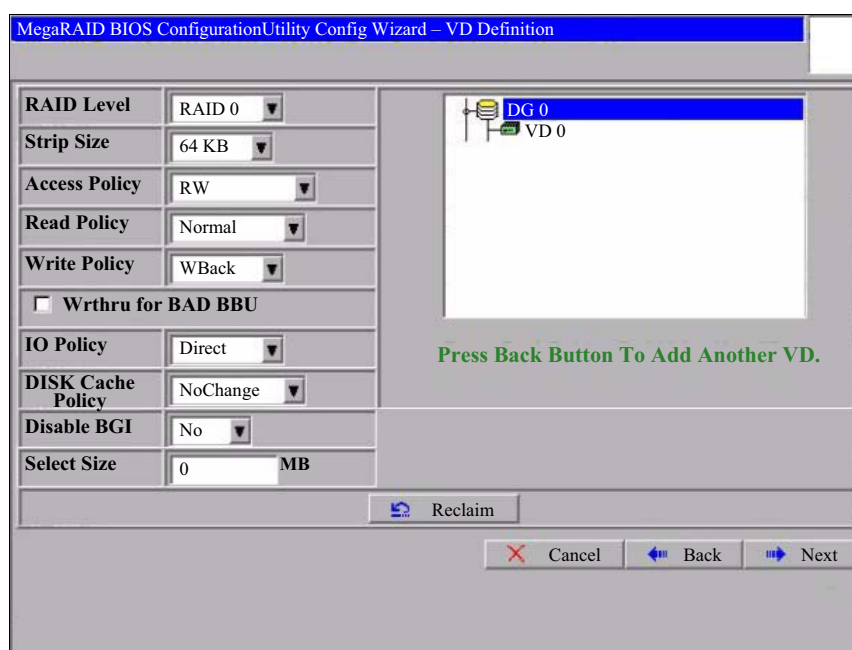


Figure 2-44 Configuring the RAID volume

3. Configure the same RAID and other disk settings from the source server.

The following tools also are available to configure RAID:

- ServerGuide:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=SERV-GUIDE>

- MegaRAID Storage Manager:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5077712>

If you need a lightweight, mouse-compatible utility, ServerGuide is the best option. This option also can help you with a Microsoft Windows OS deployment. MegaRAID Storage Manager is an advanced solution for users who require more configuration flexibility.

For the external storage configuration, you must plan Fibre Channel (FC) switch and storage settings. For switches, zone settings are the most important settings. *Zoning* is the partitioning of a Fibre Channel fabric into smaller subsets to avoid conflicts, add security, and simplify management.

Several virtual disks, or LUNs, are available through the SAN. Each system that is connected to the SAN is allowed access only to a controlled subset of the LUNs. For storage, the configuration is similar to local storage. You need to configure the RAID level, virtual disks, and so on.

Follow these steps to capture and apply the configuration settings of the Fibre Channel switch, including zone settings:

1. Log in to the AMM.
2. In the navigation bar, choose **I/O Module Tasks** → **Configuration**. As shown in Figure 2-45, there is an FC Switch in Slot (Bay) 3.

IBM BladeCenter® H Advanced Management Module

Bay 1: SN#YK17808BD1PY

- Monitors
 - System Status
 - Event Log
 - LEDs
 - Power Management
 - Hardware VPD
 - Firmware VPD
 - Remote Chassis
- Blade Tasks
 - Power/Restart
 - Remote Control
 - Firmware Update
 - Configuration
 - Serial Over LAN
 - Open Fabric Manager
- I/O Module Tasks
 - Admin/Power/Restart
 - Configuration**
 - Firmware Update
- MM Control

I/O Module Configuration

IPv6 Support Slot 1 Slot 3 Slot 9

IPv6 Support and Status

Click the checkboxes in the first column to select one or more I/O modules; then, click one of the actions in the action list below the table and click "Perform Action" to perform the desired action.

<input type="checkbox"/>	Bay	Name	IPv6 state
<input type="checkbox"/>	1	Server Conn Mod	not supported
	2	Not installed	
<input type="checkbox"/>	3	Fibre Channel SM	not supported
	4	Not installed	
	5	Not installed	
	6	Not installed	
	7	Not installed	

Figure 2-45 AMM I/O Module Configuration

3. Click the **Slot 3** tab. The I/O Module Configuration window opens, as shown in Figure 2-46.

I/O Module Configuration

IPv6 Support Slot 1 **Slot 3** Slot 9

Current IP Configuration

Configuration method:	Static
IP address:	9.125.90.164
Subnet mask:	255.255.255.0
Gateway address:	9.125.90.1

To change the IP configuration for this I/O module, fill in the fields and click "Save". This will save and enable the new configuration.

New Static IP Configuration

Configuration status	Enabled ▼
IP address	9.125.90.164
Subnet mask	255.255.255.0
Gateway address	9.125.90.1

[Advanced Options](#)

Figure 2-46 I/O Module Configuration

4. Click **Advanced Options**. Scroll down to Start CLI/Web Session, as shown in Figure 2-47.

Start CLI/Web Session ?

Choose your session parameters below, and then click Start Session. All available options for this module will be shown.

Protocol:	Web ▼
IP Address:	9.125.90.164 ▼
Security:	Unsecure ▼

[Start Session](#)

Figure 2-47 Starting the web interface to the FC switch

5. Click **Start Session** to start a management web page, as shown in Figure 2-48.

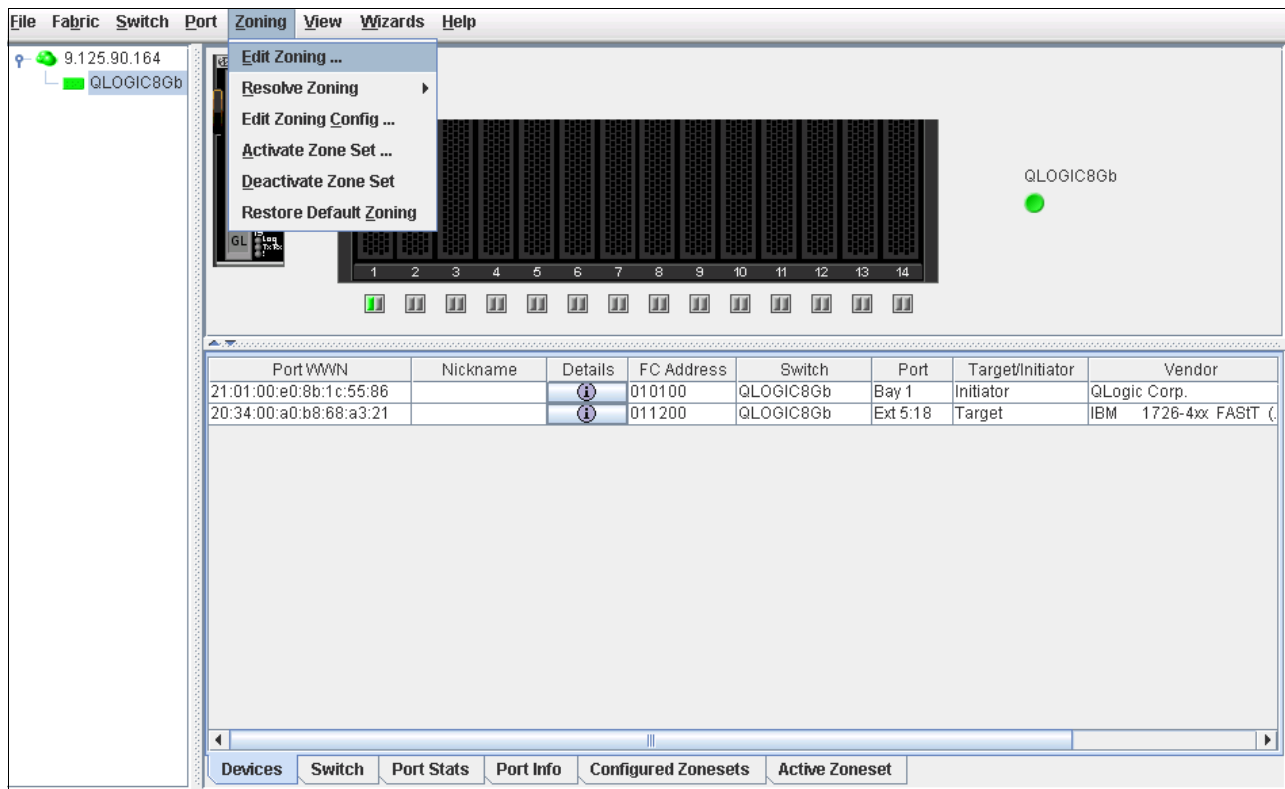


Figure 2-48 Fibre Channel switch web interface

6. From the menu bar, click **Zoning** → **Edit Zoning**. The Edit Zoning windows opens, as shown in Figure 2-49 on page 49.

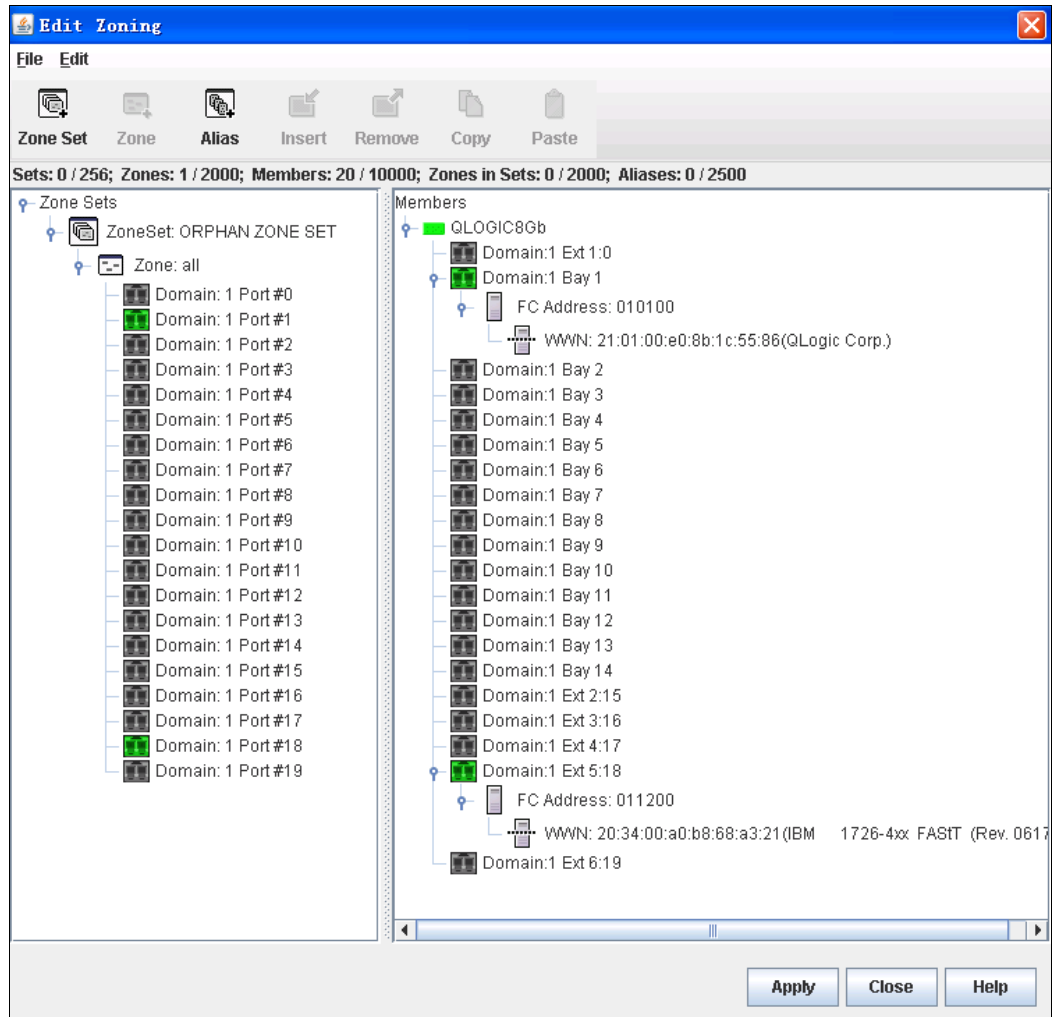


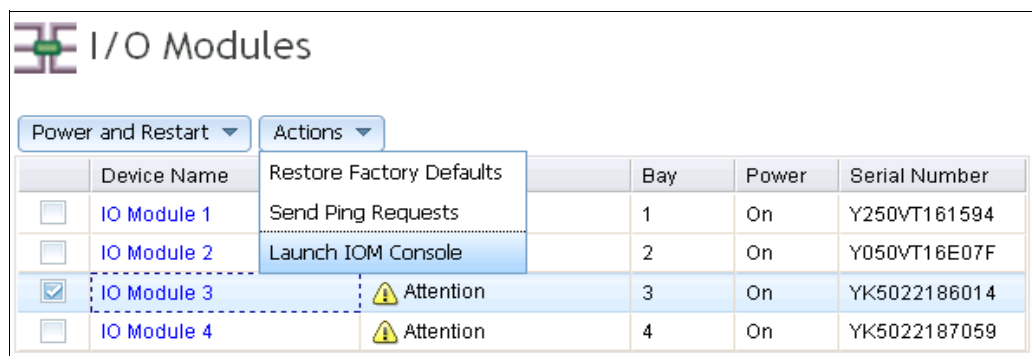
Figure 2-49 Edit Zoning

From these windows, you can see that there is one zone available, all. All the blade servers belong to this zone. You can also find details of the switch from this window.

Follow these steps to move to the Flex System chassis and migrate these settings to the FC switch:

1. Log in to the CMM.

2. From the menu bar, click **Chassis Management** → **I/O modules** to see the IO module list, as shown in Figure 2-50.



	Device Name	Bay	Power	Serial Number
<input type="checkbox"/>	IO Module 1	1	On	Y250VT161594
<input type="checkbox"/>	IO Module 2	2	On	Y050VT16E07F
<input checked="" type="checkbox"/>	IO Module 3	3	On	YK5022186014
<input type="checkbox"/>	IO Module 4	4	On	YK5022187059

Figure 2-50 Launching the I/O Module console

3. Select the SAN switch module (in this example, it is an FC5022 16Gb SAN Scalable Switch). Click **Actions** → **Launch IOM Console**. The web interface for the switch opens, as shown in Figure 2-51.

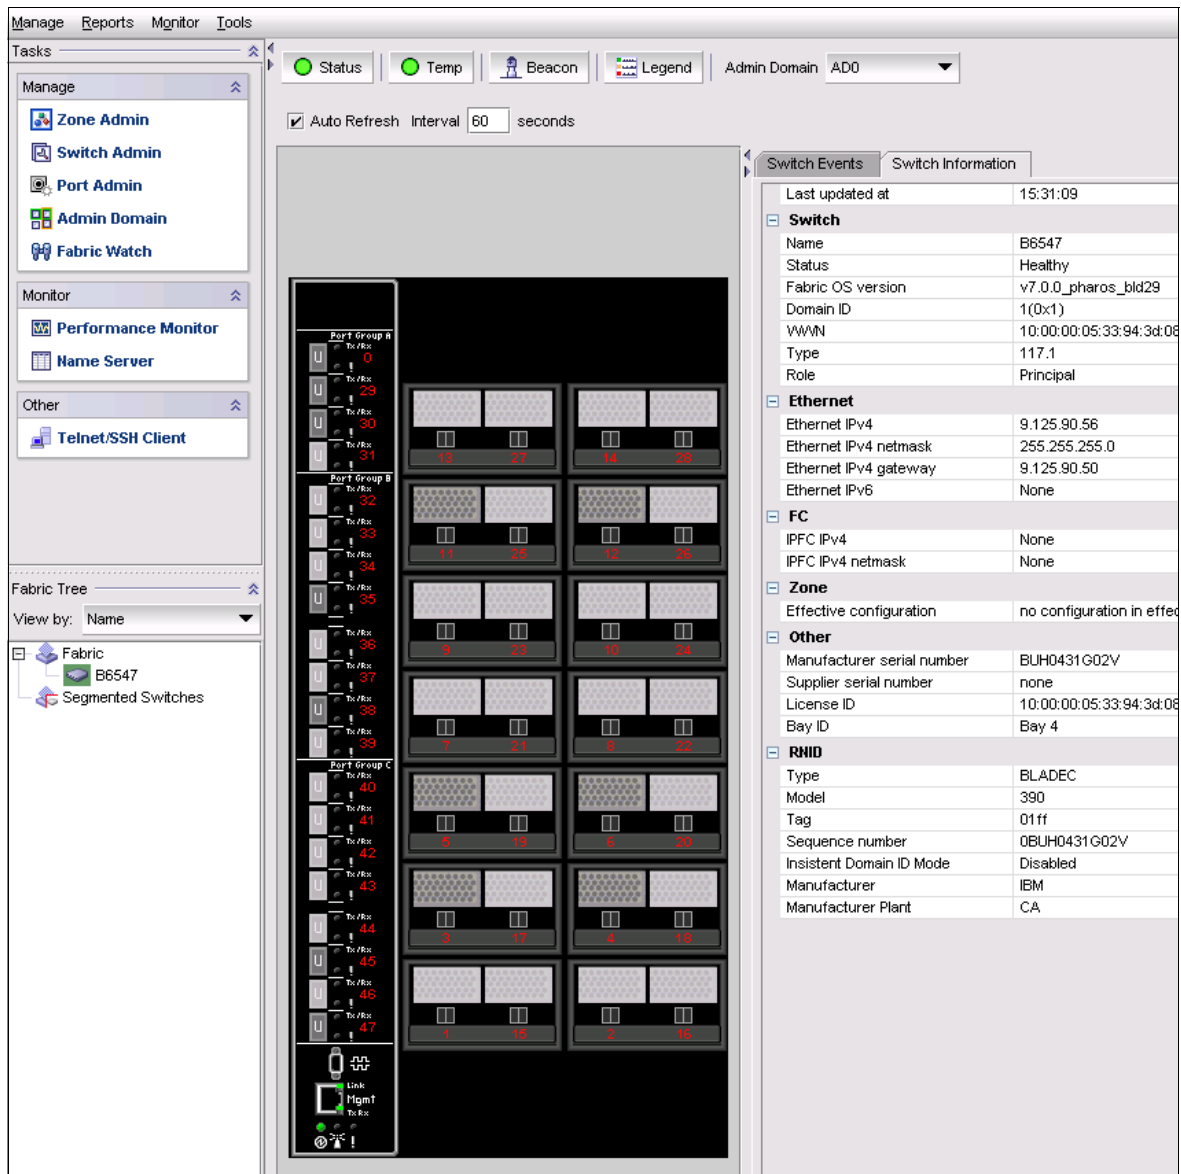


Figure 2-51 Web interface for the FC5022 16Gb SAN Scalable Switch

4. In the Task window on the left side of the interface, click **Zone Admin**. Add one new zone and place all of the ports into the zone, as shown in Figure 2-52.

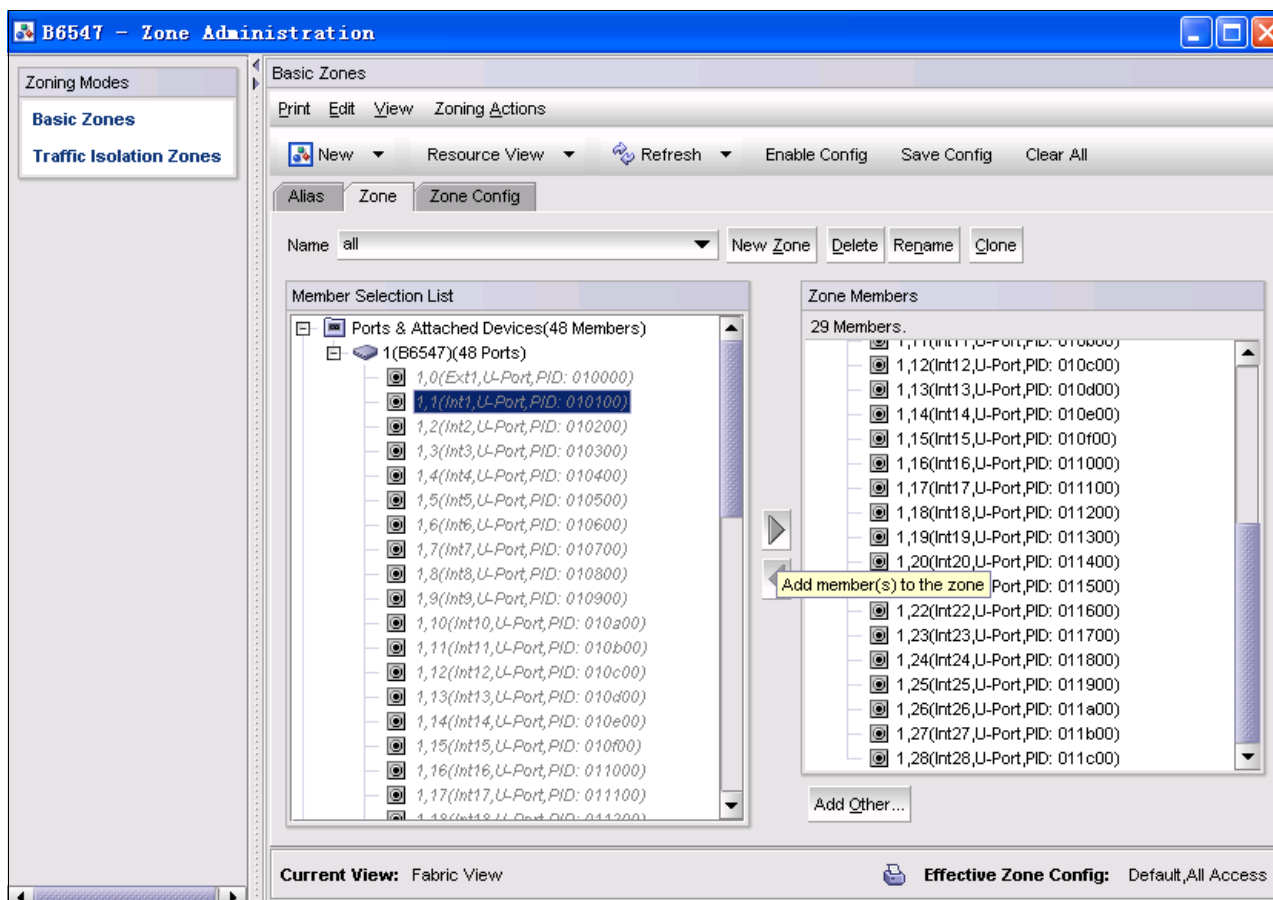


Figure 2-52 FC5022 Zone Administration

The FC/SAN switch migration can be completed, but only if the source and target are SAN switches. If the source network protocol is Fibre Channel over Ethernet (FCoE) and the target is SAN or Internet Small Computer System Interface (iSCSI), the configuration becomes more complex.

For the external storage setting, you must use the management software of the vendor to manage the third-party product. For example, to manage IBM System Storage Data Studio storage, use the IBM Data Studio Storage Manager that is available at this web page:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5077693>

There are two other complex configurations in which users must migrate network settings: moving from FCoE to SAN, or from iSCSI to SAN.

From a hardware device perspective, FCoE uses Converged Network Adapters (CNAs) and Converged switches. SAN has Fibre Channel Host Bus Adapters (FC HBAs) and SAN switches. iSCSI uses a normal NIC Adapter or iSCSI HBA and Ethernet switches. You must determine the properties of these kinds of network devices for migration.

However, FCoE and SAN share some configuration settings. These shared settings provide an opportunity for migration. For iSCSI to SAN or FCoE, there is no direct relationship between the two groups of settings, so there might be more work to complete the migration.

Another possible scenario is to migrate the configuration settings from an existing local RAID or external storage (the source) to the Flex System V7000 Storage Node (the target). Given that the V7000 Storage Node is similar to the Storwize V7000, you can access the IBM Storwize management software and duplicate the configuration settings that you defined in the source storage.

To gain access to the storage node management software, complete the following steps:

- 1. Log in to the CMM.
- 2. From the menu bar, click **Chassis Management** → **Component IP Configuration**, as shown in Figure 2-53.

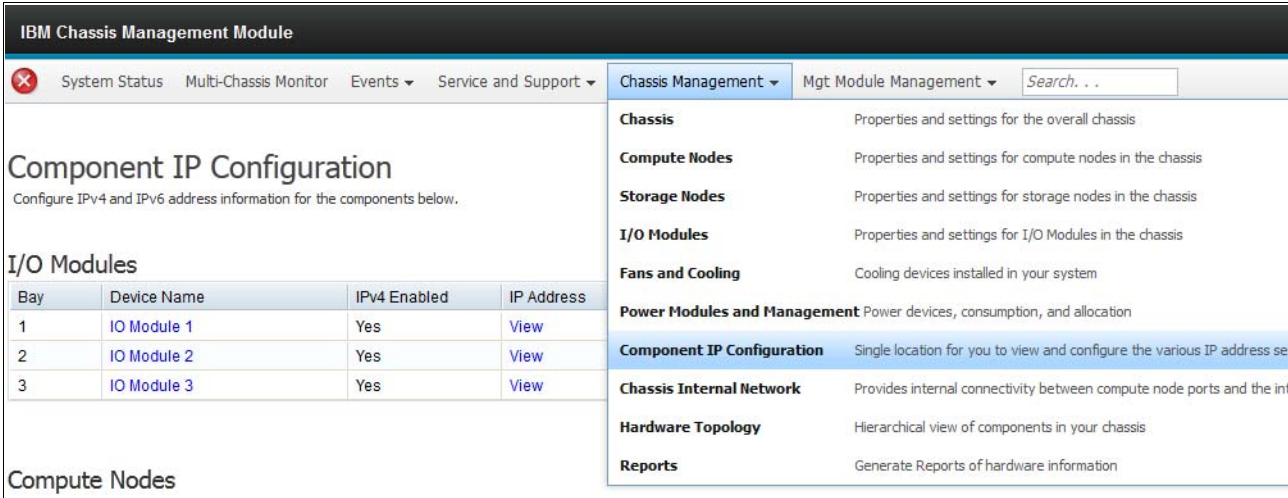


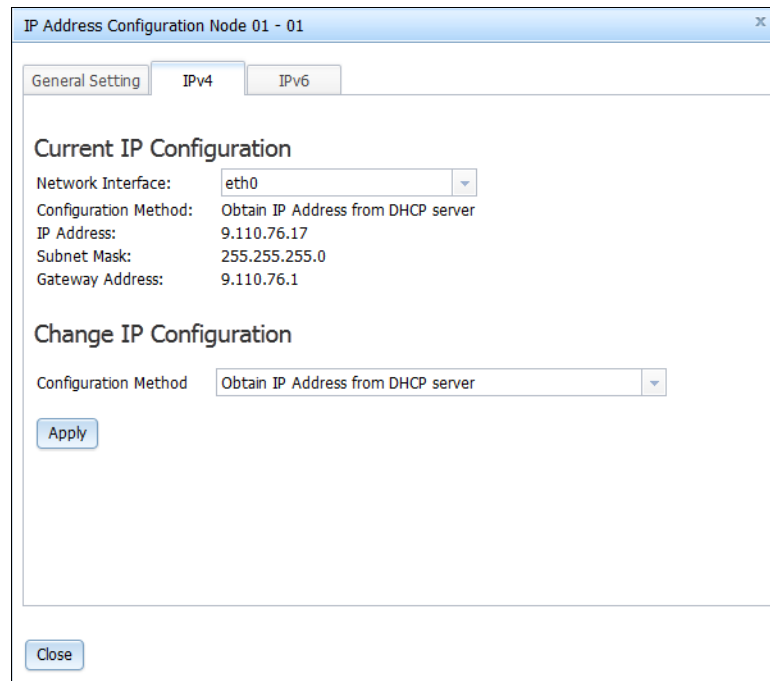
Figure 2-53 Component IP Configuration menu in the CMM

- 3. Go to the Storage Nodes section and click the storage nodes device that you want to configure, as shown in Figure 2-54.

Storage Nodes			
Bay	Device Name	IPv4 Enabled	IP Address
1-4:1	Node 01 - 01	Yes	View

Figure 2-54 Storage Node list

4. From the Tab menu, you can get or set the IP address of your storage node, as shown in Figure 2-55.



The image shows a web-based configuration window titled "IP Address Configuration Node 01 - 01". It has three tabs: "General Setting", "IPv4", and "IPv6". The "IPv4" tab is selected. The window is divided into two main sections: "Current IP Configuration" and "Change IP Configuration".

Current IP Configuration

Network Interface:	eth0
Configuration Method:	Obtain IP Address from DHCP server
IP Address:	9.110.76.17
Subnet Mask:	255.255.255.0
Gateway Address:	9.110.76.1

Change IP Configuration

Configuration Method	Obtain IP Address from DHCP server
----------------------	------------------------------------

Below the "Change IP Configuration" section is an "Apply" button. At the bottom left of the window is a "Close" button.

Figure 2-55 Storage node IP address

5. Open a new window in the web browser and enter the IP address of the storage node. You then see the welcome page of the IBM Storwize storage management software interface.

Migrating operating system images

When the hardware configurations are completed on the Flex System side, the next step is to migrate the operating system (OS) and applications from the source server to the Flex System. Before starting this migration process, it is helpful to understand the available migration options:

- ▶ **Physical-to-Physical (P2P):** Migrating the OS, applications, and data from one physical server to another physical server
- ▶ **Physical-to-Virtual (P2V):** Decoupling and migrating the OS, applications, and data from a physical server to a virtual machine that is guest-hosted on a virtualized platform. This is a typical scenario of server consolidation.
- ▶ **Virtual-to-Virtual (V2V):** Migrating an OS, application programs, and data from a virtual machine or disk partition to another virtual machine or disk partition. Typically, it is a move of virtual machine images from one physical host server to another.

Virtual-to-Physical migration: Although Virtual-to-Physical (V2P) migration is an option, it is rarely used and therefore is not covered in this document.

Review your current hardware configuration (the source and the target) to determine whether the source servers and Flex System (target servers) are connected on the same network. If the source and target servers are connected on the same network, the settings of the hardware and virtual image files can be transferred directly through the network. If the source and target servers are disconnected, the settings must be exported, transferred through removable storage, and then imported.

The migration solutions for the following scenarios are described in these sections:

- ▶ 3.1, “Source server and Flex System are disconnected” on page 56
- ▶ 3.2, “Source servers and Flex System are connected” on page 88
- ▶ 3.3, “Conclusion” on page 137

3.1 Source server and Flex System are disconnected

In this section, the following scenarios are described in which the source server and Flex System server are not connected by a local area network (LAN) or other network:

- ▶ Disconnected Physical-to-Physical
- ▶ Disconnected Physical-to-Virtual
- ▶ Disconnected Virtual-to-Virtual

3.1.1 Disconnected Physical-to-Physical

Before installing an OS to Flex System, you must check the ServerProven® list at the following website to confirm that the OS is supported:

<http://ibm.com/systems/info/x86servers/serverproven/compat/us>

In this paper, details are not provided about how to move a physical OS to the Flex System. However, the Norton Ghost backup product is one of the simplest tools to implement P2P, from which you can capture the image from the source disk and deploy the image to another disk.

For more information about Norton Ghost, see the following websites:

- ▶ <http://www.symantec.com/themes/theme.jsp?themeid=ghost>
- ▶ <http://us.norton.com/ghost>

Is it important to ensure the consistency of the drivers between the source server and the target server. The hardware for the source and target servers is different, which results in the issue of inconsistent drivers during the image migration. You must install updated drivers for the OS image of the target machine before or after transferring the drivers to the target server.

3.1.2 Disconnected Physical-to-Virtual

Before you attempt a P2V migration, the key performance of the old infrastructure must be monitored for at least one business cycle. Monitoring this cycle helps you understand more about the system usage over the entire business cycle period.

Numerous performance monitoring tools are available. For example, the Guided Consolidation module in VMware vCenter can monitor a server and report its average processing and memory quantities over a specified time. With this information, you can determine the amount of processing power and memory that you must assign to virtual machines before creating those machines.

You must also consider the type of hypervisor to deploy on the destination server. The following vendors provide hypervisors:

- ▶ VMware vSphere (ESX/ESXi)
- ▶ Microsoft Hyper-V
- ▶ KVM
- ▶ Citrix XenServer
- ▶ Qemu

These vendors have incompatible virtual disk formats such as VHD for Microsoft, VMDK for VMware, or qcow and qcow2 for QEMU, KVM, and Xen. Under these incompatibility conditions, you must use the corresponding vendor-specific tools or solutions for each vendor-specific platform. Regardless of the vendor solution you choose, the following solution implementation process is similar:

1. Capture the OS image from the source server as an OS image.
2. Copy the image to the destination hypervisor as a virtual machine.

We use Microsoft Hyper-V as the example in the following process to migrate a physical machine image to another machine as a virtual machine.

Before you begin, complete the following steps:

1. Prepare the target server that is deployed by Hyper-V and the System Center Virtual Machine Manager (SCVMM) server.

Hyper-V is included in Microsoft Windows Server 2008 R2 as one component. This hypervisor is part of the Windows Server 2008 R2 installation. The installation process can be completed from the retail Windows CD installation media or with the help of ServerGuide. You must ensure that the user enabled the Hyper-V role after installation.

An ServerGuide for your server configuration is available for download at this website:

<http://ibm.com/support/entry/portal/docdisplay?ln docid=SERV-GUIDE>

At the download website, click the corresponding server version number and browse to the guide download window. A link also is available for the ServerGuide Scripting Toolkit.

2. Identify a server to provide virtual machine management capabilities by installing and configuring System Center Virtual Machine Manager on the server.

For more information about Microsoft System Center, see this website:

<http://www.microsoft.com/en-us/server-cloud/system-center/default.aspx>

During the migration process, you must take the following measures:

- Connect the management server to the target host server in the Ethernet network.
- Ensure that the management server and target host server are in the same domain.
- Log in to the management server by using a domain user ID.

After you prepare all of the hardware and software, follow these steps to start the P2V process:

1. Open the System Center Virtual Machine Manager (SCVMM) Console on the management server. Click **Add host** to add the destination server to the VMM pool, as shown in Figure 3-1.

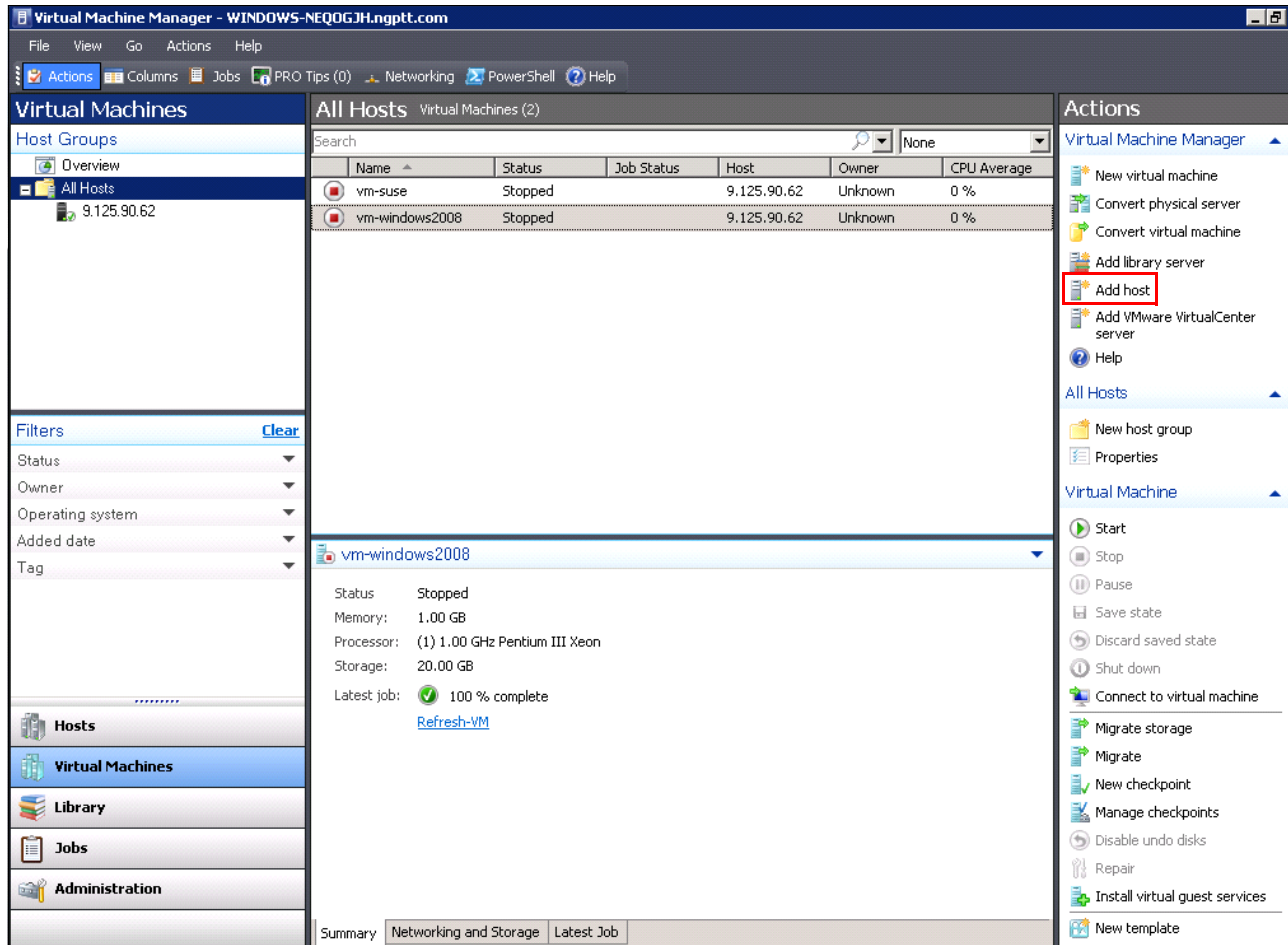


Figure 3-1 SCVMM main window

The Add Hosts window opens, as shown in Figure 3-2.

The screenshot shows the 'Add Hosts' wizard window. The title bar says 'Add Hosts'. The main heading is 'Select Host Location'. On the left is a sidebar with a tree view containing: 'Select Host Location' (selected), 'Select Host Servers', 'Configuration Settings', 'Host Properties', and 'Summary'. The main area has a blue instruction text: 'Select the host location and then enter the required credentials.' Below this are three radio button options: 'Windows Server-based host on an Active Directory domain' (selected), 'Windows Server-based host on a perimeter network', and 'VMware ESX Server host (any location)'. A horizontal line separates the options from the details section. The details section has a server icon and the text 'Windows Server-based host on an Active Directory domain'. Below this is the instruction 'Enter the credentials for connecting to the host.' followed by three text boxes: 'User name:' with 'Administrator', 'Password:' (empty), and 'Domain:' with 'NGPTT'. Below the text boxes is a checked checkbox labeled 'Host is in a trusted domain' with a note: 'Clear this option if the host does not have a two-way trust relationship with the domain of the VMM server.' At the bottom right are 'Next' and 'Cancel' buttons.

Figure 3-2 Add Hosts wizard: Select Host Location

2. Choose **Windows Server-based host on an Active Directory domain** if the management server and target host are in one domain. Enter the domain information and click **Next**.
3. Select the destination machine as shown in Figure 3-3 and click **Next**.

No hosts listed: If no domain is available, no selected hosts are listed here. The user must install the VMM Agent on the destination server. Also, the user must enter the IP and credential information of the destination server into the Host details fields in Figure 3-3. Entering this information ensures that the management server can access the destination server.

The screenshot shows the 'Add Hosts' wizard window with the title bar 'Add Hosts'. The main heading is 'Select Host Servers'. On the left is a navigation pane with the following items: 'Select Host Location', 'Select Host Servers' (which is highlighted), 'Configuration Settings', 'Host Properties', and 'Summary'. The main area is divided into two sections. The top section, 'Host details', contains four input fields: 'Computer name or IP address:', 'Encryption key:', 'Confirm encryption key:', and 'Security file path:'. The 'Security file path:' field has a 'Browse...' button next to it. Below these fields is an 'Add' button. The bottom section, 'Selected hosts:', contains a list box with one entry: '9.125.90.107'. Below the list box is a 'Remove' button. At the bottom of the window are three buttons: 'Previous', 'Next' (which is highlighted), and 'Cancel'. A link 'How to add a host on a perimeter network' is located below the 'Selected hosts' list box.

Figure 3-3 Add Hosts wizard: Select Host Servers

4. Select the host group that contains the host in the Configuration Settings window (shown in Figure 3-4).

If any of the selected hosts are managed by another Virtual Machine Management server, click **Reassociate host with this Virtual Machine Manager server** to change the associations so that you do not have to manage the association manually.

Click **Next**.

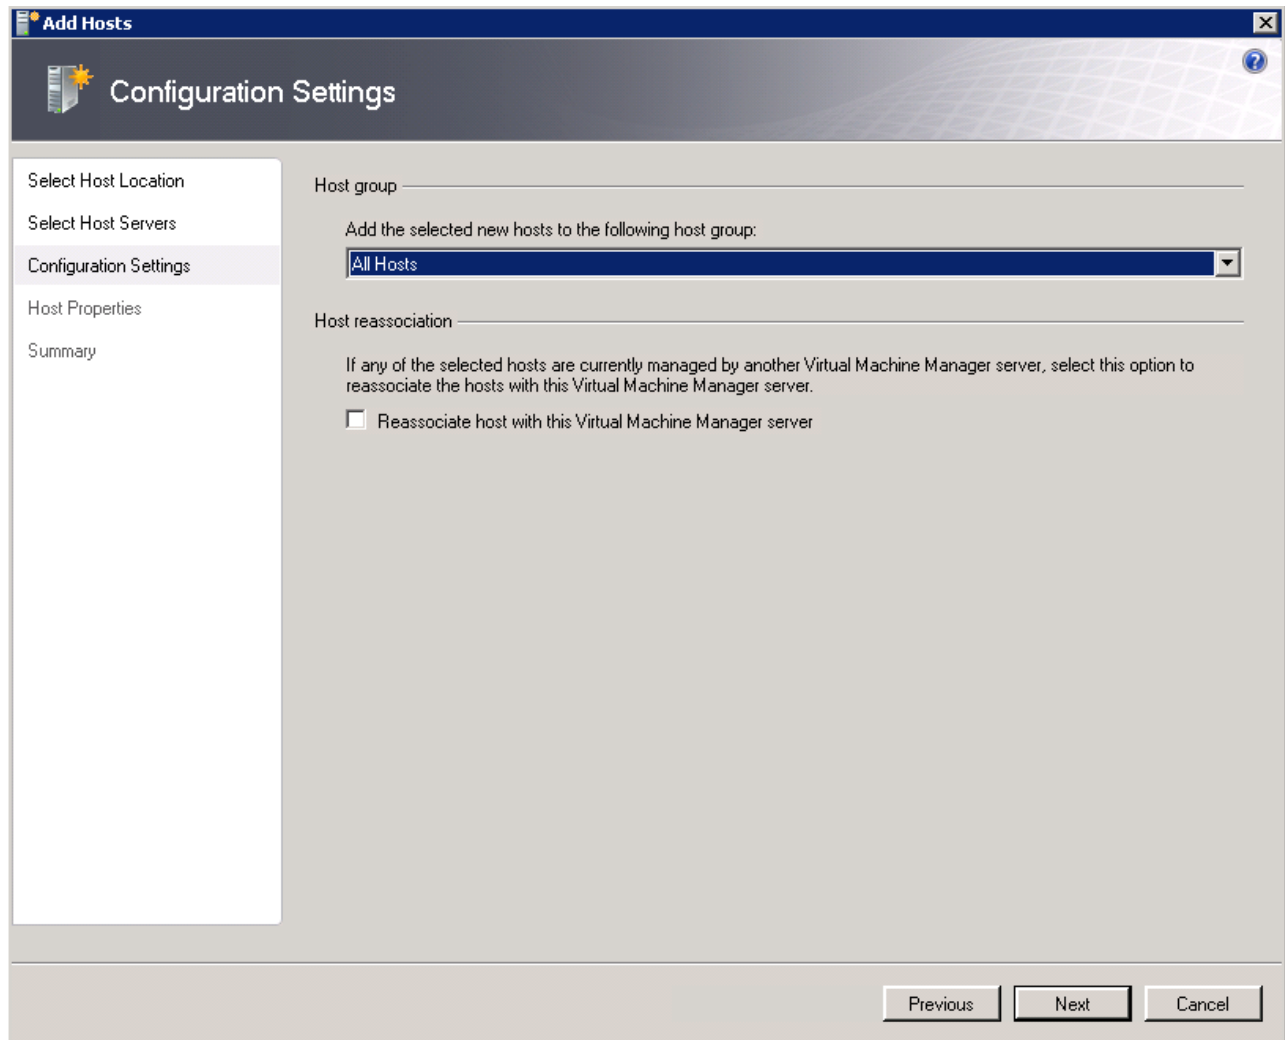


Figure 3-4 Add Hosts wizard: Configuration Settings

5. Enter the default path of the virtual machine in the Host Properties window (shown in Figure 3-5). This path can be left blank.
Click **Next**.

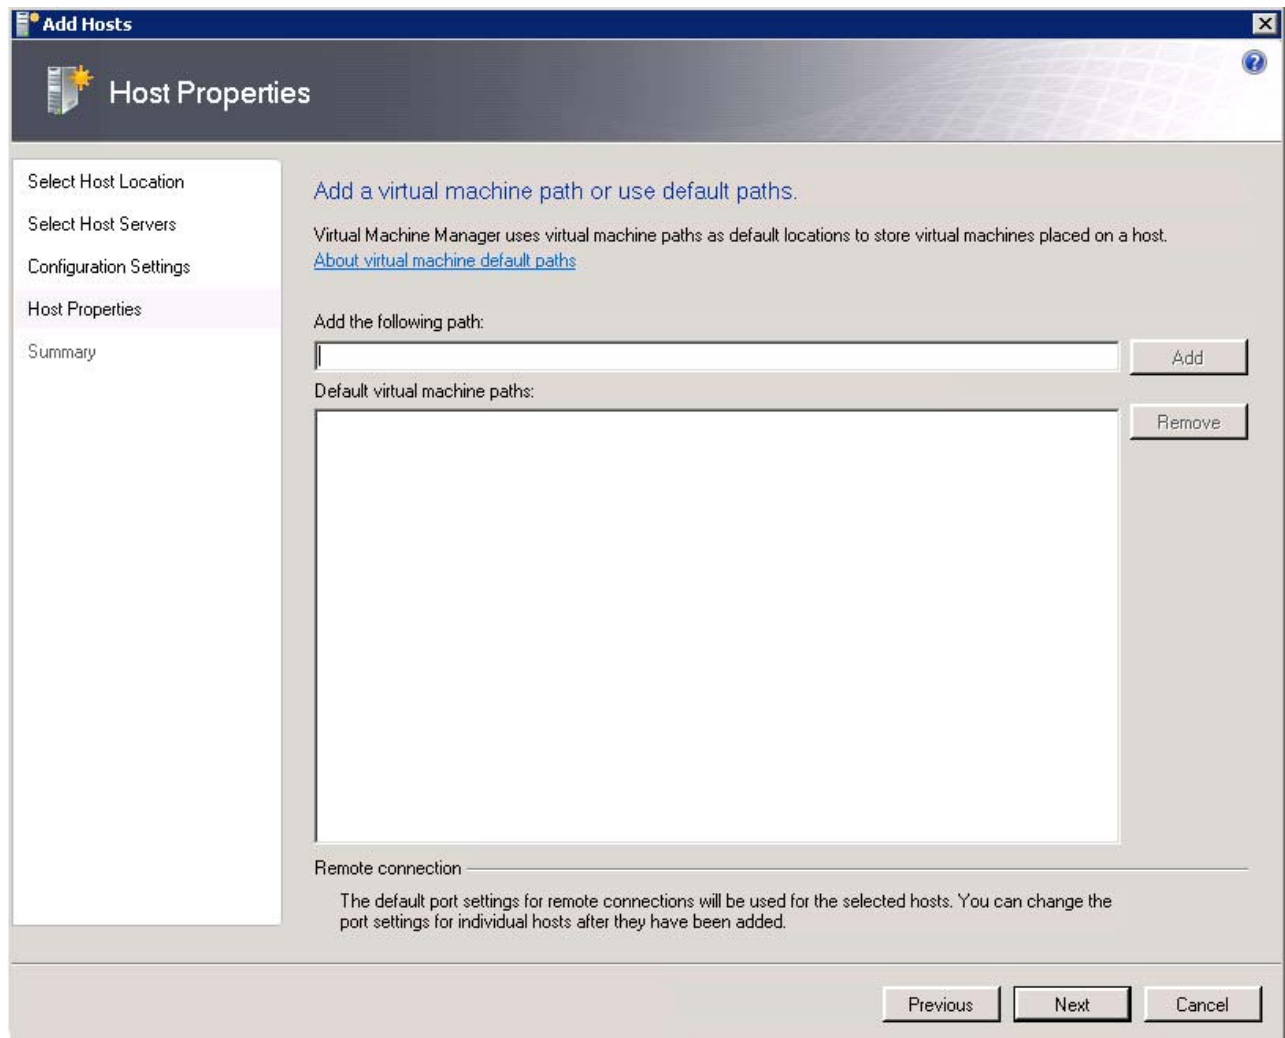


Figure 3-5 Add Hosts wizard: Host Properties

6. Review the summary and click **Add Hosts**, as shown in Figure 3-6.

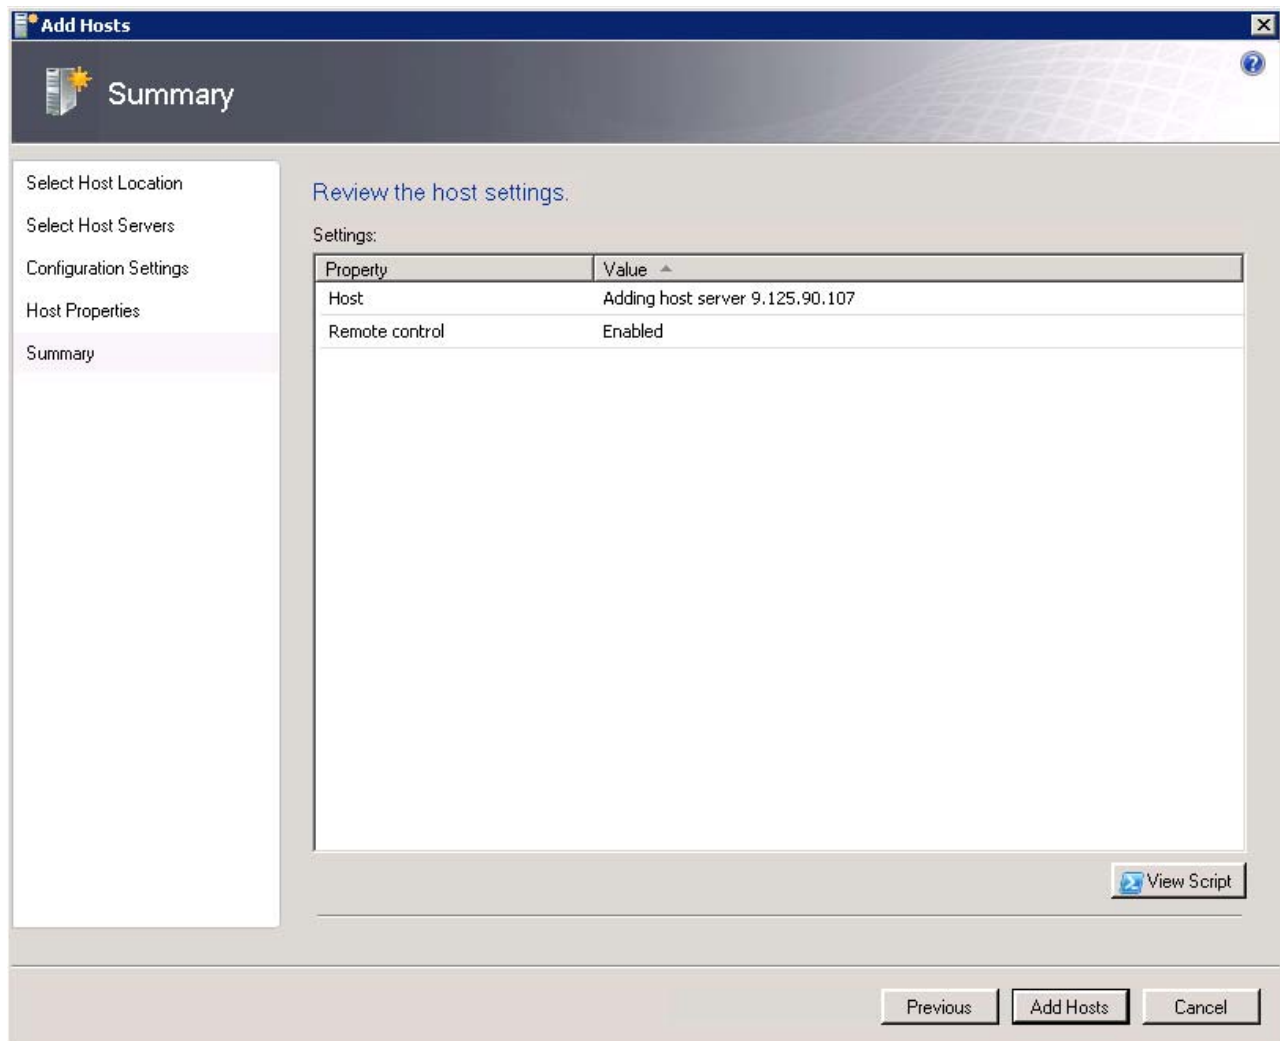


Figure 3-6 Add Hosts wizard: Summary

7. After a short period, the destination server is added into the list of managed hosts, as shown in Figure 3-7.

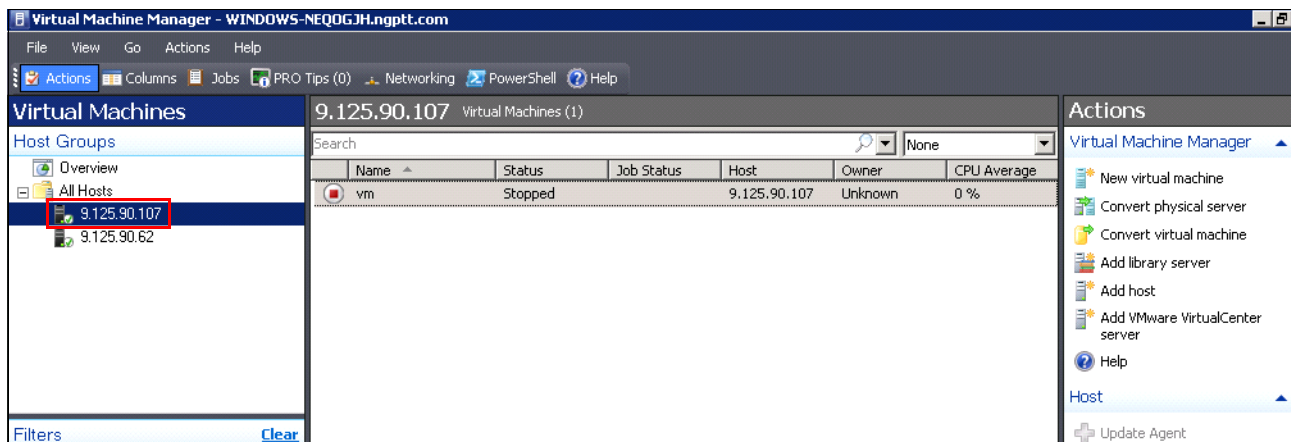


Figure 3-7 SCVMM main window after the host is added

The destination server preparation is now complete. The next step is to capture the OS image on the source server. Follow these steps to capture the OS image:

1. Download the Disk2vhd tool to capture the OS from the source server.

The Disk2vhd tool creates Virtual Hard Disk (VHD) versions of physical disks for use in Microsoft Virtual PC or Microsoft Hyper-V virtual machines. (VHD is the Microsoft Virtual Machine disk format.) For more information about the Disk2vhd tool, see this website:

<http://technet.microsoft.com/en-us/sysinternals/ee656415>

2. Start the Disk2vhd tool on the source physical machine. Select the volumes (C:\ in Figure 3-8) that you want to migrate. Document the target path to which you want to migrate the image file. Click **Create**.




Figure 3-8 Disk2vhd: Select a volume

3. The chosen volumes are converted to the virtual machine disk file, as shown in Figure 3-9.




Figure 3-9 Disk2vhd: Creating the VHD file

- After the conversion is complete, navigate to the location where the .vhd file was saved (see Figure 3-10). Copy the VHD file to a location that the target server can access.

Name	Date modified	Type	Size
disk2vhd.exe	10/11/2010 9:56 PM	Application	1,725 KB
disk2vhd-tmp.exe	1/20/2012 10:27 AM	Application	373 KB
WINDOWS-B2S1EMX.VHD	1/20/2012 10:33 AM	VHD File	14,739,097...

Figure 3-10 Location of the VHD file

- Share to the network the folder where the VHD file is located to ensure that the management server can access the VHD file.
- Return to the management server and in the SCVMM Console, click the **Library** tab (at the bottom of the left side of the window). Add the folder path to the Library tab by right-clicking the server name and then clicking **Select Path**.
- The VHD file is listed in the Library tab, as shown in Figure 3-11.

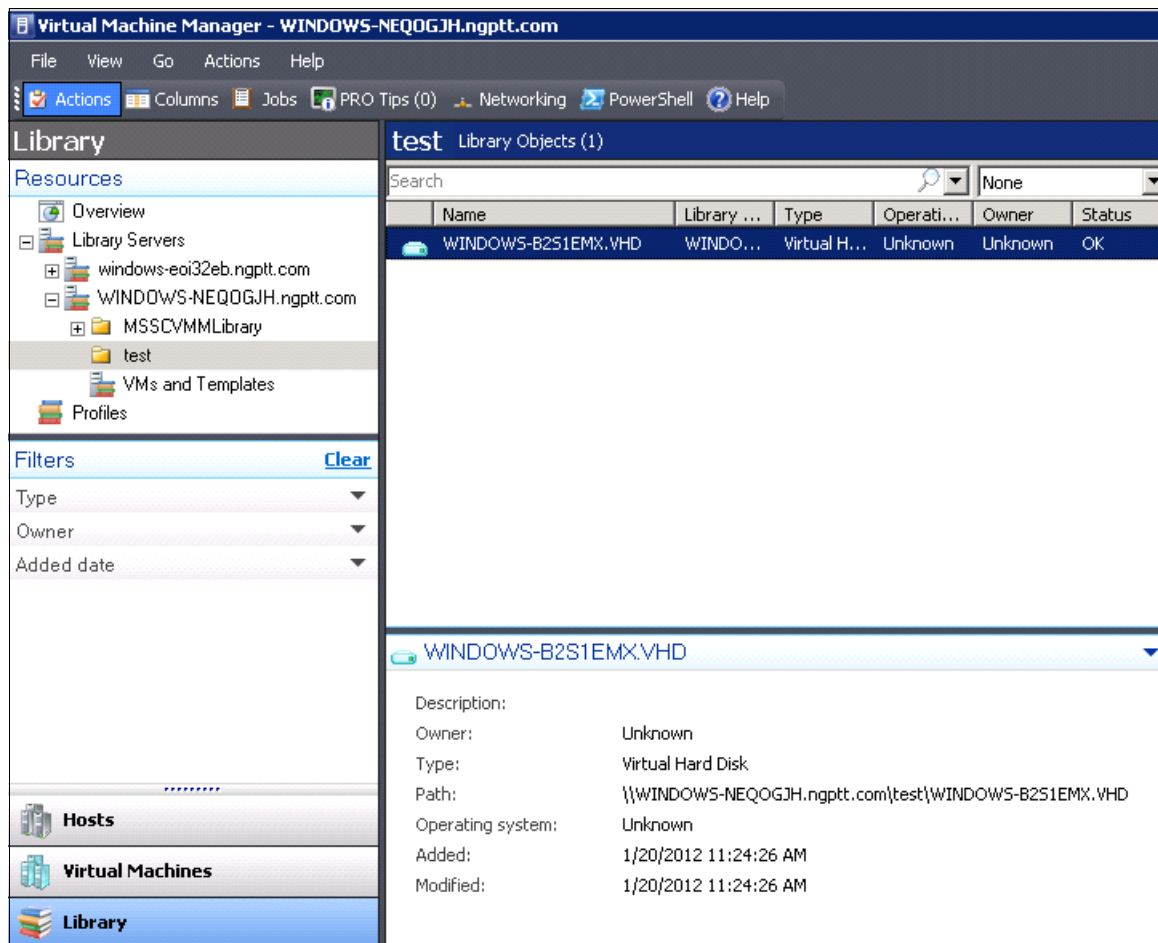


Figure 3-11 VHD added to the Library tab

- Choose the **Virtual Machines** tab in the left navigation window. Click **New virtual machine** in the Actions window on the right, as shown in Figure 3-12.

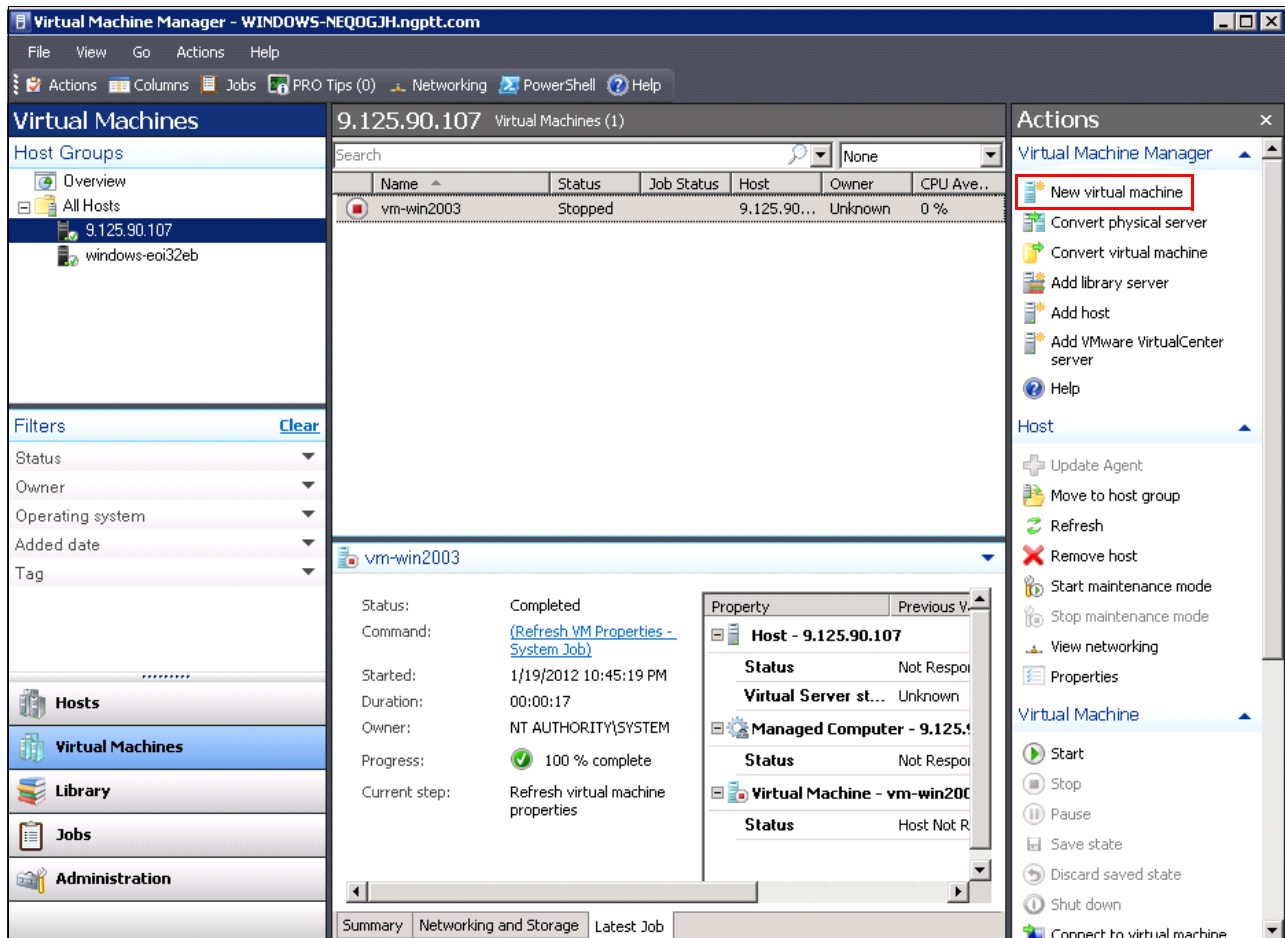


Figure 3-12 Creating a new virtual machine

9. The New Virtual Machine wizard starts. Choose the OS image file that you captured from the source server, as shown in Figure 3-13.

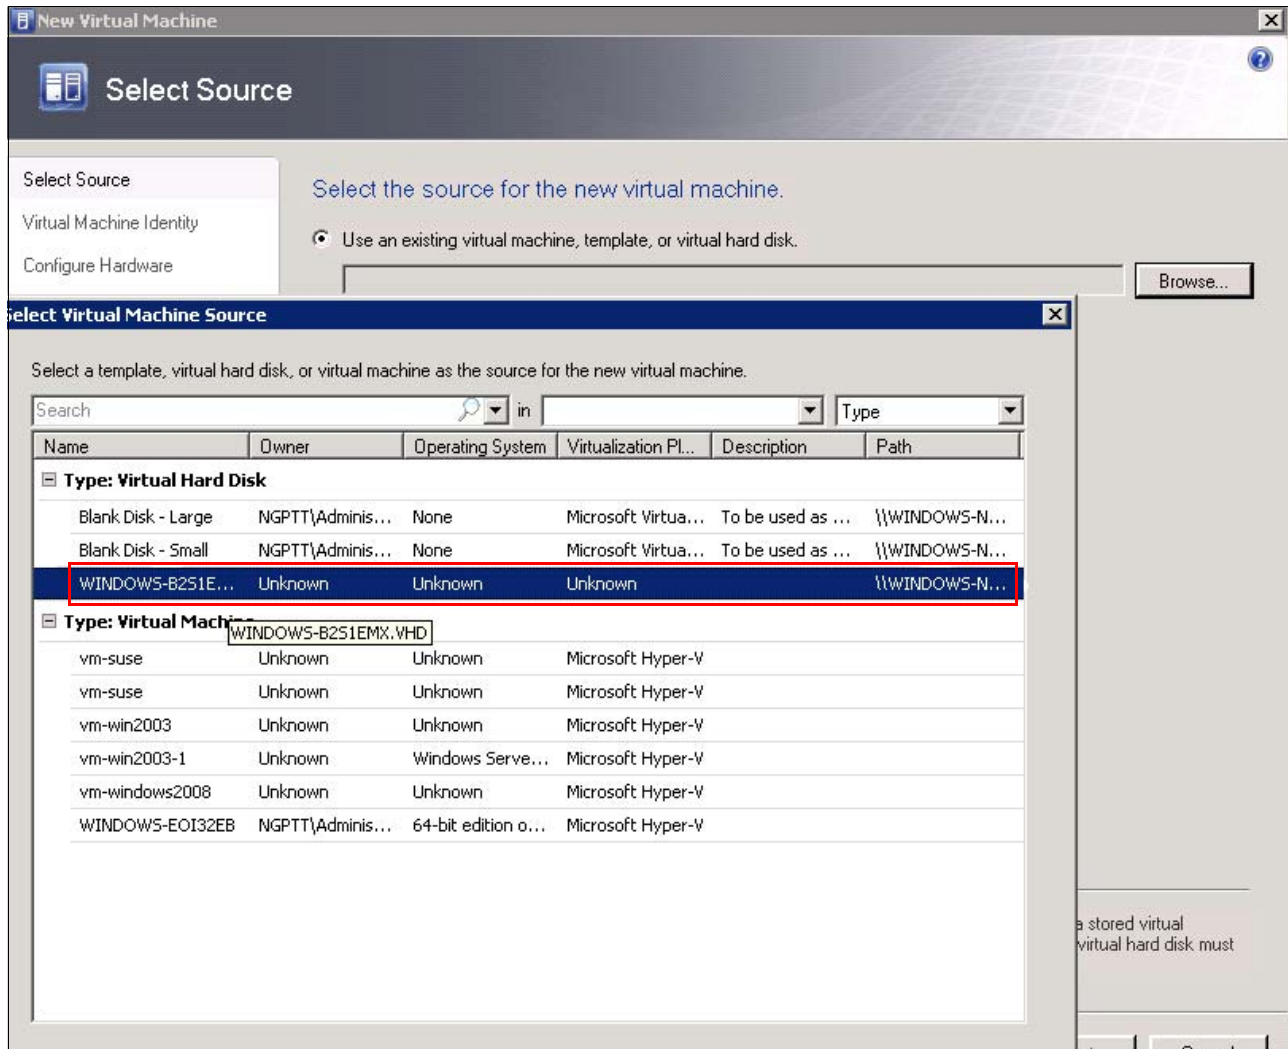


Figure 3-13 New Virtual Machine wizard: Select source

10. In the Virtual Machine Identity window, enter the virtual machine identity information, as shown in Figure 3-14. Click **Next**.

The screenshot shows the 'New Virtual Machine' wizard with the 'Virtual Machine Identity' step selected in the left-hand navigation pane. The main area contains the following fields and controls:

- Virtual machine name:** A text box containing 'vm-import-vhd'.
- Owner:** A text box containing 'NGPTT\Administrator' and a 'Browse...' button.
- Format:** A label indicating 'domain\username'.
- Description:** A large, empty text area.
- Information:** A blue information icon followed by the text: 'The virtual machine name identifies the virtual machine to VMM. The name does not have to match the computer name of the virtual machine. However, using the same name ensures consistent displays in System Center Operations Manager.'
- Navigation:** 'Previous', 'Next', and 'Cancel' buttons at the bottom right.

Figure 3-14 New Virtual Machine wizard: Virtual Machine Identity

11. In the Configure Hardware window, configure the virtual hardware information about the destination server according to the source physical machine, as shown in Figure 3-15.
- You can configure the settings of the hardware based on the performance monitoring analysis of the source server to provide the best performance for the new virtual machine. Click **Next**.

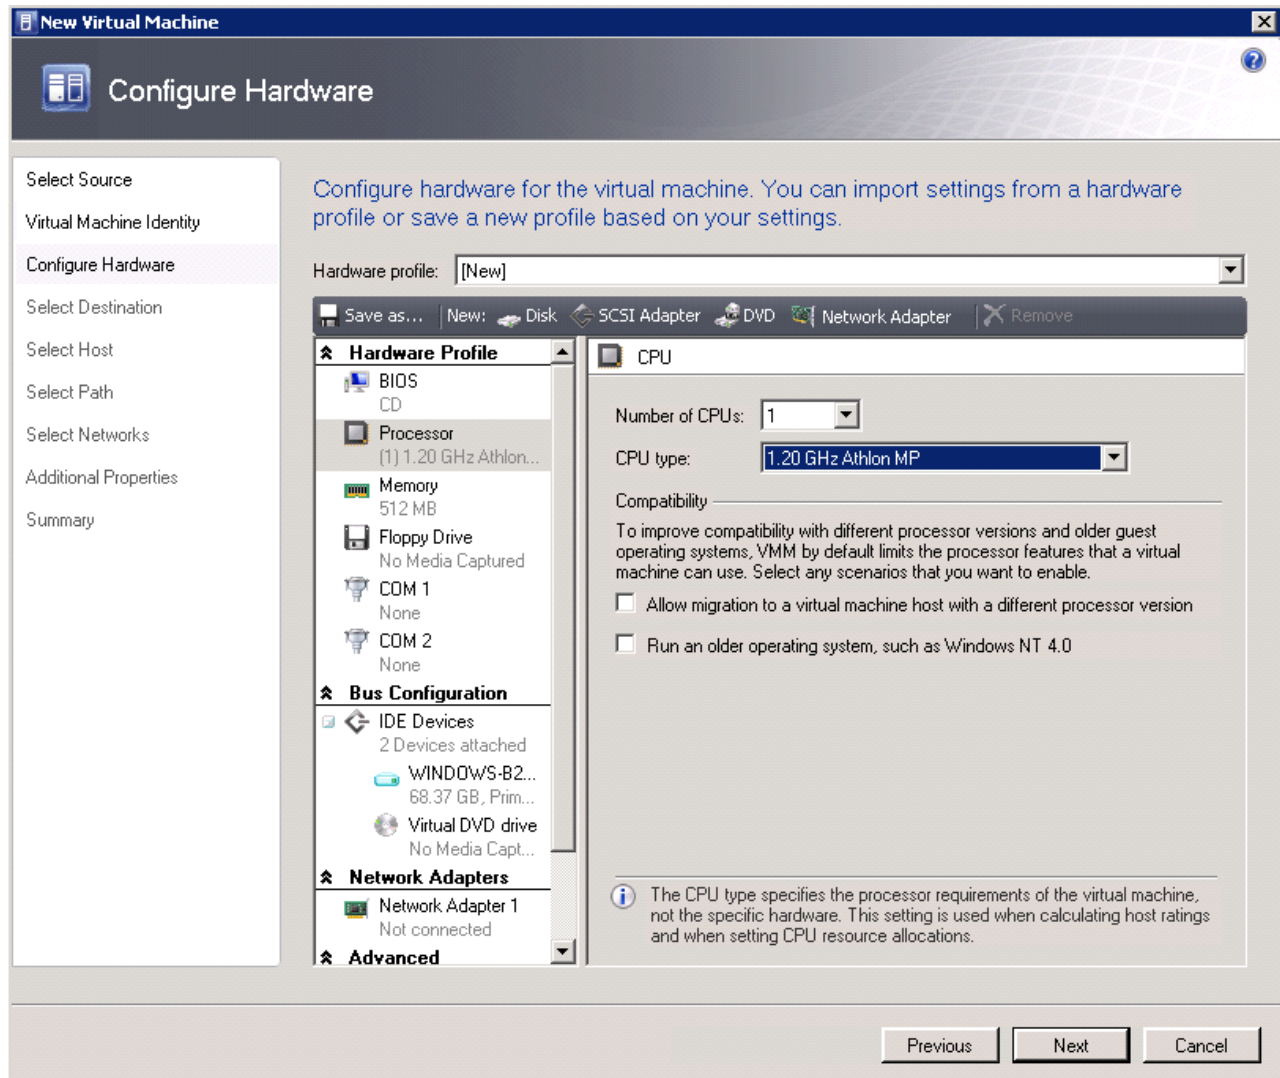


Figure 3-15 New Virtual Machine wizard: Configure Hardware

12. In the Select Destination window, select **Place the virtual machine on a host**, as shown in Figure 3-16. Click **Next**.

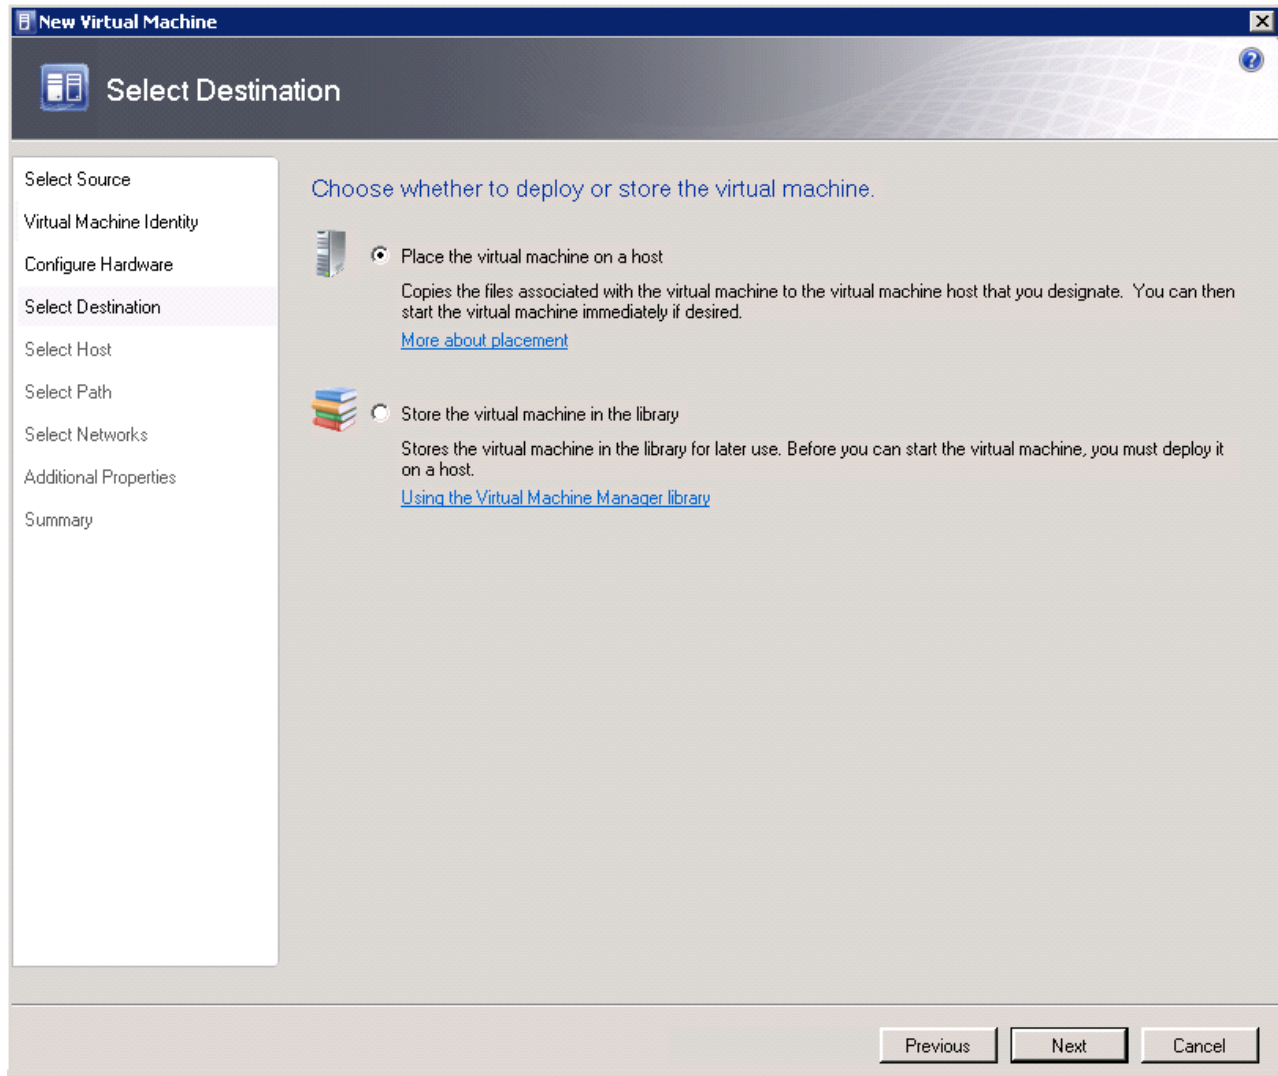


Figure 3-16 New Virtual Machine wizard: Select Destination

13. In the Select Host window, select the destination server that functions as the host, as shown in Figure 3-17. Click **Next**.

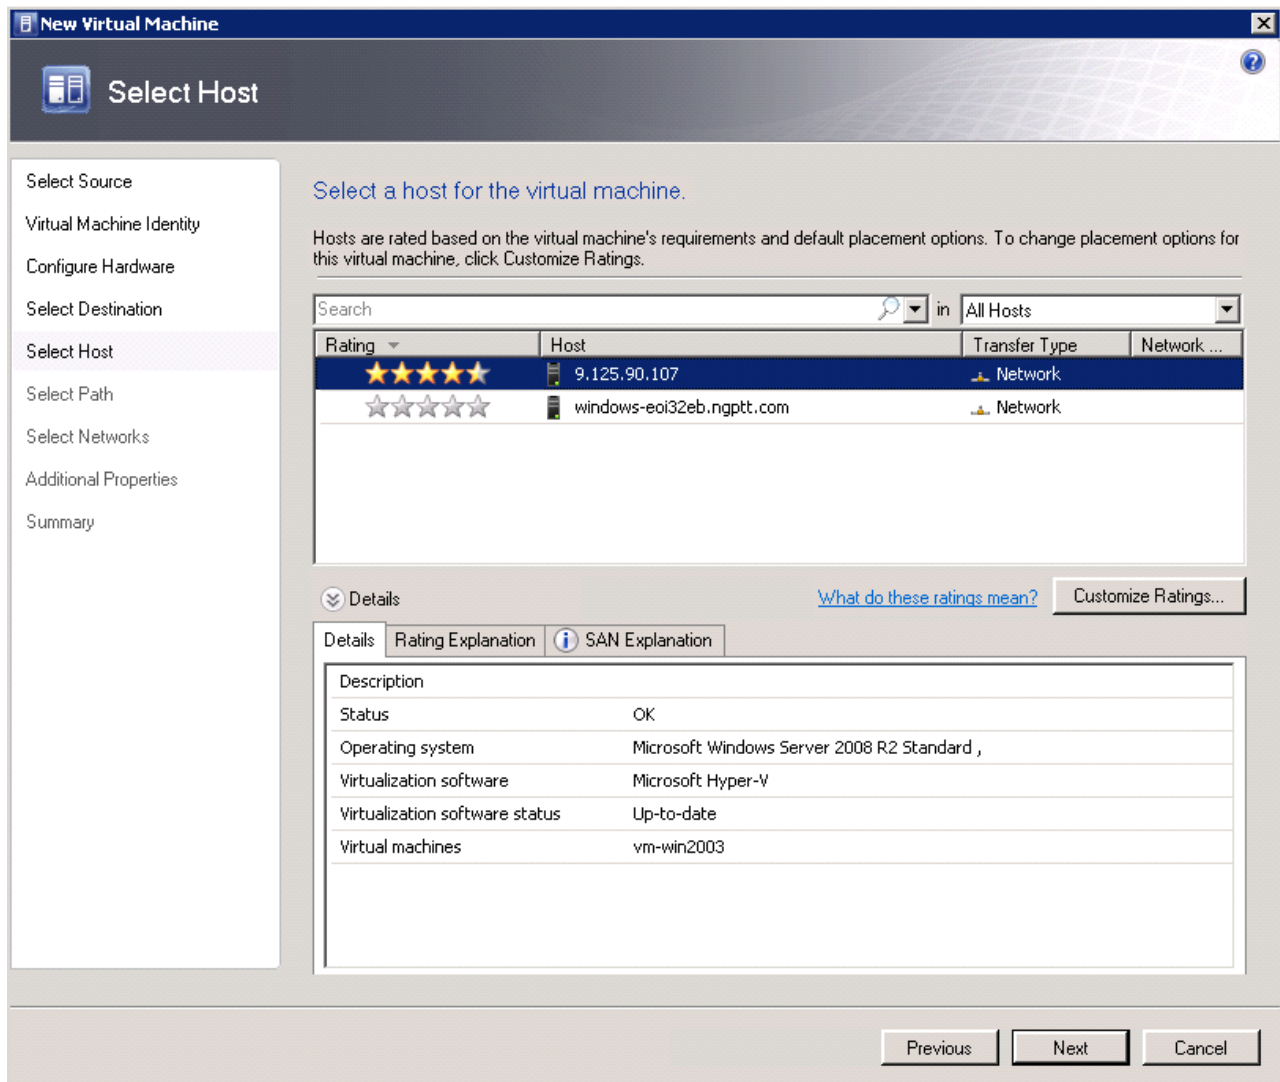


Figure 3-17 New Virtual Machine wizard: Select Host

14. In the Select Path window, enter the storage location for the virtual machine files, as shown in Figure 3-18. Click **Next**.




Figure 3-18 New Virtual Machine wizard: Select Path

15. In the Select Networks window, configure the network for the virtual machine as shown in Figure 3-19. Click **Next**.

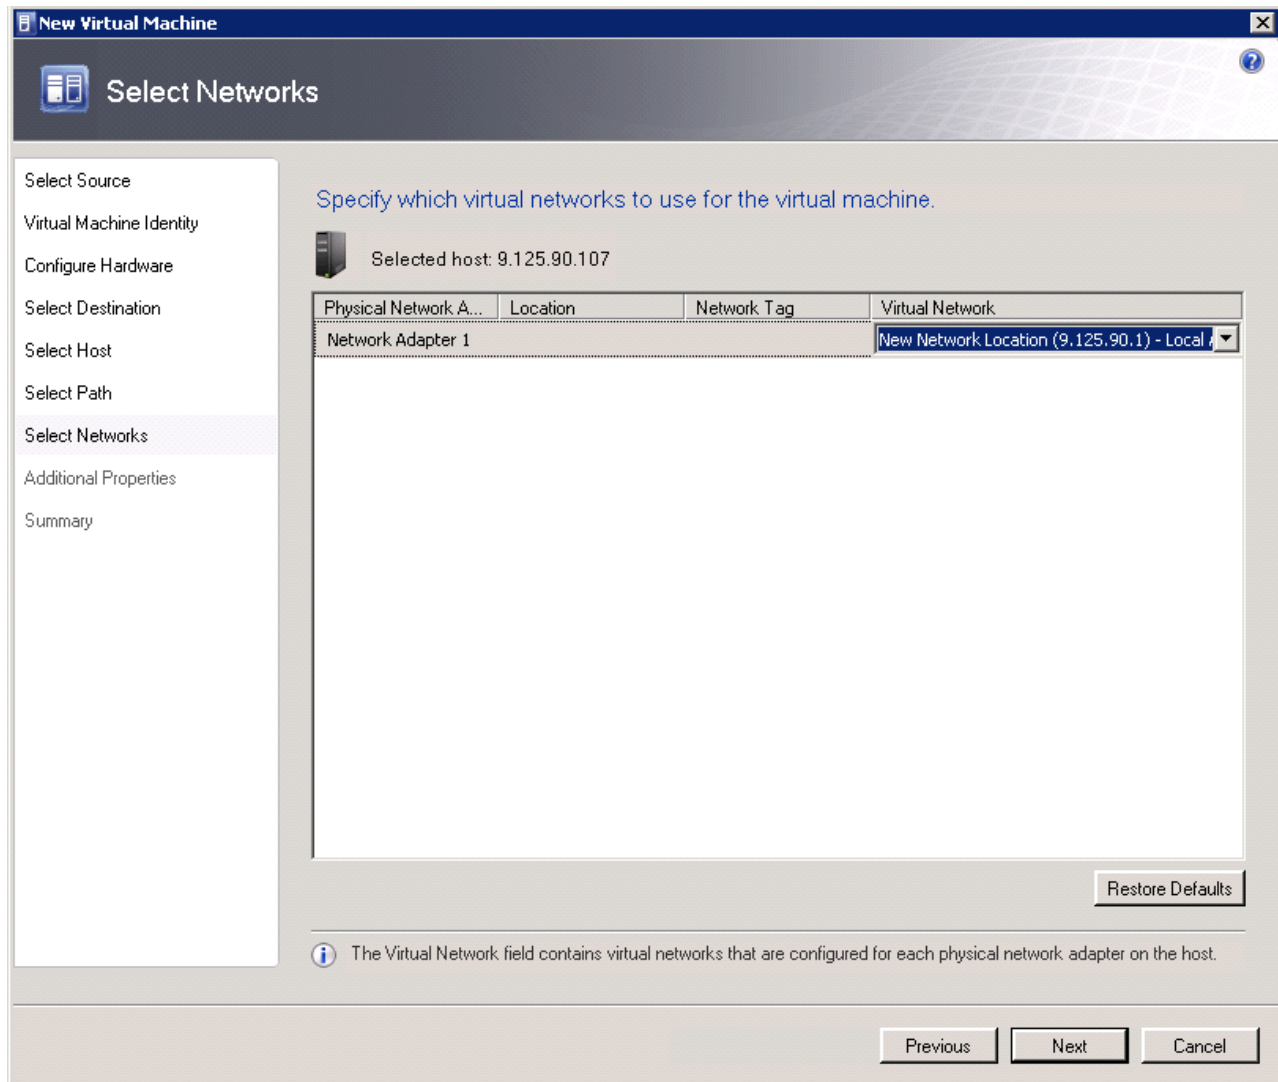


Figure 3-19 New Virtual Machine wizard: Select Networks

16. In the Additional Properties window, confirm any additional properties, as shown in Figure 3-20. Click **Next**.

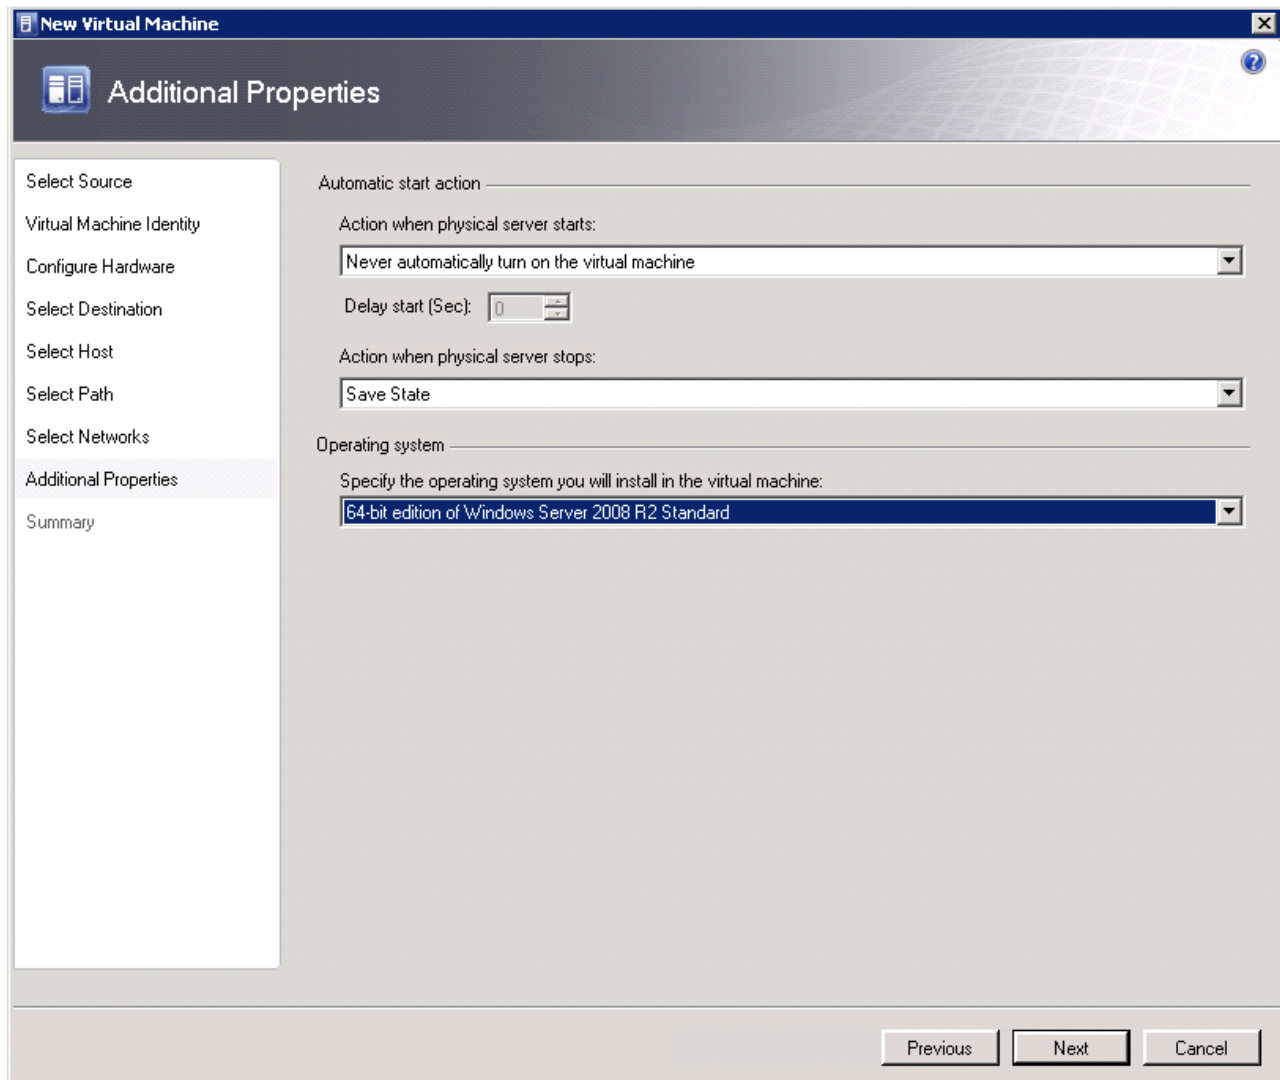


Figure 3-20 New Virtual Machine wizard: Additional Properties

17. In the Summary window, review the summary information, as shown in Figure 3-21. Click **Create**.

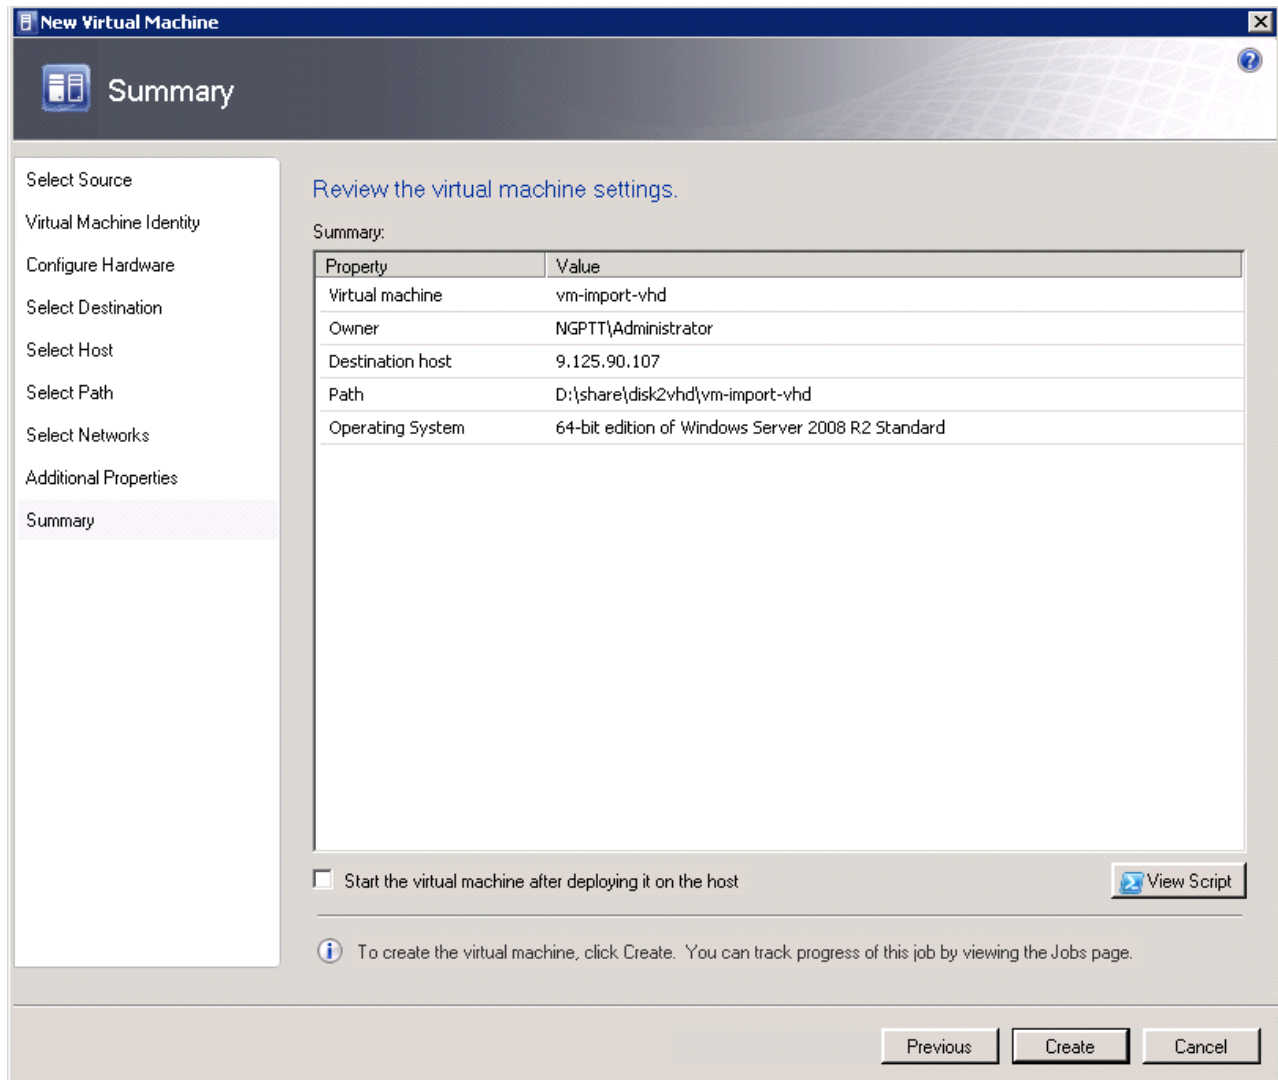


Figure 3-21 New Virtual Machine wizard: Summary

After a short period, a virtual machine is created on the target host. The virtual machine has the same configuration and disk layout as the physical source machine.

3.1.3 Disconnected Virtual-to-Virtual

Any V2V migration must start with an evaluation of the computing resource requirements of each original virtual machine (VM). This evaluation ensures that those resources are available on the destination server. This type of resource contains the processor capability, and the memory and storage capacity. If the required resources are not available, the VM must not be deployed on the destination server unless other workloads can be redistributed to free the necessary resources.

Furthermore, the hypervisor must be installed on the destination server. There are several options for hypervisors that are offered by different vendors:

- ▶ VMware vSphere (formerly ESXi)
- ▶ Microsoft Hyper-V
- ▶ KVM
- ▶ Xen
- ▶ Qemu

You must use the same hypervisor for the source and destination servers to reduce the risk of complications. By using the same hypervisor, you also save costs (particularly VM management costs) when making the transition from one hypervisor to another. This extra cost can be incurred because of the effort to convert VM images between incompatible virtual disk formats of different vendors. Those vendors that support the Open Virtual Machine Format do so only for virtual appliances. Hypervisors handle proprietary formats, such as VHD for Microsoft, VMDK for VMware, or qcow and qcow2 for QEMU, KVM, and Xen.

The following example process uses VMware vSphere to migrate a VM image to another machine:

1. Install VMware vSphere/ESX/ESXi on the destination server. For ESXi or vSphere, use the Lenovo customized image that contains the necessary drivers. To download an image and order a supported USB memory key, see this website:

<http://ibm.com/systems/x/os/vmware/esxi/index.html>

The no charge ServerGuide Scripting Toolkit can be a valuable tool to help you prepare for the installation of VMware. Alternatively, a USB Memory Key with embedded VMware can be purchased and requires no installation.

For more information and to download the ServerGuide Scripting Toolkit, see this website:

<http://ibm.com/support/entry/portal/docdisplay?lnocid=SERV-TOOLKIT>

2. Install vCenter Server (as shown in Figure 3-22 on page 77), which centralizes virtual machine management (although in this example scenario, it is not required). Additionally, you can install the vSphere client.

For more information about and to download the VMware vCenter and the VMware vSphere Client, see this website:

<http://www.vmware.com/products>



Figure 3-22 VMware vCenter Installer

3. After the vSphere client software is installed, start the vSphere client. Enter the IP address and credentials of vCenter for the source machine in the login window, as shown in Figure 3-23. Click **Login**.




Figure 3-23 VMware vSphere login window

The main vCenter console opens, as shown in Figure 3-24.

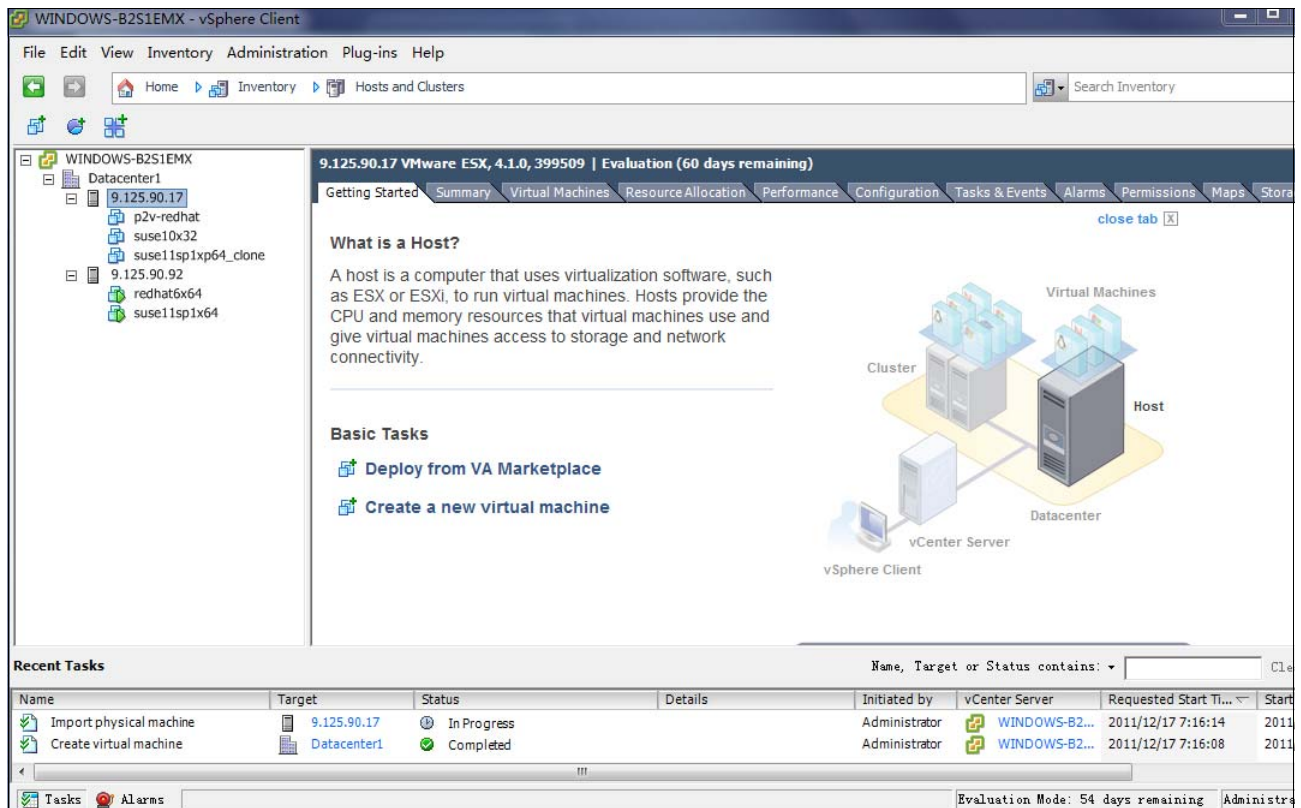


Figure 3-24 vCenter console

4. If you are using vCenter for the first time, right-click **Hosts & Clusters** and click **Add Datacenter**. Name the data center (in this example, datacenter 1).
5. Right-click **datacenter 1** and click **Add Host**. The Add Host Wizard opens.

6. As shown in Figure 3-25, in Connection Settings, enter the fully qualified Host name of the VMware source Server and the Username and Password for the host. Click **Next**.

The screenshot shows the 'Add Host Wizard' window with the title bar 'Add Host Wizard'. The main content area is titled 'Specify Connection Settings' with the instruction 'Type in the information used to connect to this host.' On the left, a sidebar lists the wizard steps: 'Connection Settings' (selected), 'Host Summary', 'Virtual Machine Location', and 'Ready to Complete'. The main area is divided into two sections: 'Connection' and 'Authorization'. The 'Connection' section has a label 'Enter the name or IP address of the host to add to vCenter.' and a text input field labeled 'Host:'. The 'Authorization' section has a label 'Enter the administrative account information for the host. vSphere Client will use this information to connect to the host and establish a permanent account for its operations.' and two text input fields labeled 'Username:' and 'Password:'. At the bottom, there are four buttons: 'Help', '< Back', 'Next >', and 'Cancel'.

Figure 3-25 Add Host Wizard: Connection Settings

7. vCenter discovers VMs that are running on the host and displays the details of each host server, as shown in Figure 3-26. Click **Next** to enter the license key. Click **Next** again.

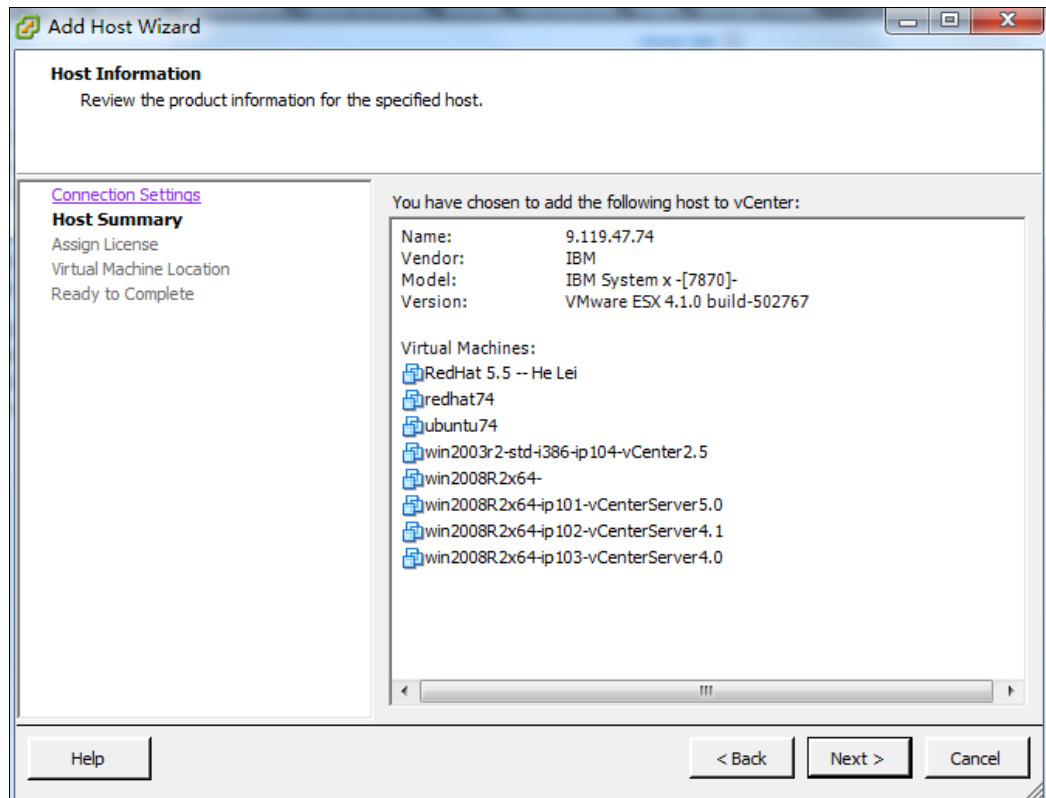


Figure 3-26 Add Host Wizard: Host summary

8. Select the data center to which to add the host, as shown in Figure 3-27. Click **Next** to review the summary. Then, click **Finish**.

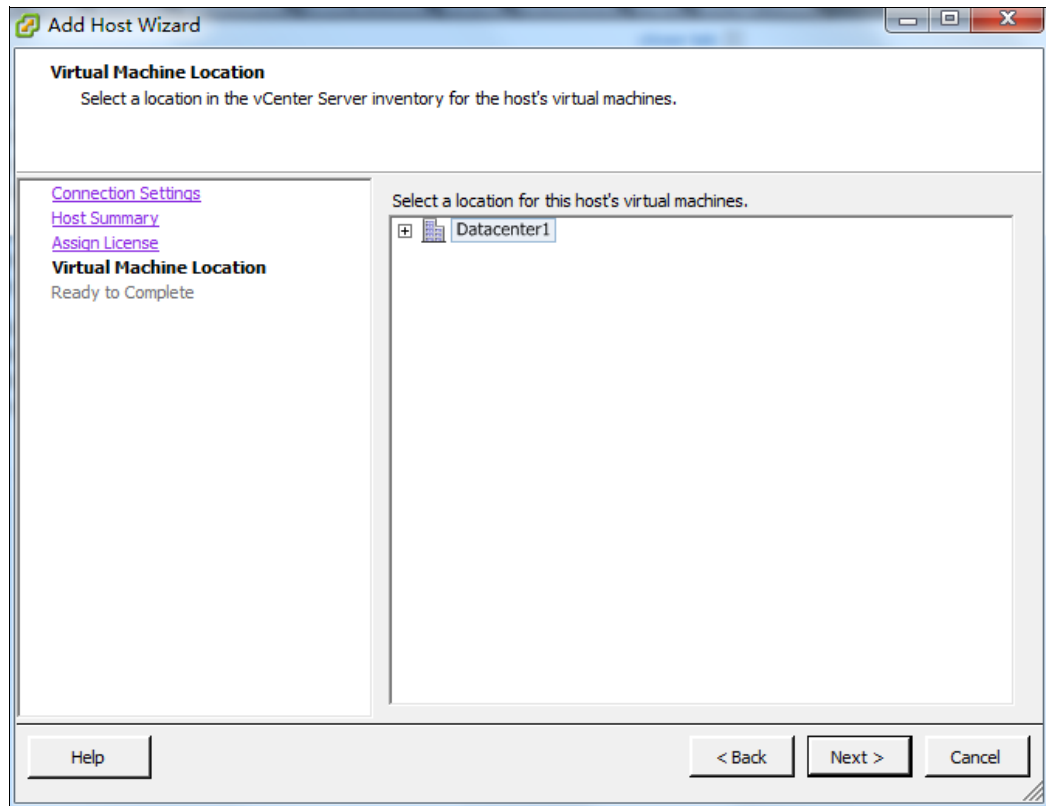


Figure 3-27 Add Host Wizard: Virtual machine location

9. The server and VMs are shown on the left side of the window. Click **Next** to complete the Add Host wizard.

10. In the vCenter Console, select the virtual machine for which you want to capture the image and click **File** → **Export** → **Export OVF Template**, as shown in Figure 3-28.

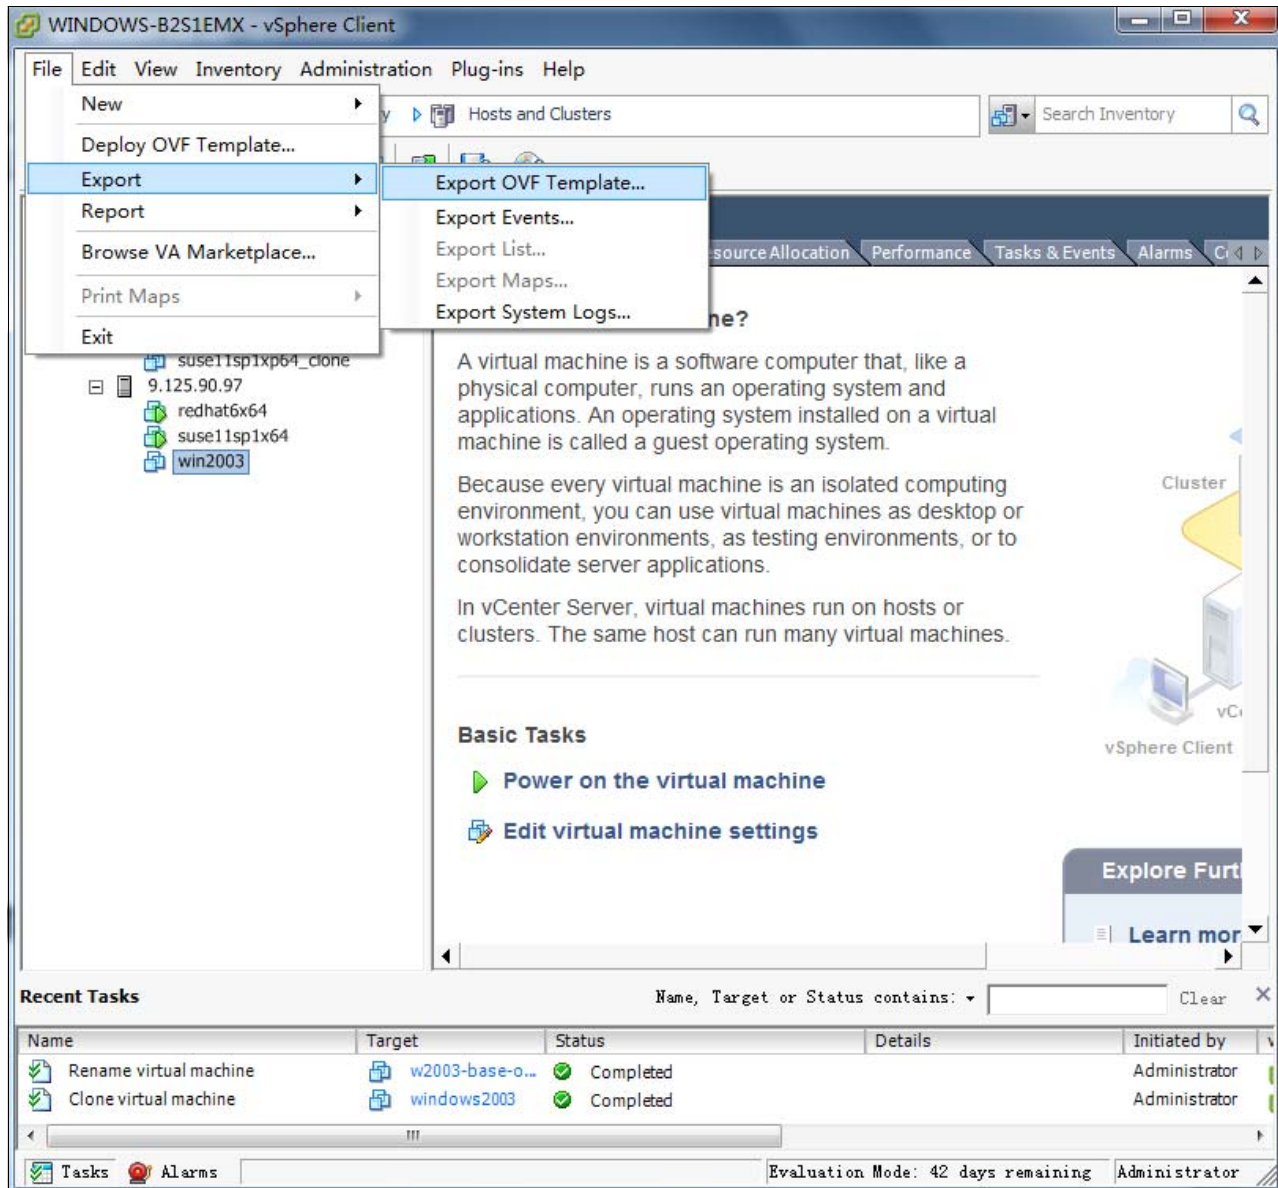


Figure 3-28 Export OVF Template

11. The Export OVF Template dialog opens. Enter the Name, Directory, and Format for the OVF template package and click **OK**, as shown in Figure 3-29.

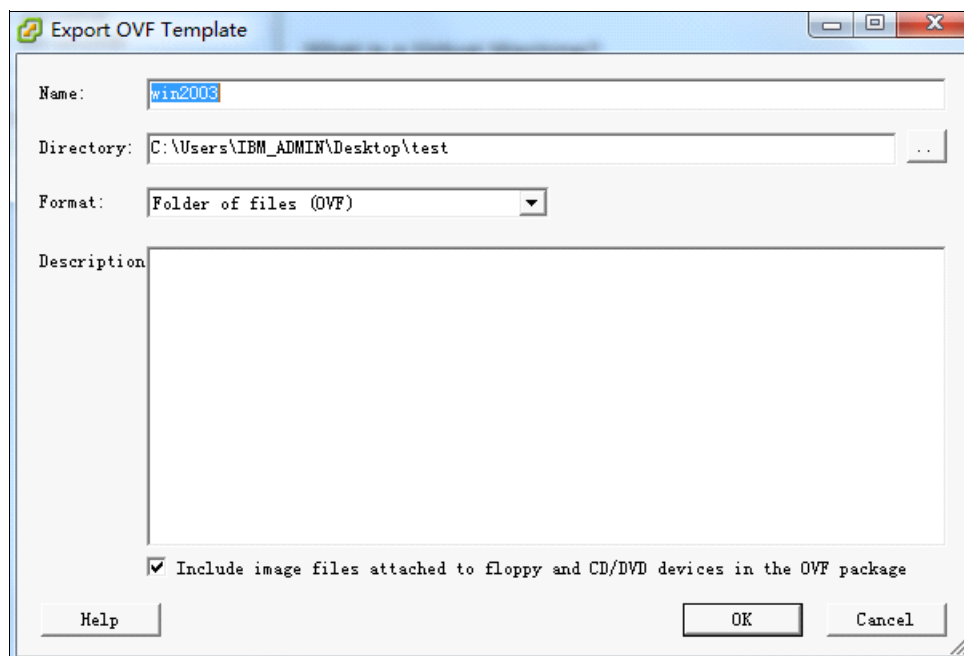


Figure 3-29 Export OVF Template

12. The time that it takes to complete the export process depends on the image size. After the process completes, the OVF package is shown in the directory that is selected for the export location, as shown in Figure 3-30.




Figure 3-30 Export process complete

13. Ensure that the OVF packages can be accessed by the vSphere Client that is referenced in the next step.
14. Restart the vSphere Client so that vCenter connects to and manages the destination compute node in the Flex System chassis. Add the compute node into the data center as shown in the process in step 1 on page 76 - step 9 on page 81.

Disconnected: The source server and destination server are not connected in this migration scenario. There are two installations of vCenter on different management servers. One installation is connected to the source server and the other is connected to the destination server.

The next step to perform on the target (Flex System) that was added to data center is to deploy a virtual OS image to the Flex System chassis.

15. Use the vCenter Console and click **File** → **Deploy OVF Template**, as shown in Figure 3-31.

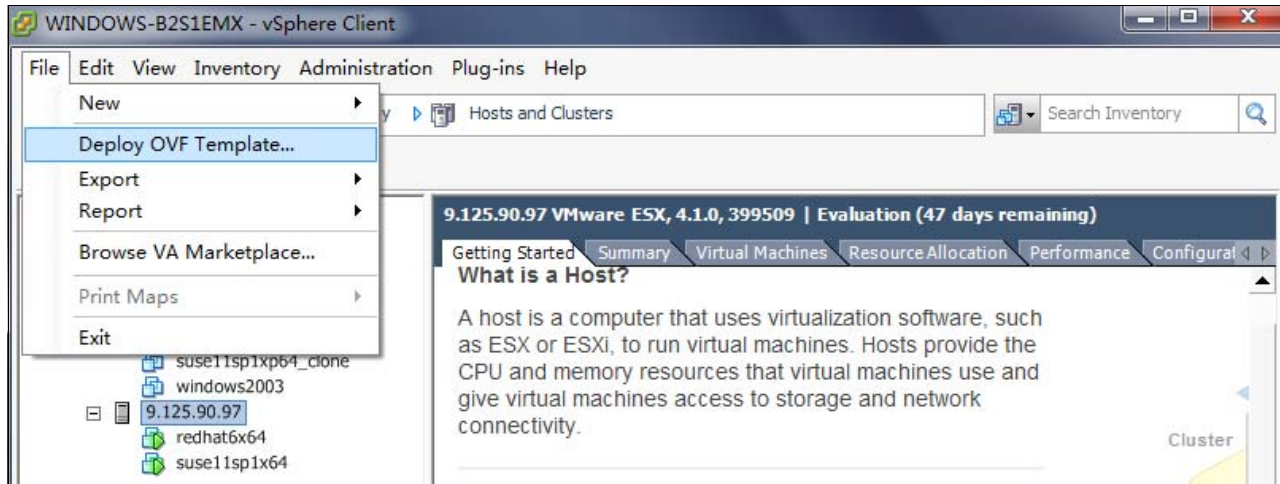


Figure 3-31 Deploy OVF Template

16. A wizard starts, as shown in Figure 3-32. Enter the path of OS images, define the virtual machine names and locations, and choose the disk format. Click **Next**.

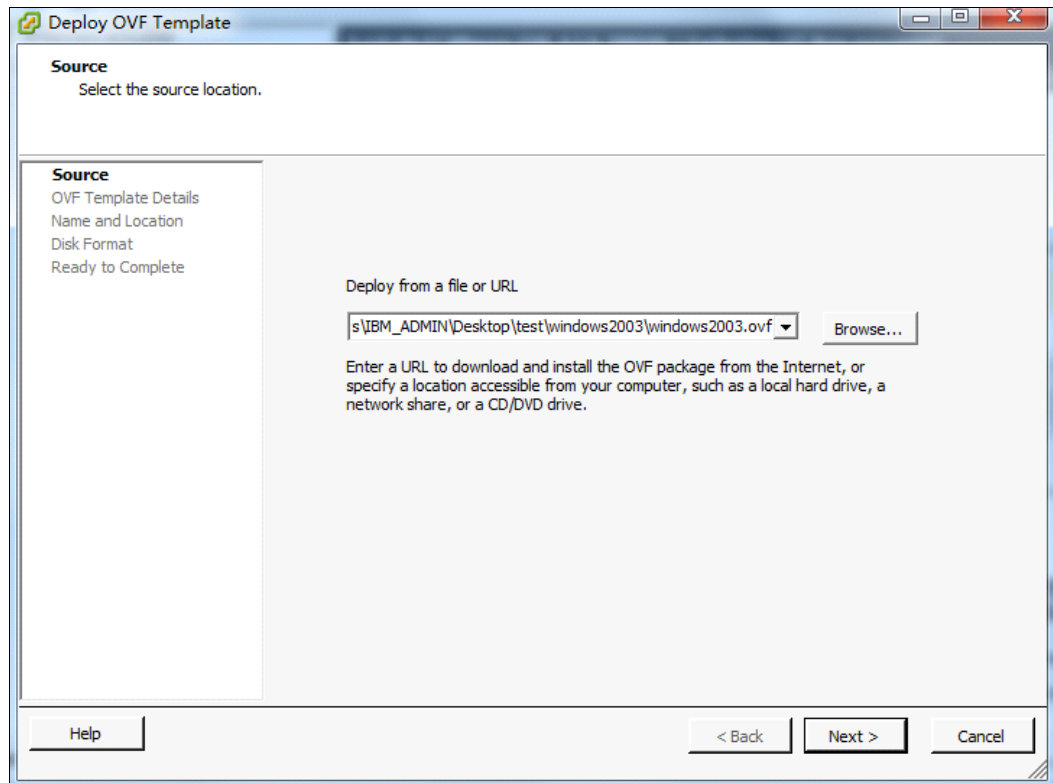


Figure 3-32 Deploy OVF Template

17. You see the summary of the information that you entered, as shown in Figure 3-33. Click **Finish** to start the deployment process.

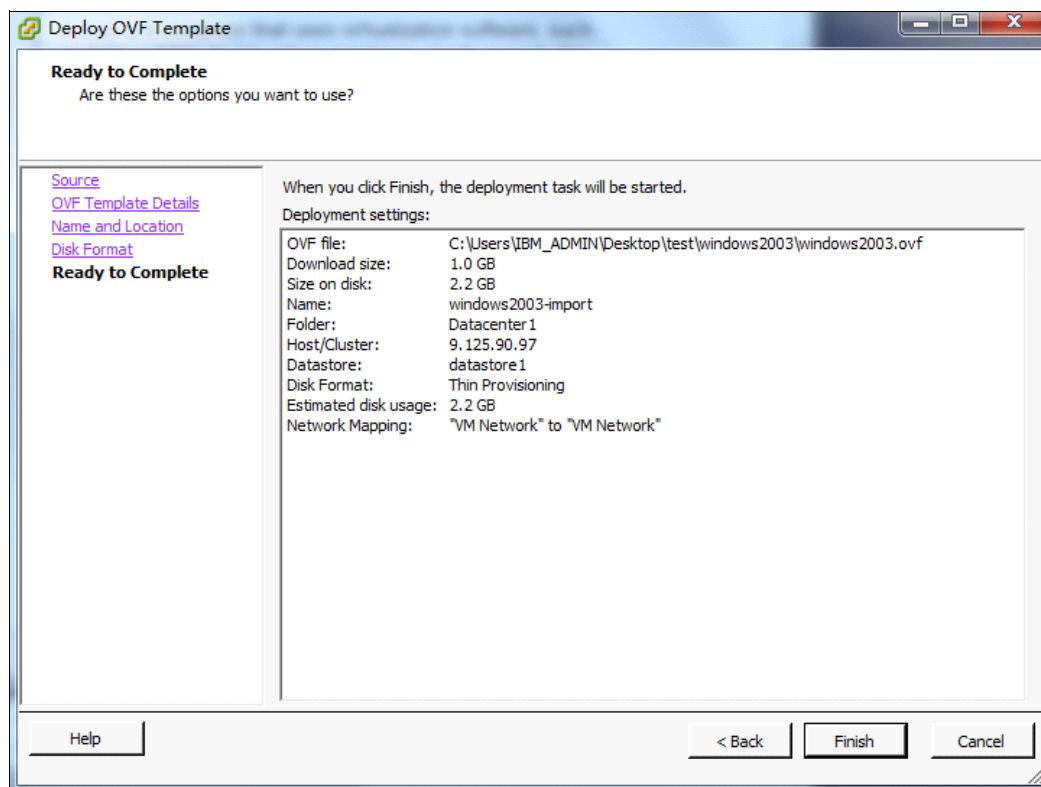


Figure 3-33 Deploy OVF Template: Summary

The deployment process takes minutes or hours, depending on the image size. After the deployment process is finished, an additional virtual machine is created and is displayed beneath the destination server, as shown in Figure 3-34.

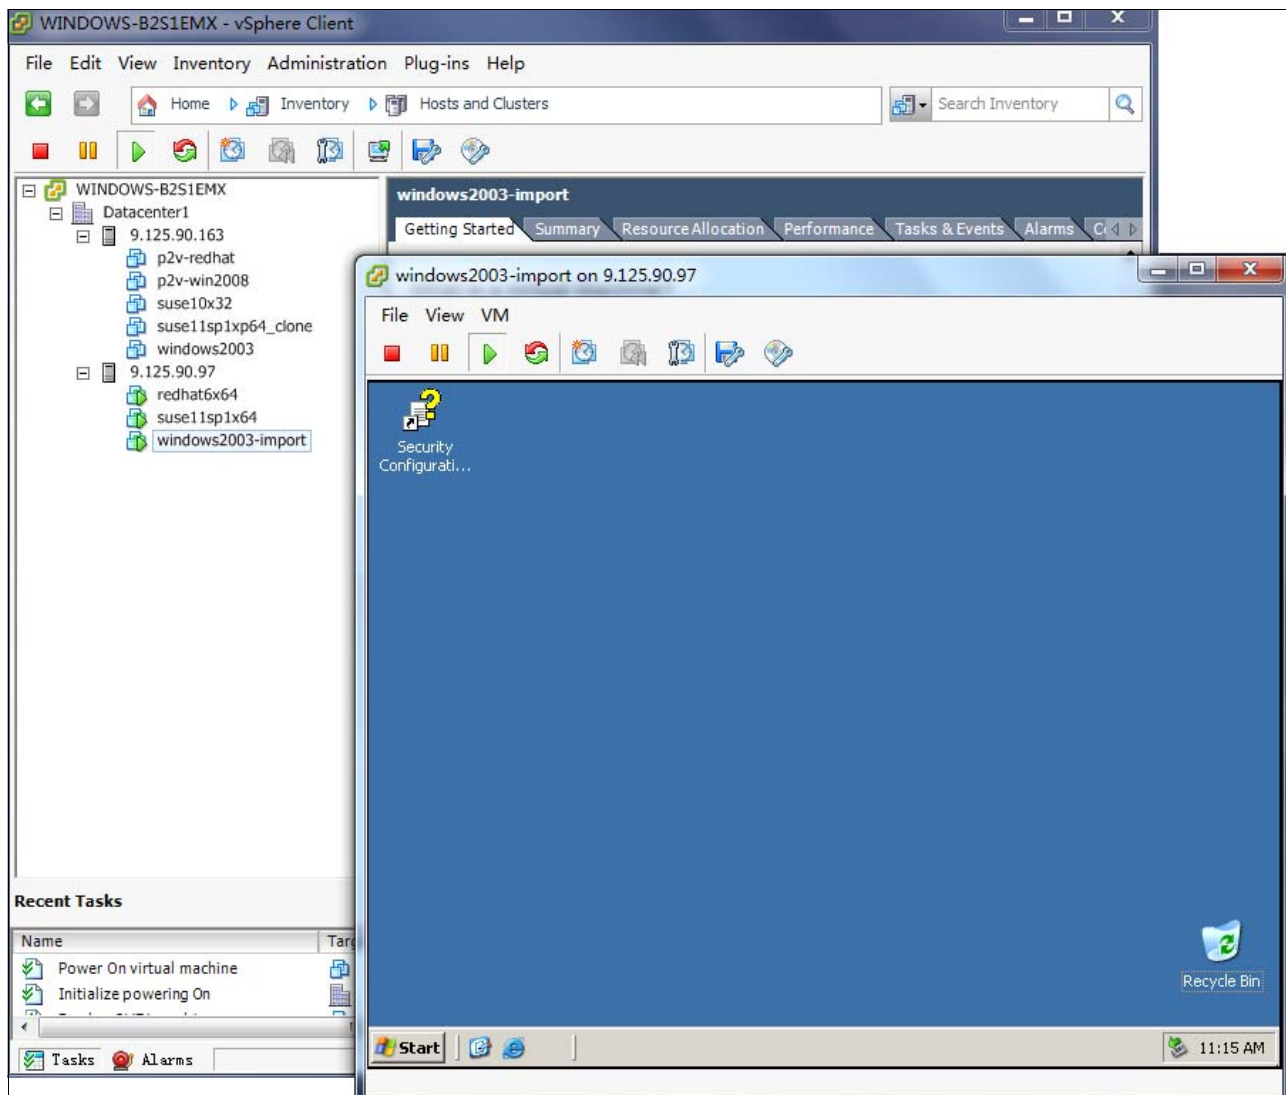


Figure 3-34 Deployed

The V2V process is now complete.

Use the System Center Virtual Machine Manager management tool when migrating a virtual machine on the Hyper-V hypervisor.

If you need to implement a V2V migration across hypervisors from different vendors, the format of the virtual disks is the key issue to consider. Ensure that the disk format of the virtual machine is supported by the source and target hypervisors. For example, the raw disk format is supported by hypervisors from nearly all the vendors. If the format is not supported, some transfer tools of the vendors are available, such as VMware vCenter Converter. For more information, see the website of the hypervisor vendor.

3.2 Source servers and Flex System are connected

In this section, the following scenarios are reviewed in which the source and target systems are connected:

- ▶ Connected Physical-to-Physical
- ▶ Connected Physical-to-Virtual
- ▶ Connected Virtual-to-Virtual

3.2.1 Connected Physical-to-Physical

Similar to Disconnected P2P as described in 3.1.1, “Disconnected Physical-to-Physical” on page 56, in addition to Norton Ghost, you can use the following end-to-end solutions:

- ▶ IBM Tivoli® Provisioning Manager:
<https://ibm.com/software/tivoli/products/prov-mgr/>
- ▶ Novell PlateSpin Migrate:
<http://www.novell.com/products/migrate/>
- ▶ Vision Solutions Double-Take Move
<http://www.visionsolutions.com/Products/DT-Move.aspx>

It is possible to plan a migration with Double-Take Move which involves near zero downtime.

3.2.2 Connected Physical-to-Virtual

As in the Disconnected P2V scenario described in 3.1.2, “Disconnected Physical-to-Virtual” on page 56, you need to understand the source system usage over the entire activity period. Understanding this usage helps you determine the amount of processing power and memory to assign to a virtual machine.

The task here is to choose the tools to implement the P2V function. The P2V migration functions in VMware vCenter or Microsoft System Center Virtual Machine Manager work fine for converting physical machines to virtual machines. In addition, some third-party companies have their own P2V conversion tools that are faster and have more features.

One important consideration for P2V conversion tools is determining the hypervisor that you want to use as the target. Some factors, such as resources, budget, time, and risks are key to the decision-making process.

With both the source and target interconnected (for example, in a LAN environment), some end-to-end solutions are available to capture the physical machine image into a virtual machine image and place it on the new target host, such as Vision Solutions Double-Take Move. For information about this tool, go to the following website:

<http://www.visionsolutions.com/Products/DT-Move.aspx>

The following scenarios illustrate the more general process of implementing a P2V migration from the source server to the Flex System solution:

- ▶ VMware vSphere scenario
- ▶ Microsoft Hyper-V scenario

VMware vSphere scenario

Some of the main components of a VMware vSphere environment, which includes vCenter server, vSphere Client, and any VMware hypervisors, must be deployed to complete the entire migration process. The Disconnected V2V scenario that is shown in 3.1.2, “Disconnected Physical-to-Virtual” on page 56 describes how to acquire, install, and configure the environment.

When all source servers are properly managed by vCenter Server, follow these steps:

1. Install the VMware OS on the needed compute nodes in the Flex System chassis.
The Flex System is managed by at least one vCenter server and the vCenter server installs the vCenter Converter plug-in.
2. After the new vCenter server is in place, use the vSphere Client to connect to the vCenter server.
3. Add the VMware hosts in the chassis to the new data center. (For more information, see 3.1.3, “Disconnected Virtual-to-Virtual” on page 75.)

4. Right-click the target server to which you want to migrate and click **Import Machine**, as shown in Figure 3-35. The Import Machine wizard starts.

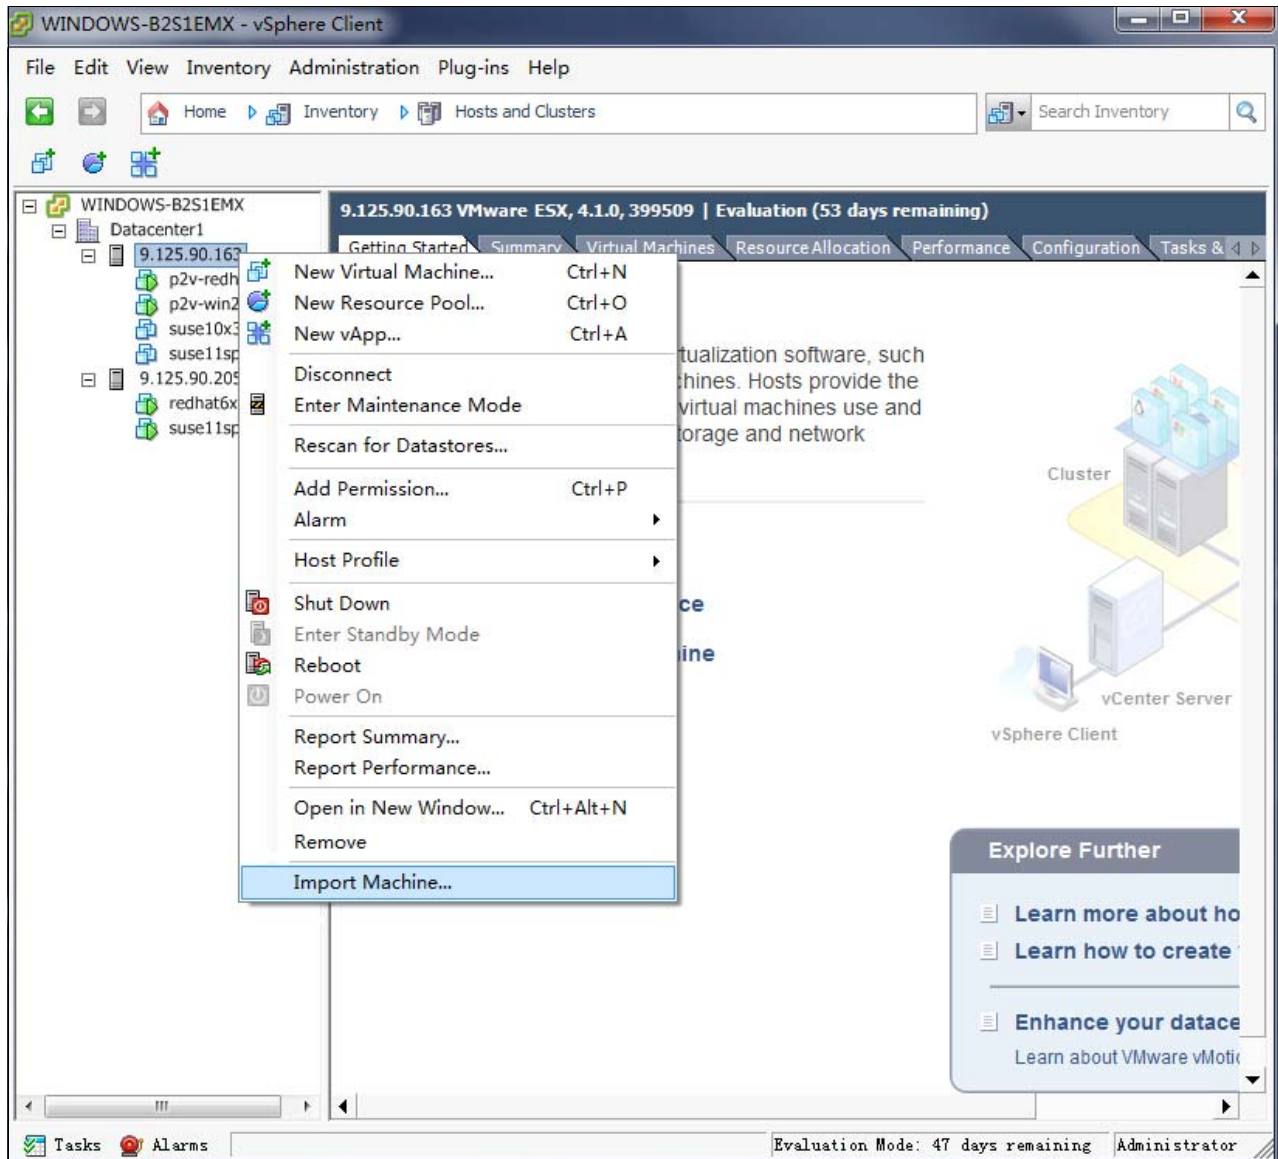


Figure 3-35 Import Machine

5. In the Source System window (see Figure 3-36), choose the Source type as **Powered-on machine**. Enter the IP address or name, User name, Password, and the OS Family of the source physical OS.

The screenshot shows the 'Import Machine' wizard window. The title bar says 'Import Machine'. Inside, the 'Source System' tab is selected. Below the title, it says 'Select the source system you want to convert'. On the left, there's a sidebar with 'Source System', 'Destination Location', 'Options', and 'Summary'. The main area shows 'Source: none' and 'Destination: 9.125.90.17'. A dropdown menu for 'Select source type:' is set to 'Powered-on machine'. Below it, a text box says 'Convert any powered-on physical or virtual machine.' Another section titled 'Specify the powered-on machine' contains fields for 'IP address or name:' (9.125.90.72), 'User name:' (root), 'Password:' (masked with dots), and 'OS Family:' (Linux). A link 'View source details...' is below these fields. At the bottom, there are buttons for 'Help', 'Export diagnostic logs...', '< Back', 'Next >', and 'Cancel'.

Figure 3-36 Import Machine wizard

Tip: This wizard also can implement the V2V connected migration, as shown in Figure 3-37.

This screenshot shows the same 'Import Machine' wizard window as Figure 3-36, but with the 'Select source type:' dropdown menu open. The menu lists several options: 'Powered-on machine', 'Powered-on machine', 'VMware Infrastructure virtual machine', 'VMware Workstation or other VMware virtual machine', 'Backup image or third-party virtual machine', and 'Hyper-V Server'. The 'Hyper-V Server' option is currently selected and highlighted. The rest of the window, including the sidebar and other fields, remains the same as in Figure 3-36.

Figure 3-37 Source type options

6. Confirm the destination location, as shown in Figure 3-38.

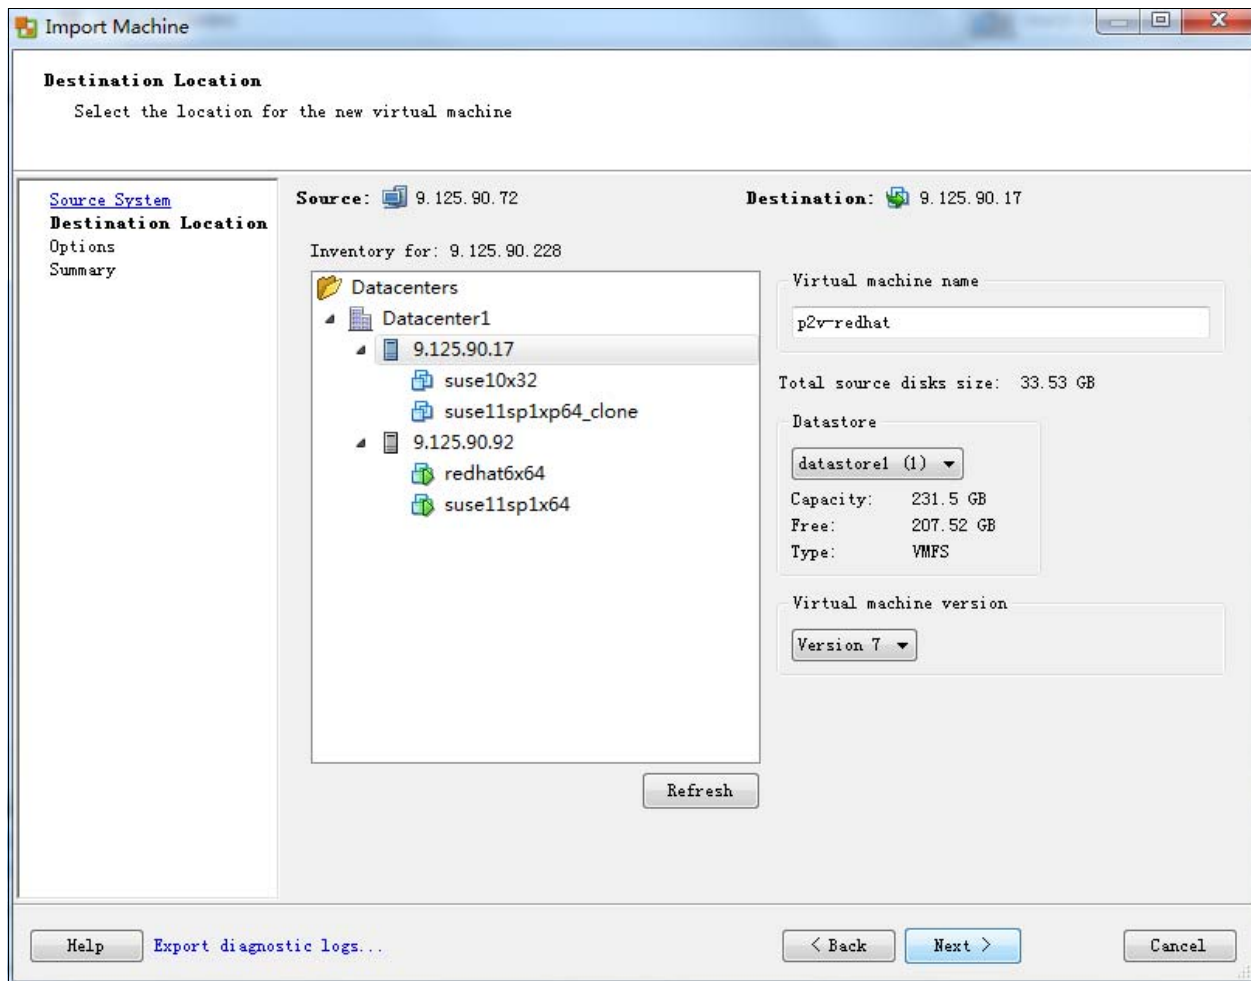


Figure 3-38 Import Machine wizard: Destination Location

7. Edit the options for the virtual machine, such as CPU, memory, or the disk settings in the Import Machine window, as shown in Figure 3-39. Use the same settings on the physical OS of the source server if the target server has the capacity.

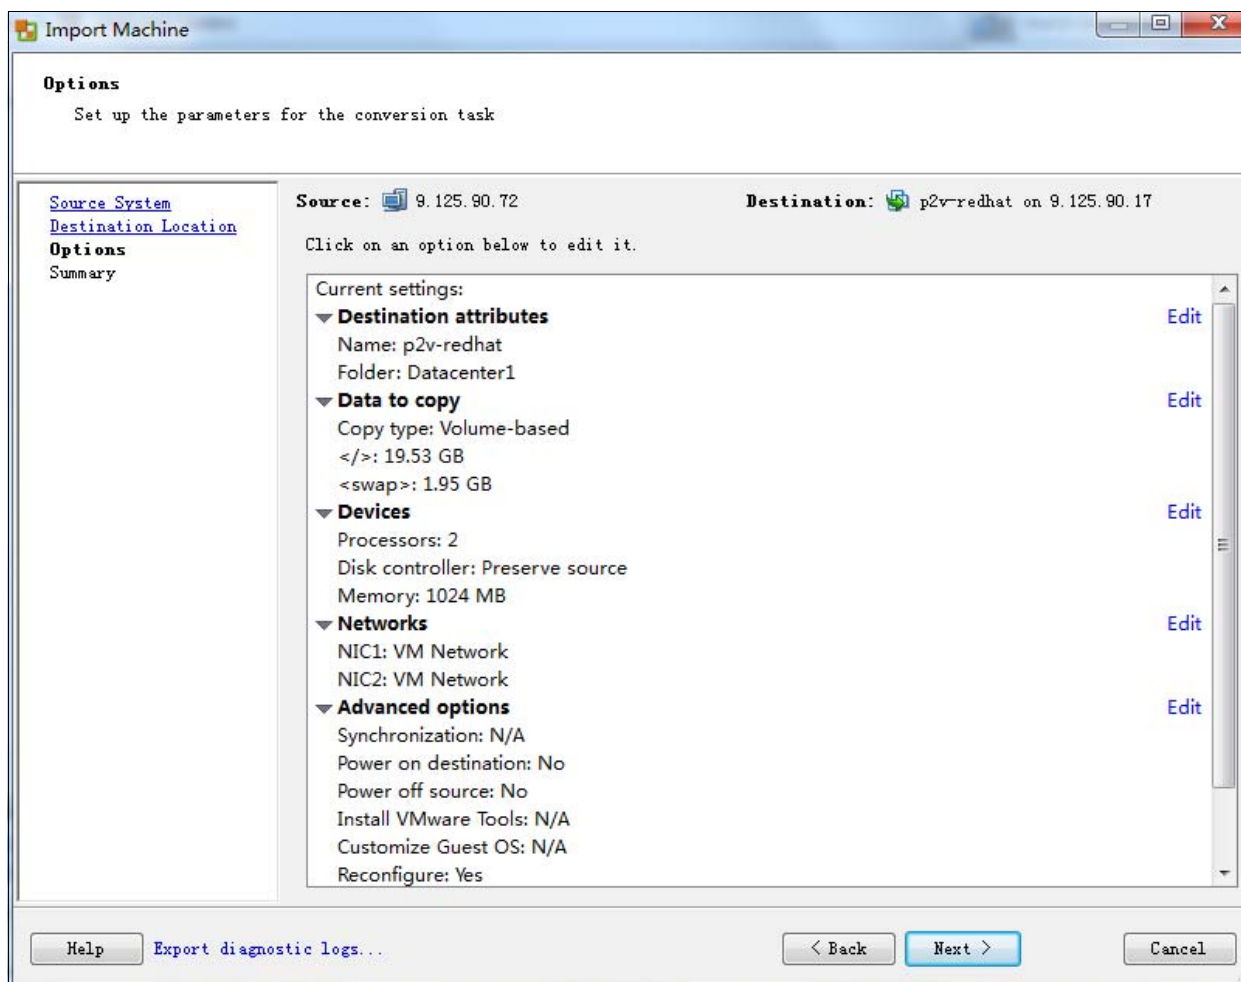


Figure 3-39 Import Machine wizard: Options

Figure 3-40 shows the dialog that opens when you click **Edit** next to Data to copy, for example.

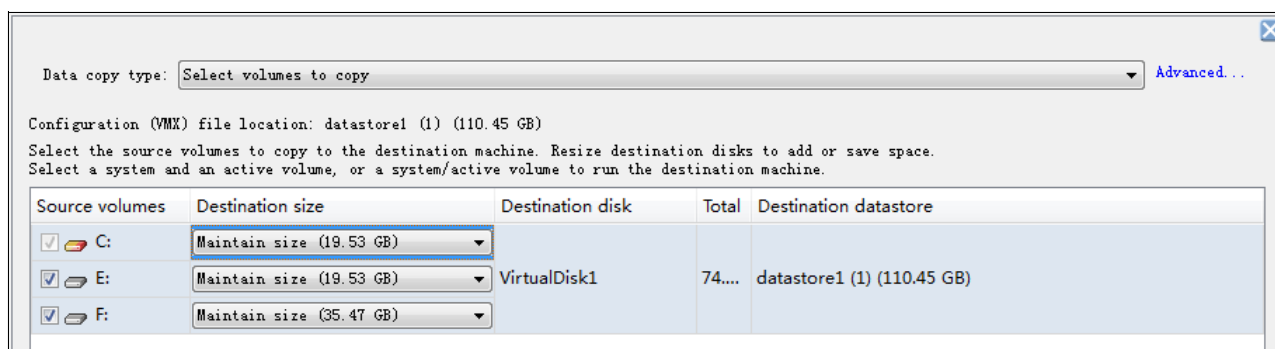


Figure 3-40 Changing the options for Data to copy

8. Click **Next** to confirm the setting (see Figure 3-41). Click **Finish**.

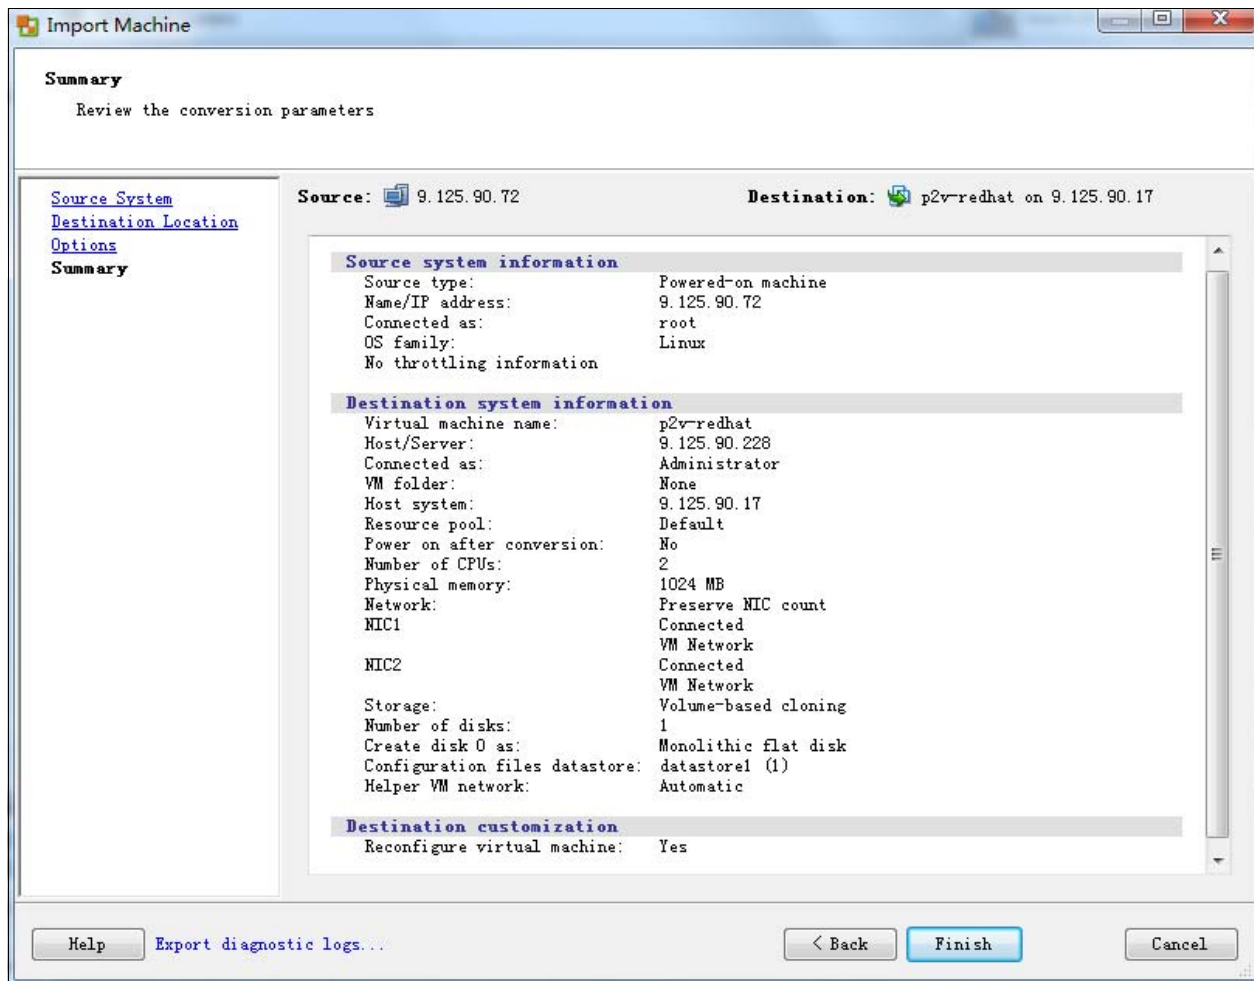


Figure 3-41 Import Machine wizard: Summary

Another task is added to the list in the Recent Tasks window at the bottom of the vSphere client interface, as shown in Figure 3-42. The task processing can take minutes or hours, depending on the disk size. After the task is finished, another virtual machine opens and is associated with the chosen host.

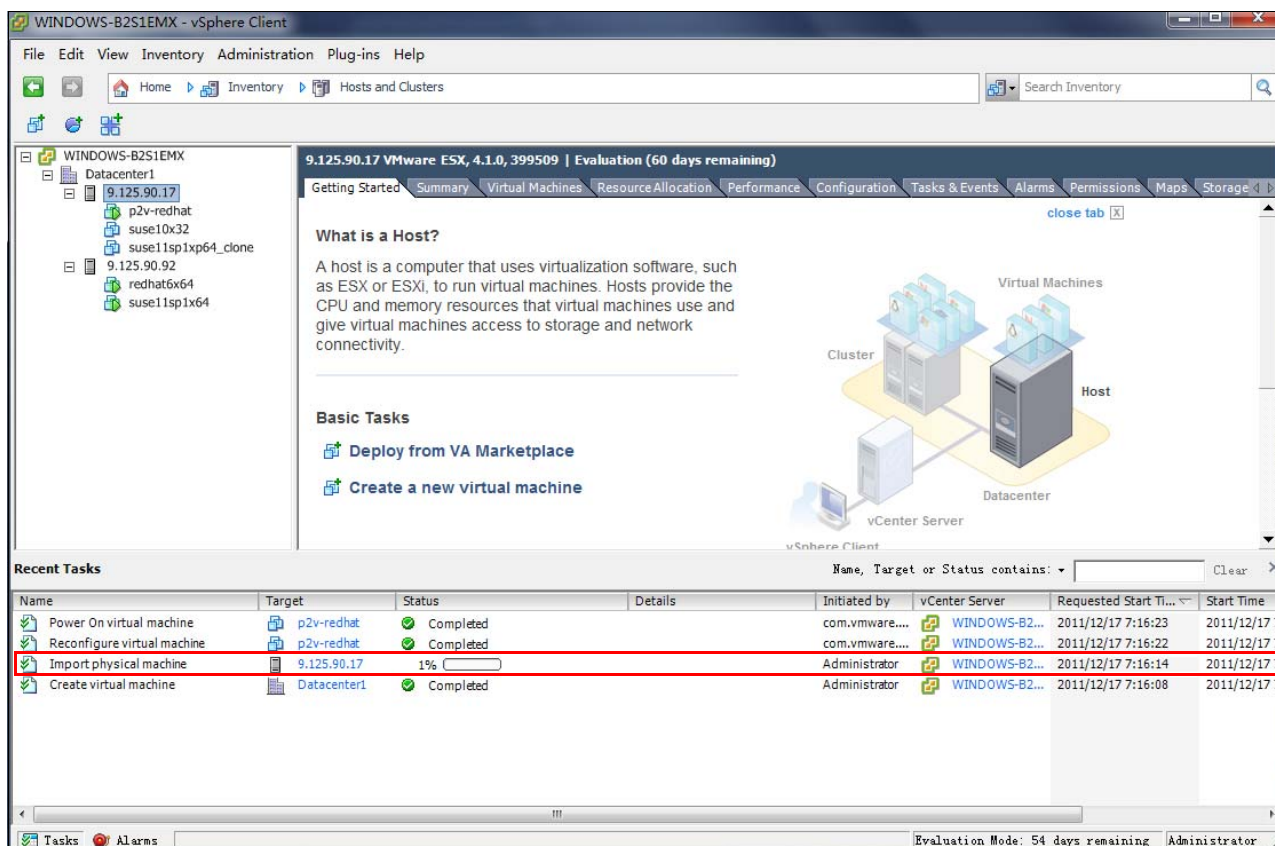


Figure 3-42 Import Machine task in progress

Microsoft Hyper-V scenario

Before you start this migration process, the following requirements and restrictions must be reviewed.

The source computer must meet the following requirements:

- ▶ No volumes larger than 2040 GB
- ▶ Must be accessible by VMM and the host computer
- ▶ Cannot be in a DMZ

A DMZ (screened subnet) is a collection of devices and subnets that are placed between an intranet and the Internet to help protect the intranet from unauthorized Internet users. The source computer for a P2V conversion can be in any other network topology in which the SCVMM management server can connect to the source machine to temporarily install an agent. The management server also must be able to make Windows Management Instrumentation (WMI) calls to the source computer.

For a list of supported physical machine operating systems for P2V conversions in VMM 2008 and VMM 2008 R2, see this website:

<http://technet.microsoft.com/en-us/library/cc764232.aspx>

P2V has the following restrictions for VMM:

- ▶ VMM does not support P2V on source computers that are running Windows NT Server 4.0. However, you can use the Microsoft Virtual Server 2005 Migration Toolkit (VSMT) or third-party solutions for converting computers that are running Windows NT Server 4.0.
- ▶ VMM 2008 R2 does not support converting a physical computer that is running Windows Server 2003 SP1 to a virtual machine that is managed by Hyper-V.

Hyper-V does not support Integration Components on computers that are running Windows Server 2003 SP1. As a result, there is no mouse control when you use Remote Desktop Protocol (RDP) to connect to the virtual machine. To avoid this issue, update the OS to Windows Server 2003 SP2 before converting the physical computer. You also can convert the computer by using VMM 2008 and then deploy the virtual machine in VMM 2008 R2.

The following requirements for the destination host server must be met:

- ▶ The destination host must be in the domain of the management server and must not be in a DMZ.
- ▶ As in any virtual machine creation or migration, the destination host for a P2V conversion must have sufficient memory for the virtual machine. The host also must have memory that is reserved for the host OS. By default, the amount of memory that is reserved for the host OS is 256 MB in VMM 2008 or 512 MB in VMM 2008 R2. If the host does not have enough memory for the virtual machine in addition to the memory reserved for the host, the user receives a placement error in the Convert Physical Server Wizard.

During a P2V conversion, the destination host can be running Windows Server 2008 with Hyper-V, Windows Server 2008 R2 with Hyper-V, or Virtual Server R2 SP1 (or later).

It is assumed that you prepared all of the hardware and software for the source and destination server. It is also assumed that the servers were added to the pool of System Center Virtual Machine Manager. (See 3.1.2, “Disconnected Physical-to-Virtual” on page 56.)

Follow these steps to implement a P2V conversion with Hyper-V:

1. From the SCVMM console, click **Convert physical server** to start the P2V process, as shown in Figure 3-43.

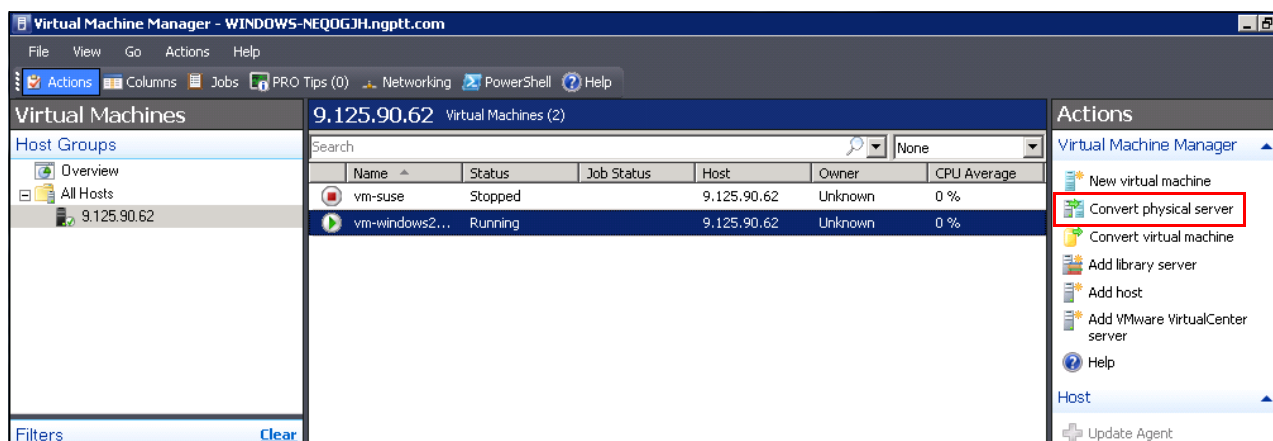


Figure 3-43 Starting the P2V process in SCVMM

Tip: The source physical server is connected to the destination server. The management server must be able to capture the source OS image and deploy this image to a virtual machine in one step.

2. The P2V wizard starts. Enter the IP address or host name and the Administrative account information for the physical computer, as shown in Figure 3-44. Click **Next**.

The screenshot shows the 'Convert Physical Server (P2V) Wizard' window. The title bar reads 'Convert Physical Server (P2V) Wizard'. The main window has a dark blue header with a 'Select Source' icon and text. On the left is a vertical navigation pane with the following items: 'Select Source' (highlighted), 'Virtual Machine Identity', 'System Information', 'Volume Configuration', 'VM Configuration', 'Select Host', 'Select Path', 'Select Networks', 'Additional Properties', 'Conversion Information', and 'Summary'. The main content area has a light gray background with the instruction 'Select the physical computer that you want to convert to a virtual machine.' Below this are several input fields: 'Computer name or IP address:' with the value '9.125.90.62' and a 'Browse...' button; 'Administrative account' section with the instruction 'Specify the administrative account to use to connect to the physical computer.'; 'User name:' with the value 'Administrator'; 'Password:' with a masked field of dots; and 'Domain or computer name:' with the value '9.125.90.62'. An information icon and text state: 'If the source machine is not in a domain, specify the source machine name or IP address.' At the bottom left is a link 'Requirements for a P2V conversion'. At the bottom right are 'Next' and 'Cancel' buttons.

Figure 3-44 SCVMM P2V wizard: Select Source

3. Enter suitable values in the Virtual Machine identity page. Click **Next**.
4. Click **Scan System** (as shown in Figure 3-45 on page 98) to scan the source physical computer and install the VMM agent on the source physical computer. (The VMM agent is installed automatically on the source server if the source server is in the same domain as the management server.)

5. The characteristics of the source system are displayed in the System Information window of the wizard, as shown in Figure 3-45. Click **Next**.

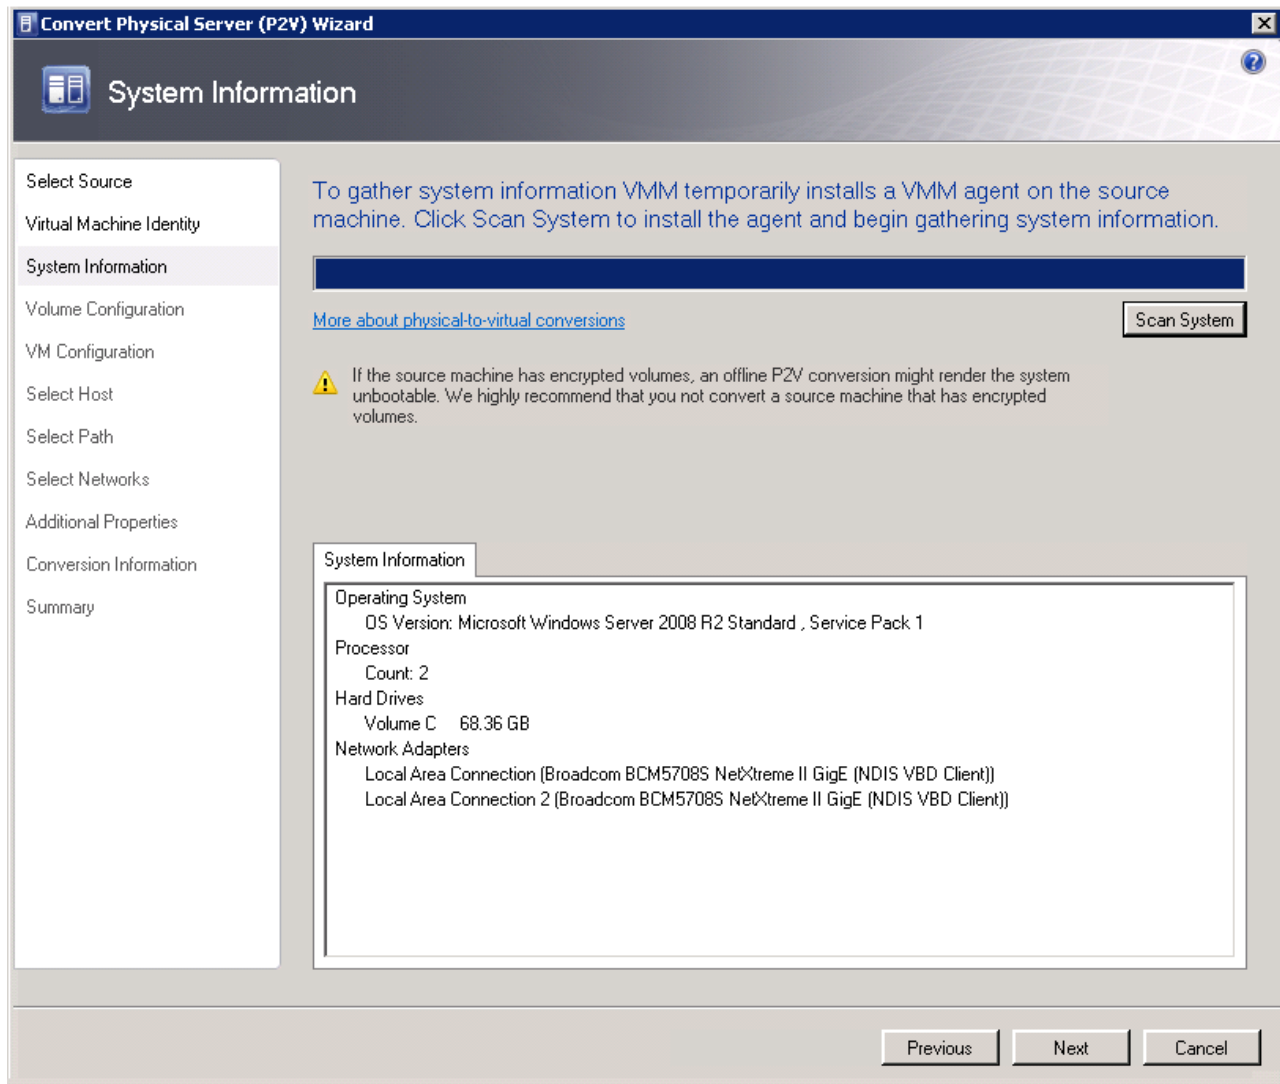


Figure 3-45 SCVMM P2V wizard: System Information

6. In the Volume Configuration window, select the volumes that are captured on the source machine, as shown in Figure 3-46. Click **Next**.

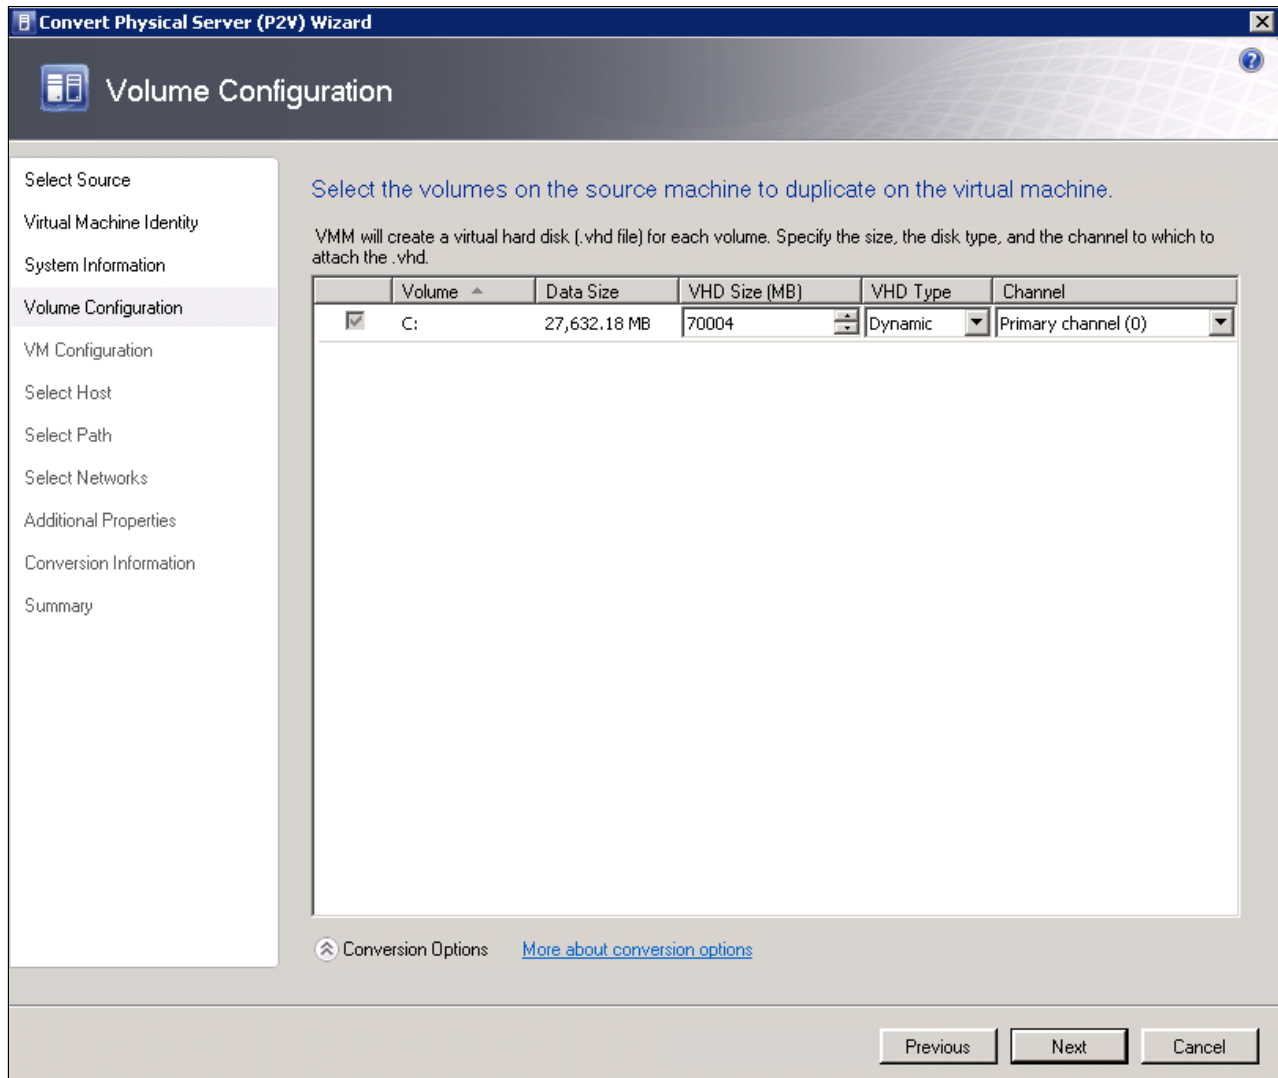


Figure 3-46 SCVMM P2V wizard: Volume Configuration

- Define the processors, memory, disk, network, and any additional custom properties for the destination virtual machine, as shown in Figure 3-47, Figure 3-48, Figure 3-49, and Figure 3-50 on page 101.

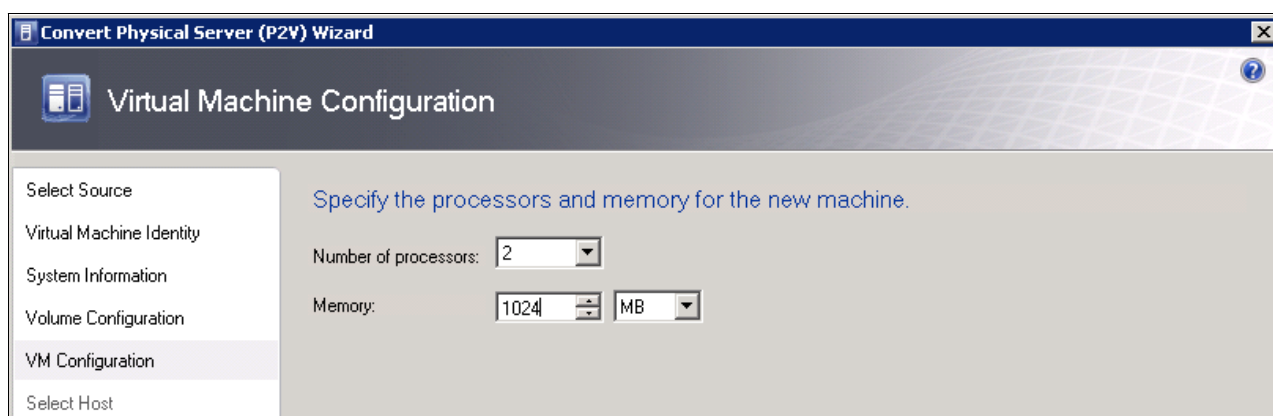


Figure 3-47 SCVMM P2V wizard: VM Configuration

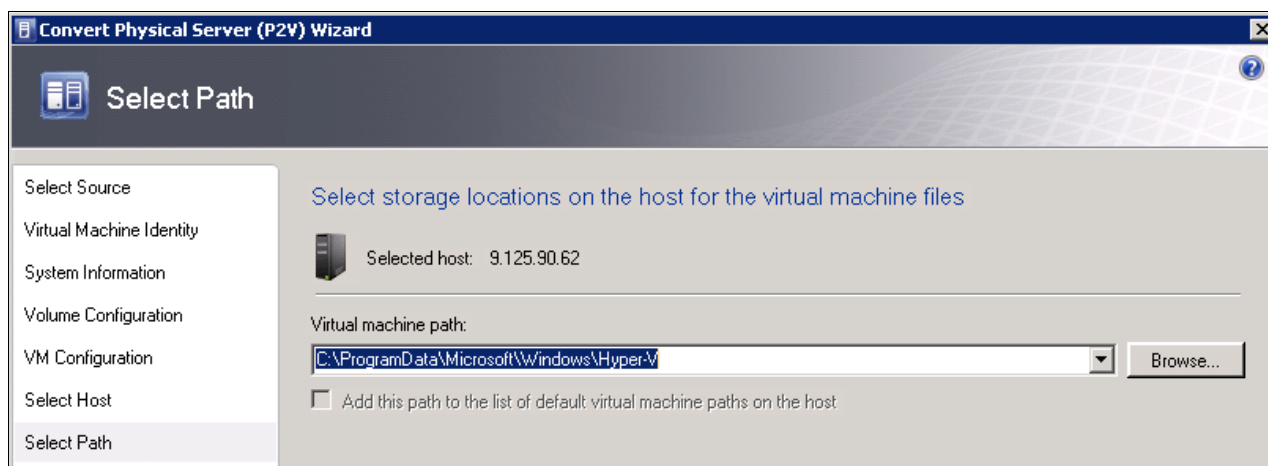


Figure 3-48 SCVMM P2V wizard: Select storage location

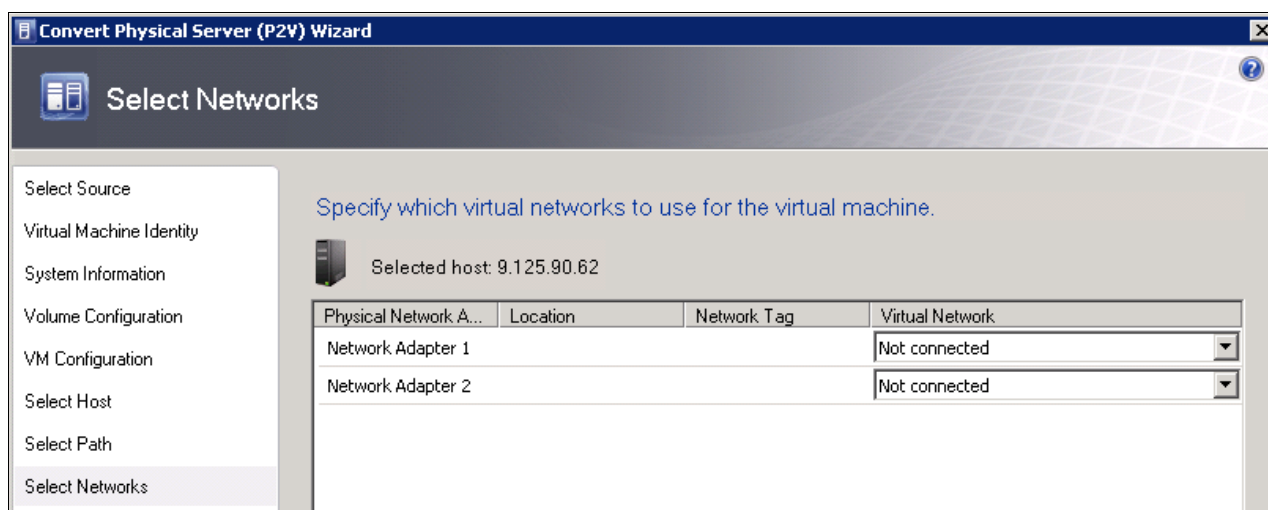


Figure 3-49 SCVMM P2V wizard: Select virtual networks

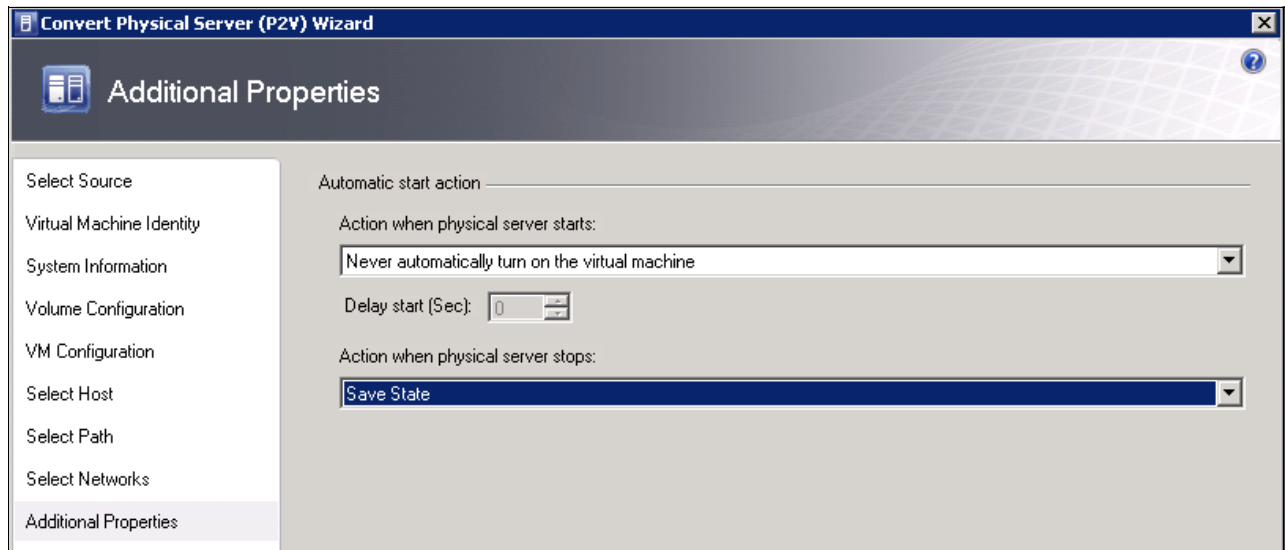


Figure 3-50 SCVMM P2V wizard: Additional Properties

8. After the configuration is complete, review the Summary page, as shown in Figure 3-51. Click **Create**.

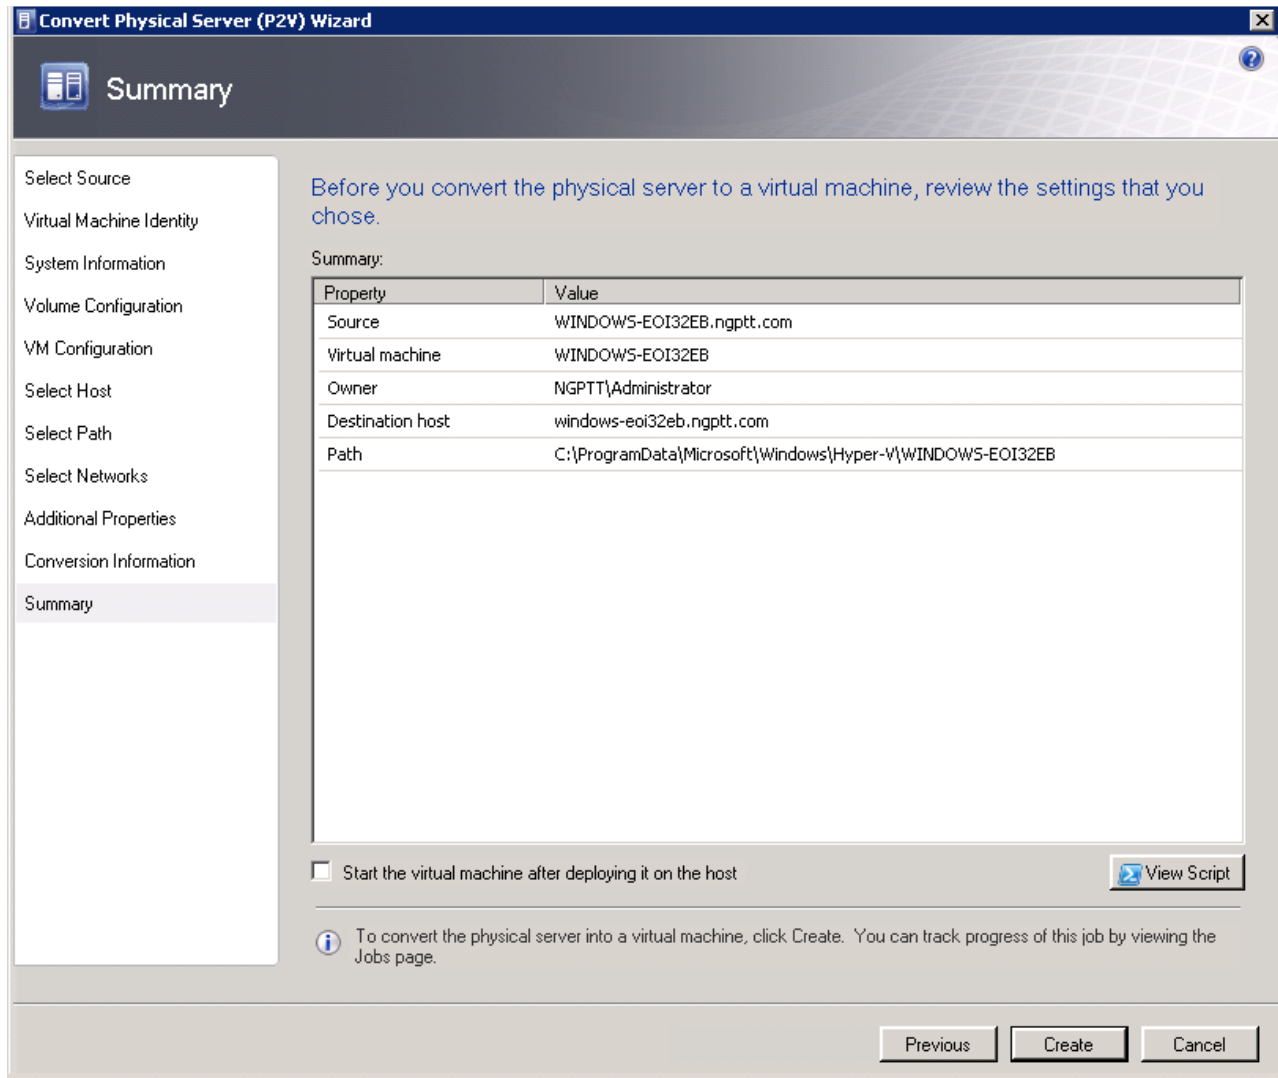


Figure 3-51 SCVMM P2V wizard: Summary

The job runs for a short time. The progress of the job is displayed, as shown in Figure 3-52.

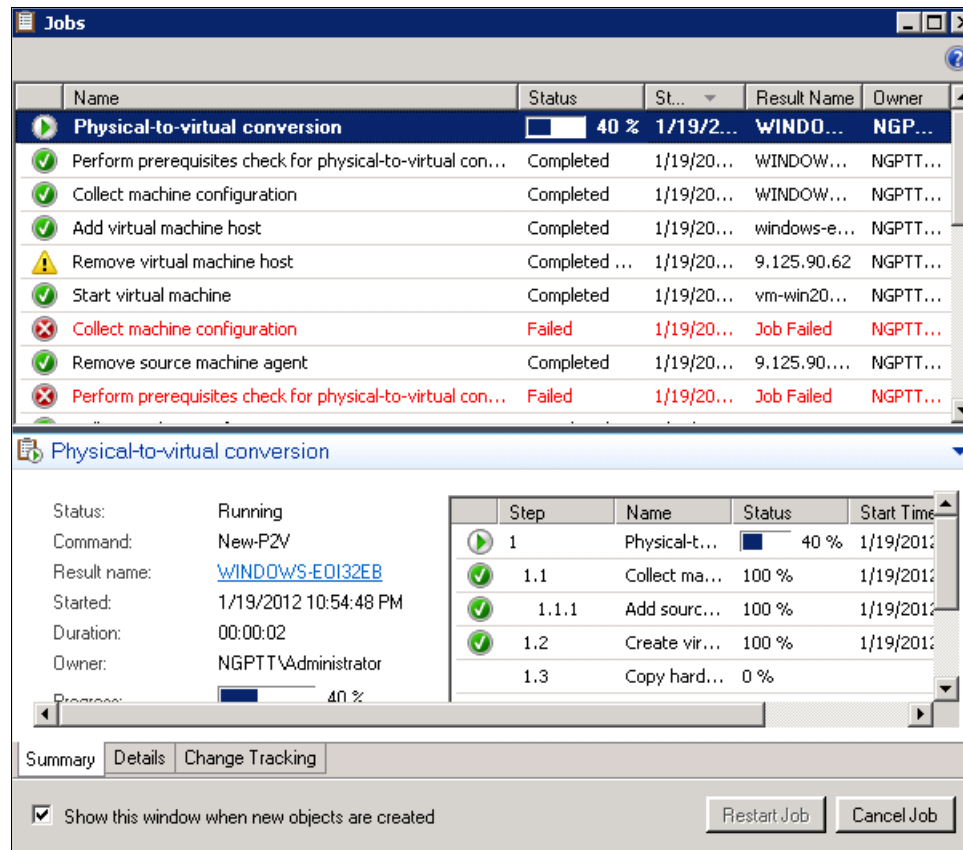


Figure 3-52 Progress of the P2V job

When the P2V conversion is complete, you see that one more virtual machine is displayed on the destination server, as shown in Figure 3-53.

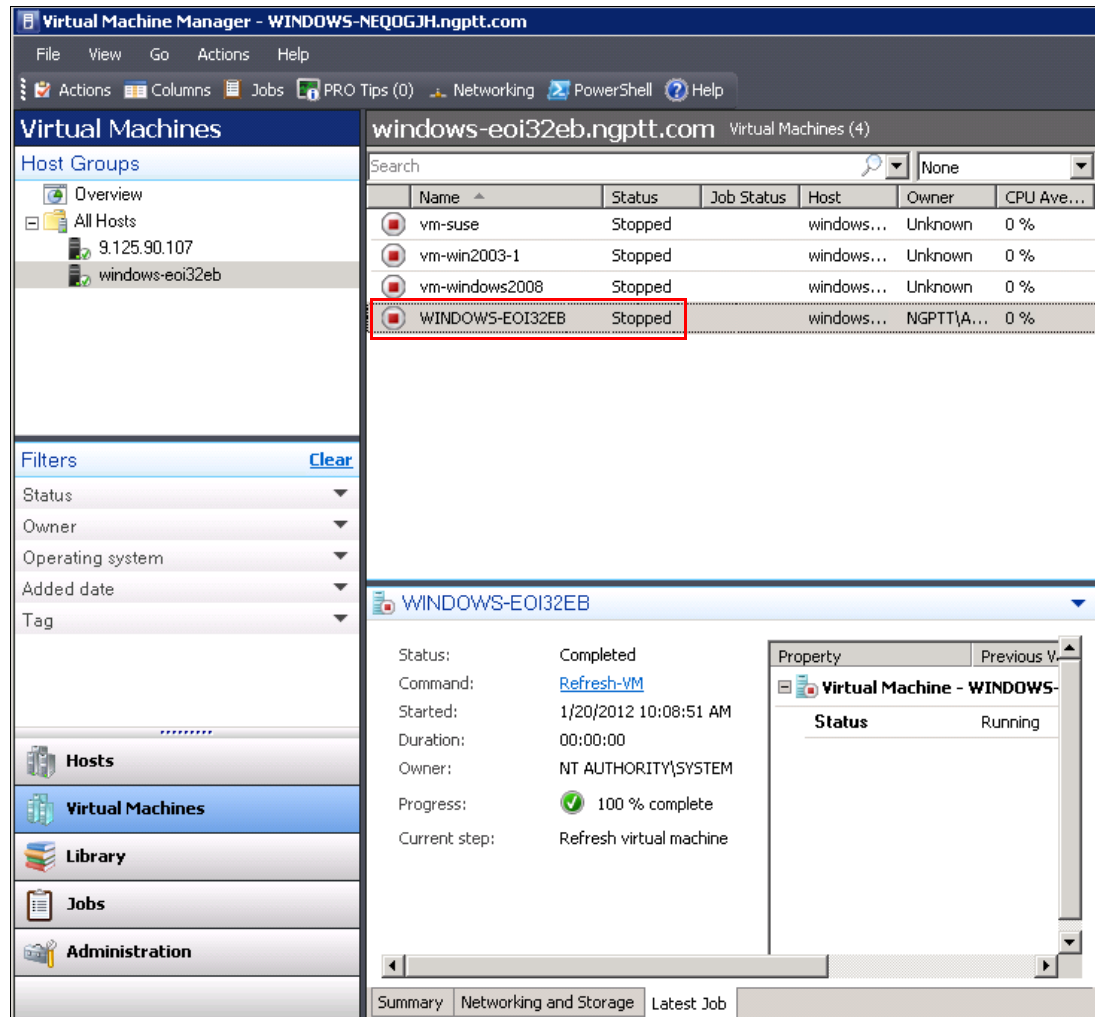


Figure 3-53 P2V conversion complete

3.2.3 Connected Virtual-to-Virtual

The Connected V2V migration scenario has the same rules and requirements as the Disconnected V2V migration scenario as described in 3.1.3, “Disconnected Virtual-to-Virtual” on page 75.

As of this writing, there are two types of V2V migrations in the industry:

- ▶ Live V2V migration
- ▶ Offline V2V migration

The Live V2V migration does not require you to take the source virtual machine offline. Normal operations are completed during the migration process. One example of a live migration solution is Vision Solutions Double-Take Move:

<http://www.visionsolutions.com/Products/DT-Move.aspx>

Offline V2V migration requires you to shut down the source virtual machine before starting the migration process.

You implement the offline migration process from the source server to the Flex System Compute Node. Implementing this type of migration in this manner is necessary because this live migration requires the source and destination servers to have the same type of processor. This requirement is not met in current Flex System transition scenarios. Also, live migration is a fee-based function of the virtualization management software.

In this section, the following migration scenarios are described:

- ▶ VMware scenario
- ▶ Hyper-V scenario
- ▶ KVM scenario

VMware scenario

The following requirements must be met before starting the VMware migration process:

- ▶ VMware must be installed on the destination Flex System compute nodes.
- ▶ A VMware vCenter server must be connected to the source server and the destination server.
- ▶ The source server and the destination server must be added to the necessary data server in vCenter.

If any of these requirements are not met, see 3.1.3, “Disconnected Virtual-to-Virtual” on page 75 for more information about completing the requirements.

Follow these steps to implement a V2V migration:

1. Use the vSphere Client to connect to the vCenter server. Right-click the source virtual machine that you are migrating and click **Clone**, as shown in Figure 3-54.

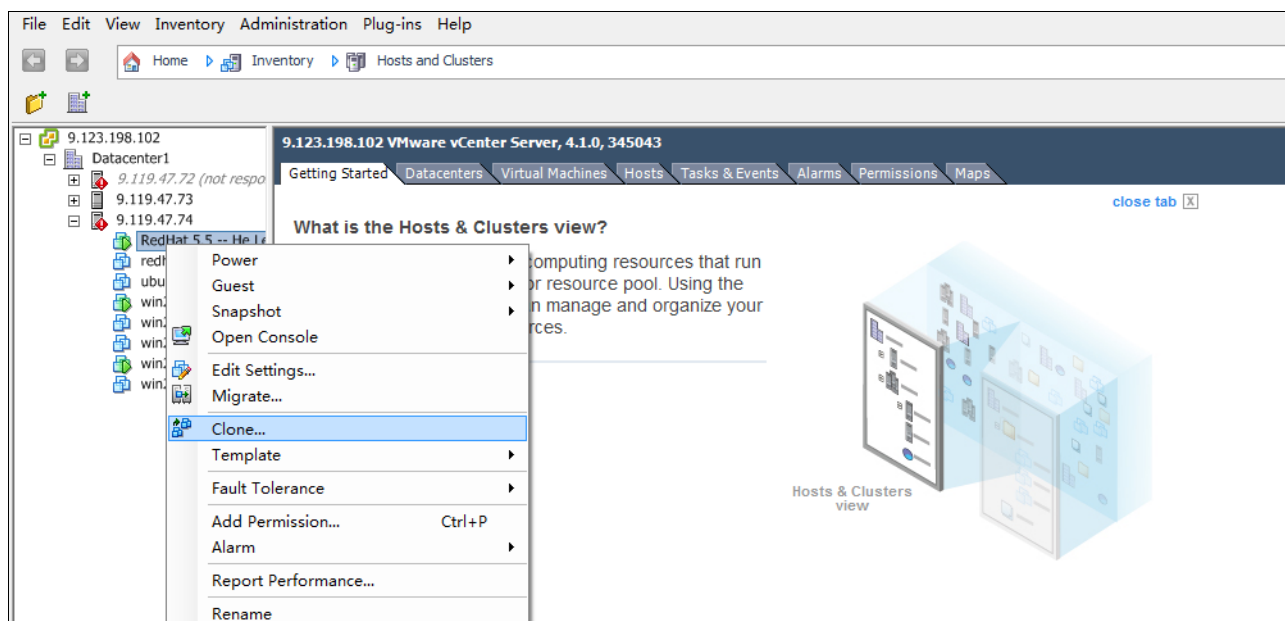


Figure 3-54 Starting the clone process

Migrate versus Clone: There are two options for this scenario that are available on the context-sensitive menu: *Migrate* and *Clone*. Migrate deletes the source VM image after copying the image from the source server to the destination server. Clone does not delete the image after the image is copied. To reduce the risk of data loss, the user must choose the Clone option.

2. The Clone Virtual Machine wizard starts. Enter the name and the location of the destination virtual machine, as shown in Figure 3-55. Click **Next**.

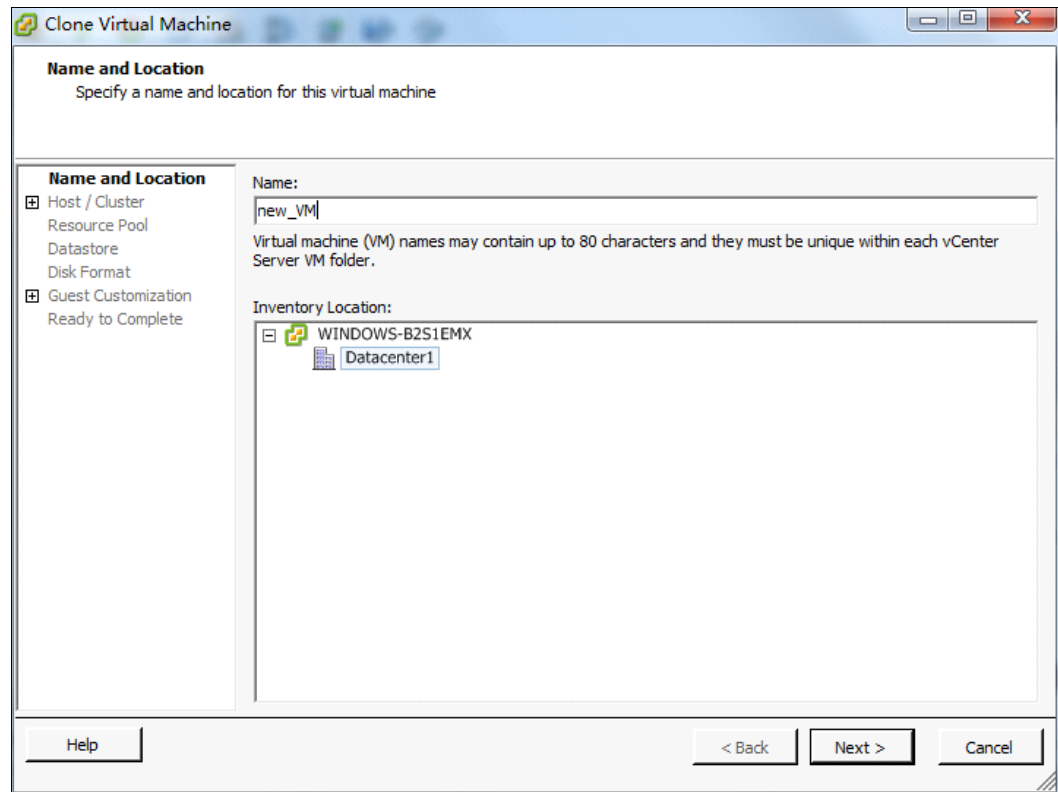


Figure 3-55 Clone Virtual Machine wizard: Name and Location

3. In the Host/Cluster window, as shown in Figure 3-56, select the destination server. Click **Next**.




Figure 3-56 Clone Virtual Machine wizard: Host/Cluster

4. Select the data store for the new virtual machine, as shown in Figure 3-57.

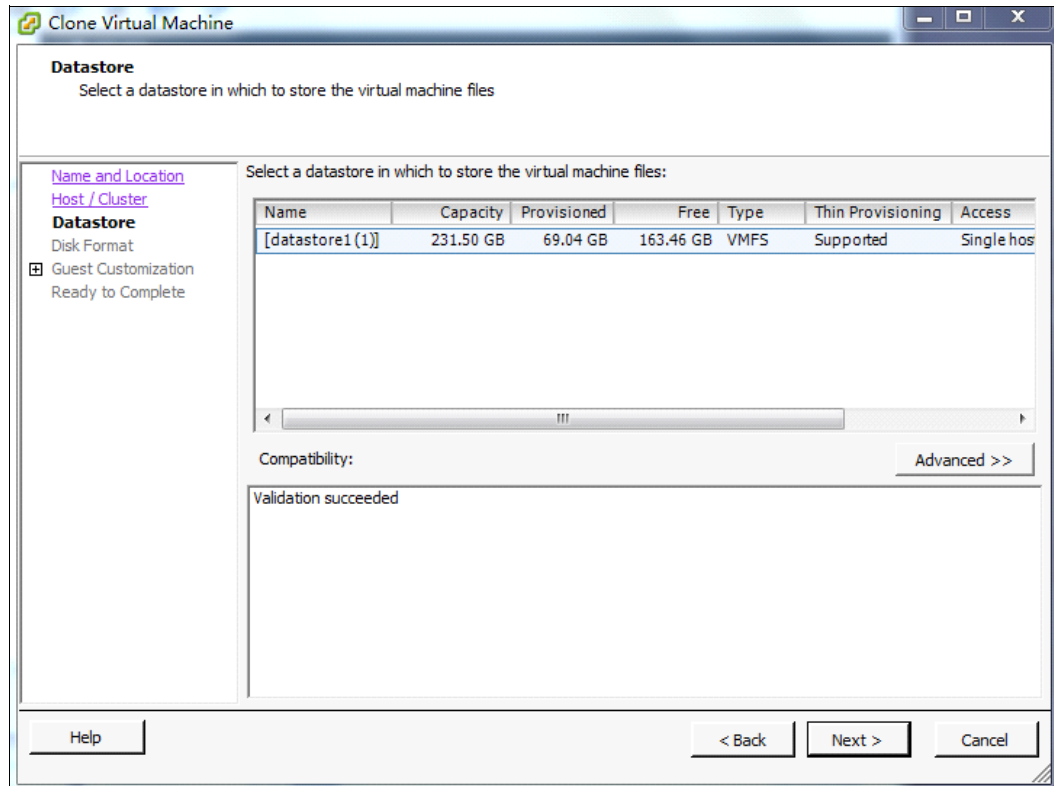


Figure 3-57 Clone Virtual Machine wizard: Datastore

5. Select the disk format for the new virtual machine, as shown in Figure 3-58. To avoid any additional risk that a disk format change presents, the user must choose the **Same format as source** option. The potential disk format change risk might include (but is not limited to) lost disk access or data corruption. Click **Next**.

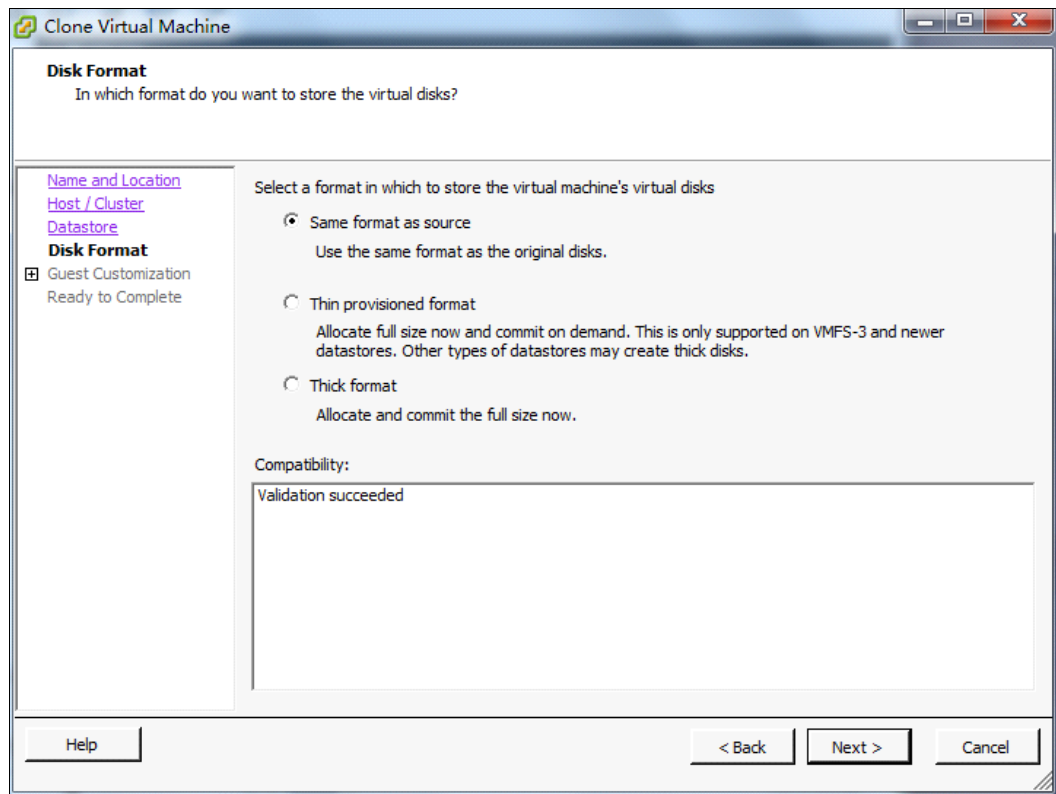


Figure 3-58 Clone Virtual Machine wizard: Disk Format

6. Customize any other settings of the virtual machine. In our example, the Guest Customization wizard was run, as shown in Figure 3-59 and Figure 3-60.

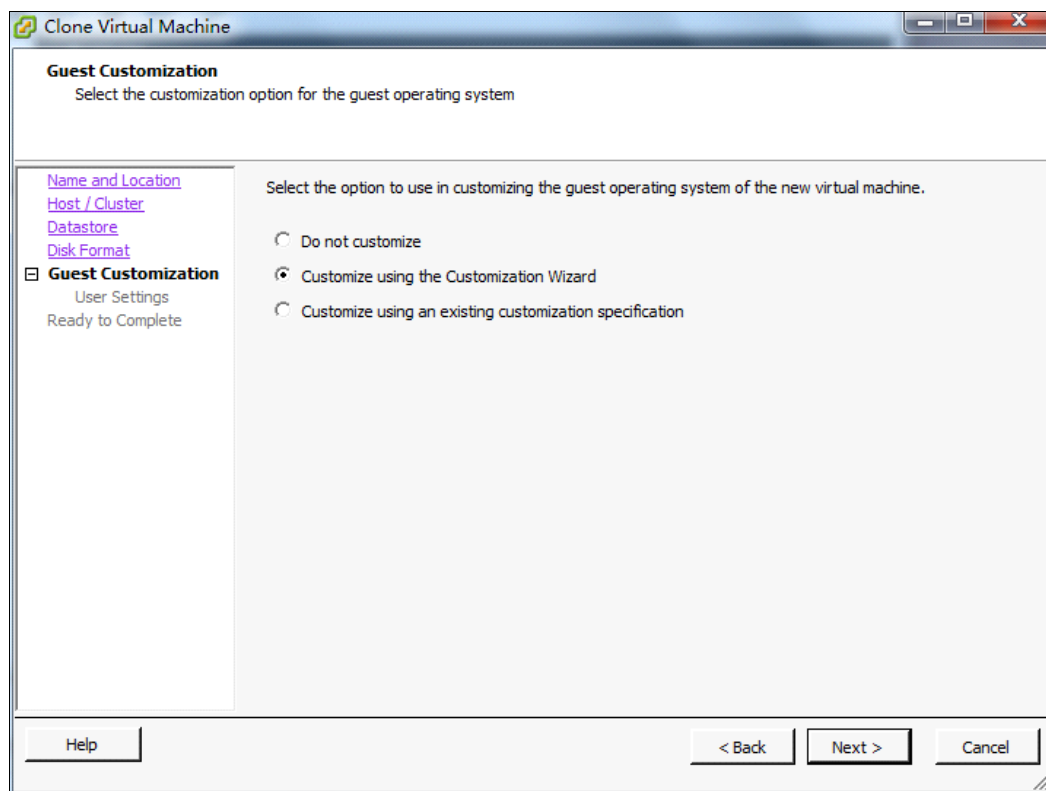


Figure 3-59 Clone Virtual Machine wizard: Guest Customization

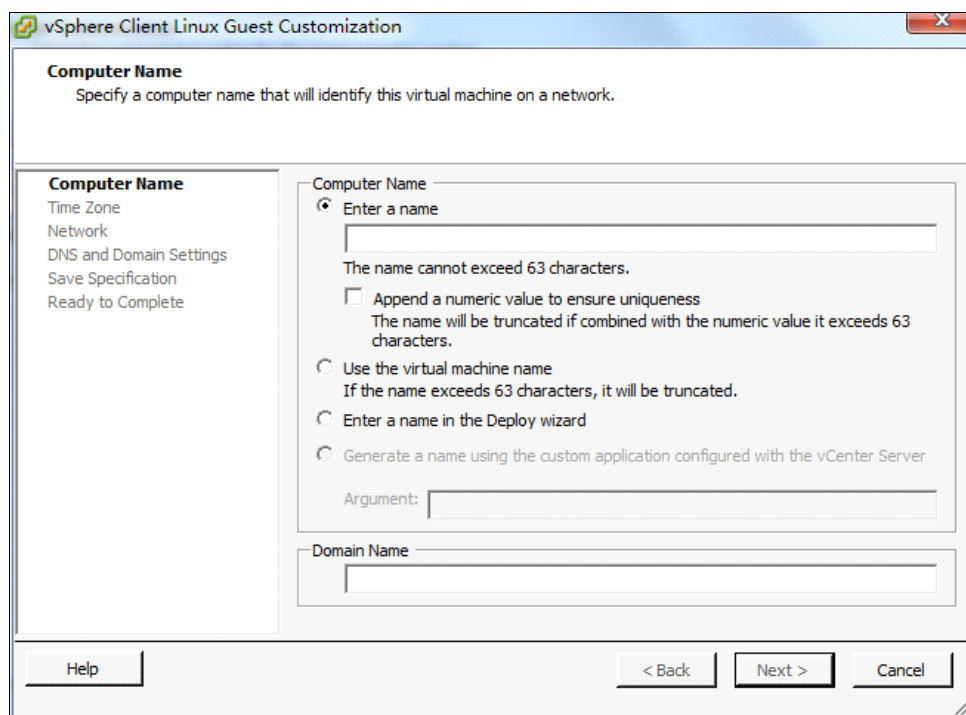


Figure 3-60 Clone Virtual Machine wizard - Computer Name

7. After all of the customization is complete, click **Finish**.

Recently completed tasks are displayed in the Recent Tasks pane, as shown in Figure 3-61. You see the newly migrated virtual machine on the target Flex System Compute Node.

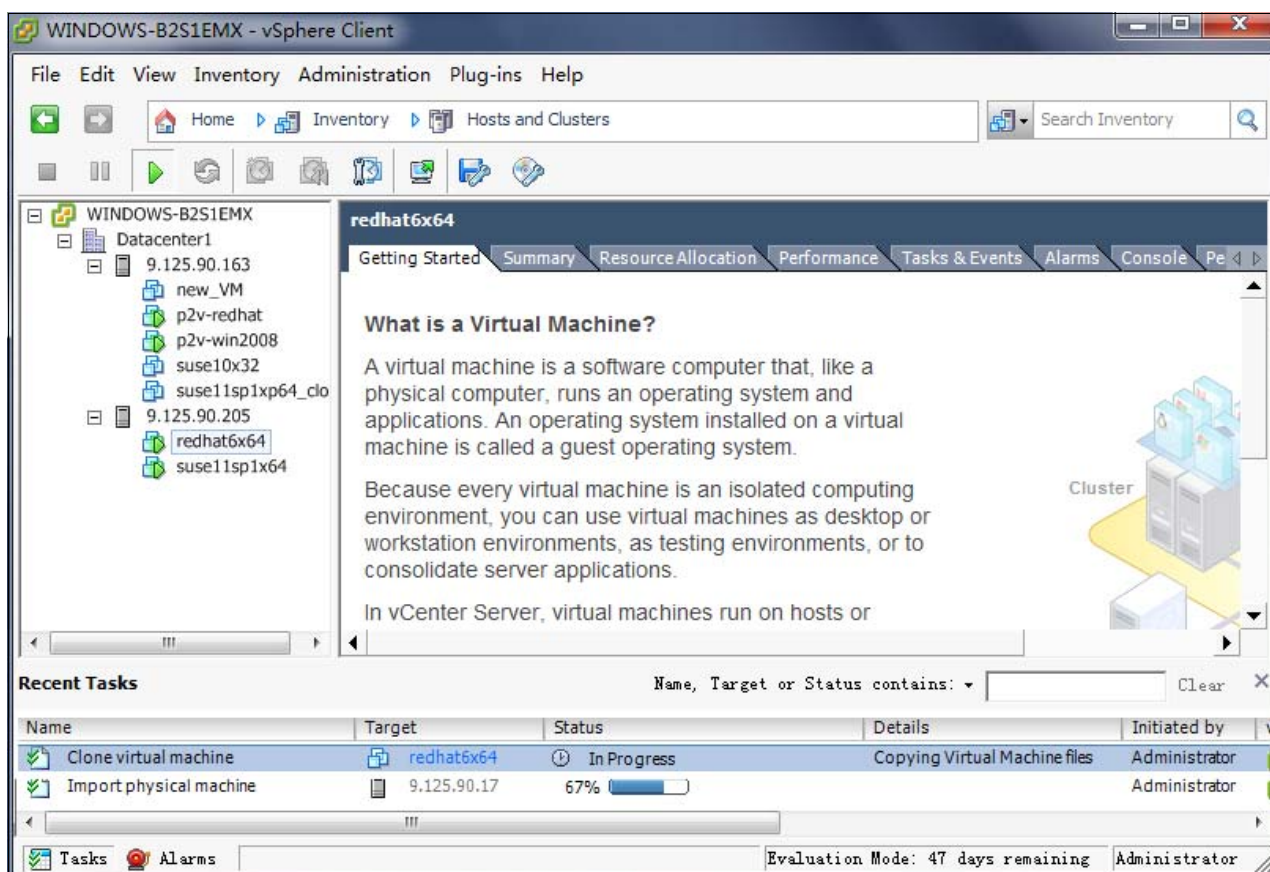


Figure 3-61 V2V migration in progress

Hyper-V scenario

The following requirements must be met before starting the Hyper-V migration process:

- ▶ Hyper-V must be installed on the Flex System compute node (destination server).
- ▶ A System Center Virtual Machine Management (SCVMM) server must be connected to the source server and the destination server.
- ▶ The source server and destination server must be added to the server list database of the SCVMM server.

If any of these requirements are not met, see 3.1.2, “Disconnected Physical-to-Virtual” on page 56 for more information about completing the requirements.

Follow these steps to implement a V2V migration with SCVMM:

1. Log in to the SCVMM console. Right-click the source virtual machine and click **Clone**, as shown in Figure 3-62.

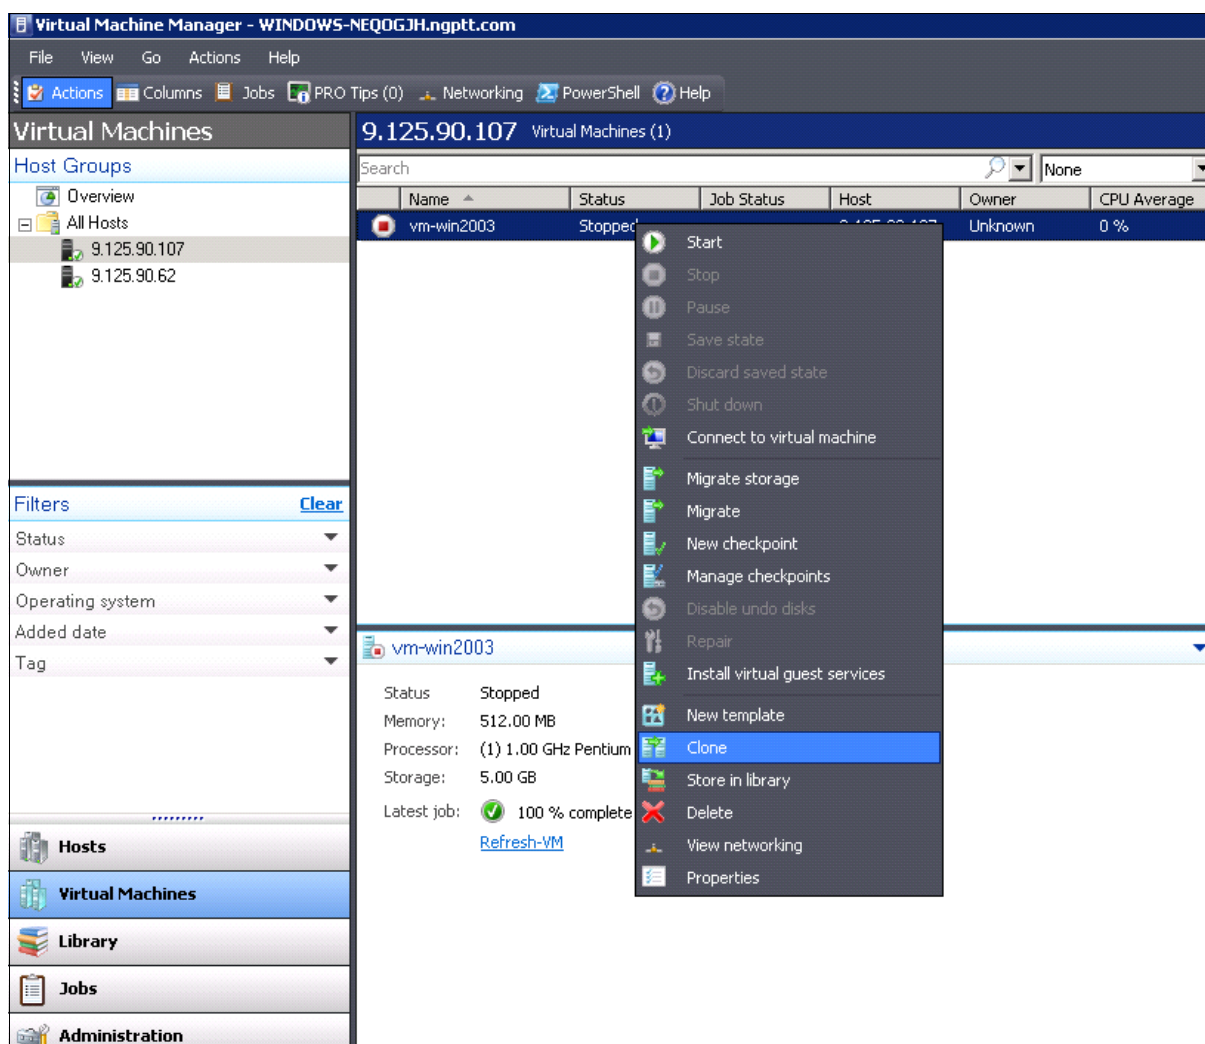


Figure 3-62 Starting the clone operation

2. The New Virtual Machine wizard starts. In the Virtual Machine Identity window, enter the name and domain information for the new virtual machine, as shown in Figure 3-63. Click **Next**.

The screenshot shows the 'New Virtual Machine' wizard with the 'Virtual Machine Identity' step selected in the left-hand navigation pane. The main area contains the following fields and controls:

- Virtual machine name:** A text box containing 'vm-win2003-1'.
- Owner:** A text box containing 'NGPTT\Administrator' and a 'Browse...' button to its right.
- Format:** A label 'Format: domain\username'.
- Description:** A large, empty text box.
- Information:** A blue information icon followed by the text: 'The virtual machine name identifies the virtual machine to VMM. The name does not have to match the computer name of the virtual machine. However, using the same name ensures consistent displays in System Center Operations Manager.'
- Navigation:** 'Next' and 'Cancel' buttons at the bottom right.

Figure 3-63 SCVMM New Virtual Machine: Virtual Machine Identity

3. In the Configure Hardware window, customize the virtual hardware configuration and virtual machine type (virtual machine or virtual appliance), as shown in Figure 3-64. Click **Next**.

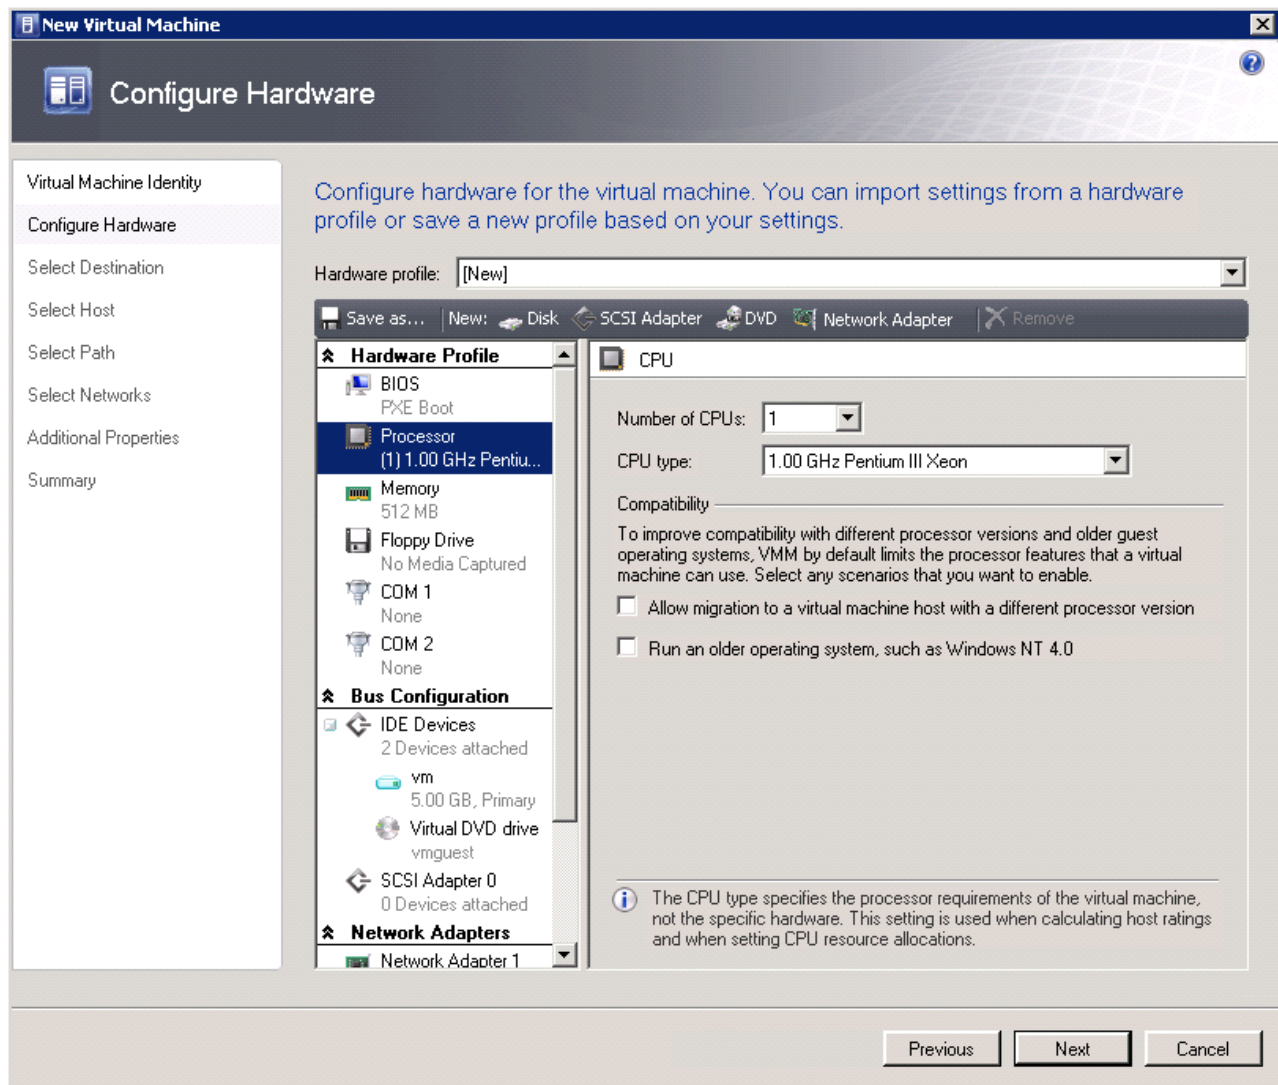


Figure 3-64 SCVMM New Virtual Machine: Configure Hardware

4. In the Select Destination window, choose the destination server and its storage locations for the virtual machine image files, as shown in Figure 3-65. Click **Next**.

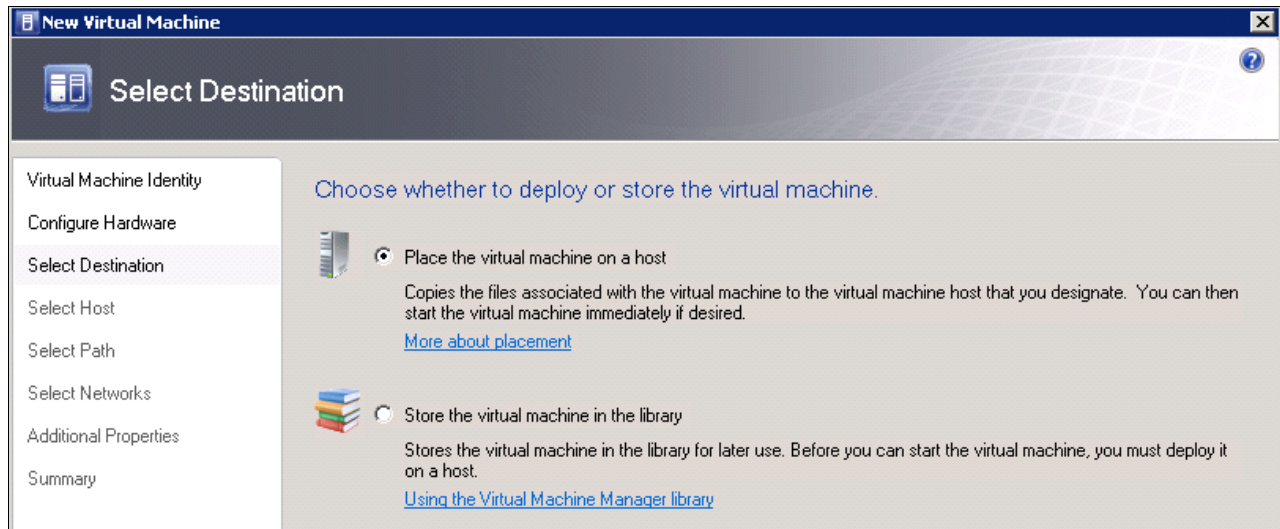


Figure 3-65 SCVMM New Virtual Machine: Select Destination

5. In the Select Host window, select the host, as shown in Figure 3-66. Click **Next**.

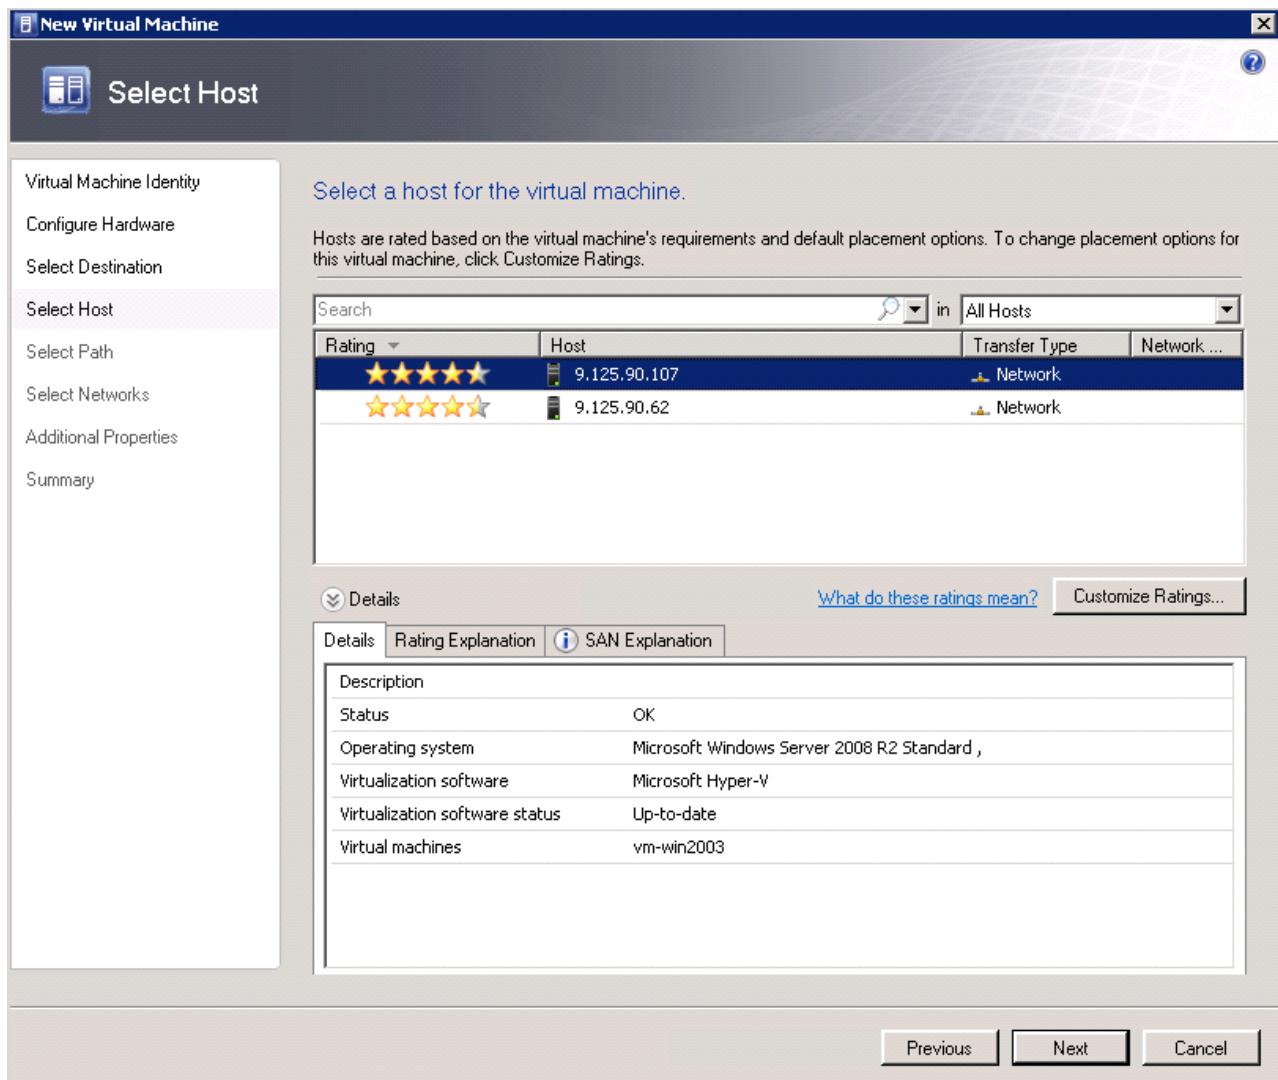


Figure 3-66 SCVMM New Virtual Machine: Select Host

6. In the Select Path window, enter the file path for virtual machine files, as shown in Figure 3-67. Click **Next**.

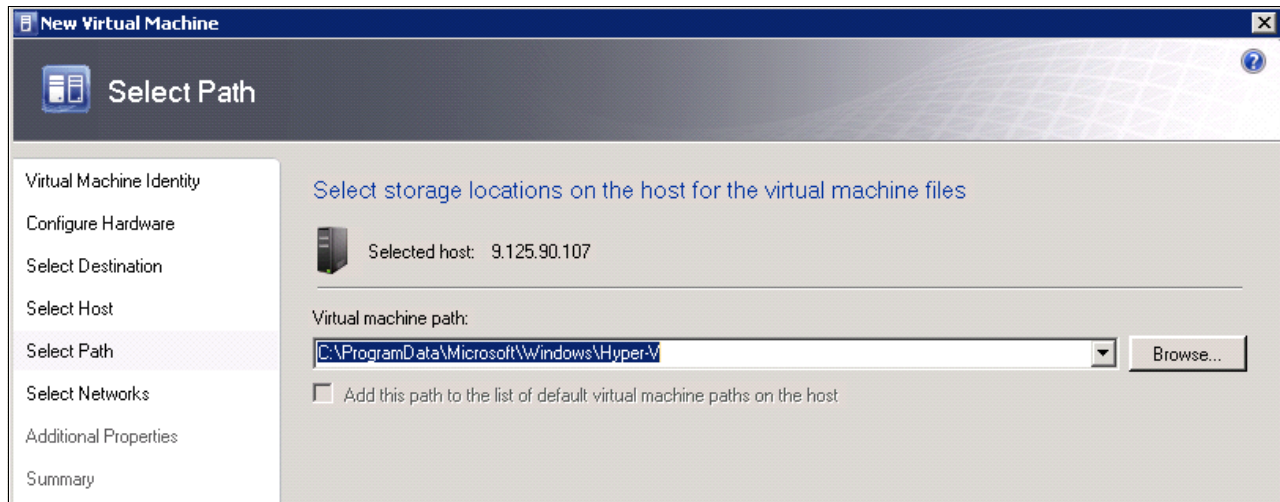


Figure 3-67 SCVMM New Virtual Machine: Select Path

7. In the Select Networks window, define the settings of the network for the new virtual machine, as shown in Figure 3-68. Click **Next**.

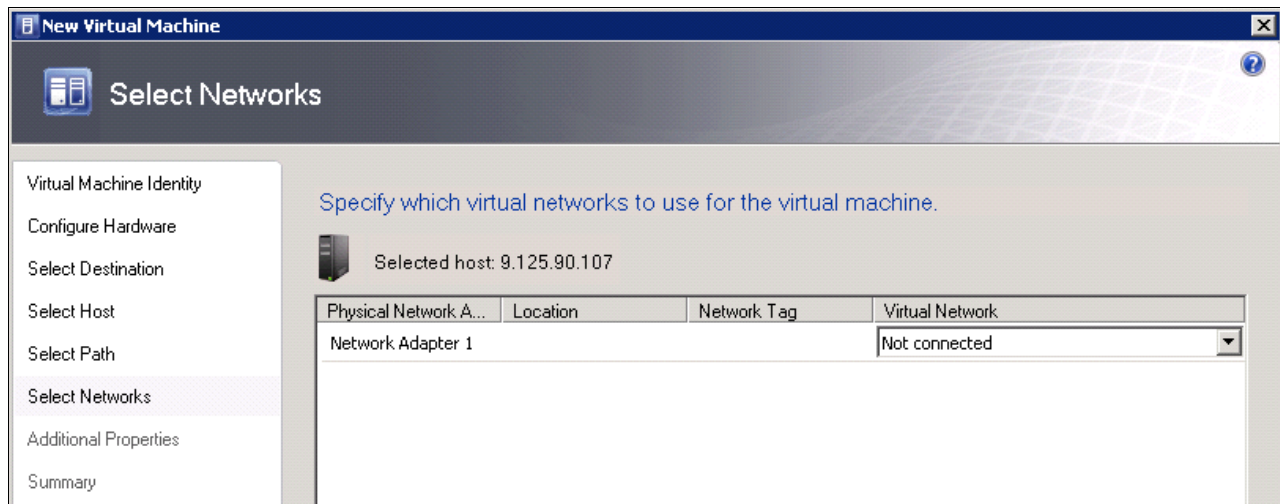


Figure 3-68 SCVMM New Virtual Machine: Select Networks

8. In the Additional Properties window, select any additional properties for the new virtual machine, as shown in Figure 3-69. Click **Next**.

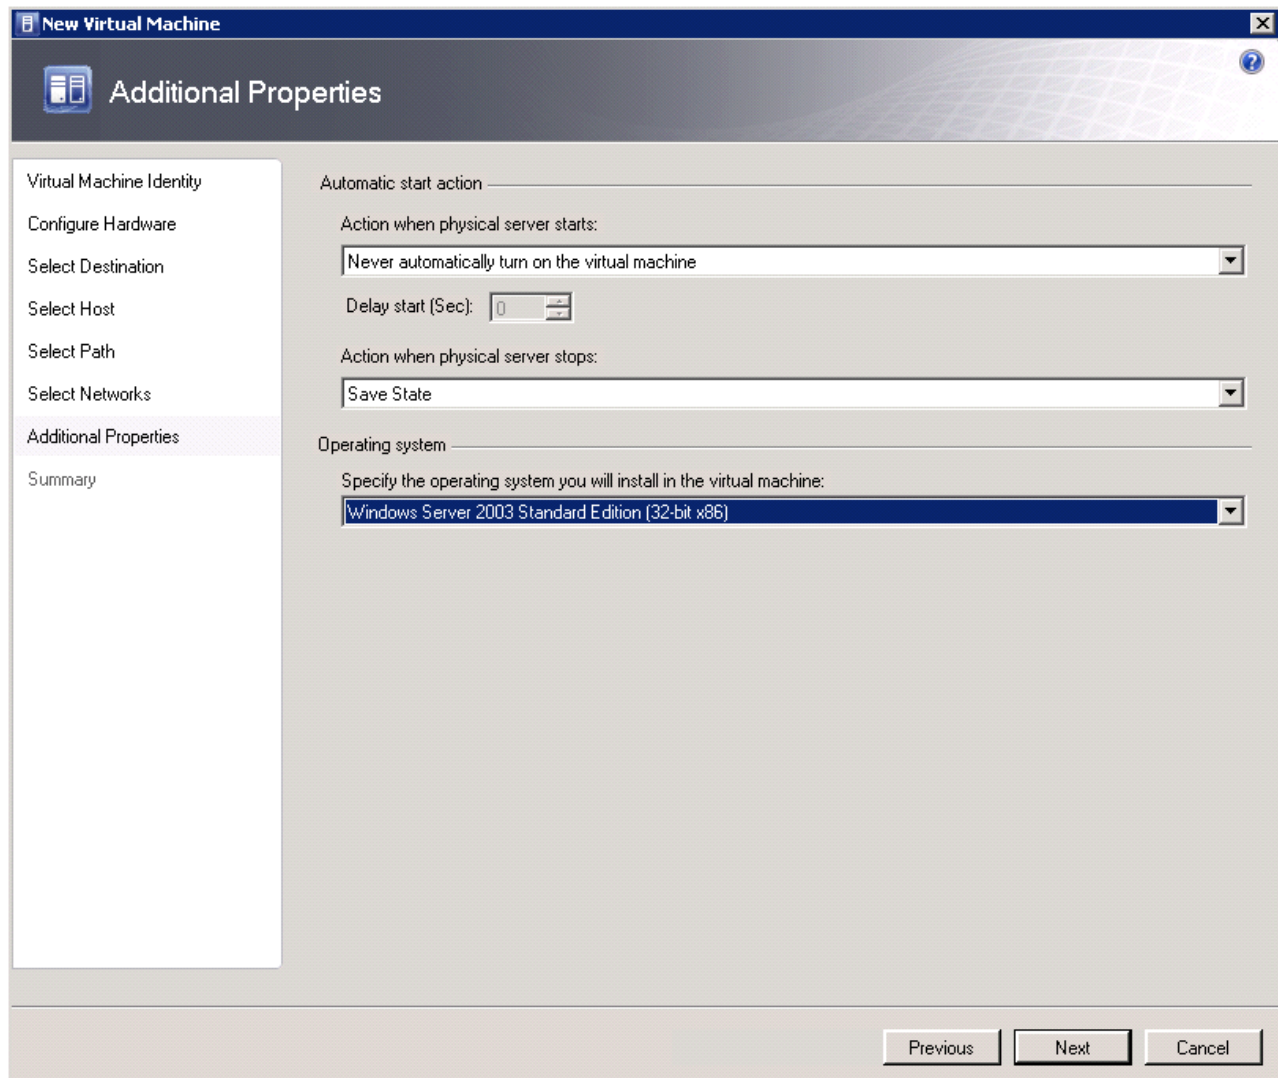


Figure 3-69 SCVMM New Virtual Machine: Additional Properties

9. Review the Summary page, as shown in Figure 3-70. Click **Create** to start the migration.

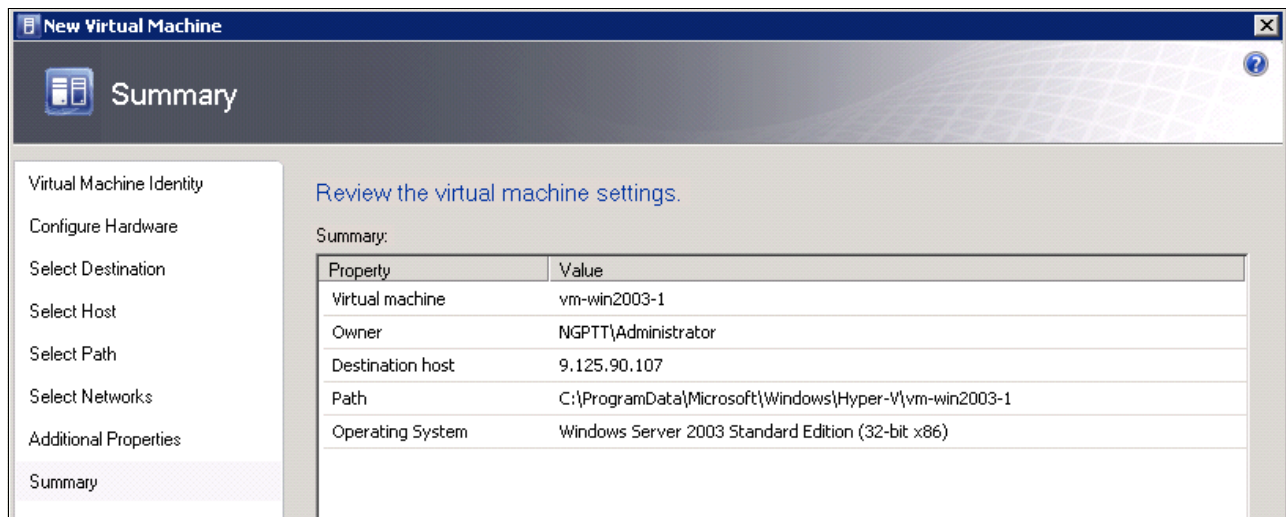


Figure 3-70 SCVMM New Virtual Machine: Summary

The migration job takes time to complete. Upon completion, you see the newly migrated virtual machine on the Flex System Compute Node, as shown in Figure 3-71.

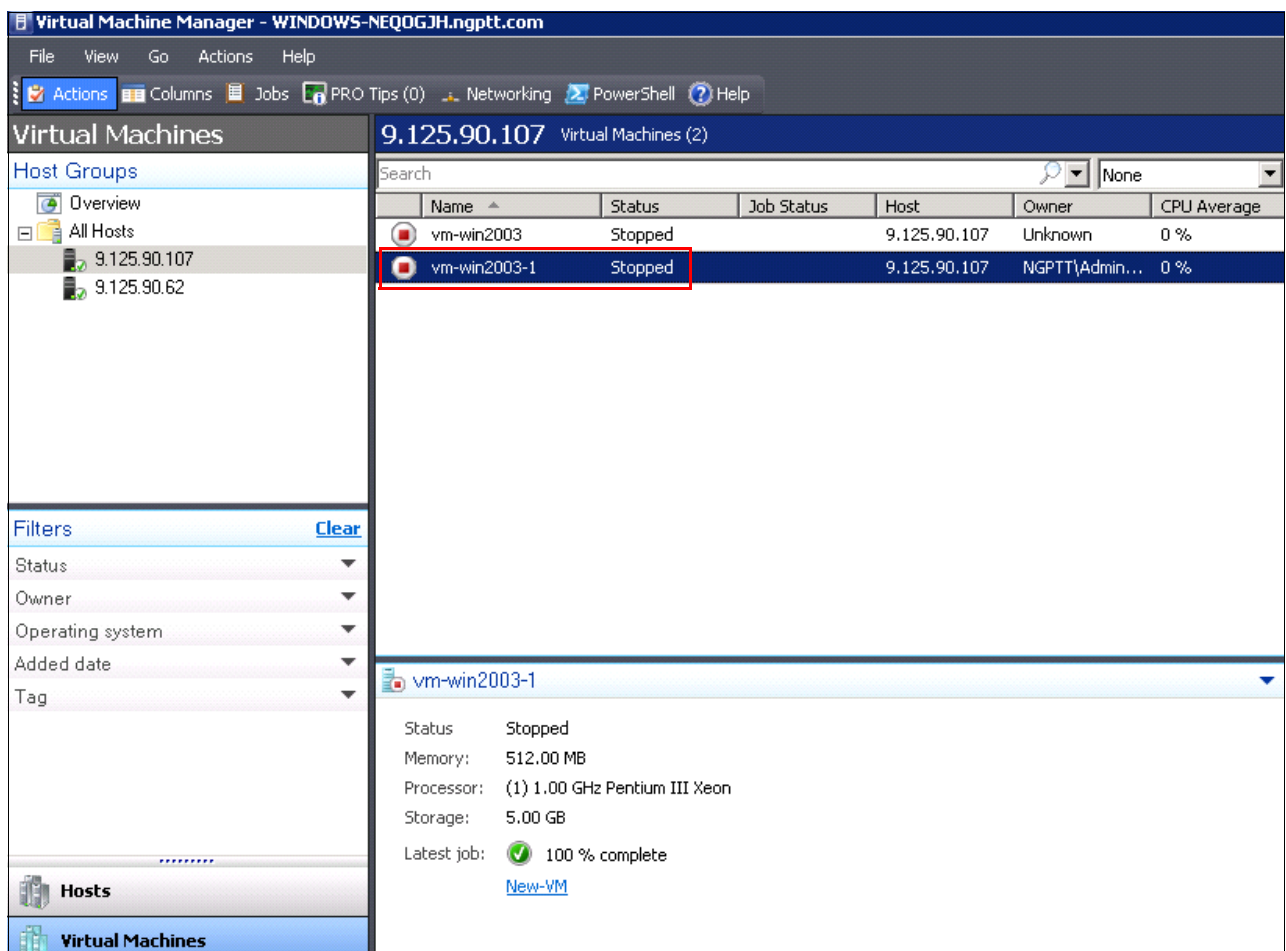


Figure 3-71 V2V migration complete

KVM scenario

Normally, a management server is needed for this migration process. This server captures the virtual machine image from the source server and deploys the image to the destination server. In this scenario, Flex System Manager is the management server for the source server and the destination server. PureFlex System includes Flex System Manager, but Flex System Manager is optional with Flex System. In the absence of an FSM, you can use any KVM management software to implement VM migration.

If you do not have Flex System Manager in your configuration, there are several open source tools that are available, such as QEMU, libvirt, virsh, and virsh-manager. For more information, go to the following website:

<http://www.linux-kvm.org/page/Migration>

The remainder of this section assumes that you are using Flex System Manager.

The following requirements must be met before you start the KVM migration process:

- ▶ Confirm that Flex System Manager VMControl Express is working because the manager is used to import captured virtual appliances.
- ▶ An RHEL x86_64 server with network file system (NFS) must be enabled or a shared SAN storage pool must be set up and configured. This configuration is used to store the virtual machine image. The KVM Platform agent also must be installed on the NFS server.
- ▶ The following image repository conditions must be met:
 - Flex System Manager Common Agent is installed.
 - VMControl Common Repository subagent is installed.
 - The shared NFS exported storage is mounted on the Image Repository server or is connected to the SAN Fibre Channel network, which is used as a repository.
 - The image repository server is discovered and inventory is collected.
 - The image repository is configured within VMControl.
- ▶ One or more RHEL KVM hosts must be set up and available. In this scenario, the user must install RHEL KVM on the Flex System Compute Node.

For more information, see the Flex System Manager product publications at this website:

<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>

Follow these steps to implement a V2V migration process by using KVM:

1. Deploy KVM to the target machines. KVM is a component module of Red Hat Enterprise Linux 6.x. Therefore, the KVM components can be installed during the RHEL installation process.

In addition, the ServerGuide Toolkit can be a useful tool to help you install Red Hat. To download the ServerGuide Scripting Toolkit, see this website:
<http://ibm.com/support/entry/portal/docdisplay?lnidocid=SERV-TOOLKIT>
2. After the KVM host is deployed, download and install the KVM Platform agent on the host. The agent can be downloaded from this website (registration is required):
https://www14.software.ibm.com/webapp/iwm/web/reg/download.do?source=dmp&S_PKG=dir_63_x86_MDagents&lang=en_US&cp=UTF-8
3. Set up an NFS and image repository server to meet the prerequisites that were listed previously in this section. After these servers are in place, you discover, authenticate, and inventory all servers, including the NFS Server, image repository server, source server, and target server.

4. Start a web browser and enter the following Flex System Manager URL:

`http://System_Name:Port_Number/ibm/console`

System_Name is the name of the system on which the management node is installed, and Port_Number is the port that is specified for the web server to use. The default ports for the web server are 8421 and 8422. If port 8422 is used, specify `https` to indicate a secure port, which is required by default. The Flex System Manager login window opens, as shown in Figure 3-72.



Figure 3-72 Flex System Manager login window

5. Enter the User ID and Password and click **Log in**. The main Flex System Manager window opens, as shown in Figure 3-73.

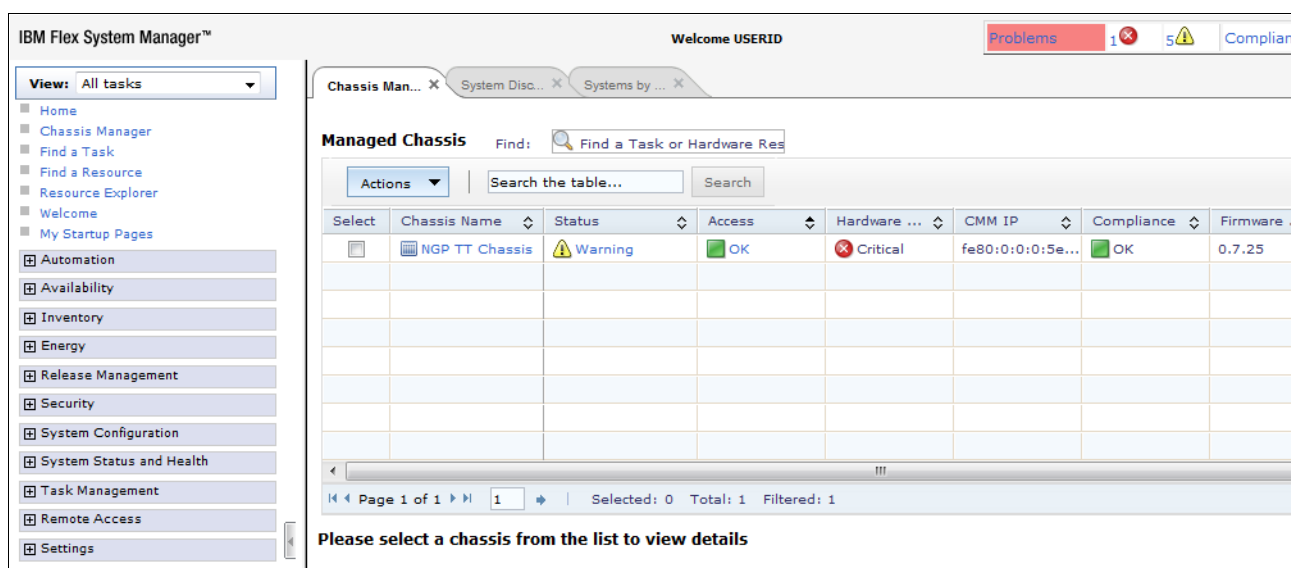


Figure 3-73 Flex System Manager main window

- From the left navigation pane, click **Inventory** → **System Discovery**. Enter the IP address of the server that you want to add to the management list, as shown in Figure 3-74. Add the target servers as needed.

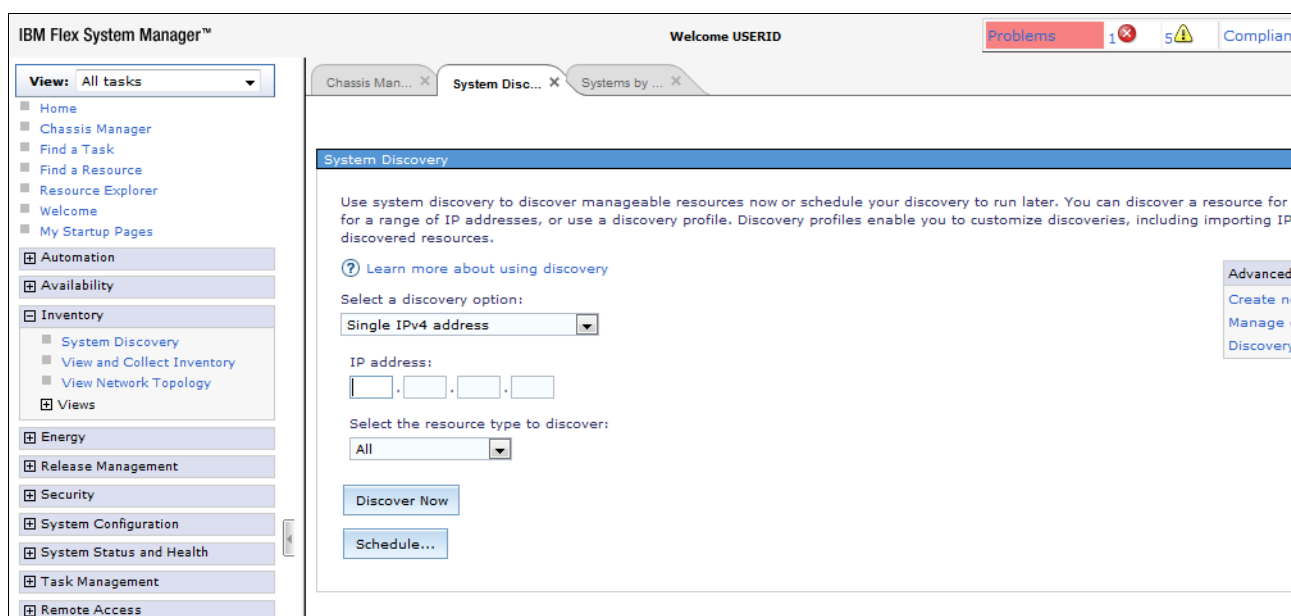


Figure 3-74 Flex System Manager: System Discovery

- When prompted, enter the User ID and Password to gain full access to the source and target server, as shown in Figure 3-75.




Figure 3-75 Request access to remote server

Important: If the source servers are managed by Flex System Manager (FSM), use the FSM page to capture the virtual images from the source servers. Although virtual images can be captured by means other than FSM, the images must be imported to the target servers by using VMControl on FSM.

The remaining steps apply only to FSM users.

8. In the left navigation pane, click **System Configuration** → **IBM System Director VMControl**, as shown in Figure 3-76. In the main part of the window, click **IBM Systems Director VMControl** to start the VMControl plug-in.

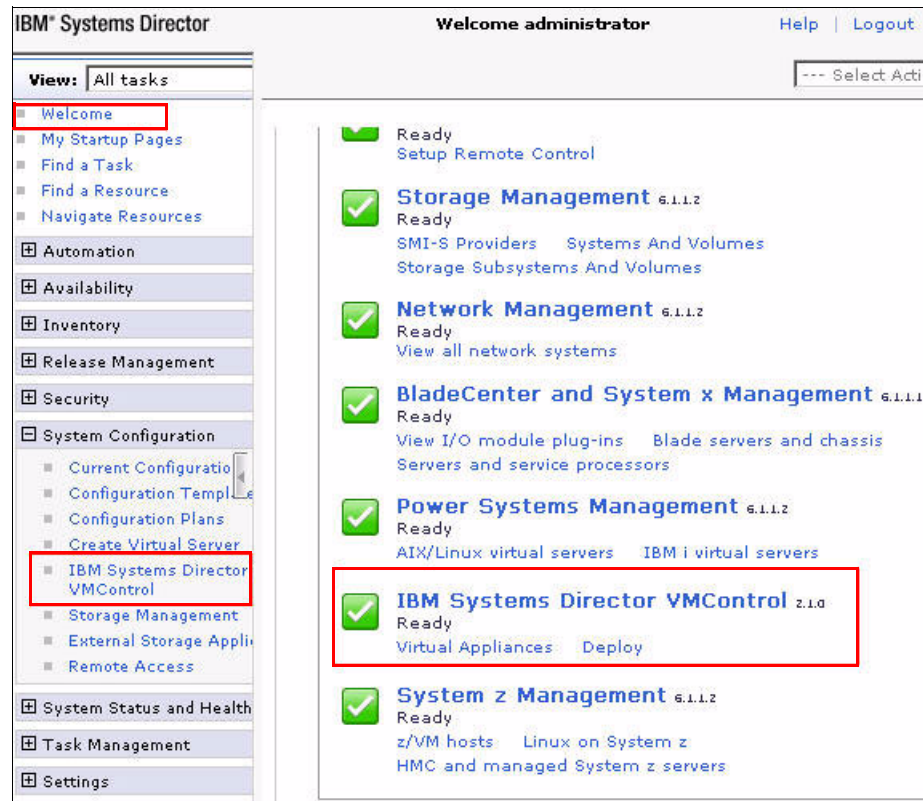


Figure 3-76 Starting the VMControl plug-in

- In the VMControl tab, select the **Virtual Appliances** tab to view virtual machine-related information, which is shown in Figure 3-77.

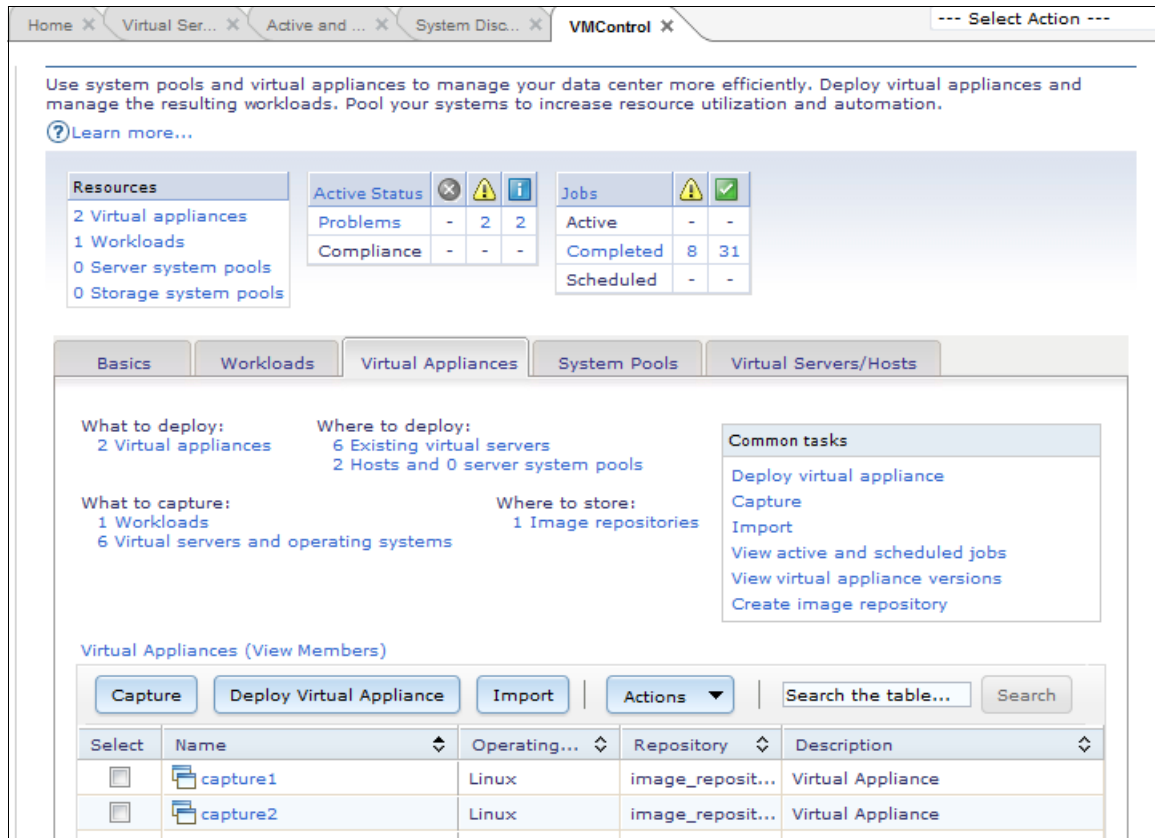


Figure 3-77 VMControl: Virtual Appliances tab

- Click the **Virtual Servers and Hosts** tab. The source and target servers that you discovered are listed, as shown in Figure 3-78.

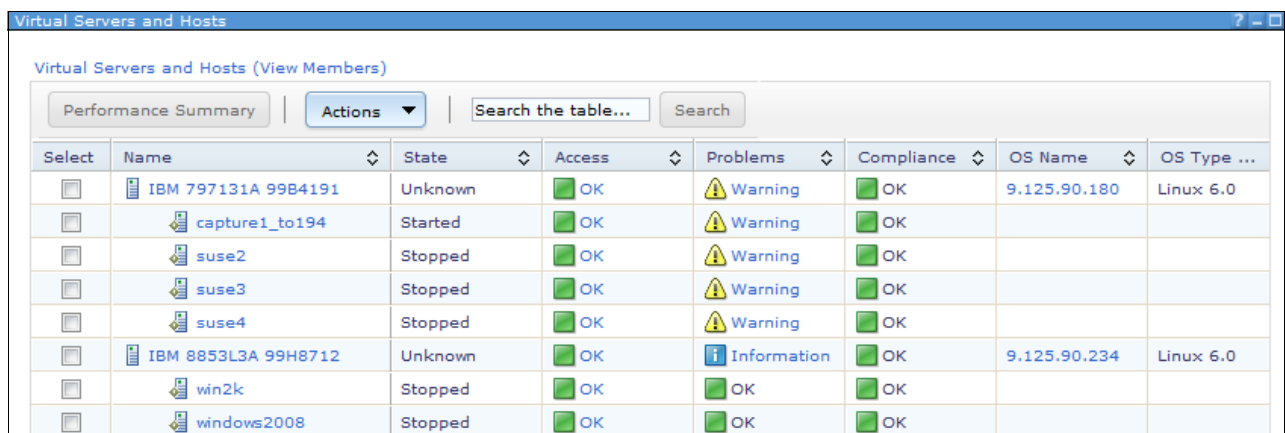


Figure 3-78 Virtual Servers and Hosts window

11. On the **Virtual Appliance** tab (see Figure 3-77 on page 124), click **Capture** in the Common tasks box.
12. The VMControl Capture wizard starts and opens the Welcome page, as shown in Figure 3-79. Click **Next**.

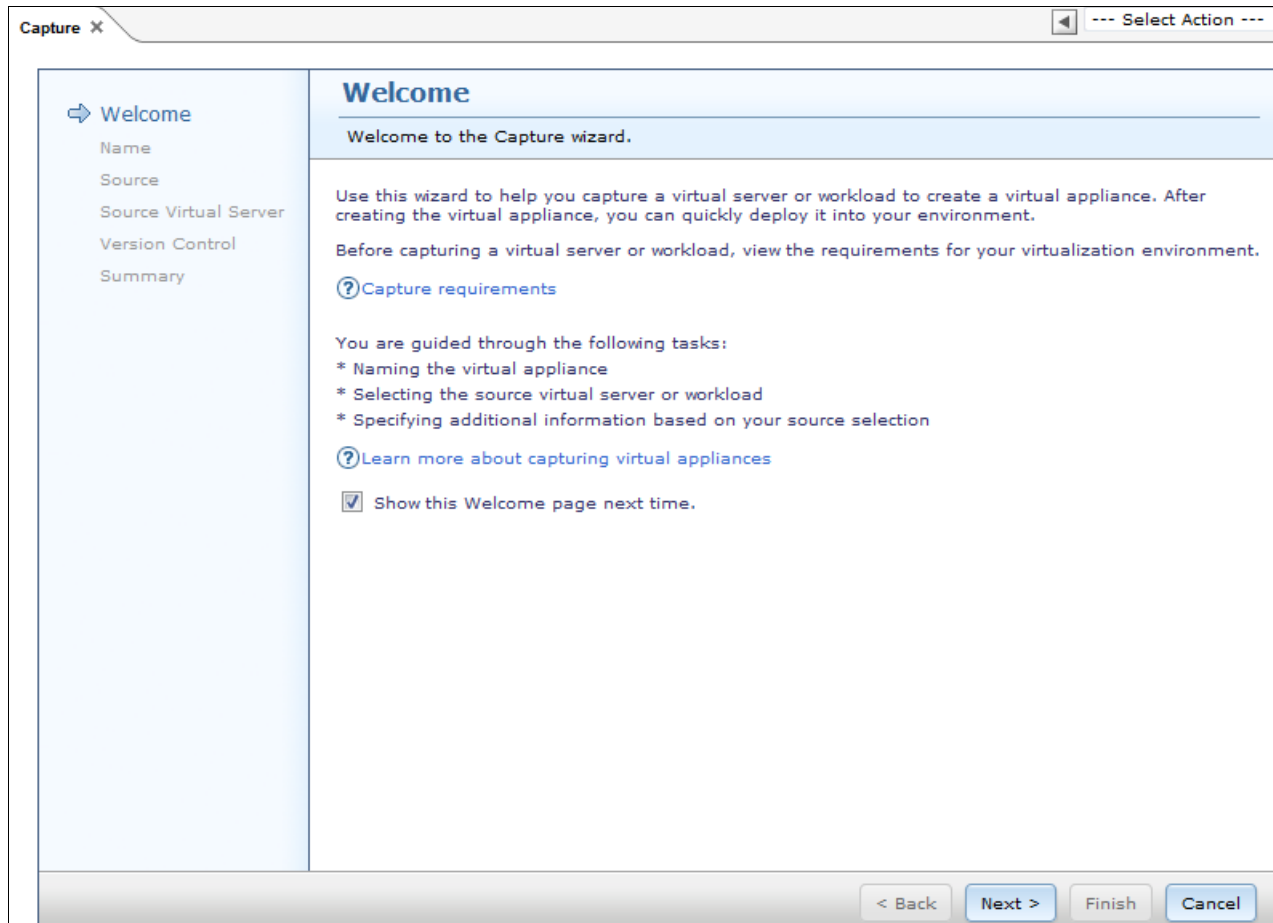
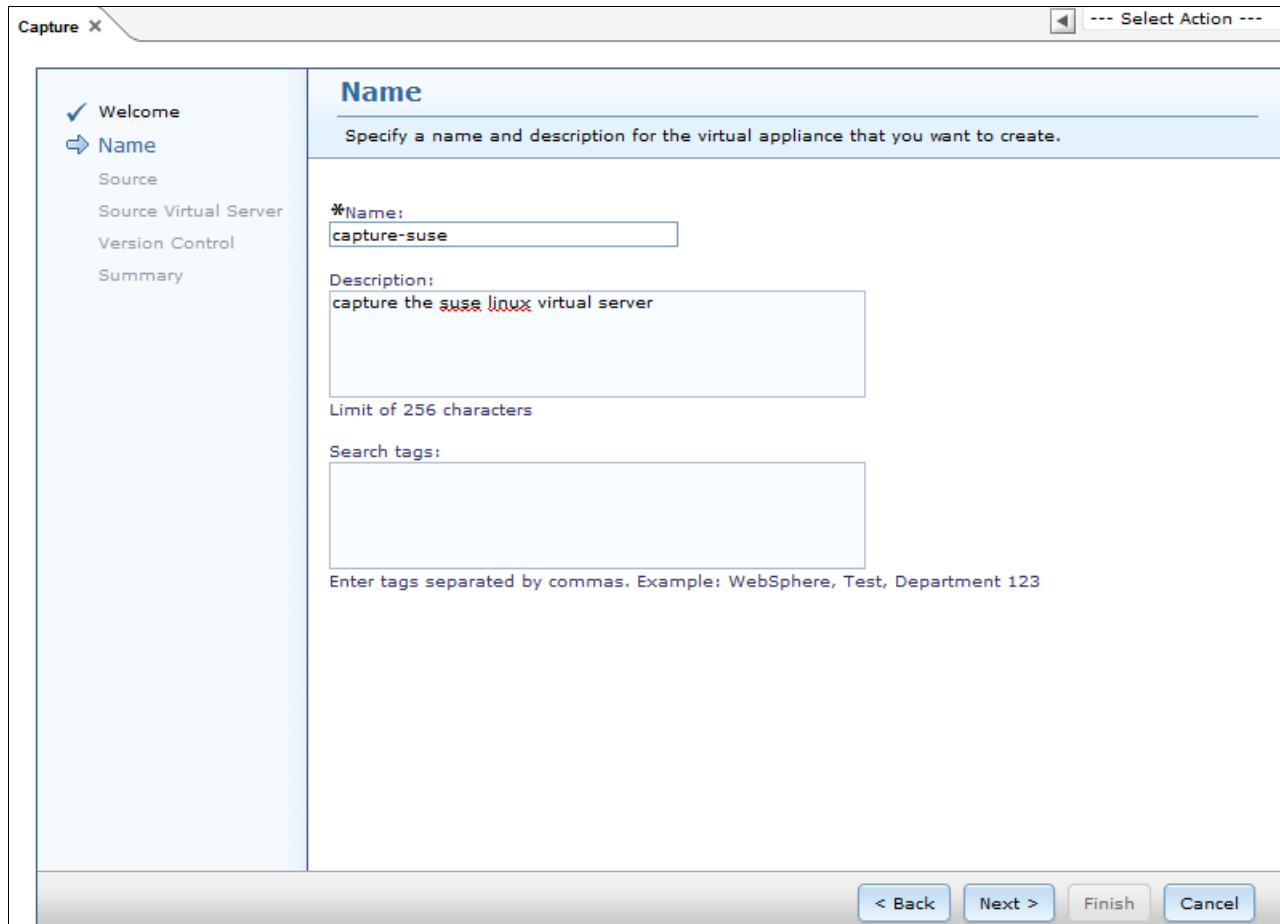


Figure 3-79 VMControl Capture wizard: Welcome

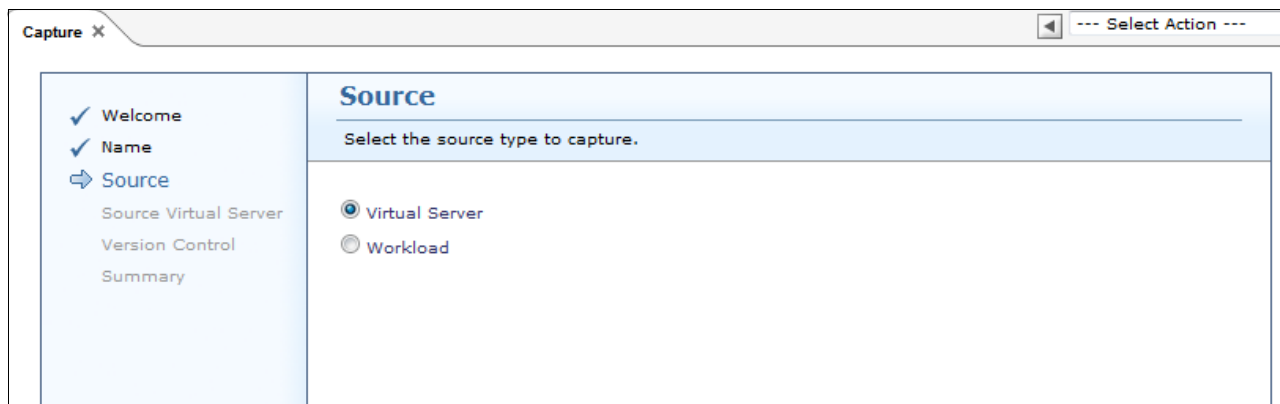
13. Enter the Name and Description for the virtual appliance (target virtual machine image package), as shown in Figure 3-80. Click **Next**.



The screenshot shows the 'Name' step of the VMControl Capture wizard. The left sidebar contains a list of steps: 'Welcome' (checked), 'Name' (active), 'Source', 'Source Virtual Server', 'Version Control', and 'Summary'. The main area is titled 'Name' and contains the instruction 'Specify a name and description for the virtual appliance that you want to create.' Below this, there are three input fields: '*Name:' with the value 'capture-suse', 'Description:' with the value 'capture the suse linux virtual server', and 'Search tags:' which is empty. A note below the search tags field says 'Enter tags separated by commas. Example: WebSphere, Test, Department 123'. At the bottom right, there are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 3-80 VMControl Capture wizard: Name

14. In the Source window, select **Virtual Server** as the source type, as shown in Figure 3-81.



The screenshot shows the 'Source' step of the VMControl Capture wizard. The left sidebar contains a list of steps: 'Welcome' (checked), 'Name' (checked), 'Source' (active), 'Source Virtual Server', 'Version Control', and 'Summary'. The main area is titled 'Source' and contains the instruction 'Select the source type to capture.' Below this, there are two radio button options: 'Virtual Server' (selected) and 'Workload'.

Figure 3-81 VMControl Capture wizard: Source

15. In the Source Virtual Server window, choose the migration source server, as shown in Figure 3-82.

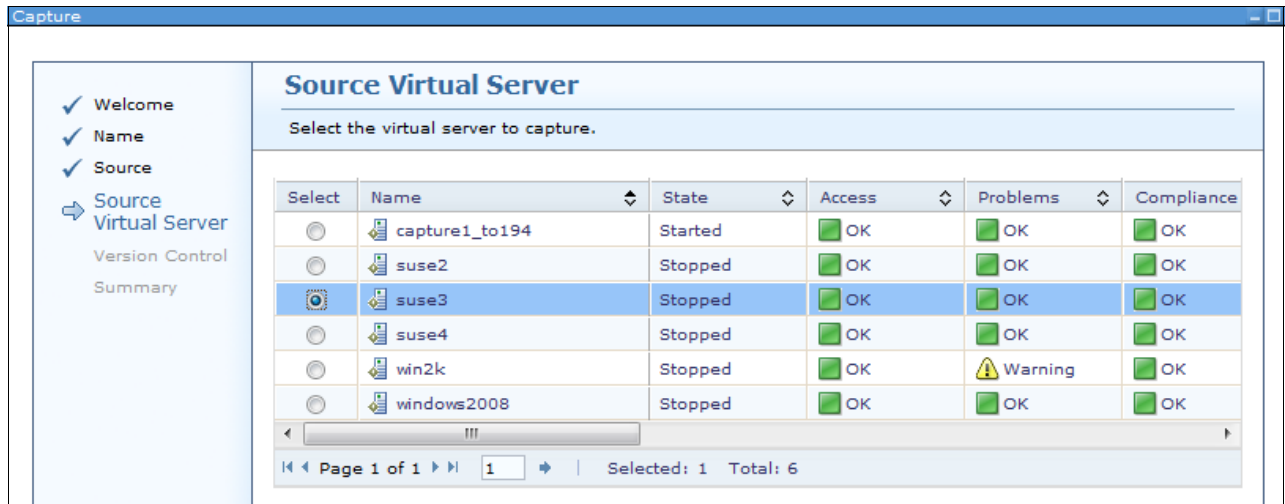


Figure 3-82 VMControl Capture wizard: Source Virtual Server

16. In the Disks window, specify the virtual appliance disk settings, as shown in Figure 3-83. Click **Next**.

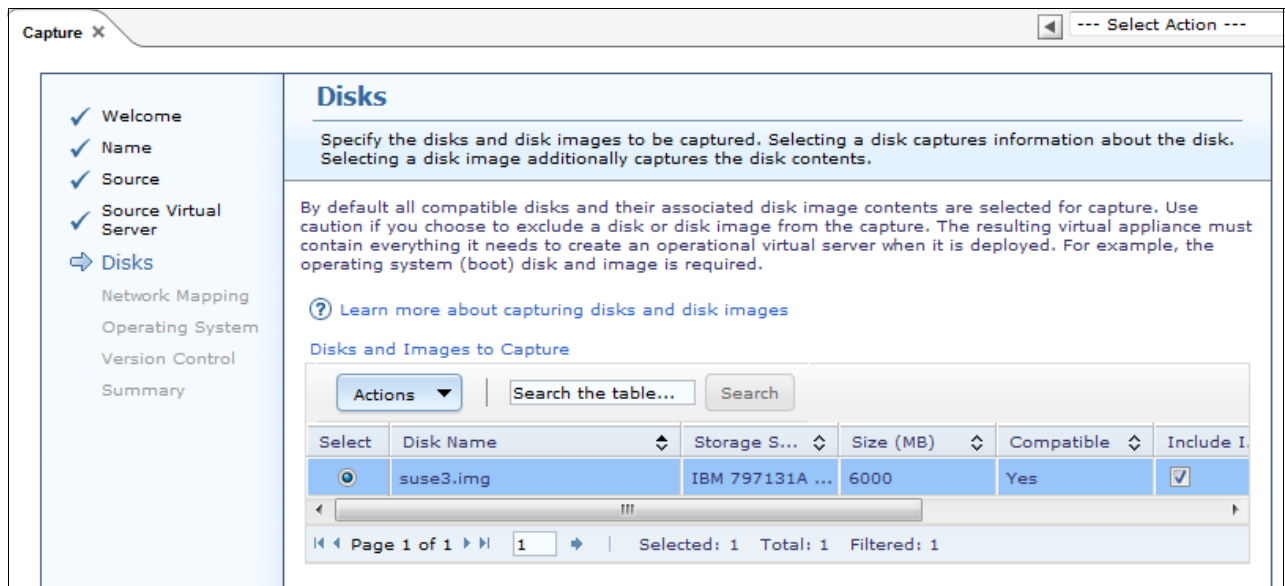


Figure 3-83 VMControl Capture wizard: Disks

17. In the Network Mapping window, specify the virtual appliance network mapping settings, as shown in Figure 3-84. Click **Next**.

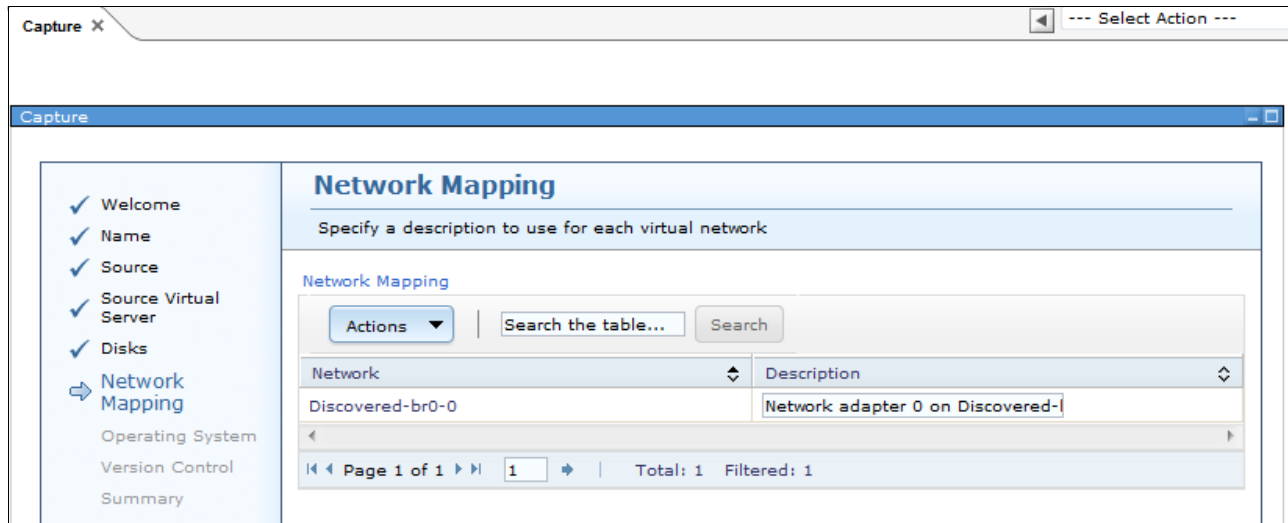


Figure 3-84 VMControl Capture wizard: Network Mapping

18. As shown in Figure 3-85, specify the OS in the Operating System window. Click **Next**.

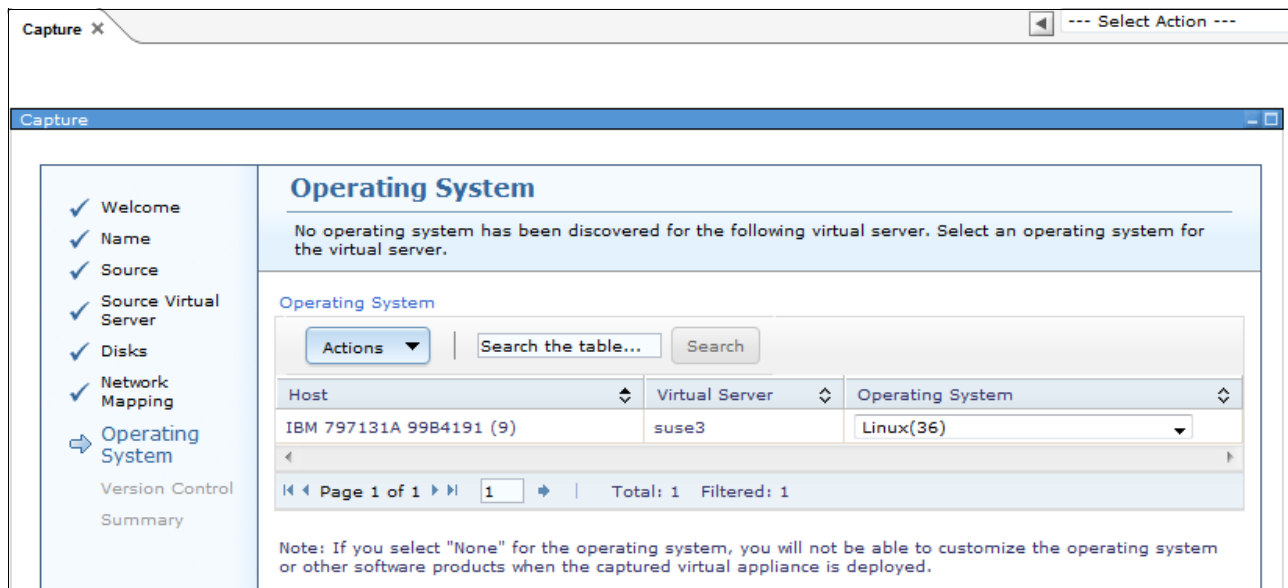


Figure 3-85 VMControl Capture wizard: Operating System

19. In the Version Control window, confirm the version control information for the new virtual appliance, as shown in Figure 3-86. You also can choose **Create a new version tree with the new virtual appliance as the root** if you want to establish a new group with the new virtual appliance. If you want to use the existing virtual appliance and assign that appliance as a parent version, select **Select a virtual appliance to be the parent version of the new virtual appliance**. Click **Next**.

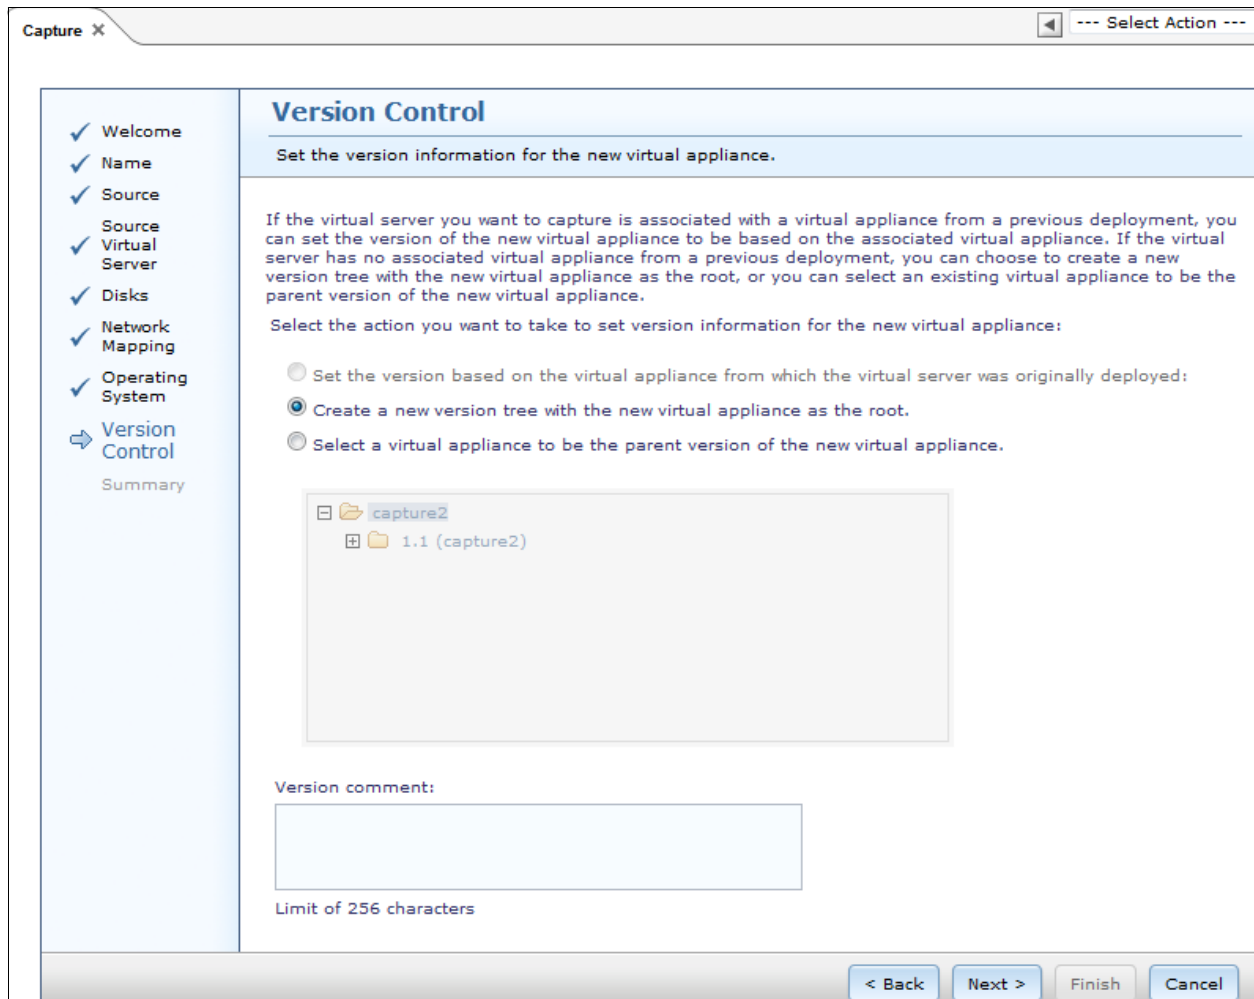


Figure 3-86 VMControl Capture wizard: Version Control

20. In the Summary page, review the summary information of the virtual appliance, as shown in Figure 3-87. Click **Finish**.

Capture X --- Select Action ---

Capture

Summary

You are now ready to capture the virtual server or workload to create a virtual appliance.

Virtual appliance or workload details:

Virtual appliance name:	capture-suse
Virtual appliance description:	
Source server:	suse3
Repository:	image_repository
Disks:	
Disk Name	suse3.img
Storage Server	IBM 797131A 99B4191
Size (MB)	6000
Compatible	Yes
Include Image	Yes
Disk Description	

Note: The virtualization manager will provide access to the target server so that it can be captured.

Click Finish to capture the virtual server or workload and create a virtual appliance. Once completed, you can deploy the virtual appliance into your environment.

< Back Next > Finish Cancel

Figure 3-87 VMControl Capture wizard: Summary

21. You are prompted to run the job immediately or schedule the run for a later time, as shown in Figure 3-88. Click **Run Now**.

The screenshot shows a dialog box titled "Launch Job" with three tabs: "Schedule", "Notification", and "Options". The "Schedule" tab is active. It contains a section titled "Job name and schedule" with a text input field for the job name. The job name is "Capture virtual appliance - January 12, 2012 12:57:08 AM GMT+08:00". Below the text field is a label "Choose when to run the job." and two radio buttons: "Run Now" (which is selected) and "Schedule". At the bottom of the dialog box are three buttons: "OK", "Cancel", and "Help".

Figure 3-88 VMControl Capture wizard: Schedule

After the job completes, you can see that another appliance is available in the Virtual Appliances tab, as shown in Figure 3-89. The virtual appliances that are listed in Figure 3-89 can be captured by FSM or imported by the user. By using either method, the user can deploy the virtual appliances into target KVM servers, as described in the next steps.

The screenshot shows the 'Virtual Appliances' tab in the VMControl interface. At the top, there are tabs for 'Basics', 'Workloads', 'Virtual Appliances', 'System Pools', and 'Virtual Servers/Hosts'. Below the tabs, the interface is divided into several sections:

- What to deploy:** 1 Virtual appliances
- Where to deploy:** 6 Existing virtual servers, 2 Hosts and 0 server system pools
- What to capture:** 1 Workloads, 6 Virtual servers and operating systems
- Where to store:** 1 Image repositories
- Common tasks:**
 - Deploy virtual appliance
 - Capture
 - Import
 - View active and scheduled jobs
 - View virtual appliance versions
 - Create image repository

Below these sections, there is a header for 'Virtual Appliances (View Members)' with buttons for 'Capture', 'Deploy Virtual Appliance', 'Import', and 'Actions'. A search bar is also present.

Select	Name	Operating Syst...	Repository	Description
<input type="checkbox"/>	capture2	Linux	image_repository	Virtual Appliance
<input type="checkbox"/>	capture	Linux	image_repository	Virtual Appliance
<input type="checkbox"/>				
<input type="checkbox"/>				
<input type="checkbox"/>				

At the bottom, there is a pagination bar showing 'Page 1 of 1', a selected count of 1, and a total count of 1.

Figure 3-89 VMControl Capture wizard: Capture complete

22. The captured virtual appliance now must be deployed by using VMControl on Flex System Manager. Click **Deploy Virtual Appliance**. A new wizard starts with a Welcome page, as shown in Figure 3-90. Click **Next**.

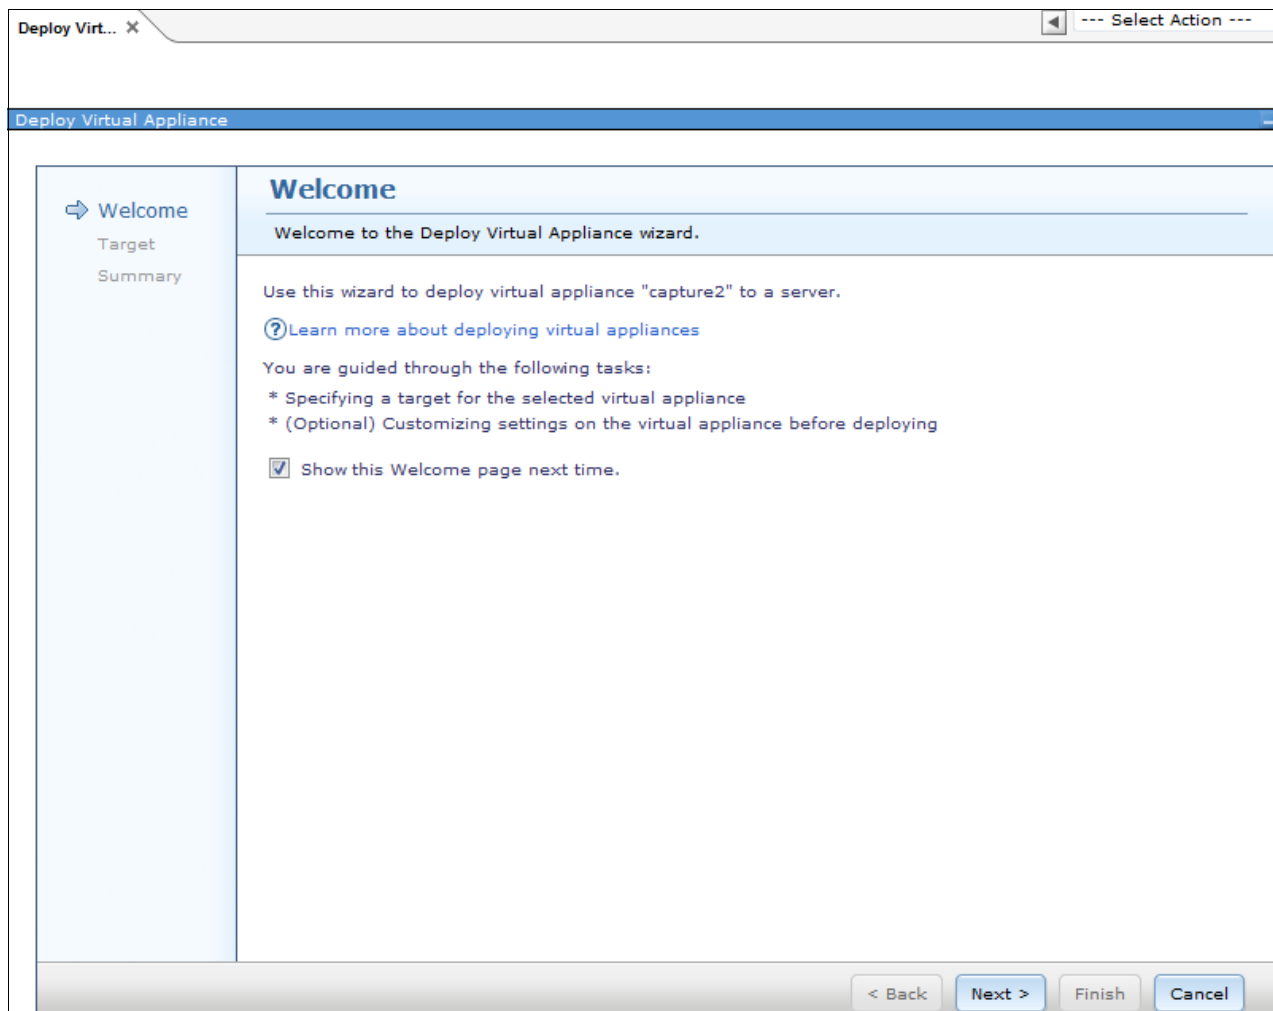


Figure 3-90 Deploy Virtual Appliance wizard: Welcome

23. In the Target window, choose **Deploy to a new virtual server on the following**, as shown in Figure 3-91. Select the target system in the pool list. Click **Next**.

Target

Select the location where you want to deploy the virtual appliance.

You can deploy the virtual appliance to create a new virtual server on an existing host system or system pool. Or, you can deploy the virtual appliance to an existing virtual server.

☒ Deploy to a new virtual server on the following:

Select	Name	State	IP Address...	Installed ...	Description
<input type="radio"/>	IBM 8853L3A 99H8712	Unknown	9.125.90.234,...	9.125.90.234	
<input checked="" type="radio"/>	IBM 797131A 99B4191	Unknown	9.125.90.194,...	9.125.90.194	

Page 1 of 1 | 1 | Selected: 1 Total: 2 Filtered: 2

☐ Deploy to an existing virtual server:

Select	Name	State	IP Address...	Installed ...	Description
<input type="radio"/>	suse3	Stopped			Virtual S
<input type="radio"/>	suse4	Stopped			Virtual S
<input type="radio"/>	capture1_to194	Started			Virtual S

Page 1 of 1 | 1 | Selected: 0 Total: 3 Filtered: 3

Note: When deploying to a server system pool, the server system pool must identify the host where the virtual appliance will be deployed. This process might take a few minutes to complete.

< Back Next > Finish Cancel

Figure 3-91 Deploy Virtual Appliance wizard: Target

24. In the Name window, enter the name for the workload (virtual machine), as shown in Figure 3-92. Click **Next**.

Name

Specify a name for the virtual server that you want to deploy.

*Type the name of the virtual server that you want to create.

capture2_vs86048

< Back Next >

Figure 3-92 Deploy Virtual Appliance wizard: Name

25. In the Storage Mapping window, assign the storage pool for the virtual disks when the virtual appliance is deployed, as shown in Figure 3-93. A storage pool must be defined for the target host before running the deployment operation.

Storage Mapping

Specify how to assign the storage for the virtual disks when you deploy the virtual appliance.

Ensure each disk in the table is assigned to either a storage volume or storage pool. To assign a disk to a storage volume, select a single disk. You can select multiple disks to assign to a storage pool. If one or more available storage locations were found, then a suggested storage pool has been assigned by default for each disk. If the default assignment(s) are adequate, you can just click Next to continue with the wizard.

[Learn more about storage mapping for deploying to a new virtual server](#)

Storage Mapping

Assign to Storage Volume... Assign to Storage Pool... Actions Search the table... Search

Select	Disk Required by Virtual ...	Assigned Storage	Size (MB)	Image
<input checked="" type="radio"/>	disk1	images (Local-Shared Storage Pool)	6,000	True

Page 1 of 1 1 Selected: 1 Total: 1 Filtered: 1

< Back Next >

Figure 3-93 Deploy Virtual Appliance wizard: Storage Mapping

26. In the Network Mapping window, select the network configuration for the new virtual server, as shown in Figure 3-94.

Network Mapping

Select a virtual network for each network defined for the appliance.

The following networks will be assigned for this virtual server.

Network Mapping

Actions Search the table... Search

VA Network Name	Description	Assigned Virtual LAN on Host	MAC Address
Network adapter 0 on Discovered-...	Network adapter ...	Discovered-br0-0	Auto set

Page 1 of 1 1 Total: 1 Filtered: 1

< Back Next >

Figure 3-94 Deploy Virtual Appliance wizard: Network Mapping

27. In the Product window, enter the OS configuration for the new virtual server, as shown in Figure 3-95. Click **Next**.

Deploy Virtual Appliance

Product
Specify the product settings you want to use when you deploy the virtual appliance.

General System Product Section
Time zone setting for the virtual system:

System Level Networking
Short host name for the system:
DNS domain name for the system:
IP addresses of DNS servers for system:
Default IPv4 gateway:

Network adapter configuration for Network adapter 0 on Discovered-br0-0

Internet Protocol Version 4
Static IP address for the network adapter "Network adapter 0 on Discovered-br0-0":
Static network mask for network adapter "Network adapter 0 on Discovered-br0-0":
Use DHCP for network adapter "Network adapter 0 on Discovered-br0-0":

Internet Protocol Version 6
Static IP address for the network adapter "Network adapter 0 on Discovered-br0-0":
Static default gateway for network adapter "Network adapter 0 on Discovered-br0-0":
Use IPv6 stateless address autoconfiguration for network adapter "Network adapter 0 on Discovered-br0-0":

Deployment use
The adapter order for network adapter "Network adapter 0 on Discovered-br0-0":
The MAC address for network adapter "Network adapter 0 on Discovered-br0-0":

Remove ISO Product Section
Remove ISO after customization (requires shutdown):

Figure 3-95 Deploy Virtual Appliance wizard: Product

28. In the Summary window, review the summary information of the deployment, as shown in Figure 3-96. Click **Finish**. You see the newly created job.

Deploy Virtual Appliance

Summary
You are now ready to deploy the virtual appliance.

Deployment details:

Virtual appliance to deploy:	capture2
Target server or system pool:	IBM 797131A 99B4191
Workload Name	deploy-workload
Name:	capture2_vs86048
Storage Mapping:	
Disk Required by Virtual Appliance	disk1
Assigned Storage	images (Local-Shared Storage Pool)
Size (MB)	6000
Image	Yes

Click Finish to deploy the virtual appliance.

Figure 3-96 Deploy Virtual Appliance wizard: Summary

29. Click **Run Now** to start the job. Click **OK**.

30. After the job is complete, you see that the virtual server was migrated to the KVM-based target machine.

3.3 Conclusion

Migration is a complex topic, the details of which cannot be covered in a single paper.

Users must understand the real migration environment and determine the specific solution. The examples in this paper are a good start to understanding more about the overall migration process. By following the guidelines in this document with a thorough understanding of the IT environment, you can conduct a successful migration.

Abbreviations and acronyms

ACPI	Advanced control and power interface	RAID	Redundant array of independent disks
AMM	Advanced Management Module	RAM	Random access memory
ATS	Advanced Technical Support	RDP	Remote Desktop Protocol
BIOS	Basic input output system	RHEL	Red Hat Enterprise Linux
CD	Compact disk	RSS	Receive-side scaling
CIM	Common Information Model	SAN	Storage area network
CMM	Chassis Management Module	SAS	Serial Attached SCSI
CPU	Central processing unit	SCVMM	System Center Virtual Machine Manager
CSTL	China Systems & Technology Labs	SLES	SUSE Linux Enterprise Server
DHCP	Dynamic Host Configuration Protocol	TOR	Top of rack
DVD	Digital Video Disc	UEFI	Unified Extensible Firmware Interface
FC	Fibre Channel	URL	Uniform Resource Locator
FDR	Fourteen data rate	USB	Universal serial bus
FSM	Flex System Manager	VHD	Virtual hard disk
GB	Gigabyte	VLAN	Virtual LAN
GUI	Graphical user interface	VM	Virtual machine
HBA	Host bus adapter	VMDK	Virtual machine disk
HDD	Hard disk drive	VMM	Virtual Machine Manager
I/O	input/Output	VSMT	Virtual Server Migration Toolkit
IBM	International Business Machines	VT	Virtualization Technology
ID	Identifier	WMI	Windows Management Instrumentation
IM	Instant messaging		
IOM	I/O Module		
IP	Internet Protocol		
IT	Information technology		
ITSO	International Technical Support Organization		
KVM	Keyboard video mouse		
LUN	Logical unit number		
MAC	Media access control		
MB	Megabyte		
MM	Management Module		
NFS	Network file system		
NIC	Network interface card		
OS	Operating system		
OVF	Open Virtualization Format		
PC	Personal computer		
PNG	Portable Network Graphics		
POST	Power-on self test		

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

Lenovo Press publications

The following Lenovo Press publications provide additional information about the topic in this document.

- ▶ *Flex System Products and Technology*, SG24-8255
- ▶ *Flex System Enterprise Chassis*, TIPS0863
- ▶ *IBM Flex System Manager*, TIPS0862
- ▶ *Flex System x240 Compute Node (E5-2600)*, TIPS0860
- ▶ *Flex System x220 Compute Node*, TIPS0885

Other publications and online resources

These publications and websites are also relevant as further information sources:

The following websites also are relevant as additional information sources:

- ▶ *Advanced Management Module User's Guide*
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5073887>
- ▶ Lenovo Customized images for VMware ESXi and vSphere
<http://ibm.com/systems/x/os/vmware/esxi/index.html>
- ▶ IBM DS Storage Manager
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5077693>
- ▶ Flex System Information Center
<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>
- ▶ ServerGuide
<http://ibm.com/support/entry/portal/docdisplay?ln docid=SERV-GUIDE>
- ▶ ServerProven
<http://ibm.com/systems/info/x86servers/serverproven/compat/us>
- ▶ MegaRAID Storage Manager
<http://ibm.com/support/entry/portal/docdisplay?ln docid=MIGR-5077712>
- ▶ ServerGuide Scripting Toolkit
<http://ibm.com/support/entry/portal/docdisplay?ln docid=SERV-T00LKIT>

