


Implementing Systems Management of IBM PureFlex System

Explores IBM PureFlex System and its systems management capabilities

Provides planning and deployment considerations

Gives step-by-step implementation instructions



Ilya Krutov
Carl Gola
Jeff Schoby

Redbooks



International Technical Support Organization

Implementing Systems Management of IBM PureFlex System

April 2014

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

Second Edition (April 2014)

This edition applies to the following products:

- ▶ IBM PureFlex System
- ▶ IBM Flex System
- ▶ IBM Chassis Management Module firmware 1.50.0D
- ▶ IBM Flex System Manager software Version 1.3.0
- ▶ IBM Flex System V7000 Storage Node software Version 7.1

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

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

Contents

Notices	ix
Trademarks	x
Preface	xi
Authors	xi
Now you can become a published author, too!	xiii
Comments welcome	xiii
Stay connected to IBM Redbooks	xiii
Summary of changes	xv
April 2014, Second Edition	xv
Part 1. Introduction	1
Chapter 1. Introduction to IBM PureFlex System and IBM Flex System	3
1.1 IBM PureFlex System	4
1.2 IBM PureFlex System capabilities	6
1.3 IBM Flex System overview	7
1.3.1 IBM Flex System Manager	8
1.3.2 IBM Flex System Enterprise Chassis	9
1.3.3 Compute nodes	10
1.3.4 Expansion nodes	11
1.3.5 Storage nodes	11
1.3.6 I/O modules	12
Chapter 2. IBM PureFlex System and IBM Flex System management devices and appliances	13
2.1 Management network	14
2.2 Chassis Management Module	15
2.3 Compute node management	17
2.3.1 Integrated Management Module II	18
2.3.2 Flexible service processor	18
2.4 I/O modules	19
2.5 IBM Flex System Manager	20
2.5.1 Hardware overview	24
2.5.2 Software features	26
Part 2. Chassis Management Module	31
Chapter 3. Planning for Chassis Management Module-based management	33
3.1 Chassis Management Module management network	34
3.2 Chassis Management Module interfaces	35
3.3 Chassis Management Module security	36
3.3.1 Security policies	36
3.3.2 User account policies	37
3.3.3 External authentication of certificates	39
3.4 Features on Demand planning	40
Chapter 4. Chassis Management Module operations	41
4.1 Initial configuration of Chassis Management Module	42

4.1.1	Connecting to Chassis Management Module	42
4.1.2	Configuring Chassis Management Module by using Initial Setup Wizard.	44
4.1.3	Preparing for Chassis Management Module redundancy.	56
4.1.4	Configuring Chassis Management Module user authority	58
4.1.5	Restoring a Chassis Management Module.	62
4.2	Chassis Management Module management tasks	62
4.2.1	Monitoring the chassis	63
4.2.2	Monitoring multiple chassis.	64
4.2.3	Event notifications.	65
4.2.4	Chassis Management Module Features on Demand	68
4.2.5	Chassis management	70
4.2.6	Using the Chassis Management Module CLI	80

Part 3. IBM Flex System Manager 83

Chapter 5. Planning for IBM Flex System Manager management 85

5.1	Planning for IBM Flex System Manager	86
5.1.1	Flex System Manager network integration architecture	86
5.1.2	Planning for security	87
5.1.3	Planning for Features on Demand	90
5.1.4	Features on Demand for components in the Chassis.	93
5.1.5	Agents and tasks supported	94
5.1.6	Planning for the management of networking infrastructure.	96
5.1.7	Planning for the management of storage infrastructure	97
5.1.8	Planning for IBM Fabric Manager	99
5.2	Planning for the management of virtualized environments	100
5.2.1	Virtualization and task supported	100
5.2.2	Planning for Linux KVM virtualization	102
5.2.3	Planning for PowerVM virtualization	111
5.2.4	Planning for VMware virtualization	115
5.2.5	Planning for Hyper-V virtualization	118

Chapter 6. IBM Flex System Manager initial configuration 121

6.1	IBM Flex System Manager Setup Wizard	123
6.2	Updating Flex System Manager	137
6.3	Selecting chassis to manage	138
6.4	Configuring centralized user management	142
6.5	Configuring chassis components	144
6.6	Configuring compute nodes using Configuration Patterns	150
6.6.1	Overview of Configuration Patterns	150
6.6.2	Creating and applying compute node Configuration Patterns.	153
6.6.3	Automating compute node failover with Configuration Patterns	161
6.7	Deploying compute node images	168
6.7.1	Importing operating system images	169
6.7.2	Deploying a new image.	171
6.8	System discovery, access, and inventory collection	176
6.8.1	Discovery basics	176
6.8.2	Operating system discovery	179
6.8.3	Requesting access to the discovered operating system.	182
6.8.4	Collecting operating system inventory	184
6.9	Updating chassis components	189
6.9.1	Acquiring updates for chassis components.	191
6.9.2	Updating the CMM firmware	199
6.9.3	Updating compute node firmware	201

6.9.4	Updating I/O module firmware	205
6.9.5	Compliance policies	211
6.10	Manage Feature-on-Demand keys	216
6.11	Flex System V7000 Storage Node initial configuration	223
6.11.1	Creating a new system on the V7000 Storage Node	224
6.11.2	Flex System V7000 Storage Node Setup wizard	228
6.12	Discover and manage external Storwize V7000	234
6.12.1	Discover an IBM Storwize V7000	234
6.12.2	Collect inventory on the discovered V7000.	237
6.13	Overview of Flex System V7000 and Storwize V7000 systems management (Storage Control).	241
6.14	External Fibre Channel SAN switch discovery	247
6.15	Configuring network parameters (Network Control)	256
Chapter 7. Managing chassis components with IBM Flex System Manager		259
7.1	Using FSM Explorer	260
7.2	Using the Chassis Map	264
7.3	Using the Event Log	269
7.4	Automating tasks with event automation plans	271
7.5	Handling problems with Service and Support Manager	279
7.6	Integrating Flex System Manager with an enterprise monitoring system	288
7.7	Monitoring system status and health.	288
7.8	Remote management	298
Chapter 8. IBM Fabric Manager		305
8.1	IBM Fabric Manager overview	306
8.2	Starting the IBM Fabric Manager interface	306
8.3	Adding devices	307
8.4	Adding a device pool.	310
8.5	Adding a boot target template.	310
8.6	Adding a profile	312
8.7	Profile deployment	313
8.8	Pushing a deployment.	313
8.9	Verifying an IBM Fabric Manager deployment	314
8.10	Adding and starting a monitor	316
Chapter 9. Managing the KVM environment with IBM Flex System Manager		319
9.1	KVM management architecture.	320
9.2	KVM platform agent installation	320
9.2.1	Preparation	321
9.2.2	KVM Platform Agent installation	322
9.2.3	KVM host discovery, granting access, and inventory collection	325
9.3	Image repository for KVM	328
9.3.1	Preparation	329
9.3.2	Common Agent installation on a KVM host image repository.	330
9.3.3	Subagent installation on a KVM image repository host.	336
9.3.4	Host mappings	339
9.3.5	Discover and manage V7000 storage system	344
9.3.6	Discover and manage SAN switches	344
9.3.7	Discover and configure an image repository server for SAN storage	345
9.4	Creating KVM storage system pools.	350
9.5	Creating KVM network system pools	352
9.6	Creating KVM server system pools.	365
9.7	Add host to an existing server system pool.	372

9.8	Operating a KVM virtual infrastructure	375
9.8.1	Importing a virtual appliance	375
9.8.2	Deploy a virtual appliance to create a virtual server	379
9.8.3	Capturing a virtual appliance	387
9.8.4	Relocate virtual servers.	396
Chapter 10. Managing the PowerVM environment with IBM Flex System Manager		401
10.1	Initial deployment of virtual machine	402
10.1.1	Solution architecture	402
10.1.2	Setting up VIOS and Network Installation Manager server.	403
10.2	Capturing virtual machines	416
10.2.1	Capturing AIX by using Network Installation Manager (NIM)	417
10.2.2	Capturing the Network Installation Manager server	435
10.2.3	Capturing AIX by using storage copy services (SCS).	444
10.3	Deploying virtual machines	471
10.3.1	Deploying virtual machines by using the LPP_source	471
10.3.2	Deploying a virtual machine by using mksysb	478
10.3.3	Deploying a virtual machine by using Storage Copy Services (SCS).	481
10.4	Relocating virtual machines	486
10.4.1	Manual relocation	486
10.4.2	Automatic relocation	486
10.4.3	Relocating virtual servers manually	487
Chapter 11. Managing the VMware environment with IBM Flex System Manager		495
11.1	Environment overview.	496
11.2	Deploying a VM.	498
11.3	Relocating a VM	514
11.4	Relocating all VMs from a host and saving a relocation plan	519
11.5	Modifying the Virtual Server resource allocation.	523
11.6	Enabling VMware Distributed Resource Scheduler (DRS)	529
11.7	Putting a host in maintenance mode.	534
11.8	Topology view	539
11.9	Automating preventive actions in response to hardware alerts.	544
Chapter 12. Managing the Hyper-V environment with IBM Flex System Manager		555
12.1	Initial setup tasks for a Hyper-V node	556
12.1.1	Discovering your Hyper-V server	556
12.1.2	Importing the Common Agent for Windows	557
12.1.3	Granting access and collecting inventory on a Hyper-V node	558
12.1.4	Installing the Common Agent on a Hyper-V host	559
12.2	Managing Hyper-V with IBM Flex System Manager	562
12.2.1	Deploying virtual servers.	562
12.2.2	Editing a virtual server	567
12.2.3	Deleting a virtual server	568
12.2.4	Viewing the virtual server network topology	568
Chapter 13. Mobile management.		571
13.1	Obtaining the mobile application.	572
13.2	Configuring secure communications to the FSM	572
13.2.1	Generating a Java keystore and Certificate Signing Request.	572
13.2.2	Installing the keystore into the IBM Flex System Manager.	574
13.2.3	Installation on Android	574
13.2.4	Installation on BlackBerry	575
13.2.5	Installation on iOS.	576

13.3 Using the Flex System Manager mobile application	576
Abbreviations and acronyms	587
Related publications	589
IBM Redbooks	589
Online resources	589
Help from IBM	589

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

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


COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

AIX®	Power Systems™	Storwize®
BladeCenter®	POWER7®	System Storage®
DS8000®	POWER7+™	System x®
Electronic Service Agent™	PowerVM®	Tivoli®
FlashCopy®	PureApplication®	Tivoli Enterprise Console®
IBM®	PureFlex®	X-Architecture®
IBM Flex System®	PureSystems®	XIV®
IBM Flex System Manager™	Real-time Compression™	z/OS®
NetView®	Redbooks®	
POWER®	Redbooks (logo)  ®	

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.

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

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

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

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

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

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

Preface

To meet today's complex and ever-changing business demands, you need a solid foundation of compute, storage, networking, and software resources. This system must be simple to deploy and be able to quickly and automatically adapt to changing conditions. You also need to be able to take advantage of broad expertise and proven guidelines in systems management, applications, industry solutions, and more.

IBM® PureFlex® System combines no-compromise system designs along with built-in expertise and integrates them into complete, optimized scalable solutions. With IBM Flex System® Manager, multiple solution components that include compute nodes, network and storage infrastructures, storage systems, and heterogeneous virtualization environments can be managed from a single panel.

This IBM Redbooks® publication introduces IBM PureFlex System and IBM Flex System and their management devices and appliances. It provides implementation guidelines for managing Linux kernel-based virtual machine (KVM), IBM PowerVM®, VMware vSphere, and Microsoft Hyper-V virtualization environments.

This book is intended for the IT community of clients, IBM Business Partners, and IBM employees who are interested in planning and implementing systems management of the IBM PureFlex System.

Authors

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



Ilya Krutov is a Project Leader at the ITSO Center in Raleigh and has been with IBM since 1998. Before he joined the ITSO, Ilya served in IBM as a Run Rate Team Leader, Portfolio Manager, Brand Manager, Technical Sales Specialist, and Certified Instructor. Ilya has expertise in IBM System x®, BladeCenter®, and PureFlex System products; server operating systems; and networking solutions. He has authored over 150 books, papers, and Product Guides. He has a Bachelor's degree in Computer Engineering from the Moscow Engineering and Physics Institute.



Carl Gola is an IT Specialist with IBM, and works in the Boulder PureFlex Center of Competency supporting both internal and external clients. His areas of expertise include the IBM PureFlex System, IBM Power Systems™, AIX®, Linux, IBM X-Architecture®, and virtualization. He is certified in Red Hat Linux and is proficient in security policies, auditing, performance tuning, and backup/restore technologies. He has worked for IBM since 1997 and is based in Boulder, CO.



Jeff Schoby has been with IBM since 2012 as a Subject Matter Expert for a UNIX team in the Columbia, Missouri, Delivery Center. He has 20 years of experience in System x and Power server operating systems and administration, networking, storage, and virtualization technologies. In 2013, he was chosen to assist with IBM PureSystems® support. His end-to-end knowledge of IT systems has enabled him to quickly familiarize himself with much of the IBM PureSystems environment and help multiple clients and IBM teams with a wide range of configuration scenarios.

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

From the ITSO:

- ▶ Kevin Barnes
- ▶ Tamikia Barrow
- ▶ Ella Buslovich
- ▶ Mary Comianos
- ▶ Cheryl Gera
- ▶ David Watts

From IBM:

- ▶ Hunter Cook
- ▶ Daniel Daley
- ▶ Kevin Hoff
- ▶ Chris Long
- ▶ Josh Niemeyer
- ▶ Meleata Pinto
- ▶ David Tareen
- ▶ Grant Taylor
- ▶ Erica St. John

Thanks to the authors of the first edition, *Implementing Systems Management of IBM PureFlex System*, SG24-8060-00, published in November 2012:

- ▶ Ilya Krutov
- ▶ Frederik Aouizerats
- ▶ Brandon Harrell
- ▶ MinChul Kim
- ▶ Stanimir Markov

Now you can become a published author, too!

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

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

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

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review Redbooks form found at this website:

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

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

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

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-8060-01
for Implementing Systems Management of IBM PureFlex System
as created or updated on April 22, 2014.

April 2014, Second Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information

- ▶ IBM Flex System Manager™ (FSM) Explorer
- ▶ Configuration patterns
- ▶ Operating system deployment
- ▶ IBM Fabric Manager
- ▶ Mobile management
- ▶ Integrated V7000 Storage Node management
- ▶ Network System Pools

Changed information

- ▶ PureFlex System overview
- ▶ FSM software features
- ▶ FSM ordering information
- ▶ Flex System Feature on Demand (FoD) upgrades
- ▶ Centralized FSM security



Part 1

Introduction

This book is divided into multiple parts. This part introduces IBM PureFlex System and IBM Flex System, and describes their management architecture, devices, and appliances.

This part includes the following chapters:

- ▶ Chapter 1, “Introduction to IBM PureFlex System and IBM Flex System” on page 3
- ▶ Chapter 2, “IBM PureFlex System and IBM Flex System management devices and appliances” on page 13



Introduction to IBM PureFlex System and IBM Flex System

In today's ever-changing marketplace, IT solutions are driving business – making it faster, more nimble, and ultimately, more successful. Users want access to data and applications on their terms. They want full service from mobile devices. They rely on social networks to drive ever increasing numbers of transactions.

Cloud computing solutions from IBM can help you meet these demands. IBM PureFlex System and IBM Flex System offerings and solutions are optimized for cloud deployments. With the latest IBM PureFlex System technology enhancements, you can accelerate system deployment while simplifying cloud service delivery and improving efficiency in your IT environment.

IBM PureFlex System is a complete, flexible cloud infrastructure system with integrated expertise. The system integrates and optimizes all compute, storage, and networking resources to deliver infrastructure-as-a-service (IaaS) out of the box. To simplify acquisition of your solution, you can choose one of the predefined and fully integrated, optimized configurations as the starting point.

IBM Flex System offers a broad range of x86 and IBM POWER® compute nodes in an innovative chassis design that goes beyond blade servers with advanced networking and system management, to support extraordinary simplicity, flexibility, and upgradeability.

If you want a pre-configured, pre-integrated infrastructure with integrated management and cloud capabilities, that is factory tuned from IBM with an x86 and Power hybrid solution, IBM PureFlex System is the answer.

If you want to build and tune a custom configuration with efficient x86 compute and memory performance, or on the latest IBM POWER and POWER7+™ processors for maximum performance and efficiency, IBM Flex System is the right fit.

This chapter covers the following topics:

- ▶ 1.1, “IBM PureFlex System” on page 4
- ▶ 1.2, “IBM PureFlex System capabilities” on page 6
- ▶ 1.3, “IBM Flex System overview” on page 7

1.1 IBM PureFlex System

To meet today's complex and ever-changing business demands, you need a solid foundation of server, storage, networking, and software resources. Furthermore, it must be simple to deploy and able to quickly and automatically adapt to changing conditions. You also need access to, and the ability to take advantage of, broad expertise and proven guidelines in systems management, applications, hardware maintenance, and more.

IBM PureFlex System is a comprehensive infrastructure system that provides an expert-integrated computing system. It combines servers, enterprise storage, networking, virtualization, and management into a single structure. Its built-in expertise enables organizations to manage and flexibly deploy integrated patterns of virtual and hardware resources through unified management. These systems are ideally suited for clients who want a system that delivers the simplicity of an integrated solution with the capability to tune middleware and the runtime environment.

IBM PureFlex System uses workload placement that is based on virtual machine compatibility and resource availability. By using built-in virtualization across servers, storage, and networking, the infrastructure system enables automated scaling of resources and true workload mobility.

IBM PureFlex System has undergone significant testing and experimentation so that it can mitigate IT complexity without compromising the flexibility to tune systems to meet business demands. By providing flexibility and simplicity, IBM PureFlex System can provide extraordinary levels of IT control, efficiency, and operating agility. This combination enables businesses to rapidly deploy IT services at a reduced cost. Moreover, the system is built on decades of expertise. This expertise enables deep integration and central management of the comprehensive, open-choice infrastructure system. It also dramatically cuts down on the skills and training that are required for managing and deploying the system.

IBM PureFlex System combines advanced IBM hardware and software along with patterns of expertise. It integrates them into three optimized configurations that are simple to acquire and deploy so that you get fast time to value.

IBM PureFlex System is built and integrated before shipment so it can be quickly deployed into the data center. PureFlex System is shipped complete, integrated within a rack incorporating all the required power, networking, and SAN cabling together with all the associated switches, compute nodes, and storage.

Figure 1-1 shows an IBM PureFlex System 42U rack, complete with its distinctive PureFlex door.



Figure 1-1 IBM PureFlex System

The PureFlex System includes the following configurations:

- ▶ IBM PureFlex System Express, which is designed for small and medium businesses and is the most affordable entry point for PureFlex System.
- ▶ IBM PureFlex System Standard, which is optimized for application servers with supporting storage and networking, and is designed to support your key ISV solutions.
- ▶ IBM PureFlex System Enterprise, which is optimized for transactional and database systems. It has built-in redundancy for highly reliable and resilient operation to support your most critical workloads.

These configurations are summarized in Table 1-1.

Table 1-1 IBM PureFlex System configurations

Component	IBM PureFlex System Express	IBM PureFlex System Standard	IBM PureFlex System Enterprise
IBM PureFlex System 42U Rack	1	1	1
IBM Flex System Enterprise Chassis	1	1	1
IBM Flex System Fabric EN4093 10 Gb Scalable Switch	1	1	2 with both port-count upgrades
IBM Flex System FC3171 8 Gb SAN Switch ^a	1	2	2
IBM Flex System FC5022 24-port 16Gb ESB SAN Scalable Switch ^a	1	2	2
IBM Flex System Manager Node	1	1	1
IBM Flex System Manager software license	IBM Flex System Manager with 1-year service and support	IBM Flex System Manager Advanced with 3-year service and support	Flex System Manager Advanced with 3-year service and support
Chassis Management Module	2	2	2
Chassis power supplies (standard /maximum)	2/6	4/6	6/6
Chassis 80 mm fan modules (standard /maximum)	4/8	6/8	8/8
IBM Flex System V7000 Storage Node ^b	Yes (redundant controller)	Yes (redundant controller)	Yes (redundant controller)
IBM Storwize V7000 Disk System ^b	Yes (redundant controller)	Yes (redundant controller)	Yes (redundant controller)
IBM Storwize V7000 Software	<ul style="list-style-type: none"> ▶ Base with 1-year software maintenance agreement ▶ Optional Real Time Compression 	<ul style="list-style-type: none"> ▶ Base with 3-year software maintenance agreement ▶ Real Time Compression 	<ul style="list-style-type: none"> ▶ Base with 3-year software maintenance agreement ▶ Real Time Compression

a. Select the IBM Flex System FC3171 8 Gb SAN Switch or IBM Flex System FC5022 24-port 16Gb ESB SAN Scalable Switch module.

b. Select the IBM Flex System V7000 Storage Node that is installed inside the Enterprise chassis or the external IBM Storwize® V7000 Disk System.

1.2 IBM PureFlex System capabilities

The PureFlex System offers these advantages:

- ▶ Configurations that ease acquisition experience and match your needs
- ▶ Optimized to align with targeted workloads and environments

- ▶ Designed for cloud with SmartCloud Entry included on Standard and Enterprise
- ▶ Choice of architecture, operating system, and virtualization engine
- ▶ Designed for simplicity with integrated, single-system management across physical and virtual resources
- ▶ Simplified ordering that accelerates deployment into your environments
- ▶ Ships as a single integrated entity directly to you
- ▶ Includes factory integration and lab services optimization

IBM PureFlex System has three preintegrated offerings that support compute, storage, and networking requirements. You can select from these offerings, which are designed for key client initiatives and help simplify ordering and configuration. As a result, PureFlex System reduces the cost, time, and complexity of system deployments.

The IBM PureFlex System is offered in these configurations:

- ▶ Express: The infrastructure system for small-sized and midsized businesses, and the most cost-effective entry point.
- ▶ Standard: The infrastructure system for application servers with supporting storage and networking.
- ▶ Enterprise: The infrastructure system that is optimized for scalable cloud deployments. It has built-in redundancy for highly reliable and resilient operation to support critical applications and cloud services.

A PureFlex System configuration has these main components:

- ▶ Preinstalled and configured IBM Flex System Enterprise Chassis
- ▶ Compute nodes with either IBM POWER or Intel Xeon processors
- ▶ IBM Flex System Manager, preinstalled with management software and licenses for software activation
- ▶ IBM Flex System V7000 Storage Node or IBM Storwize V7000 external storage unit
- ▶ All hardware components that are preinstalled in an IBM PureFlex System 42U rack
- ▶ Choice of the following options:
 - Operating system: IBM AIX, IBM i, Microsoft Windows, Red Hat Enterprise Linux, or SUSE Linux Enterprise Server
 - Virtualization software: IBM PowerVM, KVM, VMware vSphere, or Microsoft Hyper V
 - SmartCloud Entry
- ▶ Complete pre-integrated software and hardware
- ▶ Onsite services included to get you up and running quickly

The fundamental building blocks of the three IBM PureFlex System solutions are the compute nodes, storage nodes, and networking of the IBM Flex System Enterprise Chassis.

1.3 IBM Flex System overview

IBM Flex System is a full system of hardware that forms the underlying strategic basis of IBM PureFlex System and IBM PureApplication® System and forms the underlying hardware basis of other IBM PureSystems offerings. IBM Flex System optionally includes a management appliance, known as *Flex System Manager*.

IBM Flex System is the next generation blade chassis offering from IBM that features the latest innovations and advanced technologies.

The major components of the IBM Flex System are described next.

1.3.1 IBM Flex System Manager

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

It is an ideal solution that allows you to reduce administrative expense and focus your efforts on business innovation.

A single user interface controls the following features:

- ▶ Intelligent automation
- ▶ Resource pooling
- ▶ Improved resource usage
- ▶ Complete management integration
- ▶ Simplified setup

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

An increased focus on optimizing time-to-value is evident in the following features:

- ▶ Setup wizards, including initial setup wizards, provide intuitive and quick setup of the Flex System Manager.
- ▶ The Chassis Map provides multiple view overlays to track health, firmware inventory, and environmental metrics.
- ▶ Configuration management for repeatable setup of compute, network, and storage devices.
- ▶ Remote presence application for remote access to compute nodes with single sign-on.
- ▶ Quick search provides results as you type.

Beyond the physical world of inventory, configuration, and monitoring, IBM Flex System Manager enables virtualization and workload optimization for a new class of computing:

- ▶ Resource usage: Detects congestion, notification policies, and relocation of physical and virtual machines that include storage and network configurations within the network fabric.
- ▶ Resource pooling: Pooled network switching, with placement advisors that consider virtual machine (VM) compatibility, processor, availability, and energy.
- ▶ Intelligent automation: Automated and dynamic VM placement that is based on usage, hardware predictive failure alerts, and host failures.

Figure 1-2 shows the IBM Flex System Manager appliance.



Figure 1-2 IBM Flex System Manager

1.3.2 IBM Flex System Enterprise Chassis

The IBM Flex System Enterprise Chassis is the foundation of the Flex System offering, which features 14 standard (half-width) Flex System form factor compute node bays in a 10U chassis that delivers high-performance connectivity for your integrated compute, storage, networking, and management resources.

Up to a total of 28 independent servers can be accommodated in each Enterprise Chassis, if double-dense x222 compute nodes are deployed.

The chassis is designed to support multiple generations of technology and offers independently scalable resource pools for higher usage and lower cost per workload.

With the ability to handle up to 14 nodes, supporting the intermixing of IBM Power Systems and Intel x86, the Enterprise Chassis provides flexibility and tremendous compute capacity in a 10U package. Additionally, the rear of the chassis accommodates four high-speed I/O bays that can accommodate up to 40 GbE high-speed networking, 16 Gb Fibre Channel or 56 Gb InfiniBand. With interconnecting compute nodes, networking, and storage that uses a high-performance and scalable mid-plane, the Enterprise Chassis can support latest high-speed networking technologies.

The “ground-up” design of the Enterprise Chassis reaches new levels of energy efficiency through innovations in power, cooling, and air flow. Simpler controls and futuristic designs allow the Enterprise Chassis to break free of “one size fits all” energy schemes.

The ability to support the workload demands of tomorrow’s workloads is built in with a new I/O architecture, which provides choice and flexibility in fabric and speed. With the ability to use Ethernet, InfiniBand, Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), and iSCSI, the Enterprise Chassis is uniquely positioned to meet the growing and future I/O needs of large and small businesses.

Figure 1-3 shows the IBM Flex System Enterprise Chassis.



Figure 1-3 The IBM Flex System Enterprise Chassis

1.3.3 Compute nodes

IBM Flex System offers compute nodes that vary in architecture, dimension, and capabilities.

Optimized for efficiency, density, performance, reliability, and security, the portfolio includes a range of IBM POWER and Intel Xeon-based nodes that are designed to make full use of the full capabilities of these processors that can be mixed within the same Enterprise Chassis.

Power Systems nodes are available in either a two and four socket variety that uses the IBM POWER7® and IBM POWER7+ processors. Also available is a POWER7 node that is optimized for cost-effective deployment of Linux.

Compute nodes that use Intel processors are available that range from the two-socket Intel Xeon E5-2400 product family and the two-socket Intel E5-2600 product family to the four-socket Intel E5-4800 product family.

Up to 28 two-socket Intel Xeon E5-2400 servers can be deployed in a single Enterprise Chassis where high-density cloud, virtual desktop, or server virtualization is wanted.

Figure 1-4 shows a four-socket IBM POWER7 compute node, the p460.



Figure 1-4 IBM Flex System p460 Compute Node

The nodes are complemented with leadership I/O capabilities of up to 16 channels of high-speed I/O lanes per standard wide node bay and 32 lanes per full wide node bay. Various I/O adapters and matching I/O modules are available.

1.3.4 Expansion nodes

Expansion nodes can be attached to certain standard form factor (half-width) Flex System compute nodes. This attachment allows the expansion of the nodes' capabilities with locally attached storage or PCIe adapters.

The IBM Flex System Storage Expansion Node provides locally attached disk expansion to the x240 and x220. SAS and SATA disks are supported.

With the attachment of the IBM Flex System PCIe Expansion Node, an x220 or x240 can have up to four PCIe adapters attached. High-performance graphics processing units (GPUs) can also be installed within the PCIe Expansion Node from companies, such as Intel and NVIDIA.

1.3.5 Storage nodes

The storage capabilities of IBM Flex System give you advanced functionality with storage nodes in your system and make full use of your existing storage infrastructure through advanced virtualization.

Storage is available within the chassis by using the IBM Flex System V7000 Storage Node that integrates with the Flex System Chassis or externally with the IBM Storwize V7000.

IBM Flex System simplifies storage administration with a single user interface for all your storage. The management console is integrated with the comprehensive management system. These management and storage capabilities allow you to virtualize third-party storage with nondisruptive migration of your current storage infrastructure. You can also make use of intelligent tiering so that you can balance performance and cost for your storage needs. The solution also supports local and remote replication and snapshots for flexible business continuity and disaster recovery capabilities.

Flex System can also be connected to various external storage systems.

1.3.6 I/O modules

The range of available modules and switches to support key network protocols allows you to configure IBM Flex System to fit in your infrastructure. However, you can do so without sacrificing the ability to be ready for the future. The networking resources in IBM Flex System are standards-based, flexible, and fully integrated into the system. This combination gives you no-compromise networking for your 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.

IBM Flex System gives you the following key networking capabilities:

- ▶ Supports the networking infrastructure that you have today, including Ethernet, FC, FCoE, and InfiniBand
- ▶ Offers industry-leading performance with 1 Gb, 10 Gb, and 40 Gb Ethernet, 8 Gb and 16 Gb Fibre Channel, and quad data rate (QDR) and fourteen data rate (FDR) InfiniBand
- ▶ Provides pay-as-you-grow scalability so you can add ports and bandwidth when needed


Networking in data centers is undergoing a transition from a discrete traditional model to a more flexible, optimized model. The network architecture in IBM Flex System was designed to address the key challenges clients are facing today in their data centers. The key focus areas of the network architecture on this platform are unified network management, optimized and automated network virtualization, and simplified network infrastructure.

Providing innovation, leadership, and choice in the I/O module portfolio uniquely positions IBM Flex System to provide meaningful solutions to address client needs.

Figure 1-5 shows the IBM Flex System Fabric EN4093R 10Gb Scalable Switch.



Figure 1-5 IBM Flex System Fabric EN4093R 10Gb Scalable Switch



IBM PureFlex System and IBM Flex System management devices and appliances

The IBM Flex System hardware and software features can help you accomplish these tasks:

- ▶ Optimize your resource and power utilization
- ▶ Track and deploy your assets
- ▶ Maintain a secure environment
- ▶ Simplify the overall management of your data center

The Chassis Management Module (CMM), integrated compute node management controllers, and the IBM Flex System Manager (FSM) management node are designed to help simplify the overall management of your IBM Flex System resources.

This chapter includes the following sections:

- ▶ 2.1, “Management network” on page 14
- ▶ 2.2, “Chassis Management Module” on page 15
- ▶ 2.3, “Compute node management” on page 17
- ▶ 2.4, “I/O modules” on page 19
- ▶ 2.5, “IBM Flex System Manager” on page 20

2.1 Management network

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

The management network is shown in Figure 2-1 as the blue line. It connects the Chassis Management Module (CMM) to the compute nodes, the switches in the I/O bays, and the Flex System Manager (FSM). The FSM connection to the management network is through a special Broadcom 5718-based management network adapter (Eth0). The management networks in multiple chassis can be connected together through the external ports of the CMMs in each chassis by using a GbE top-of-rack switch.

The yellow line in the Figure 2-1 shows the production data network. The FSM also connects to the production network (Eth1) so that it can access the Internet for product updates and other related information.

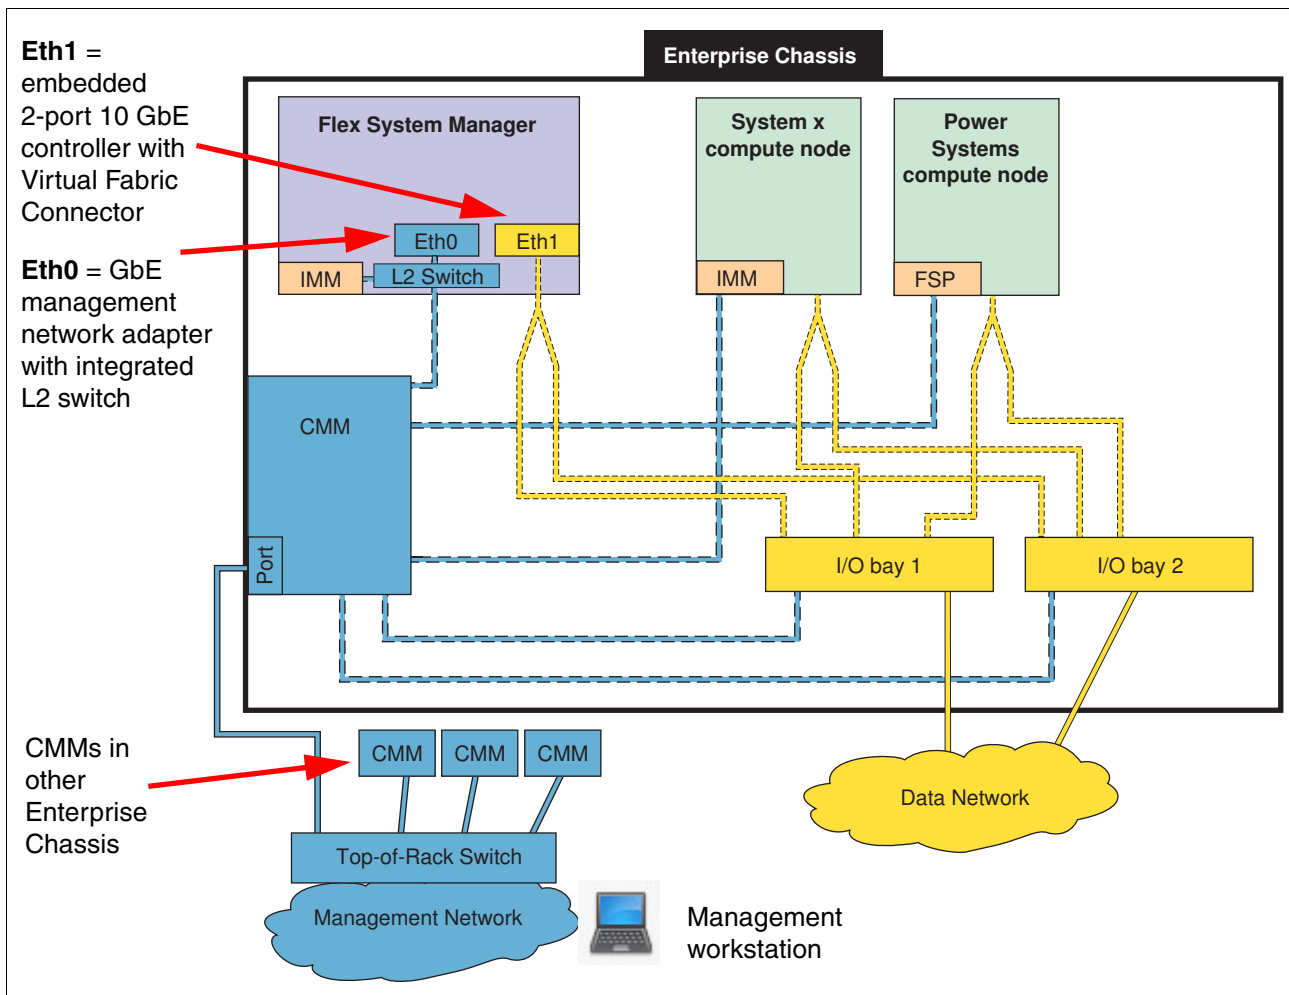


Figure 2-1 Separate management and production data networks

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

2.2 Chassis Management Module

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

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

An Enterprise chassis comes with at least one CMM installed. Table 2-1 lists the ordering information for the second CMM if required.

Table 2-1 Chassis Management Module ordering information

Part number	Feature code ^a	Description
68Y7030	A0UE/3592	IBM Flex System Chassis Management Module

a. The first feature code listed is for x-config configurations. The second feature code is for e-config configurations.

Figure 2-2 shows the location of the CMM bays on the back of the Enterprise Chassis.




Figure 2-2 CMM Bay 1 and Bay 2

Figure 2-3 shows the CMM connectors and LEDs.




Figure 2-3 Chassis Management Module

The CMM has the following connectors:

- ▶ USB connection: Can be used for insertion of a USB media key for tasks such as firmware updates.
- ▶ 10/100/1000 Mbps RJ45 Ethernet connection: For connection to a management network. The CMM can be managed through this Ethernet port.
- ▶ Serial port (mini-USB): For local serial (command-line interface (CLI)) access to the CMM. Use the cable kit that is listed in Table 2-2 for connectivity.

Table 2-2 Serial cable specifications

Part number	Feature code ^a	Description
90Y9338	A2RR/None	IBM Flex System Management Serial Access Cable Contains two cables: <ul style="list-style-type: none"> ▶ Mini-USB-to-RJ45 serial cable ▶ Mini-USB-to-DB9 serial cable

a. The first feature code listed is for x-config configurations. The second feature code is for e-config configurations.

The CMM has the following LEDs that provide status information:

- ▶ Power-on LED
- ▶ Activity LED
- ▶ Error LED
- ▶ Ethernet port link and port activity LEDs

The CMM also incorporates a reset button. It has two functions, dependent upon how long the button is held in:

- ▶ When pressed for less than 5 seconds, the CMM restarts.
- ▶ When pressed for more than 5 seconds (for example 10-15 seconds), the CMM configuration is reset to manufacturing defaults and then restarts.

Through an embedded firmware stack, the CMM implements functions to monitor, control, and provide external user interfaces to manage all chassis resources. The CMM allows you to perform these functions:

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

2.3 Compute node management

Each node in the Enterprise Chassis has a management controller that communicates upstream through the CMM-enabled 1 GbE private management network that enables management capability. Different chassis components that are supported in the Enterprise Chassis can implement different management controllers. Table 2-3 details the different management controllers that are implemented in the chassis components.

Table 2-3 Chassis components and their respective management controllers

Chassis components	Management controller
Intel Xeon processor-based compute nodes	Integrated Management Module II (IMM2)
Power Systems compute nodes	Flexible service processor (FSP)

The management controllers for the various Enterprise Chassis components have the following default IPv4 addresses:

- ▶ CMM: 192.168.70.100
- ▶ Compute nodes: 192.168.70.101-114 (corresponding to the slots 1-14 in the chassis)
- ▶ I/O Modules: 192.168.70.120-123 (sequentially corresponding to chassis bay numbering)

In addition to the IPv4 address, all I/O modules also support link-local IPv6 addresses and configurable external IPv6 addresses.

2.3.1 Integrated Management Module II

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

The IMM2 provides the following major features as standard:

- ▶ IPMI v2.0-compliance
- ▶ Remote configuration of IMM2 and UEFI settings without the need to power on the server
- ▶ Remote access to system fan, voltage, and temperature values
- ▶ Remote IMM and UEFI update
- ▶ UEFI update when the server is powered off
- ▶ Remote console by way of a serial over LAN
- ▶ Remote access to the system event log
- ▶ Predictive failure analysis and integrated alerting features (for example, by using Simple Network Management Protocol (SNMP))
- ▶ Remote presence, including remote control of server by using a Java or Active x client
- ▶ Operating system failure window (blue screen) capture and display through the web interface
- ▶ Virtual media that allow the attachment of a diskette drive, CD/DVD drive, USB flash drive, or disk image to a server
- ▶ Syslog alerting mechanism that provides an alternative to email and SNMP traps
- ▶ Support for Features On Demand (FoD) enablement of server functions, option card features, and System x solutions and applications

For more information, see these resources:

- ▶ *Integrated Management Module II User's Guide*
<http://ibm.com/support/entry/portal/docdisplay?lnocid=MIGR-5086346>
- ▶ *IMM and IMM2 Support on IBM System x and BladeCenter Servers*, TIPS0849:
<http://www.redbooks.ibm.com/abstracts/tips0849.html>

2.3.2 Flexible service processor

Several advanced system management capabilities are built into POWER7-based compute nodes. The flexible service processor (FSP) handles most of the server-level system management. The FSP used in Enterprise Chassis compatible POWER-based nodes is the same service processor that is used on POWER rack servers. It has system alerts and Serial over LAN (SOL) capability.

The FSP provides out-of-band system management capabilities, such as system control, runtime error detection, configuration, and diagnostic procedures. Generally, you do not interact with the FSP directly. Rather, you interact by using tools such as IBM Flex System Manager and Chassis Management Module.

The FSP provides an SOL interface, which is available by using the CMM and the console command. The POWER7-based compute nodes do not have an on-board video chip, and do not support keyboard, video, and mouse (KVM) connections. Server console access is obtained by a SOL connection only.

SOL provides a means to manage servers remotely by using a CLI over a Telnet or SSH connection. SOL is required to manage servers that do not have KVM support or that are attached to the FSM. SOL provides console redirection for both Software Management Services (SMS) and the server operating system.

The SOL feature redirects server serial-connection data over a LAN without requiring special cabling by routing the data through the CMM network interface. The SOL connection enables POWER7-based compute nodes to be managed from any remote location with network access to the CMM.

SOL offers the following functions:

- ▶ Remote administration without KVM
- ▶ Reduced cabling and no requirement for a serial concentrator
- ▶ Standard Telnet/SSH interface, eliminating the requirement for special client software

The Chassis Management Module command-line interface (CLI) provides access to the text-console command prompt on each server through a SOL connection. This configuration allows the POWER7-based compute nodes to be managed from a remote location.

2.4 I/O modules

The I/O modules have the following base functions:

- ▶ Initialization
- ▶ Configuration
- ▶ Diagnostic tests (both power-on and concurrent)
- ▶ Status reporting

In addition, the following set of protocols and software features are supported on the I/O modules:

- ▶ Supports configuration method over the Ethernet management port.
- ▶ A scriptable SSH CLI, a web server with SSL support, Simple Network Management Protocol v3 (SNMPv3) Agent with alerts, and a sFTP client.
- ▶ Server ports that are used for Telnet, HTTP, SNMPv1 agents, TFTP, FTP, and other insecure protocols are DISABLED by default.
- ▶ LDAP authentication protocol support for user authentication.
- ▶ For Ethernet I/O modules, 802.1x enabled with policy enforcement point (PEP) capability to allow support of Trusted Network Connect (TNC).
- ▶ The ability to capture and apply a switch configuration file and the ability to capture a first-failure data capture (FFDC) data file.
- ▶ Ability to transfer files by using URL update methods (HTTP, HTTPS, FTP, TFTP, sFTP).
- ▶ Various methods for firmware updates are supported including FTP, sFTP, and TFTP. In addition, firmware updates by using a URL that includes protocol support for HTTP, HTTPS, FTP, sFTP, and TFTP are supported.
- ▶ Supports SLP discovery in addition to SNMPv3.

- ▶ Ability to detect firmware/hardware hangs, and ability to pull a 'crash-failure memory dump' file to an FTP (sFTP) server.
- ▶ Supports selectable primary and backup firmware banks as the current operational firmware.
- ▶ Ability to send events, SNMP traps, and event logs to the CMM, including security audit logs.
- ▶ IPv4 and IPv6 on by default.
- ▶ The CMM management port supports IPv4 and IPv6 (IPv6 support includes the use of link local addresses).
- ▶ Port mirroring capabilities:
 - Port mirroring of CMM ports to both internal and external ports.
 - For security reasons, the ability to mirror the CMM traffic is hidden and is available only to development and service personnel
- ▶ Management virtual local area network (VLAN) for Ethernet switches: A configurable management 802.1q tagged VLAN in the standard VLAN range of 1 - 4094. It includes the CMM's internal management ports and the I/O modules internal ports that are connected to the nodes.

2.5 IBM Flex System Manager

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

- ▶ Reduce the number of interfaces, steps, and clicks it takes to manage IT resources
- ▶ Intelligently manage and deploy workloads that are based on resource availability and predefined policies
- ▶ Manage events and alerts to increase system availability and reduce downtime
- ▶ Reduce operational costs

The IBM Flex System Manager management appliance is shown in Figure 2-4.



Figure 2-4 IBM Flex System Manager management appliance

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

IBM Flex System Manager has the following key features:

- ▶ Optimizing your workload management through built-in expertise

With a workload-optimized approach, you can decrease infrastructure costs and improve service levels. You can create and modify system pools using virtual workloads, make dynamic virtual workload adjustments, and move workloads within system pools. These features result in an optimized virtual environment with increased resilience to cope with planned or unplanned downtime. A system pool is a group of virtualized system components that are managed as a single entity. This configuration allows you to manage the pools as easily as managing a single system, which is an essential capability for moving to cloud computing and a dynamic infrastructure.

- ▶ Managing all of your resources with one solution

IBM Flex System Manager is designed to provide all of the key management functions for your integrated IT resources from a single, easy to use interface. This support begins with deployment through maintenance, upgrades, and problem resolution. From your office or remotely through a secure connection, you can manage your compute, storage, network, and virtualized resources:

- Compute

Auto discovery and setup wizards make deploying compute nodes quick and easy using the IBM Flex System Manager. After it is deployed, IBM Flex System Manager provides real-time updates for compute node “health” summaries. With the ability to define performance thresholds to trigger alerts, you can automate responses to potential problems help keep your critical business applications running at peak performance. IBM Flex System Manager can detect many problems with essential system resources and recover automatically. IBM Flex System Manager can also run trend analysis to forecast and prevent future problems that otherwise might lead to expensive system outages.

- Storage

IBM Flex System Manager helps you address storage management challenges from device deployment and through the data lifecycle. Storage deployment capabilities in the IBM Flex System Manager include storage device discovery and simple logical and physical device configuration from a single interface. IBM Flex System Manager can provide physical and logical storage topology views, and show relationships between storage and server resources. These features give you the ability to track key resources based on their business usage. Provisioning capabilities include image management for simple virtual machine creation, deployment, and cloning. You can also manage storage system pools for data lifecycle management and storage placement based on business policies.

– Networking

Networking resources allow your virtualized compute and storage resources to communicate and function in the cloud. IBM Flex System Manager delivers end-to-end network management for your PureFlex System from a single tool. IBM Flex System Manager supports automated network discovery to speed deployments. It also offers a graphical view of the network from the integrated user interface. Network resources are pooled and virtualized. With logical network profiles, you can quickly and easily specify the network connectivity characteristics of a virtual machine.

IBM Flex System Manager supports automatic provisioning and simple movements of virtual LANs for virtual machines. You can manage MAC addresses for virtual network interface cards. IBM Flex System Manager provides detailed network usage and performance statistics for virtual machines and physical compute nodes. These statistics allow you to track valuable network resources and manage them based on your business needs.

– Virtualization

The basic virtualization functions in the IBM Flex System Manager begin with the ability to create and manage virtual servers from pooled resources. IBM Flex System Manager takes this capability further through the application of built-in expertise to make provisioning and deployment of virtual machines fast and easy. After virtual machines are deployed, the virtualization features of IBM Flex System Manager are designed to help you manage these virtualized resources efficiently. Automation features such as dynamic virtual machine placement, automated optimization, and resource balancing simplify virtualization management. IBM Flex System Manager also helps keep your virtual machines up and running with support for nondisruptive updates, virtual machine mobility, and a range of other resilience features.

The IBM Flex System Manager appliance is based on an x86 compute node that comes with preloaded management software. The software contains a set of components that are responsible for running certain management functions. These components must be activated by using the available IBM FoD software entitlement licenses. They are licensed on a per-chassis basis, so you need one license for each chassis you plan to manage. The management node comes standard without any entitlement licenses, so you must purchase a license to enable the required FSM functions.

The part number to order the management node is shown in Table 2-4.

Table 2-4 Ordering information for IBM Flex System Manager node

Part number	Description
8731A1x ^a	IBM Flex System Manager node

a. The x in the Part number represents a country-specific letter (for example, the EMEA part number is 8731A1G, and the US part number is 8731A1U). Ask your local IBM representative for specifics.

The part numbers to order FoD software entitlement licenses are shown in the following tables. The part numbers for the same features are different in different countries. Ask your local IBM representative for specifics. Table 2-5 shows the information for the United States, Canada, Asia Pacific, and Japan.

Table 2-5 Ordering information for FoD licenses (United States, Canada, Asia Pacific, and Japan)

Part number	Description
Base feature set	
90Y4217	IBM Flex System Manager Per Managed Chassis with 1-Year SW S&S
90Y4222	IBM Flex System Manager Per Managed Chassis with 3-Year SW S&S
Advanced feature set upgrade ^a	
90Y4249	IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 1-Year SW S&S
00D7554	IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 3-Year SW S&S

a. The base feature set is a prerequisite for the Advanced Upgrade.

Table 2-6 shows the ordering information for Latin America and Europe/Middle East/Africa.

Table 2-6 Ordering information for FoD licenses (Latin America and Europe/Middle East/Africa)

Part number	Description
Base feature set	
95Y1174	IBM Flex System Manager Per Managed Chassis with 1-Year SW S&S
95Y1179	IBM Flex System Manager Per Managed Chassis with 3-Year SW S&S
Advanced feature set upgrade ^a	
94Y9219	IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 1-Year SW S&S
94Y9220	IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 3-Year SW S&S

a. The base feature set is a prerequisite for the Advanced Upgrade.

IBM Flex System Manager base feature set offers the following functions:

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

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

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

Requirement: IBM Flex System Manager base license is a prerequisite for the Advanced Upgrade license.

2.5.1 Hardware overview

The IBM FSM Manager Node has the following fixed hardware specifications:

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

The FSM Manager Node ships with a preinstalled software management stack based on RHEV-H.

Figure 2-5 shows the internal layout of the FSM.




Figure 2-5 Internal view that shows the major components of IBM Flex System Manager

Front controls

The diagram in Figure 2-6 shows the front of an FSM with the location of the controls and LEDs.




Figure 2-6 FSM front panel showing controls and LEDs

Local storage

The FSM ships with 2 x IBM 200 GB SATA 1.8" MLC SSD and 1 x IBM 1 TB 7.2K 6 Gbps NL SATA 2.5" SFF HS HDD drives. The 200 GB SSD drives are configured in an RAID-1 pair that provides roughly 200 GB of usable space. The 1 TB SATA drive is not part of a RAID group.

The partitioning of the disks is listed in Table 2-7.

Table 2-7 Detailed SSD and HDD disk partitioning

Physical disk	Virtual disk size	Description
SSD	50 MB	Boot disk
SSD	60 GB	OS/Application disk
SSD	80 GB	Database disk
HDD	40 GB	Update repository
HDD	40 GB	Dump space
HDD	60 GB	Spare disk for OS/Application
HDD	80 GB	Spare disk for database
HDD	30 GB	Service Partition

Management network adapter

The management network adapter is a standard feature of the FSM, and provides a physical connection into the private management network of the chassis. The adapter contains a Broadcom 5718 Dual 1GbE adapter and a Broadcom 5389 8-port L2 switch. This card is one of the features that makes the FSM unique compared to all other nodes supported by the Enterprise Chassis. The management network adapter provides a physical connection into the private management network of the chassis. The connection allows the software stack to have visibility into both the data and management networks. The L2 switch on this card is automatically set up by the IMM2. It connects the FSM and the onboard IMM2 into the same internal private network.

2.5.2 Software features

The IBM Flex System Manager management software has these main features:

- ▶ Monitoring and problem determination:
 - A real-time, multichassis view of hardware components with overlays for more information.
 - Automatic detection of issues in your environment through an event setup that triggers alerts and actions.
 - Identification of changes that might affect availability.
 - Server resource utilization by virtual machine or across a rack of systems.
- ▶ Hardware management:
 - Automated discovery of physical and virtual servers and interconnections, applications, and supported third-party networking.
 - Inventory of hardware components.
 - Chassis and hardware component views:
 - Hardware properties.
 - Component names/hardware identification numbers.
 - Firmware levels.
 - Utilization rates.

- ▶ Network management:
 - Management of network switches from various vendors.
 - Discovery, inventory, and status monitoring of switches.
 - Graphical network topology views.
 - Support for KVM, pHyp, VMware virtual switches, and physical switches.
 - VLAN configuration of switches.
 - Integration with server management.
 - Per-virtual machine network usage and performance statistics that are provided to VMControl.
 - Logical views of servers and network devices that are grouped by subnet and VLAN.
- ▶ Storage management:
 - Discovery of physical and virtual storage devices.
 - Support for virtual images on local storage across multiple chassis.
 - Inventory of the physical storage configuration.
 - Health status and alerts.
 - Storage pool configuration.
 - Disk sparing and redundancy management.
 - Virtual volume management.
 - Support for virtual volume discovery, inventory, creation, modification, and deletion.
- ▶ Virtualization management (base feature set):
 - Support for VMware, Hyper-V, KVM, and IBM PowerVM.
 - Creates virtual servers.
 - Edits virtual servers.
 - Manages virtual servers.
 - Relocates virtual servers.
 - Discovers virtual server, storage, and network resources and visualize the physical-to-virtual relationships.
- ▶ Virtualization management (advanced feature set):
 - Creates new image repositories for storing virtual appliances and discover existing image repositories in your environment.
 - Imports external, standards-based virtual appliance packages into your image repositories as virtual appliances.
 - Captures a running virtual server that is configured the way that you want, complete with a guest operating system, running applications, and virtual server definition.
 - Imports virtual appliance packages that exist in the Open Virtualization Format (OVF) from the Internet or other external sources.
 - Deploys virtual appliances quickly to create new virtual servers that meet the demands of your ever-changing business needs.
 - Creates, captures, and manages workloads.
 - Creates server system pools, which enable you to consolidate your resources and workloads into distinct and manageable groups.

- Deploys virtual appliances into server system pools.
- Manages server system pools, including adding hosts or more storage space, and monitoring the health of the resources and the status of the workloads in them.
- Groups storage systems together by using storage system pools to increase resource utilization and automation.
- Manages storage system pools by adding storage, editing the storage system pool policy, and monitoring the health of the storage resources.
- ▶ I/O address management:
 - Manages assignments of Ethernet MAC and Fibre Channel WWN addresses.
 - Monitors the health of compute nodes, and automatically replaces a failed compute node from a designated pool of spare compute nodes without human intervention.
 - Preassigns MAC addresses, WWN addresses, and storage boot targets for the compute nodes.
 - Creates addresses for compute nodes, saves the address profiles, and deploys the addresses to the slots in the same or different chassis.
- ▶ Additional features:
 - Resource-oriented chassis map provides an instant graphical view of chassis resources, including nodes and I/O modules:
 - A fly-over provides an instant view of an individual server's (node) status and inventory.
 - A chassis map provides an inventory view of chassis components, a view of active statuses that require administrative attention, and a compliance view of server (node) firmware.
 - Actions can be taken on nodes, such as working with server-related resources, showing and installing updates, submitting service requests, and starting the remote access tools.
 - Remote console:
 - Open video sessions and mount media, such as DVDs with software updates, to the servers from local workstation.
 - Remote KVM connections.
 - Remote Virtual Media connections (mount CD/DVD/ISO/USB media).
 - Power operations against servers (Power On/Off/Restart).
 - Hardware detection and inventory creation.
 - Firmware compliance and updates.
 - Automatic detection of hardware failures:
 - Provides alerts.
 - Takes corrective action.
 - Notifies IBM of problems to escalate problem determination.
 - Health status (such as processor utilization) on all hardware devices from a single chassis view.
 - Administrative capabilities, such as setting up users within profile groups, assigning security levels, and security governance.

For more information, see the IBM Flex System Manager product publications available from the IBM Flex System Information Center at this website:

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



Part 2

Chassis Management Module

This part describes how to implement systems management of IBM PureFlex System using the Chassis Management Module.

This part includes the following chapters:

- ▶ Chapter 3, “Planning for Chassis Management Module-based management” on page 33
- ▶ Chapter 4, “Chassis Management Module operations” on page 41



Planning for Chassis Management Module-based management

This chapter describes the systems management capabilities of the Chassis Management Module (CMM). It includes things that you need to take into account when you are planning for CMM-based management in your infrastructure.

Topics that are covered include CMM tasks, management network, CMM configuration interfaces, and options for securing your chassis management components.

This chapter includes the following sections:

- ▶ 3.1, “Chassis Management Module management network” on page 34
- ▶ 3.2, “Chassis Management Module interfaces” on page 35
- ▶ 3.3, “Chassis Management Module security” on page 36
- ▶ 3.4, “Features on Demand planning” on page 40

3.1 Chassis Management Module management network

The internal chassis management network topology for CMM-based deployments is shown in Figure 3-1 as the blue line. It connects CMM to the compute nodes and the switches in the I/O bays. The CMM interfaces with the integrated management module (IMM) or flexible service processor (FSP) integrated in each compute node in the chassis through the management network. The management networks in multiple chassis can be connected together through the external ports of the CMMs in each chassis by using a GbE top-of-rack switch. The yellow line in Figure 3-1 shows the production data network.




Figure 3-1 CMM-based management network topology

The CMM-based management function is limited to the hardware management and alerting capabilities that are provided by the CMMs themselves and IMMs or FSPs. For more information, see 2.2, "Chassis Management Module" on page 15, and 2.3, "Compute node management" on page 17.

3.2 Chassis Management Module interfaces

The Chassis Management Module supports a web-based graphical user interface and command-line interface (CLI). Both the web-based and CLI interfaces are accessible through the single RJ-45 Ethernet connector on the CMM, or from any other system that is connected to the same (management) network.

The CMM has the following default IPv4 settings:

- ▶ IP address: 192.168.70.100
- ▶ Subnet: 255.255.255.0
- ▶ User ID: USERID (all capital letters)
- ▶ Password: PASSWORD (all capital letters, with a zero instead of the letter O)

The CMM does not have a fixed static IPv6 IP address by default. Initial access to the CMM in an IPv6 environment can be done by either using the IPv4 IP address or the IPv6 link-local address. The IPv6 link-local address is automatically generated based on the MAC address of the CMM. By default, the CMM is configured to respond to DHCP first before it uses its static IPv4 address. If you do not want this operation to take place, connect locally to the CMM and change the default IP settings. You can connect locally, for example, by using a mobile computer.

Requirement: Network interfaces on the devices that are connected to the management network must be on a same IP subnet. These devices include CMMs, IMMs, FSPs, and I/O modules.

The web-based GUI brings together all the functions that are needed to manage the chassis elements in an easy-to-use fashion with consistency across all System x IMM2 based platforms.

For more information about how to use the web-based interface to connect to the default CMM address, see 4.1.1, “Connecting to Chassis Management Module” on page 42. The CMM CLI provides direct access to IBM Flex System management functions as an alternative to using the web-based user interface.

Using the CLI, you can issue commands to control the power and configuration of the CMM and other components in an IBM Flex System Enterprise Chassis. The command-line interface also provides access to the text-console command prompt on each compute node through a Serial over LAN (SOL) connection.

You can access the CMM CLI through these connections:

- ▶ A direct serial or Ethernet connection to the CMM
- ▶ A Telnet connection to the IP address of the CMM
- ▶ A Secure Shell (SSH) connection to the CMM

You can initiate connections from the client system by using standard remote communication software. No special programs are required.

You do not need any special hardware to use the CMM command-line interface.

3.3 Chassis Management Module security

Unsecured systems management tools can represent a threat to the hardware and software, and place your data at risk. The CMM offers advanced security capabilities and user management settings to help you secure your environment.

Important: All security and user-management settings are replicated to the IMM that CMM manages.

The following security enhancements and features are provided in the chassis:

- ▶ Single sign-on (centralized user management)
- ▶ End to end audit logs
- ▶ Secure boot: TPM and CRTM
- ▶ Intel TXT technology (Intel Xeon-based compute nodes)
- ▶ Signed firmware updates to ensure authenticity
- ▶ Secure communications
- ▶ Certificate authority and management
- ▶ Chassis and compute node detection and provisioning
- ▶ Role-based access control
- ▶ Security policy management
- ▶ The same management protocols that are supported on BladeCenter AMM for compatibility with earlier versions
- ▶ Insecure protocols are disabled by default in CMM, with “Locks” settings to prevent user from inadvertently or maliciously enabling them
- ▶ Supports up to 84 local CMM user accounts
- ▶ Supports up to 32 simultaneous sessions
- ▶ Planned support for DRTM

3.3.1 Security policies

A CMM security policy is a set of security-related characteristics that define a particular level of protection from security exposures. The CMM security policies include hardware-related communication-protocol controls and account-related access controls.

The CMM offers two levels of security policy: *Legacy* and *Secure*. Security policies are not customizable. However, you can modify the user account policies that the security policies access. For more information, see 3.3.2, “User account policies” on page 37.

An administrator or a user with administrative privileges can use the CMM CLI or web interface to change the security policy settings. For more information about how to configure Security Policies, see “Security policies” on page 58.

Remember: If the security policy settings are changed after the compute nodes are up and running, the security policy status will remain in Pending state. This state persists until the compute nodes in the chassis are restarted.

Secure security policy

The CMM Secure security policy is the most secure and least flexible setting that is available for your configuration. The Secure security policy establishes a more restrictive chassis infrastructure with a higher level of control over users and chassis configuration. It helps secure the chassis environment and enforces the following conditions:

- ▶ Complex password policies for CMM user accounts.
- ▶ Mandatory change of password for all user accounts at first login.
- ▶ Disabling of communication protocols that are not secure: FTP, SNMPv1, Telnet, TFTP, FTP, and non-secure TCP command mode. Only secure communication protocols such as SSH, SSL, and HTTPS are allowed.
- ▶ Certificates to establish secure, trusted connections for applications that run on the management processors.

Restriction: You cannot access the CMM CLI through Telnet while you are using the Secure security policy setting.

Legacy security policy

The CMM Legacy security policy is the least secure and most flexible setting that is available for your configuration.

The Legacy level of management software security policy provides flexibility in managing the chassis infrastructure. It allows the use of the following conditions:

- ▶ Weaker password policies for CMM user accounts
- ▶ No requirement that passwords for user accounts be changed at first login
- ▶ Availability of all communication protocols, both secure and unencrypted (Telnet, SNMP v1, TCP command mode, CIM-XML, FTP, and TFTP).

3.3.2 User account policies

A CMM user account policy is a set of criteria that determines how CMM user account security, including passwords, is implemented.

User account policy conditions affect all users of the CMM. They help enforce the security policy that is chosen for the IBM Flex System Enterprise Chassis environment. For more information, see 3.3.1, “Security policies” on page 36.

The CMM offers two initial user account policy choices: *Legacy* and *High*. You can customize the default values of each of these choices to create a *Custom* user account policy for your IBM Flex System Enterprise Chassis chassis environment.

You can change individual user account policy settings from the default values for each user account policy type. However, the security policy of the CMM might require that specific user account policy settings have secure values. For example, if you attempt to change the CMM security policy level from Legacy to Secure, the CMM might require that you change some user account policy settings. However, if you change the CMM security policy from Secure to Legacy but do not manually modify the user account policy settings, some of them retain their previous secure values.

High user account policy

The CMM user account policy must have a High setting to be used with a CMM that has a Secure security policy.

The High user account policy establishes a higher level of control over users. It provides a more secure chassis environment than the Legacy setting. If the High user account policy is selected, you can override its default values to create a Custom policy. You can do so by using the CMM web interface or the CMM CLI.

Legacy user account policy

When the CMM password policy is configured for use with the Legacy security policy, it allows more flexible, and less secure, accounts.

The Legacy user account policy establishes a lower level of control over users. It provides a less secure chassis environment than the High setting. If the Legacy user account policy is selected, you can override its default values to create a Custom policy. You can do so by using the CMM web interface or the CMM CLI.

Table 3-1 provides some examples of user account policy settings.

Table 3-1 Account policy settings

User account policy setting	Description
User authentication method	The method for authenticating CMM users (local, LDAP, or both)
Maximum simultaneous user sessions	The number of concurrent login sessions that are allowed for each user through all CMM interfaces
Maximum login failures	The maximum number of failed login attempts by a user before the account is locked out
Lockout period login failure	The amount of time a user account is locked out after the maximum number of unsuccessful login attempts is reached
Complex password	Whether the CMM follows more secure complex password rules
Password change on first access	The requirement that users change their password the first time they log in to the CMM
Password expiration period	The amount of time a user password remains valid before it must be changed
Minimum password change interval	The minimum amount of time between user password changes
Password reuse cycle	The number of password changes before a password can be reused

Depending on the initial user account policy you selected, the user account policy settings are configured with different values. For example, the user account policy settings “Complex password” and “Password change on first access” are **On** if you select High user account policy. They are **Off** if you select Legacy user account policy.

For more information about user account policies, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/cmm_password_policies.html

3.3.3 External authentication of certificates

Certificates are used to establish secure, trusted connections to the CMM and from the CMM to other servers.

For an application to trust the server that it is connecting to, it must have the correct certificate in its trust store. This certificate must be a copy of either the server certificate or the certificate of the *certificate authority* (CA) that signed the server certificate. The CMM has a CA that signs certificates for the LDAP, HTTPS, and CIM servers of all systems management processors in the IBM Flex System Enterprise Chassis. You can create trust between your web browser and the HTTPS servers on the management processors in the chassis by importing the CA certificate into your web browser. When you work with an external LDAP server, you can use the CMM web interface or CLI to configure either non-mutual (server only) or mutual certificate authentication.

The CA certificate in each IBM Flex System Enterprise Chassis is unique. Download CA certificates through the primary CMM in each chassis by using the CMM web interface or CLI.

After you download each CA certificate, import it into your web browser. This configuration ensures that the web browser trusts websites that have a certificate that is signed by the CA. If there are multiple users who access the management processors in the IBM Flex System Enterprise Chassis, you can share the CA certificates with the other users. Each user that receives a CA certificate must also import it into their web browser. If your organization has a process for pushing trusted authority certificates to users, you can also use that process.

If you change a CA certificate, you must download the new certificate and import it into the following locations:

- ▶ Your web browser
- ▶ The Certificate Trust Store of your IBM Flex System Manager management software
- ▶ Any IBM Systems Director servers that might be in your network
- ▶ Any external LDAP servers that might be configured for mutual authentication

This process applies for all activities that can change a CA certificate: Manual changes, resetting the CMM to defaults, or restoring a CMM configuration from a backup image.

If your web browser advises you that a connection is untrusted or a security certificate is invalid, or has any other issue that indicates a certificate exception issue relating to a certificate exception, download and import the CA certificate. Make sure to clear all old certificates from the IBM Flex System Enterprise Chassis on all tabs in the certificate pages. You can also try clearing the browser cache. Because some certificate issues affect only certain web browsers, you might be able to correct the condition by switching to a different web browser.

Importing an LDAP certificate with non-mutual authentication

Import a certificate by using non-mutual external authentication when you need to authenticate only the LDAP server with the CMM. You can authenticate the LDAP server with the CMM by using the CMM command-line interface (CLI) or web interface.

Requirement: Certificates must be signed using SHA-1 hashes, SHA-2 hashes are not supported.

Importing an LDAP certificate with mutual authentication

Import certificates for mutual authentication when you need the external LDAP server to authenticate the CMM and the CMM to authenticate the external LDAP server.

There are two ways to establish mutual authentication between the CMM and an external LDAP server. When you use either method, you must also perform the steps for non-mutual authentication.

- ▶ Export the chassis CA certificate and import it into the trust store for your external LDAP server. This process allows mutual authentication between the LDAP server and all elements in the chassis that have their security configuration automatically provisioned.
- ▶ Export a certificate signing request (CSR) from the CMM and have it signed by a certificate authority that the LDAP server already trusts. This method provides mutual authentication between the CMM and the LDAP server.

For more information about external authentication of certificates, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.cmm.doc/cli_ext_cert_authentication.html

3.4 Features on Demand planning

Features on Demand (FoD) is the capability to activate or “unlock” features that are integrated in IBM products. The feature is in the firmware or software, but is “locked” until the activation key is installed.

You can purchase activation keys to activate the FoD for your CMMs, I/O modules, and compute nodes, if your components support these features.

You can use CMM to view activated license keys. For more information, see 4.2.4, “Chassis Management Module Features on Demand” on page 68.

To activate the keys, perform these actions:

- ▶ Do it directly on your I/O modules
- ▶ Do it directly on the compute nodes by using the IMM or Advanced Settings Utility (ASU).
- ▶ Use Flex System Manager to manage all your FoD keys. For more information, see 5.1.3, “Planning for Features on Demand” on page 90.

For a list of available FoD keys for IBM Flex System, see 5.1.3, “Planning for Features on Demand” on page 90.



Chassis Management Module operations

This chapter describes the steps that are required for initial configuration of the Chassis Management Module, and shows how to manage the Enterprise chassis with the CMM.

This chapter includes the following sections:

- ▶ 4.1, “Initial configuration of Chassis Management Module” on page 42
- ▶ 4.2, “Chassis Management Module management tasks” on page 62

4.1 Initial configuration of Chassis Management Module

This section describes how to initially configure the Chassis Management Module to enable chassis management tasks. For more information about CMM capabilities, see 2.2, “Chassis Management Module” on page 15.

The following tasks are described:

- ▶ 4.1.1, “Connecting to Chassis Management Module” on page 42
- ▶ 4.1.2, “Configuring Chassis Management Module by using Initial Setup Wizard” on page 44
- ▶ 4.1.3, “Preparing for Chassis Management Module redundancy” on page 56
- ▶ 4.1.4, “Configuring Chassis Management Module user authority” on page 58
- ▶ 4.1.5, “Restoring a Chassis Management Module” on page 62

4.1.1 Connecting to Chassis Management Module

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

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

By default, the CMM is configured to respond to DHCP first before it uses its static IP address.

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

To connect to the CMM, perform the following steps:

1. Make sure that the subnet of the client computer is set to the same value as in the CMM (the default CMM subnet is 255.255.255.0). The IP address of the CMM must also be in the same local domain as the client computer. To connect to the CMM for the first time, you might have to change the Internet Protocol properties on the client computer.
2. Open a web browser on the client computer, and direct it to the CMM IP address. For the first connection to the CMM, use the default IP address of the CMM, as shown in Figure 4-1.




Figure 4-1 Log in to the Chassis Management Module with the default IP address

Clarification: The Chassis Management Module has the following default settings:

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

3. Log in to the CMM by using the default credentials: USERID/PASSWORD. Click **Log In** as shown in Figure 4-2.



Figure 4-2 CMM login

The Chassis Management Module main window opens as shown in Figure 4-3.

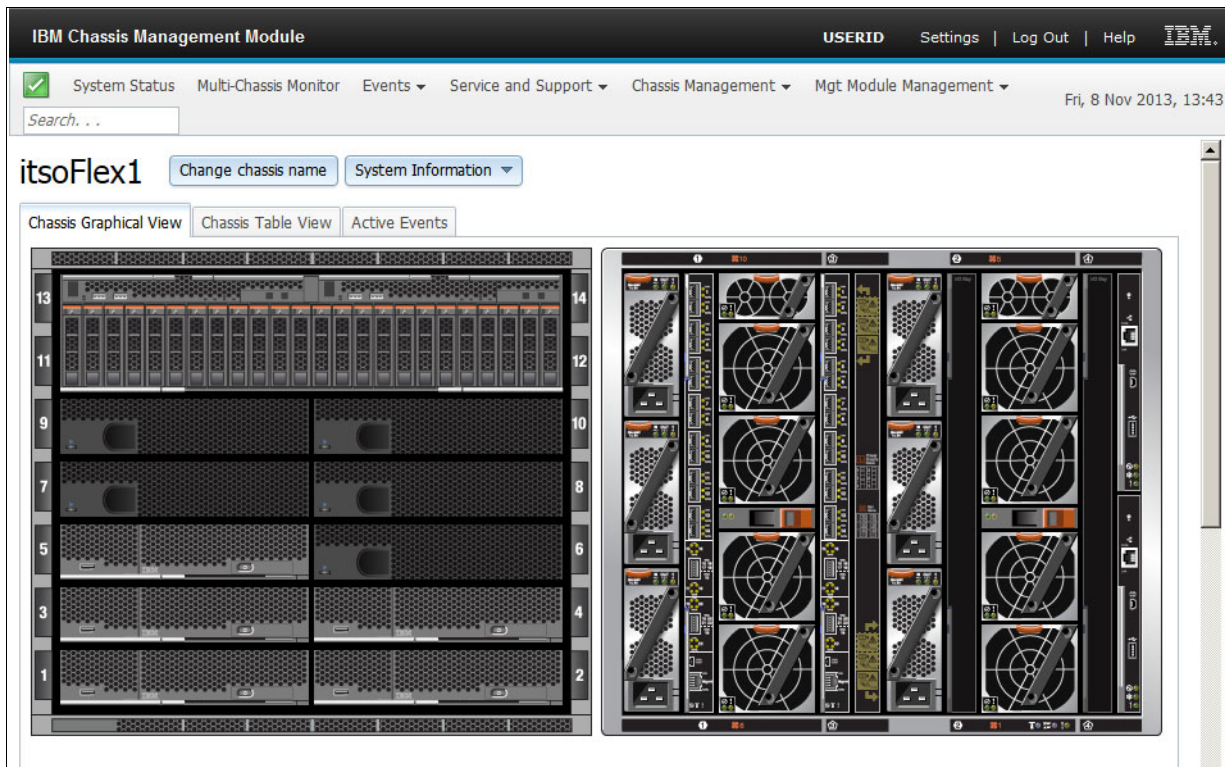


Figure 4-3 CMM main window

4.1.2 Configuring Chassis Management Module by using Initial Setup Wizard

The next step is the initial configuration of the Chassis Management Module. The initial setup wizard can help you configure the CMM through a web interface. The wizard starts automatically when you first access the web interface of a new CMM or a CMM that has been reset to its default settings.

Follow these steps to manually start the Initial Setup Wizard and perform the initial configuration:

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

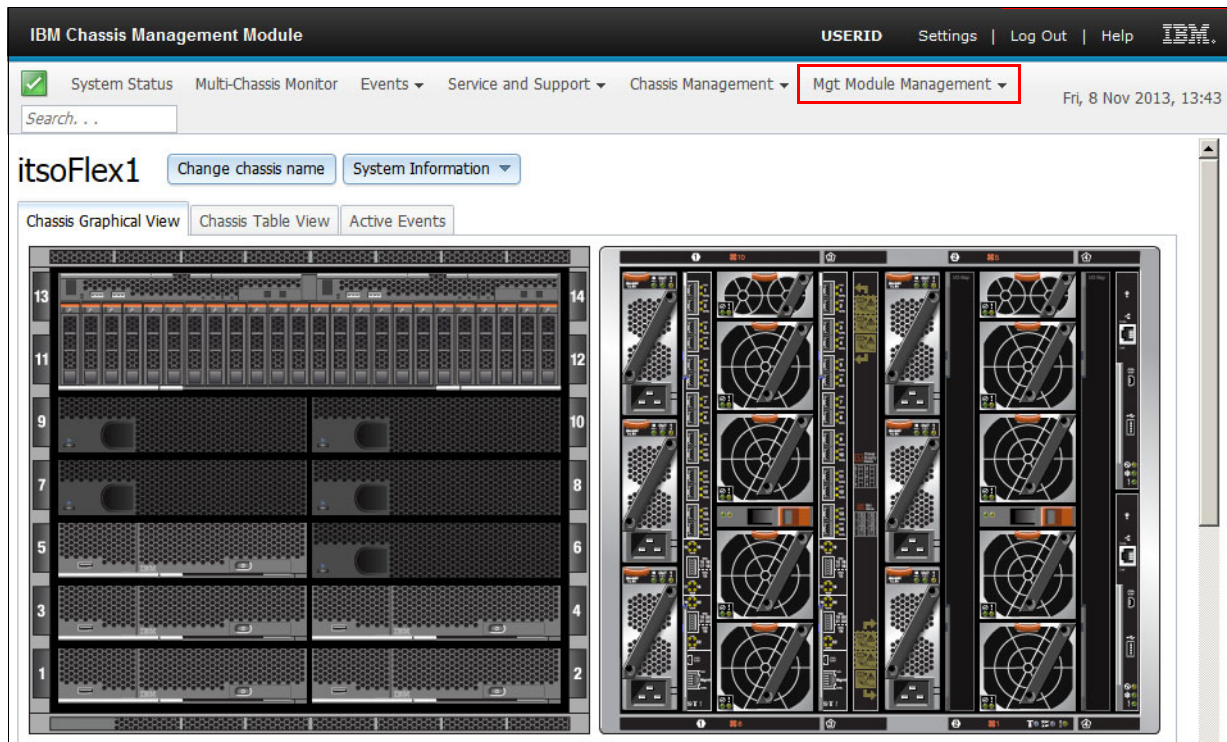
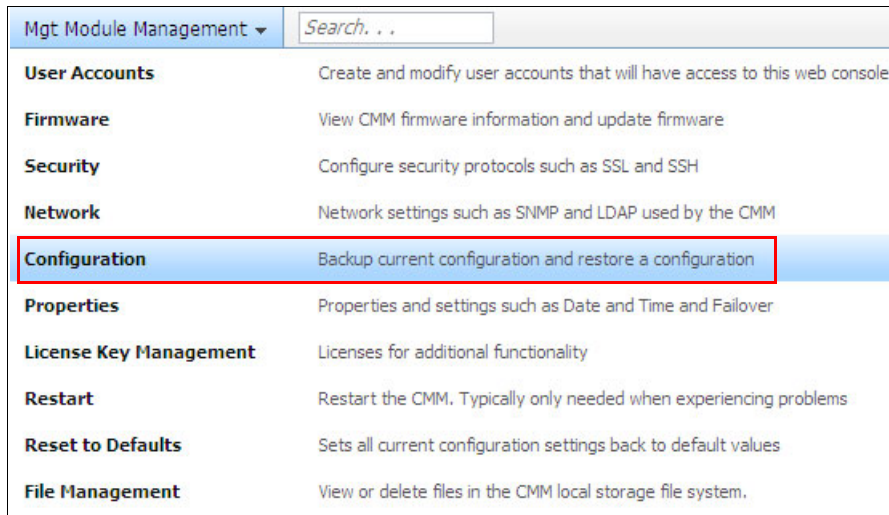


Figure 4-4 CMM main window: Mgt Module Management

The initial setup wizard is contained in the **Configuration** menu, as shown in Figure 4-5.



Mgt Module Management ▾	Search. . .
User Accounts	Create and modify user accounts that will have access to this web console
Firmware	View CMM firmware information and update firmware
Security	Configure security protocols such as SSL and SSH
Network	Network settings such as SNMP and LDAP used by the CMM
Configuration	Backup current configuration and restore a configuration
Properties	Properties and settings such as Date and Time and Failover
License Key Management	Licenses for additional functionality
Restart	Restart the CMM. Typically only needed when experiencing problems
Reset to Defaults	Sets all current configuration settings back to default values
File Management	View or delete files in the CMM local storage file system.

Figure 4-5 Mgt Module Management window

Several options are displayed for managing the Chassis Management Module configuration.

2. For the first time connection, click **Initial Setup Wizard** as shown in Figure 4-6.



Figure 4-6 Manage Configuration window

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

Figure 4-7 shows the Welcome window of the Initial Setup Wizard. Navigation buttons for the wizard are in the lower left corner of each window. Click **Next**.




Figure 4-7 Welcome window

a. The Inventory and Health window shows the detected components and their current health status, as shown in the Figure 4-8. Click **Next**.




Figure 4-8 Inventory and Health window

- b. If you have saved a configuration file, the Import Existing Configuration window allows you to upload this file to the CMM, as shown in Figure 4-9.

Import Existing Configuration

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

If this is your first time setting up a chassis, you will not have a configuration file to import. These files are useful as a backup of your management module settings, or for configuring multiple chassis. To create a configuration file, you can use the main console under Mgt Module Management -> Configuration.

i Some restore operations may cause a temporary loss of web connectivity. Under these circumstances, the final confirmation popup and restore log may not be available. If web connectivity is lost, clear the browser cache (Ctrl+F5) and restart your session. At this point, check the event log for messages related to the configuration restore operation.

Passphrase:

Confirm pass:

Upload configuration file: **Browse for file**

< Back Next > Finish Cancel

Figure 4-9 Import Existing Configuration window

4. The General Settings window prompts you to enter some descriptive information about the chassis, including location and contact person, as shown in the Figure 4-10. Click **Next**.

General Settings

General settings for the chassis and management module

Management module name:

Chassis description:

Contact person:

Chassis location:

Room ID:

Rack ID:

Lowest U-position:

Unit height of chassis:

< Back Next > Finish Cancel

Figure 4-10 General Settings window

- c. Set the date and time for the CMM in the Date and Time window, as shown in the Figure 4-11. There are two options to sync the time: Using NTP or setting manually. Click **Next**.

The screenshot shows the 'Date and Time' configuration window. On the left is a navigation pane with options: Welcome, Inventory and Health, Import Existing Configuration, General Settings, Date and Time (selected), IP Configuration, IO Modules, Security Policy, DNS, Event Recipients, and Confirm. The main area is titled 'Date and Time' and contains the following settings:

- Select method:** Synchronize with an NTP server (dropdown)
- NTP server host name and/or IP address:** 9.42.170.223 (text input)
- Synchronization frequency (minutes):** 20 (spin box)
- Enable NTP v3 Authentication:**
- NTP v3 Authentication key index:** 2 (text input)
- NTP v3 Authentication key (M - MD5):** 3291FC94 (text input)
- Status:** NTP last updated the clock on 10/03/2013 19:35:05 by 0 s. The last 1704 update attempt(s) have failed.
- GMT Offset:** -5:00 - Eastern Standard Time (Eastern USA, Ontario, Quebec) (dropdown)
- Unable to automatically determine the daylight saving time to use. Please provide the DST scheme.**
- Selected GMT offset:** -5:00 - Eastern Standard Time (Eastern USA, Ontario, Quebec)
- Available schemes:** USA and Canada (dropdown)
- Automatically adjust for daylight savings time (DST):**

At the bottom are four buttons: < Back, Next >, Finish, and Cancel.

Figure 4-11 Date and Time window

- d. Each CMM is configured with the same static IP address. You must create a unique static IP address for each CMM. If DHCP is not used, only one CMM at a time can be added onto the network for discovery. Adding more than one CMM to the network without a unique IP address assignment for each results in IP address conflicts.

Figure 4-12 shows the IP configuration window.

The screenshot shows the 'IP Configuration' window for the management module. The left sidebar contains a list of configuration categories: Welcome, Inventory and Health, Import Existing Configuration, General Settings, Date and Time, IP Configuration (selected), IO Modules, Security Policy, DNS, Event Recipients, and Confirm. The main content area is titled 'IP Configuration' and 'IP configuration for the management module'. It includes fields for 'Host name' (CMM) and 'Domain name'. A checkbox for 'Register this interface with DNS' is present. Below this, there are tabs for 'IPv4' and 'IPv6'. The 'IPv4' tab is active, showing 'Currently assigned IPv4 address information' with fields for IP address (9.42.170.215), Subnet mask (255.255.254.0), and Default gateway (9.42.170.1). A dropdown menu for 'IP address assignment methods' is set to 'Use static IP address'. Below that, 'Static IP Address Settings' are shown with fields for Static address (9.42.170.215), Subnet mask (255.255.254.0), and Default gateway (9.42.170.1). A note states '*Changing settings requires a CMM restart.' At the bottom, there are buttons for '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 4-12 IPv4 configuration window

e. If you need to set up IPv6, you can use IPv6, as shown in the Figure 4-13. Click **Next**.

The screenshot shows the 'IP Configuration' window for the management module, similar to Figure 4-12. The 'IPv6' tab is active. It shows 'Enable IPv6' checked. Below this, it displays 'Link local address: fe80::5ef3:fcff:fe25:ed85', 'Stateless address: None assigned', and 'Default gateway: 0::0'. The 'Stateful address' field is empty. Under 'IP address assignment methods', there are three options: 'Use stateless address autoconfiguration' (unchecked), 'Use stateful address configuration (DHCPv6)' (unchecked), and 'Use statically assigned IP address' (unchecked). The bottom navigation buttons are '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 4-13 IPv6 configuration window

- f. You can view the status and configure the options for the I/O modules that are connected to the CMM, as shown in Figure 4-14. Click **Next**.




Figure 4-14 I/O Modules window

- g. Select the security policy for your CMM as shown in Figure 4-15. Click **Next**.

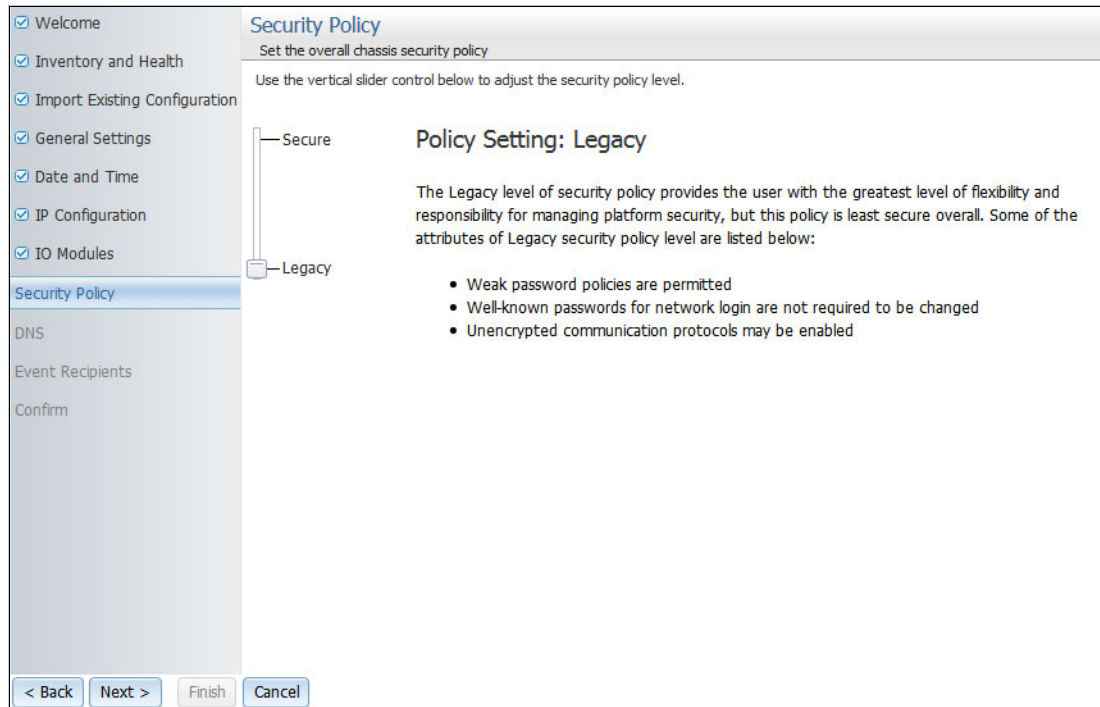


Figure 4-15 Security Policy window

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

For more information about security policies, see 3.3, “Chassis Management Module security” on page 36.

5. Select the appropriate Domain Name Server (DNS) options for your CMM, as shown in the Figure 4-16. Click **Next**.

The screenshot shows the DNS configuration window. The left sidebar contains a list of configuration steps, with 'DNS' selected. The main area is titled 'DNS' and 'Relevant IP addresses for Domain Name Server (DNS)'. It features three settings: 'Enable DNS' (unchecked), 'Preferred DNS address type: IPv4' (dropdown menu), and 'Send DDNS updates to these servers' (unchecked). At the bottom, there are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Figure 4-16 DNS setup window

- h. Enter the email addresses where notifications are to be sent as CMM events occur, as shown in Figure 4-17. Click **Next**.

Figure 4-17 Event Recipients window

- i. Confirm all of the information that has been entered in the setup wizard, as shown in the Figure 4-18. Click **Finish**.

Figure 4-18 Confirm window

Updating the Chassis Management Module firmware

This section explains how to manage the Chassis Management Module firmware.

Some IBM Flex System solutions require specific code levels or coordinated code updates. If the CMM is part of one of these solutions, verify that the level of code is supported for the solution before you update the code.

If your IBM Flex System Enterprise Chassis is configured for redundant operations and the second CMM is installed, both will have the same level of firmware after the primary CMM pushes the updates to the standby CMM. The latest level of CMM firmware is available at this website:

<http://www-947.ibm.com/support/entry/portal/Downloads?lnk=mhsd>

To update the CMM firmware, follow these steps:

1. In the CMM web interface, select **Firmware** from the **Mgt Module Management** menu, as shown in Figure 4-19.

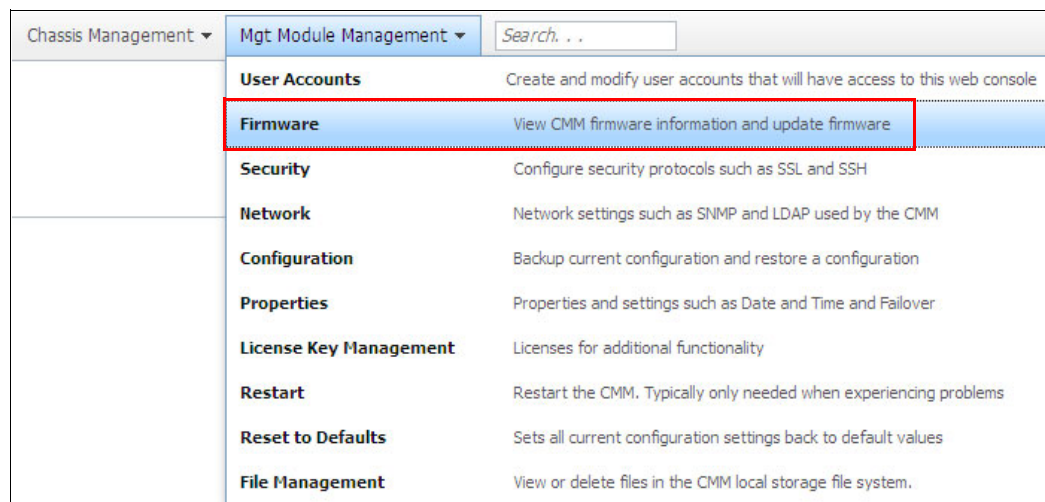


Figure 4-19 Mgt Module Management menu

2. Check the current firmware level on the CMMs as shown in Figure 4-20.

The screenshot shows the 'Firmware' page with the following table:

Bay	Name	Firmware Type	Build ID	File Name	Release Date	Revision	Role
1	cmm1	CMM firmware	2PET12D	cmefs.uxp	08/28/2013	12	Primary
2	Standby MM	CMM firmware	2PET12D	cmefs.uxp	08/28/2013	12	Standby

Figure 4-20 Checking the current firmware level

3. Click **Update Firmware**.

4. The update Firmware window opens as shown in Figure 4-21. Proceed through each step of the wizard by clicking **Next** and entering the information as required.




Figure 4-21 Select File Source window

5. Click **Browse**, and select CMM firmware file in your local directory, as shown in the Figure 4-22. Click **Next**.




Figure 4-22 Select Firmware File window

6. Select the behavior that you want after updating the CMM firmware, as shown in Figure 4-23:
- Restart CMM manually
 - Restart CMM automatically after updating




Figure 4-23 Choose Post Update Behavior window

Click **Next**. Monitor the firmware update progress as shown in Figure 4-24.




Figure 4-24 Performing Firmware Flash update progress window

Verify that update is completed, as shown in the Figure 4-25. Click **Finish** to restart the CMM or to go back to the CMM management interface if you chose to restart the CMM manually.

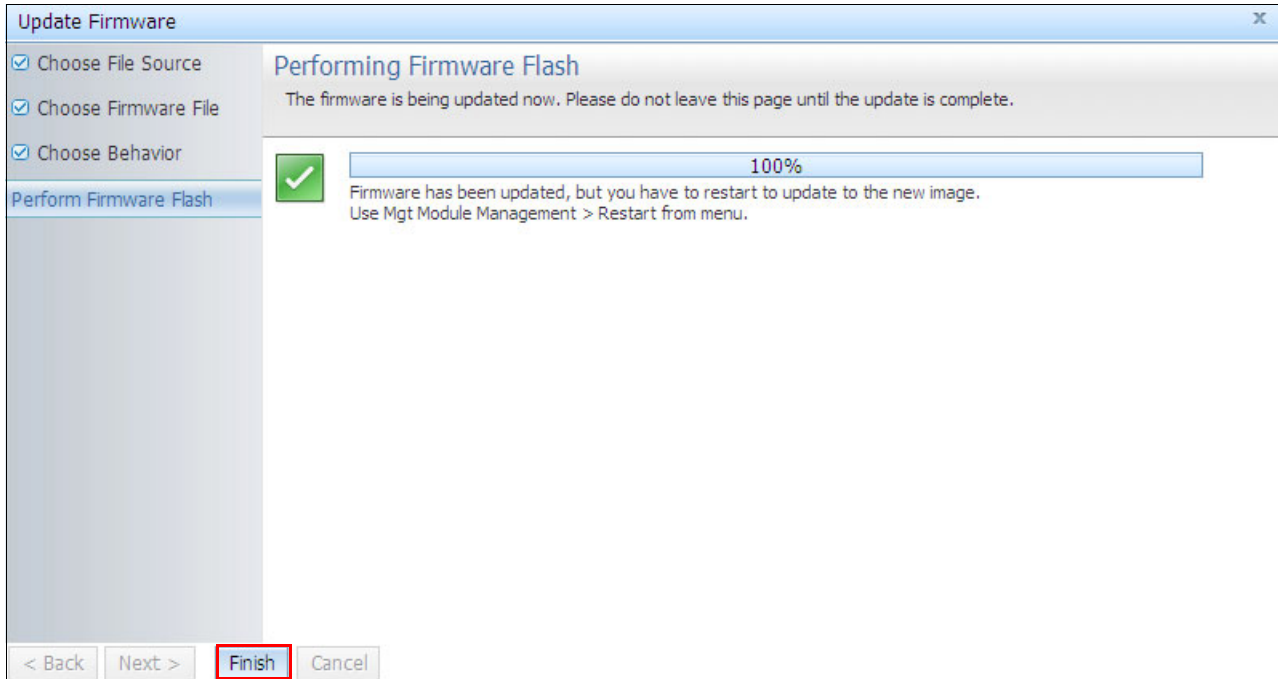


Figure 4-25 Performing Firmware Flash update status window

4.1.3 Preparing for Chassis Management Module redundancy

This section addresses how to set up fail over between two Chassis Management Modules.

To prepare your CMM for redundancy, complete the following steps:

1. Install the standby CMM in the available CMM bay.
2. Wait approximately 2 minutes while the primary CMM transfers the firmware and configuration information to the standby CMM.

Clarification: Whenever power is restored to an IBM Flex System Enterprise Chassis that has two functional CMMs, the CMM in CMM bay 1 is designated as the primary CMM. This process occurs even if the CMM in CMM bay 2 was the primary CMM before power was removed.

3. Configure Chassis Management Module failover response.

There are two options to configure CMM failover. The first case is that the hardware failure or some malfunction resulted in failure of the primary CMM. The second case is the primary CMM operates properly, but there is a network problem. For example, a network switch might go down, resulting in loss of connectivity to the primary CMM.

To configure failover response for the loss of the primary CMM:

1. Select **Properties** from the **Mgt Module Management** menu to reach the Management Module Properties window.

2. Click the **Advanced Failover** tab in the Management Module Properties window, as shown in Figure 4-26.

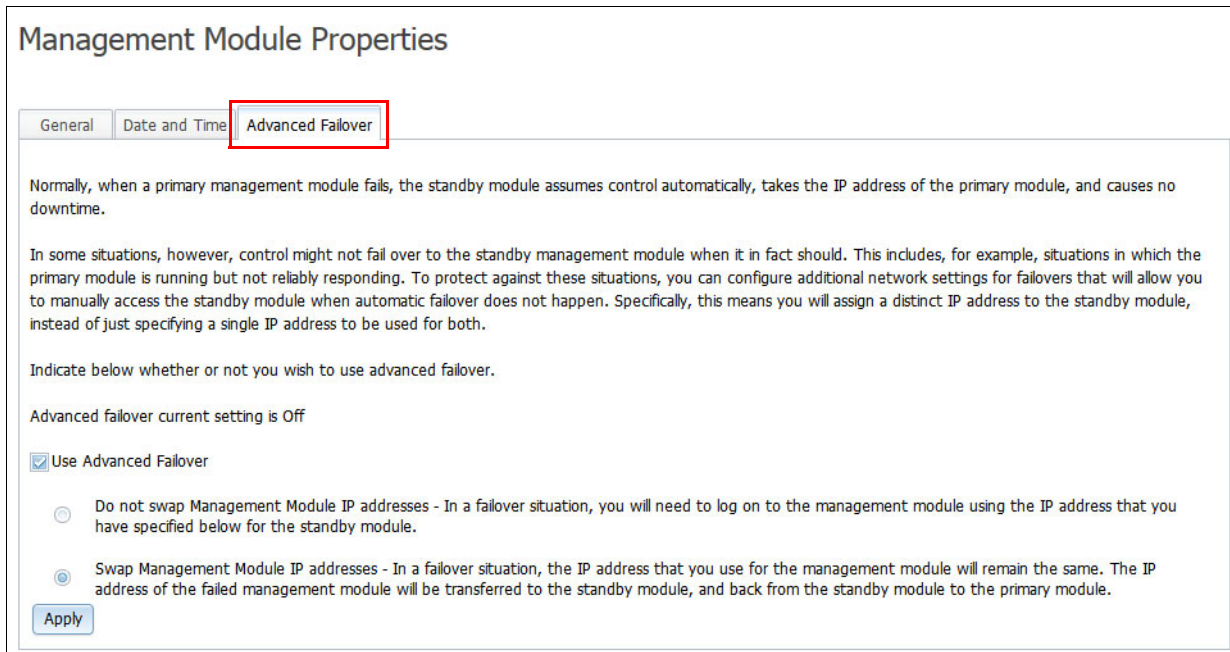


Figure 4-26 Failover menu

3. Check **Use Advanced Failover** check box, specify CMM IP address behavior, and click **Apply**.

To configure failover response for loss of the management network (uplink) connection to the primary CMM:

1. In the CMM web interface, click the **Ethernet** tab in the **Network Protocol Properties** window.
2. Select **Network** from the **Mgt Module Management** menu, then click **Advanced Ethernet** tab, as shown in Figure 4-27.

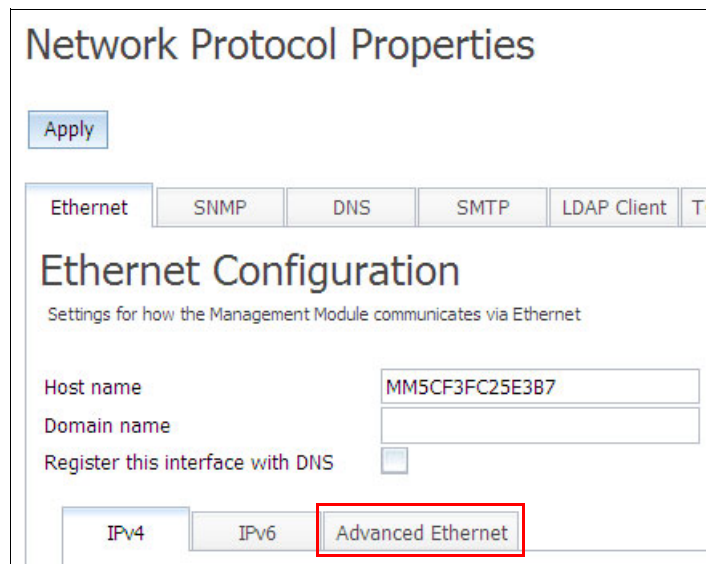


Figure 4-27 Network Protocol Properties window

Define the failover condition based on your requirements, as shown in Figure 4-28.

The screenshot shows the 'Advanced Ethernet' configuration window. It includes tabs for IPv4, IPv6, and Advanced Ethernet. The Advanced Ethernet tab is active. The configuration options are as follows:

Property	Value
Data rate	Auto
Duplex	Auto
MAC address type	Use burned in MAC address
Burned-in MAC address	5C:F3:FC:25:E3:B7
Maximum Transmission units	1,500
Failover on loss of physical network link	<input checked="" type="checkbox"/>
Failover delay for physical link loss (seconds)	60
Failover on loss of logical network link	<input type="checkbox"/>
Failover delay for logical link loss (seconds)	1,800
IPv4 address for logical link check	0.0.0.0
IPv6 address for logical link check	0::0

Use Alert and failover if:

- either IPv4 or IPv6 link check fails
- both IPv4 and IPv6 link checks fail

Figure 4-28 Network Failover setup menu

4.1.4 Configuring Chassis Management Module user authority

Configuring Chassis Management Module user authority requires you to set security policies and user account policies.

Security policies

The IBM Flex System Enterprise Chassis takes a new approach to security with a ground-up Chassis management design to meet the new Trusted Computing Group (TCG) security standards.

The Enterprise Chassis ships with secure settings by default, with two security policy settings supported: *Secure* and *Legacy*. For more information about each security policy, see 3.3.1, “Security policies” on page 36.

The centralized security policy makes the Enterprise Chassis easy to configure. In essence, all components run the same security policy that is provided by the Chassis Management Module.

To set the specific security policy level, perform these steps:

1. In the CMM web interface, select **Security** from the **Mgt Module Management** menu, as shown in Figure 4-29.




Figure 4-29 Mgt Module Management: Security

2. In the Security Policies window, use the slider bar to select **Secure**, and click **Apply** as shown in Figure 4-30.

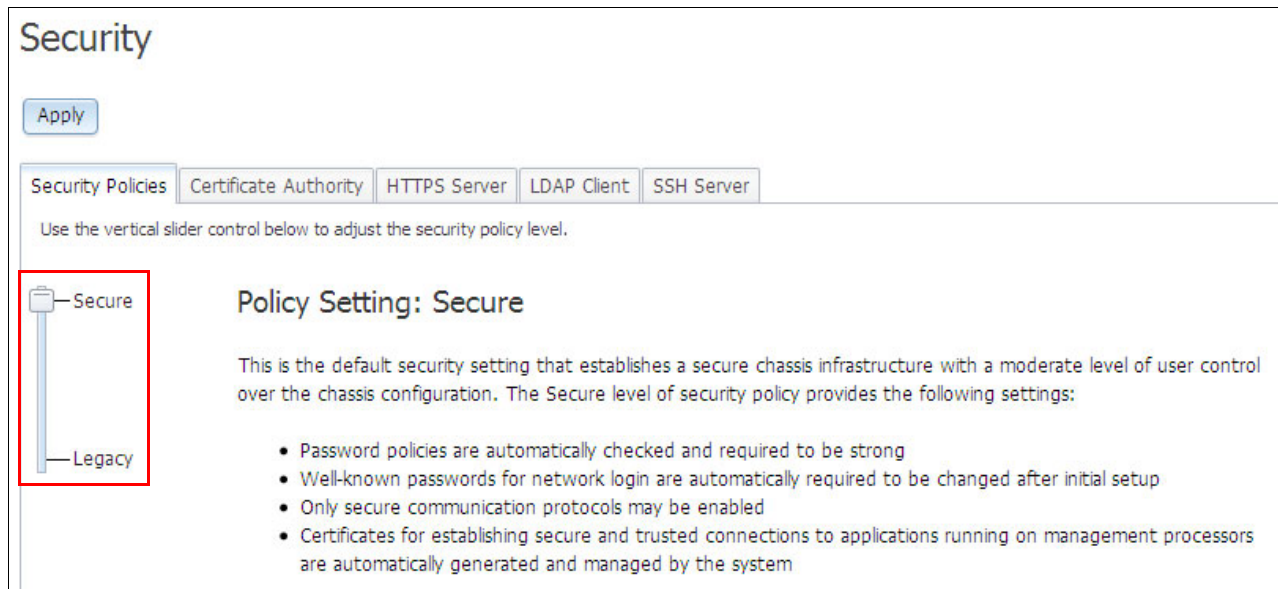


Figure 4-30 Security window

User account policies

A CMM user account policy is a set of criteria that determine how CMM user account security, including passwords, is implemented.

User account policy conditions affect all users of the CMM. They help enforce the security policy that is chosen for the IBM Flex System Enterprise Chassis environment.

The CMM offers two initial user account policy choices: Legacy and High. You can customize the default values of each of these choices to create a Custom user account policy for your IBM Flex System Enterprise Chassis environment.

Remember: You can change individual user account policy settings from the default values for each user account policy type. However, the security policy of the CMM might require that specific user account policy settings have secure values. For example, if you attempt to change the CMM security policy level from Legacy to Secure, the CMM might require that you change some user account policy settings to secure values. However, if you change the CMM security policy from Secure to Legacy but do not manually modify the user account policy settings, some retain their previous secure values.

In the CMM web interface, user account security policy settings are on the **General** tab of the Account Security Level window in the Global Login Settings window. To set the user account policies, perform these steps:

1. Select **User Accounts** from the Mgt Module Management menu, as shown in Figure 4-31.

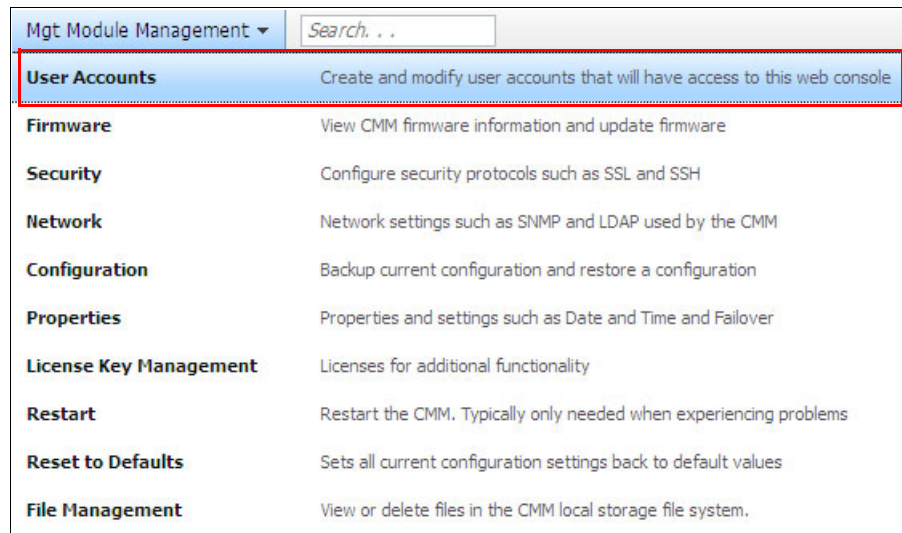


Figure 4-31 Mgt Module Management: User Accounts

2. Click **Global Login Settings** on the Accounts tab in the User Accounts window, as shown in Figure 4-32.

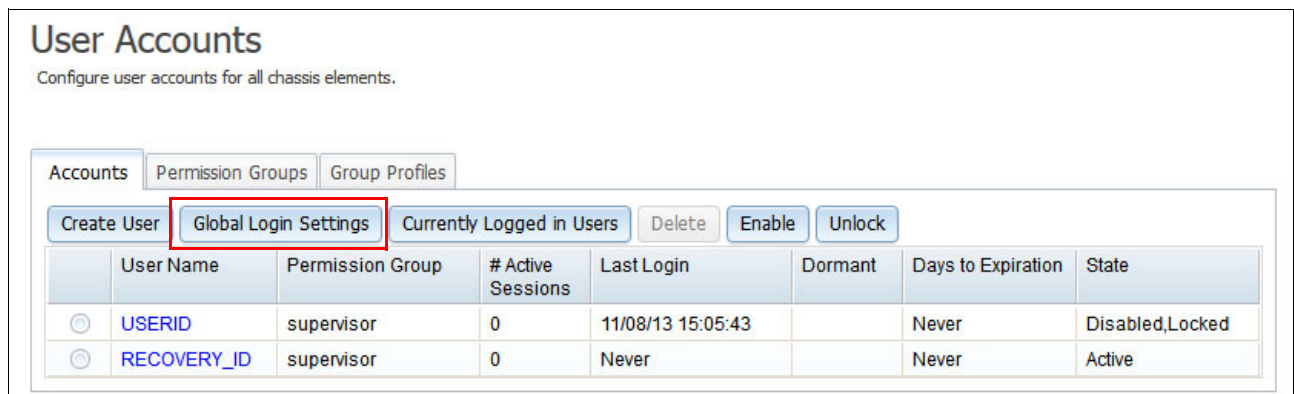


Figure 4-32 User account window

- Click the **Account Security Level** tab in the Global Login Settings window, as shown in Figure 4-33.

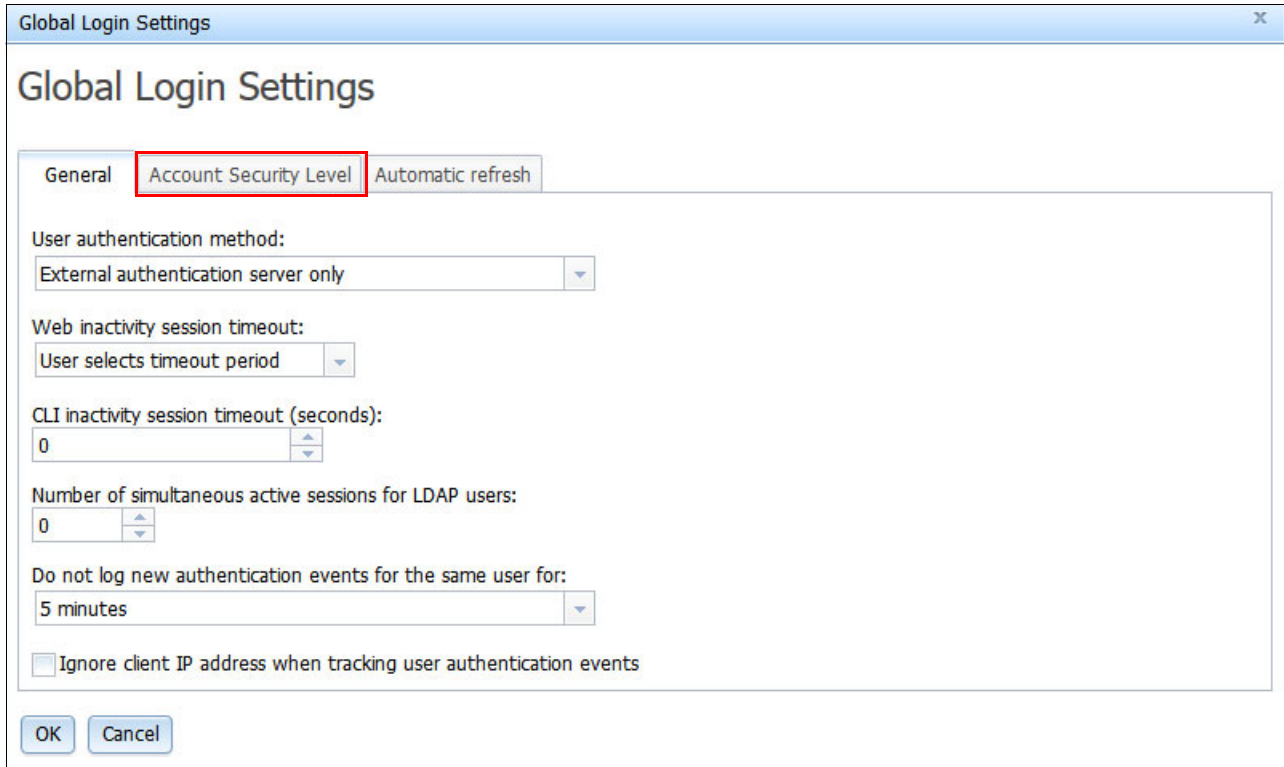


Figure 4-33 Global Login Settings window

- Select **Custom Security Settings, High Security Settings, or Legacy Security Settings** from the menu and click **OK**, as shown in Figure 4-34.

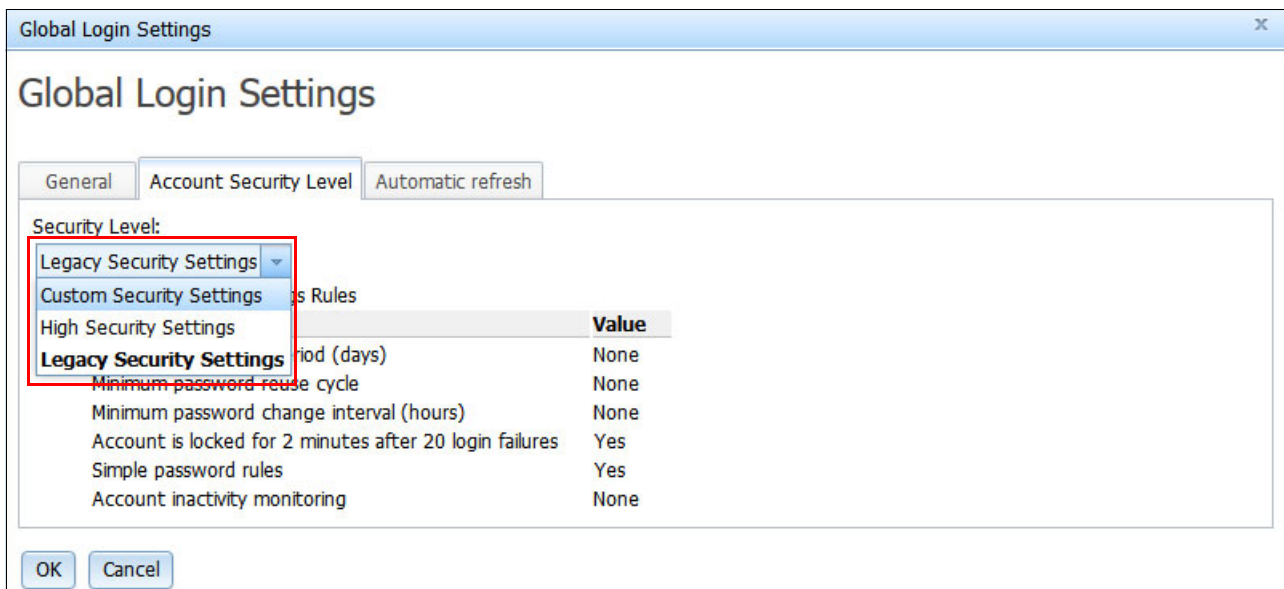


Figure 4-34 Global Login Settings: Security Settings

For more information about the CMM security, see 3.3, “Chassis Management Module security” on page 36.

4.1.5 Restoring a Chassis Management Module

Restoring a CMM differs depending on whether you are restoring the default configuration or restoring from a saved configuration.

Restoring the CMM manufacturing default configuration

You can restore the CMM to its manufacturing default configuration in these ways:

- ▶ In the CMM web interface, select **Reset to Defaults** from the Mgt Module Management menu, as shown in Figure 4-35.




Figure 4-35 Mgt Module Management: Reset to Defaults

- ▶ If you have physical access to the CMM, push the reset button and hold it for approximately 10 seconds.

Restoring a saved CMM configuration

In the CMM web interface, a saved configuration is applied from the Manage Configuration window, as shown in Figure 4-36. Select **Configuration** from the Mgt Module Management menu to open the Manage Configuration window.




Figure 4-36 Manage Configuration window

4.2 Chassis Management Module management tasks

This section addresses how to use the CMM to run specific systems management tasks:

- ▶ 4.2.1, “Monitoring the chassis” on page 63
- ▶ 4.2.2, “Monitoring multiple chassis” on page 64
- ▶ 4.2.3, “Event notifications” on page 65

- ▶ 4.2.4, “Chassis Management Module Features on Demand” on page 68
- ▶ 4.2.5, “Chassis management” on page 70
- ▶ 4.2.6, “Using the Chassis Management Module CLI” on page 80

4.2.1 Monitoring the chassis

The System Status window is the default window when you enter the CMM web interface, as shown in Figure 4-37. You can also access it by clicking **System Status** on the menu bar.




Figure 4-37 Chassis status

The graphical view of the Chassis is active, so the changes are reflected immediately. The following selections are available (matching callouts in Figure 4-37):

1. **Selected Active Resource:** All major components of the Chassis can be clicked for more information. Select a component of interest (in Figure 4-37, the IBM Flex System p260 Compute Node), and a pop-up displays information about that component, such as serial number, name, or bay. You can power a component on or off from this window by right-clicking, or view other details about the component.
2. **Selected Active Resource Details:** With a component selected, this box displays several tabs for additional information, such as events, hardware, firmware, and LEDs.
3. **Selected Resource Actions:** I/O modules and compute nodes activate the Actions menu, from which you can power on/off, restart, and perform other tasks.

4.2.2 Monitoring multiple chassis

You can view multiple networked Chassis from the CMM web interface, as shown in the Figure 4-38.


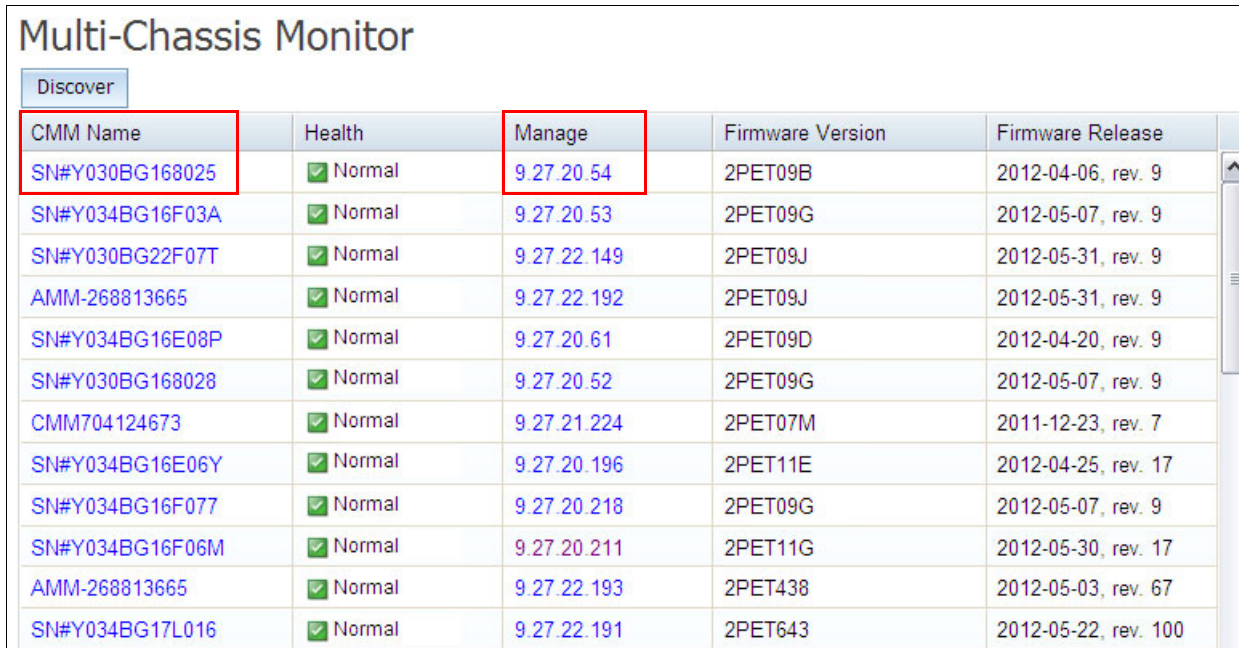


Figure 4-38 Selecting Multi-Chassis Monitor

Use the Multi-Chassis Monitor window to view the state of all compute nodes and management nodes in multiple networked chassis from one location. Monitor the status of installed compute nodes and management nodes in multiple chassis, and discover recently connected ones in the Multi-Chassis Monitor window, as shown in Figure 4-39.



The image shows the 'Multi-Chassis Monitor' window. At the top left is a 'Discover' button. Below it is a table with the following columns: 'CMM Name', 'Health', 'Manage', 'Firmware Version', and 'Firmware Release'. The first row of the table is highlighted with a red box. The table contains 13 rows of data.

CMM Name	Health	Manage	Firmware Version	Firmware Release
SN#Y030BG168025	✓ Normal	9.27.20.54	2PET09B	2012-04-06, rev. 9
SN#Y034BG16F03A	✓ Normal	9.27.20.53	2PET09G	2012-05-07, rev. 9
SN#Y030BG22F07T	✓ Normal	9.27.22.149	2PET09J	2012-05-31, rev. 9
AMM-268813665	✓ Normal	9.27.22.192	2PET09J	2012-05-31, rev. 9
SN#Y034BG16E08P	✓ Normal	9.27.20.61	2PET09D	2012-04-20, rev. 9
SN#Y030BG168028	✓ Normal	9.27.20.52	2PET09G	2012-05-07, rev. 9
CMM704124673	✓ Normal	9.27.21.224	2PET07M	2011-12-23, rev. 7
SN#Y034BG16E06Y	✓ Normal	9.27.20.196	2PET11E	2012-04-25, rev. 17
SN#Y034BG16F077	✓ Normal	9.27.20.218	2PET09G	2012-05-07, rev. 9
SN#Y034BG16F06M	✓ Normal	9.27.20.211	2PET11G	2012-05-30, rev. 17
AMM-268813665	✓ Normal	9.27.22.193	2PET438	2012-05-03, rev. 67
SN#Y034BG17L016	✓ Normal	9.27.22.191	2PET643	2012-05-22, rev. 100

Figure 4-39 Monitoring multiple chassis

Clicking the name of the CMM in the **CMM Name** column shows the managed resources in the selected chassis, as shown in Figure 4-40.

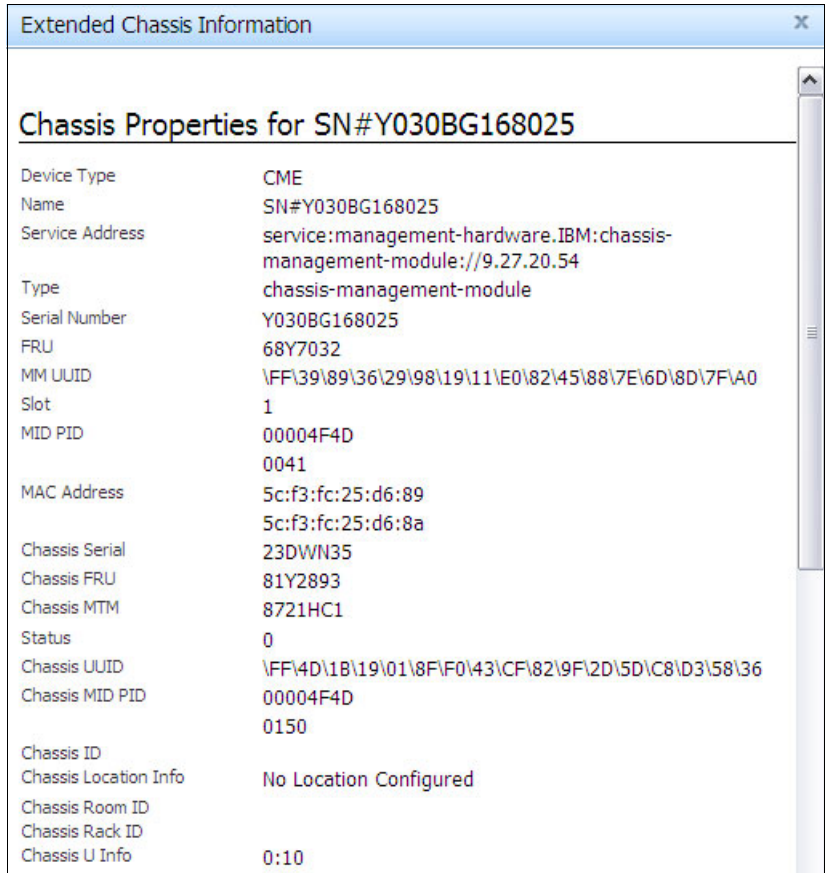


Figure 4-40 Extended Chassis Information window

Clicking the IP address of the chassis in the **Manage** column (Figure 4-39 on page 64) redirects you to another Chassis Management Module’s web interface.

4.2.3 Event notifications

Click **Event Log** in the Events menu to check events, as shown in the Figure 4-41.




Figure 4-41 Events: Event Log

The CMM event log contains a list of all events that are received from all devices in the chassis, as shown in Figure 4-42. These events are also sent by the CMM to the IBM Flex System Manager, if one is installed.

The screenshot shows the 'Event Log' window with the following table of events:

Severity	Source	Sequence #	Date	Event ID	Message
Informational	Audit	00001EB9	Today 02:53:02 PM	00285000	The name of Chassis Management Module in CMM Bay 01 was changed to USERID from Web at IP address 9.44.157.218.
Informational	Audit	00001EB8	Today 02:40:38 PM	00285000	The name of Chassis Management Module in CMM Bay 01 was changed to USERID from Web at IP address 9.44.157.218.
Informational	Audit	00001EB7	Today 01:43:33 PM	0000007A	Login successful. User ID USERID from Web at IP address 9.44.157.218.
Informational	Audit	00001EB6	Today 01:38:53 PM	0001601A	Logoff successful. User ID USERID from Web at IP address 9.44.157.218.
Informational	SERVPROC	00001EB5	Nov 5, 2013 04:46:55 PM	00016805	Service data collection completed on CMM SN#Y011BG24H0BB.
Informational	IOMod_03	00001EB4	Nov 5, 2013 04:45:41 PM	00038F03	I/O module 3 communication is online.
Informational	Audit	00001EB3	Nov 5, 2013 04:45:37 PM	00016804	Service data collection initiated on CMM SN#Y011BG24H0BB by user ID US address 9.42.170.223.

Figure 4-42 Event Log window

You can see general information about the event, including severity, source, sequence, date, and event message, as shown in Figure 4-42. In addition, several options are available to manage logs:

- ▶ The Export option allows you to export your event log in various formats (csv, XML, or pdf).
- ▶ Use the Delete Events option to delete all selected items, with the additional option of selecting audit, systems, or both.
- ▶ With the Settings option, you can add a log event when a log is 75% full.
- ▶ The Open Service Request option is enabled when you select one of the events from the table.

To configure event recipient notifications in the CMM web interface, open the **Events** menu and click **Event Recipients** as shown in Figure 4-43.

The screenshot shows the 'Event Recipients' window with the following table structure:

Name	Notification Method	Events to Receive	Status
No Data Available			

Figure 4-43 Event Recipients window

Select the **E-mail Recipient** as shown in Figure 4-44.

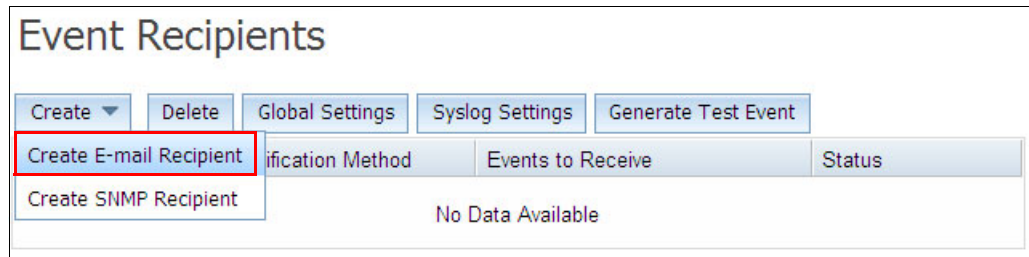


Figure 4-44 Event Recipients window continued

Enter the recipient name and email address, select which events to receive, and click **OK**, as shown in Figure 4-45.

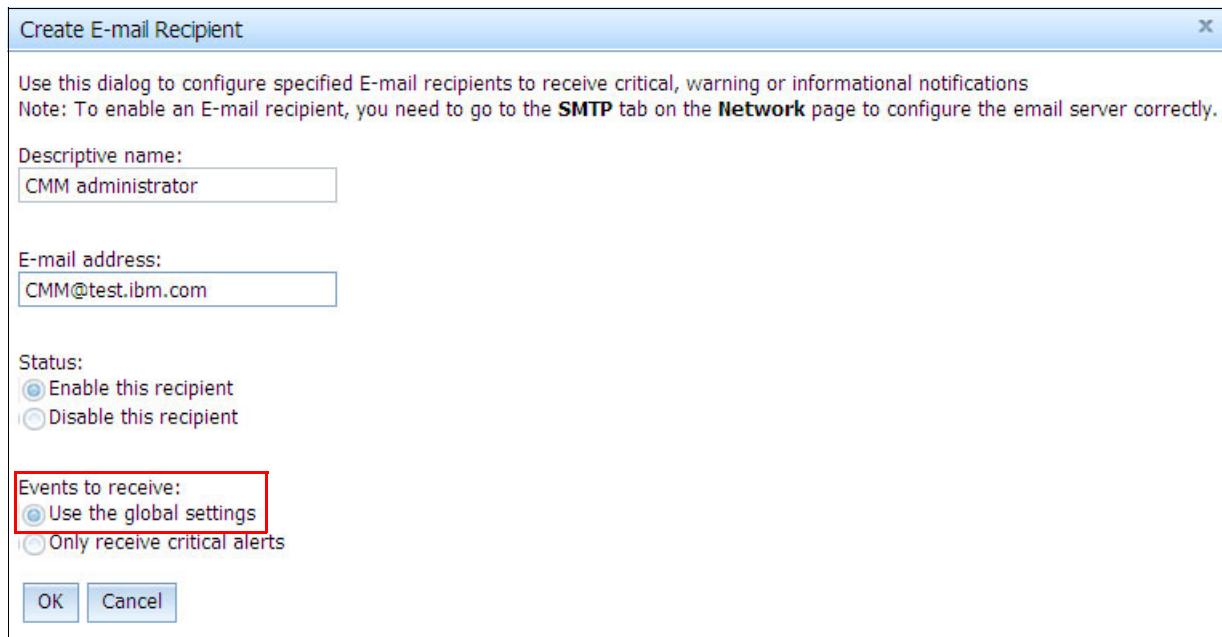


Figure 4-45 Create E-mail Recipient window

Click **Global Settings** in the Event recipients window and select the type and severity of the alerts to be sent, as shown in Figure 4-46.




Figure 4-46 Event Recipient Global Settings window

4.2.4 Chassis Management Module Features on Demand

You can check Feature on Demand (FoD) features activated on your CMM and I/O modules.

Check current Feature on Demand in Chassis Management Module. Select **License Key Management** from the **Mgt Module Management** menu, as shown in Figure 4-47.

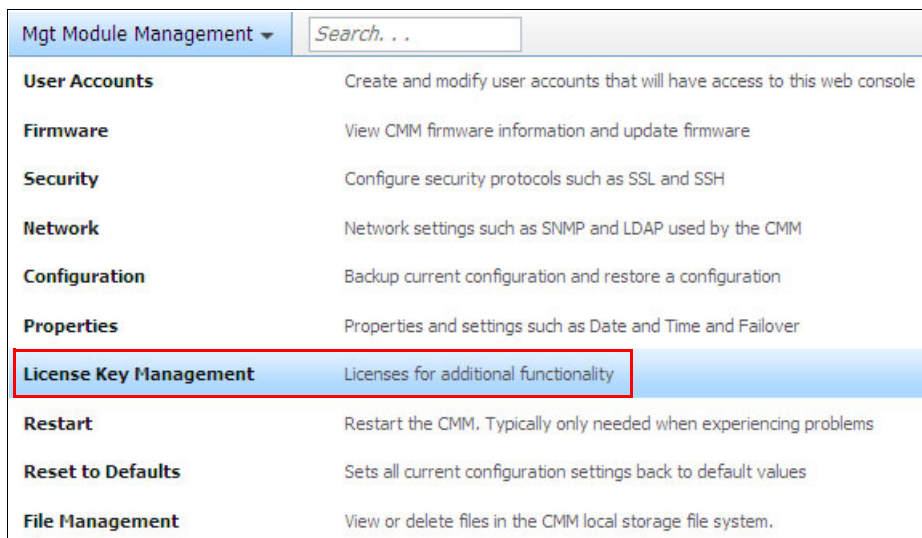


Figure 4-47 License Key Management

Figure 4-48 shows the activated licensed features for I/O modules.

License Key Management				
IOM License Keys Management		Chassis License Keys Management		
Cert	Index	Bay	Valid Through	Description
1	1	1	N/A	IBM Flex System EN4093 10Gb ScSE 24-46 Port Upgrade
2	2	1	N/A	IBM Flex System EN4093 10Gb ScSE 46-64 Port Upgrade

Figure 4-48 License Key Management: I/O Module

In the previous example, two IOM features are activated on IBM Flex System EN4093 Fabric 10Gb Scalable Switch.

Chassis License Keys Management tab shows the activated licenses on CMM itself, for example, IBM Fabric Manager, as shown in Figure 4-49.

License Key Management				
IOM License Keys Management		Chassis License Keys Management		
Index	Feature	Feature Type	Description	System
1	IBM Fabric Manager	0014	IBM SYSTEM X FEATURE ON DEMAND ACTIVATION KEY	IBM Flex System Chassis

Figure 4-49 License Key Management: Chassis

For more information about Features on Demand, see the IBM Features on Demand website at this website:

<https://www-304.ibm.com/systems/x/fod/index.wss>

Figure 4-50 shows Features on Demand main http window. You need an IBM ID to get the features that you want to activate.

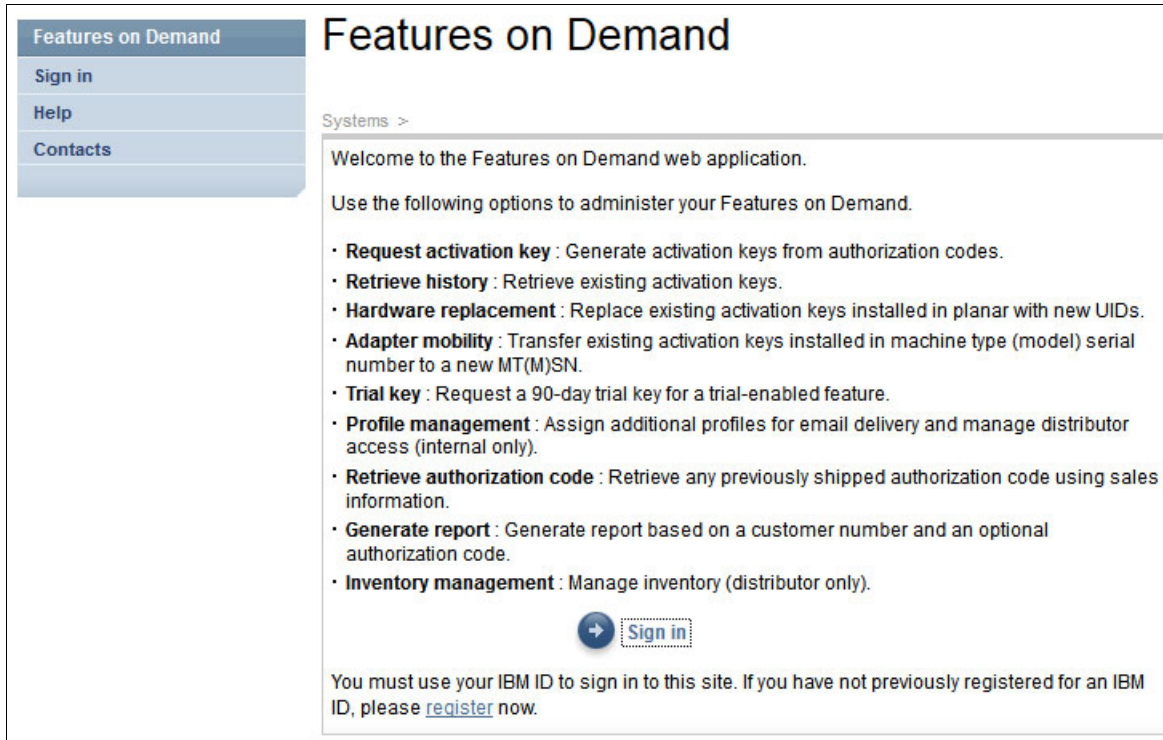


Figure 4-50 Features on Demand main window

4.2.5 Chassis management

This section describes the management of the configured resources in the chassis.

The **Chassis Management** menu is used for reviewing or changing the properties of the components in the chassis. The menu is shown in Figure 4-51. Click **Chassis**.




Figure 4-51 Chassis Management menu

Operating the chassis

You can check the hardware addition or removal history through **Hardware Activity** tab, as shown in Figure 4-52.

Chassis

Identification | LEDs | **Hardware Activity** | Temperature | Air Flow | Air Filter

Shows a summary of chassis hardware activity.

Bay	Module Name	FRU Number	Serial Number	Manufacturer ID	Action
3	4 X86 CPU Blade Server	88Y6237	Y130BG1CM003	IBM Corporation	Added
3	4 X86 CPU Blade Server	88Y6237	Y130BG1CM003	IBM Corporation	Removed
7	2 X86 CPU Blade Server	81Y5128	Y030BG18X00K	IBM Corporation	Added
7	2 X86 CPU Blade Server	81Y5128	Y030BG18X00K	IBM Corporation	Removed
7	2 X86 CPU Blade Server	81Y5128	Y030BG18X00K	IBM Corporation	Added
7	2 X86 CPU Blade Server	81Y5128	Y030BG18X00K	IBM Corporation	Removed
7	2 X86 CPU Blade Server	81Y5128	Y030BG18X00K	IBM Corporation	Added
7	2 X86 CPU Blade Server	81Y5128	Y031BG19R01E	IBM Corporation	Removed

Figure 4-52 Hardware Activity tab

Also, you can see chassis temperature and cooling status by clicking the **Temperature** tab, as shown in Figure 4-53.


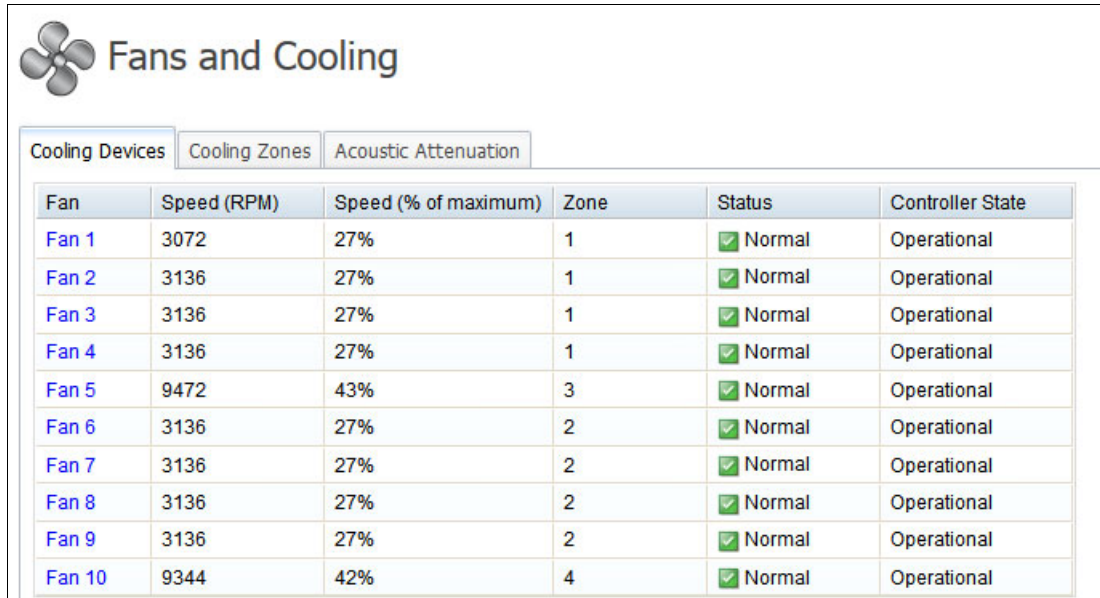


Figure 4-53 Temperature tab

You can check fan and cooling status by selecting the **Chassis Management** → **Fans and Cooling** menu item in the CMM menu as shown in Figure 4-54.



The interface shows a fan status table with columns for Fan, Speed (RPM), Speed (% of maximum), Zone, Status, and Controller State. All fans are listed as 'Normal' and 'Operational'.

Fan	Speed (RPM)	Speed (% of maximum)	Zone	Status	Controller State
Fan 1	3072	27%	1	✓ Normal	Operational
Fan 2	3136	27%	1	✓ Normal	Operational
Fan 3	3136	27%	1	✓ Normal	Operational
Fan 4	3136	27%	1	✓ Normal	Operational
Fan 5	9472	43%	3	✓ Normal	Operational
Fan 6	3136	27%	2	✓ Normal	Operational
Fan 7	3136	27%	2	✓ Normal	Operational
Fan 8	3136	27%	2	✓ Normal	Operational
Fan 9	3136	27%	2	✓ Normal	Operational
Fan 10	9344	42%	4	✓ Normal	Operational

Figure 4-54 Fans and Cooling window

Selecting the individual fan shows you events and power usage statistics that are associated with it, as shown in Figure 4-55.

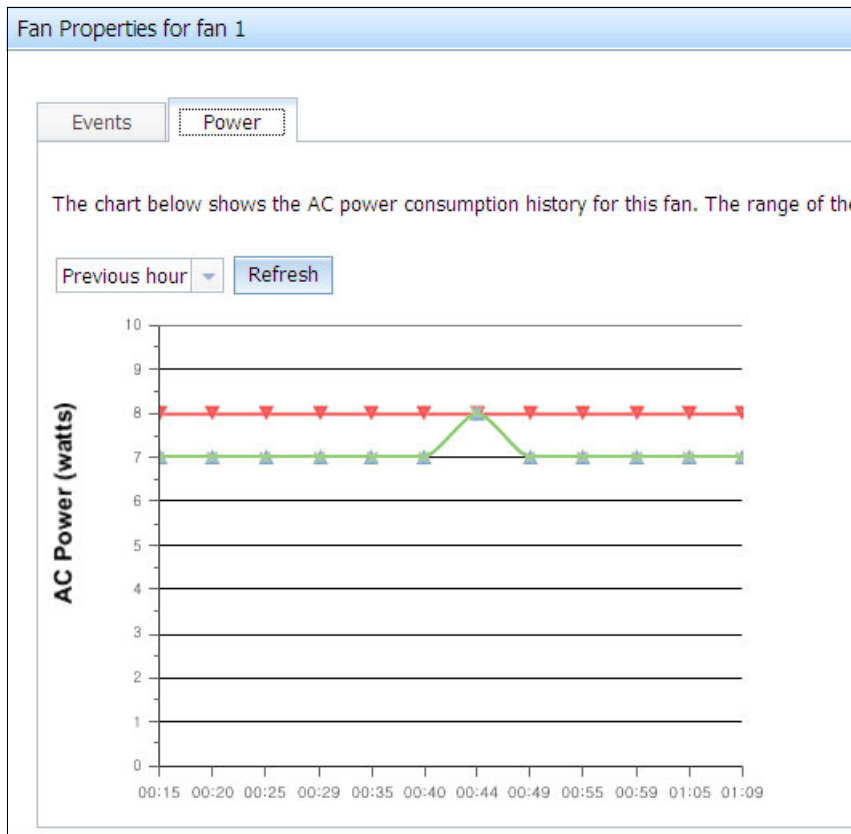


Figure 4-55 Power consumption window

Cooling zone status can also be checked from the Fans and Cooling window by clicking the **Cooling Zone** tab and then clicking the zone number, as shown in Figure 4-56 (Zone 1 is shown).

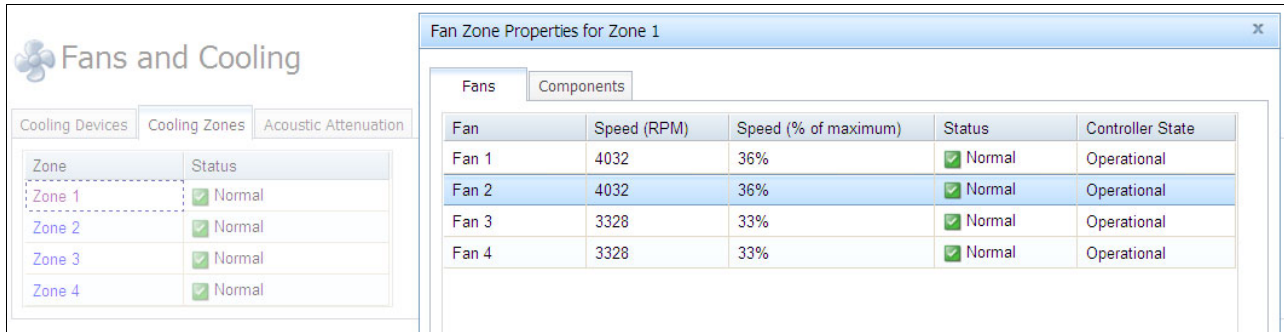


Figure 4-56 Cooling Zones status

You can select the power module policy that meets your specific needs by selecting **Power Modules and Management** menu item from the Chassis Management menu. On the Policies tab, click **Change** near the Current policy, as shown in Figure 4-57.




Figure 4-57 Power Modules and Management window

To monitor power allocation and power consumption history, click the **Input Power and Allocation** tab in the Power Modules window, as shown in Figure 4-58.

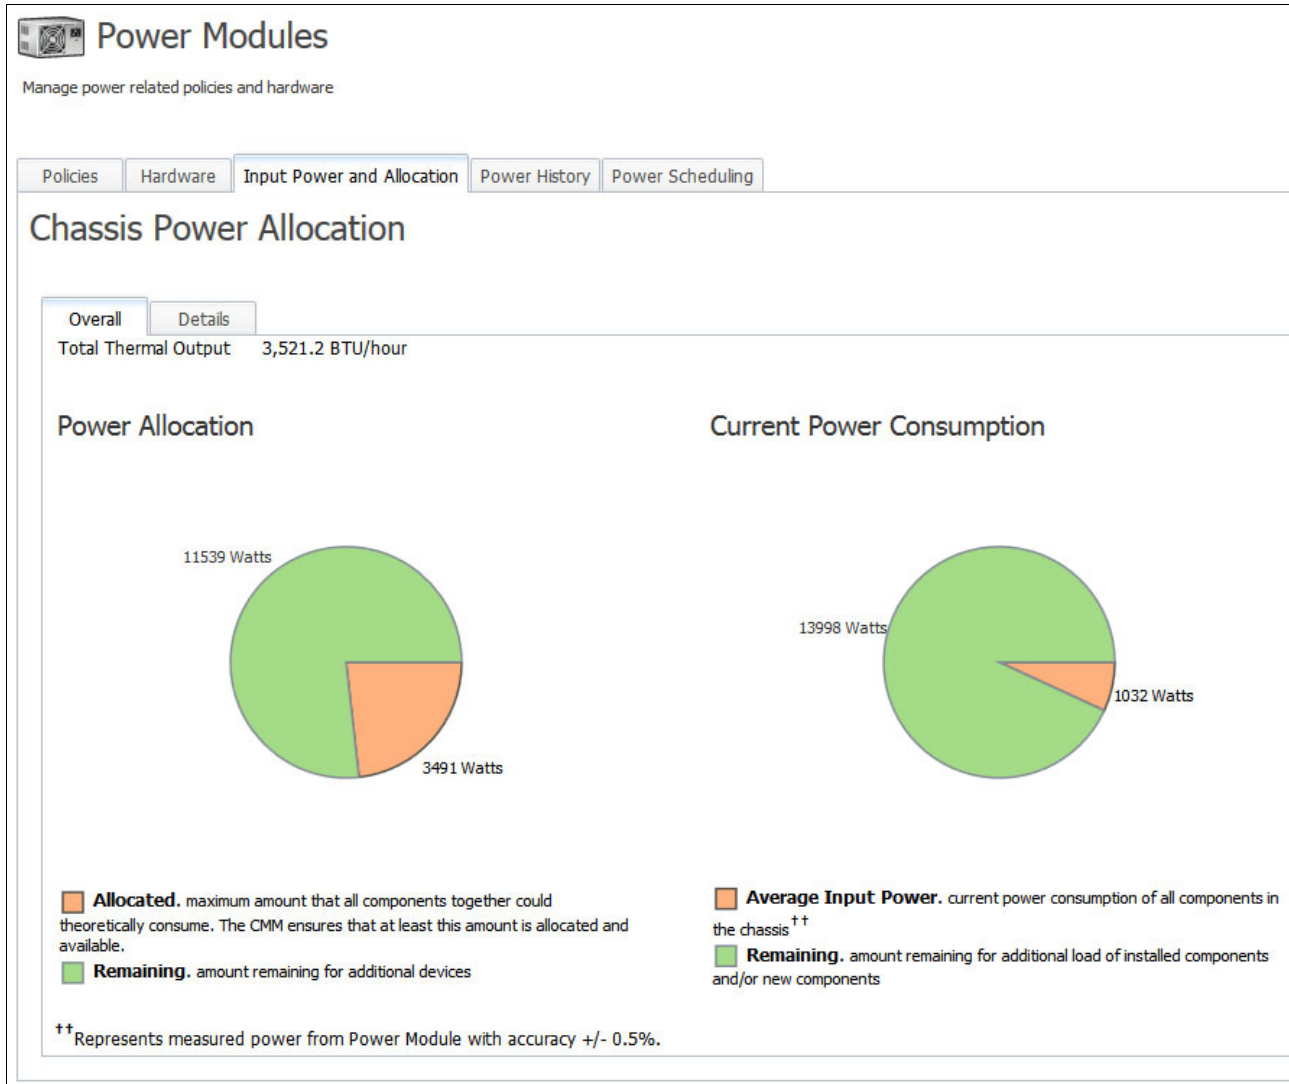


Figure 4-58 Power consumption

Component IP Configuration menu allows you to set the IP parameters on I/O modules and compute nodes, as shown in Figure 4-59.

Component IP Configuration

Configure IPv4 and IPv6 address information for the components below.

I/O Modules

Bay	Device Name	IPv4 Enabled	IP Address
1	IO Module 1	Yes	View
3	IO Module 3	Yes	View

Compute Nodes

Bay	Device Name	IPv4 Enabled	IP Address
1	Node 01 (node01-x240)	Yes	View
2	Node 02 (node02-x240)	Yes	View
3	Node 03 (node03-x240)	Yes	View
4	Node 04 (node04-x240)	Yes	View
5	Node 05 (node05-FSM)	Yes	View

Storage Nodes

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

Figure 4-59 Component IP Configuration window

Click the I/O module or compute node link to open its IP properties window, then click the **IPv4** or **IPv6** tab to verify or configure IP parameters, as shown in Figure 4-60.

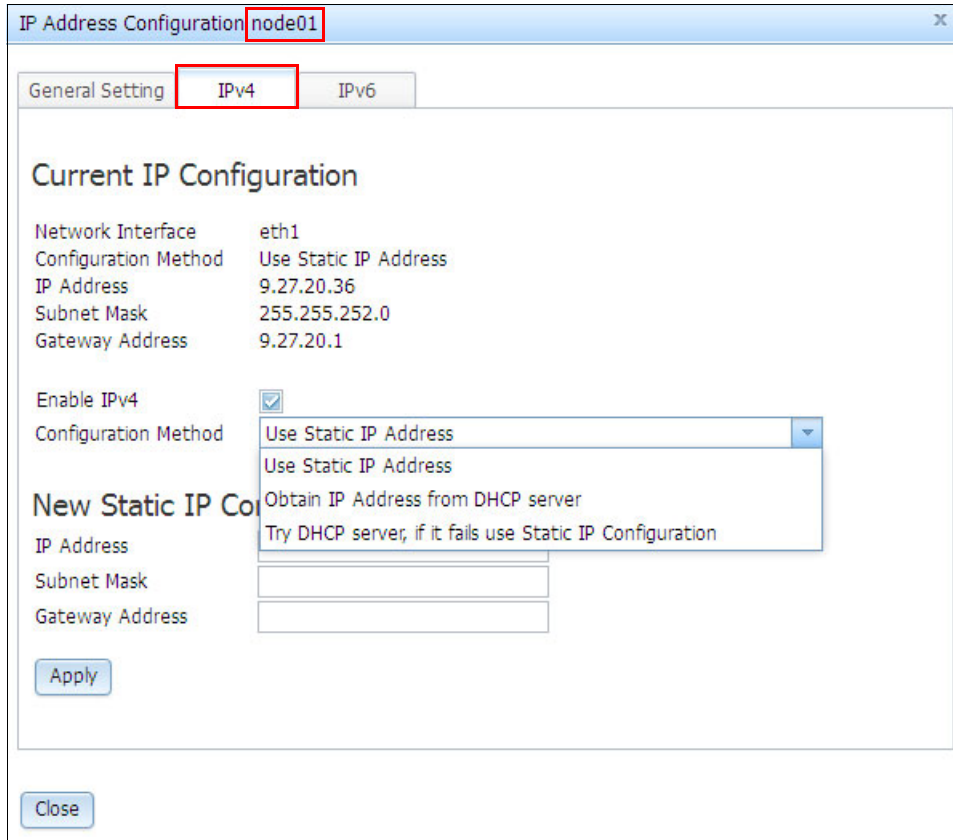


Figure 4-60 IP Address Configuration node01 window

Click **Hardware Topology** in the Chassis Management menu to check all the components in the chassis and their hierarchy, as shown in Figure 4-61.

The screenshot shows the 'Chassis Hardware Topology' window. On the left is a tree view of the chassis components, with 'Processor 1' selected. On the right, the details for 'Processor 1' are displayed:

Processor 1

- Bay: 1
- Bay Type: Processor
- Type: Processor
- Device Name: Processor 1
- Bay Width: 1
- Module Description: CPU 1
- Manufacturer: Intel(R) Corporation
- Speed: 2.00 GHz
- Product Version: Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz
- Part ID: D706 0200 FFFB EBBF
- Type: CENTRAL
- Family: Intel Xeon
- Cores: 8
- Threads: 16
- L1 ICache Size: 64K
- L2 Cache Size: 256K
- L3 Cache Size: 20480K
- Voltage: 1.2 V
- External Clock: 100 MHz
- Maximum Data Width: 64-bit Capable

Figure 4-61 Chassis Hardware Topology window

Click **Reports** in the Chassis Management menu to see hardware information in the chassis, as shown in Figure 4-62.

The screenshot shows the 'Reports' window with the 'MAC Addresses' tab selected. A note states: 'If a device contains more than 8 MAC addresses, please click the "Module name" link to see all of them'. The table below lists the MAC addresses for various modules:

Type	Module name	Description	MAC 1	MAC 2	MAC 3	MAC 4
Compute Nodes	[1] Node 01 (node01-x240)	IBM Flex System x240 with 10Gb	34:40:B5:BE:7D:00	34:40:B5:BE:7D:04	34:40:B5:BE:7D:01	34:40:B5:BE:7D:05
	Expansion Card	FC3052 8Gb FC Adapter	----	----	----	----
	[2] Node 02 (node02-x240)	IBM Flex System x240 with 10Gb	34:40:B5:BE:8E:90	34:40:B5:BE:8E:94	34:40:B5:BE:8E:91	34:40:B5:BE:8E:95
	Expansion Card	FC3052 8Gb FC Adapter	----	----	----	----
	[3] Node 03 (node03-x240)	IBM Flex System x240 with 10Gb	34:40:B5:BE:83:D0	34:40:B5:BE:83:D4	----	----
	Expansion Card	FC3052 8Gb FC Adapter	----	----	----	----
	[4] Node 04 (node04-x240)	IBM Flex System x240 with 10Gb	34:40:B5:BE:9D:58	34:40:B5:BE:9D:5C	----	----
	Expansion Card	FC3052 8Gb FC Adapter	----	----	----	----
I/O Modules	[5] Node 05 (node05-FSM)	IBM Flex System Manager Node	5C:F3:FC:5F:5E:86	----	----	----
	Add-in Card	Network Adapter	5C:F3:FC:5F:09:90	5C:F3:FC:5F:09:91	----	----
	[1] I/O Module	EN4093 10Gb Ethernet Switch	34:40:B5:34:AD:EF	----	----	----
	[3] I/O Module	FC3171 8Gb SAN Switch	00:C0:DD:24:41:1C	----	----	----

Figure 4-62 Reports window

Operating the compute node

Selecting **Compute Nodes** from the Chassis Management menu shows a window that lists the servers that are installed in the chassis (Figure 4-63). You can power on/off compute nodes, access them through a remote console, set properties such as Wake on LAN, and perform other actions.

Compute Nodes

i If specifying a power action for multiple nodes, please be aware that in case of an error you will only be informed failed executing the action. Successful nodes are ignored.
 Different node types may take different amounts of time to complete the power action, so in some cases, the power sta immediately reflected on the page. In this case, the user may have to perform a refresh (F5) one or more times to see th reflected on the page.

Power and Restart ▾ Actions ▾ Global Settings Columns ▾

	Device Name	Device Type	Health Status	Power	Bay	Bay Type
<input type="checkbox"/>	Node 01 (node01-x240)	Compute Node	✓ Normal	On	1	Node
<input type="checkbox"/>	Node 02 (node02-x240)	Compute Node	✓ Normal	On	2	Node
<input type="checkbox"/>	Node 03 (node03-x240)	Compute Node	✓ Normal	On	3	Node
<input type="checkbox"/>	Node 04 (node04-x240)	Compute Node	✓ Normal	On	4	Node
<input type="checkbox"/>	Node 05 (node05-FSM)	Compute Node	✓ Normal	On	5	Node

Figure 4-63 Compute Nodes window

Operating the I/O module

The I/O Modules window is similar to the Compute Nodes window. After clicking **I/O Modules** in the Chassis management, a table is displayed that shows the I/O modules. Clicking the module name opens a pop-up window with the properties of that module as shown in Figure 4-64.

I/O Modules

Power and Restart ▾ Actions ▾

	Device Name	Health Status	Bay	Power	Serial Number	Part Number
<input checked="" type="checkbox"/>	IO Module 1	✓ Normal	1	On	Y250VT1BW111	49Y4272
<input type="checkbox"/>	IO Module 3	✓ Normal	3	On	Y251NY26K054	69Y1932

Details for I/O Module 'IO Module 1' - IBM Flex System Fabric EN4093 10Gb Scalable Switch

Events | General | Hardware | Firmware | Power | IO Connectivity | Port Info | LEDs

Severity	Source	Sequence #	Date	Event ID
0 items				
No Data Available				

Apply Cancel

Figure 4-64 I/O Modules window

Start the switch module remote console by clicking **Launch IOM Console**, as shown in Figure 4-65.

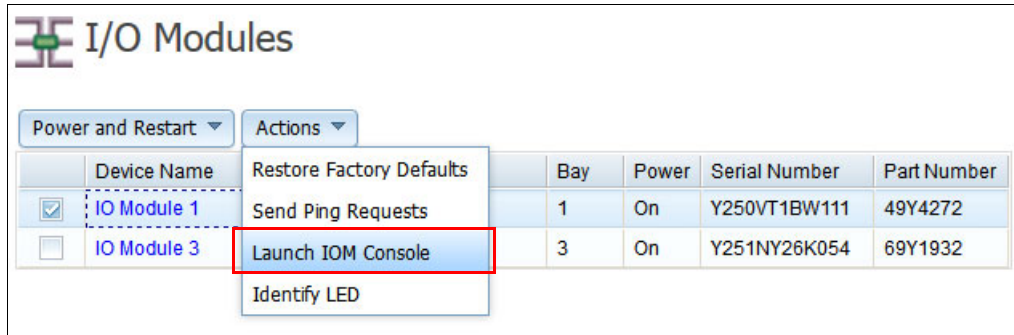


Figure 4-65 Launch I/O module console menu

Check or type the IP address of the managed I/O module, select the protocols, and click **Launch** as shown in Figure 4-66.

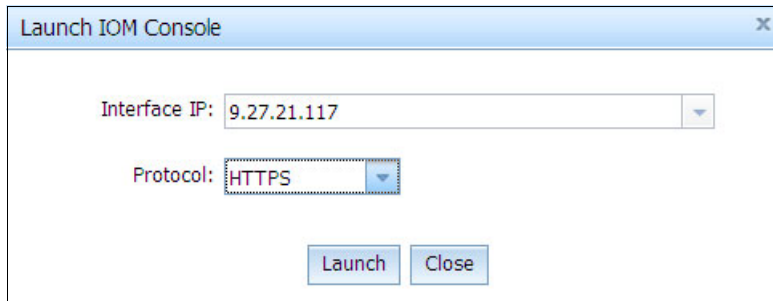


Figure 4-66 Launch IOM Console window

This process prompts for user ID and password and then displays the selected I/O module’s management user interface (UI), in this example the web UI for the network switch, as shown in Figure 4-67.

Switch Dashboard	
Switch Name	
Switch Location	
Switch Type	IBM Flex System Fabric EN4093 10Gb Scalable Switch(Upgrade1)
Switch Up Time	92 days, 6 hours, 3 minutes and 26 seconds.
Last Boot Time	15:24:00 Thu Aug 8, 2013 (power cycle)
Time and date	21:19:50 , 11/8/2013
Timezone Location	
Daylight Savings Time Status	disabled
MAC Address	34:40:b5:34:ad:00
IP Address	9.42.171.8
PCBA Part Number	BAC-00072-01
Hardware Part Number	49Y4272
Serial Number	Y250VT1BW111
Manufacturing Date	48/11

Figure 4-67 Network I/O module main window

4.2.6 Using the Chassis Management Module CLI

This section addresses configuring the CMM and managing components that are installed in an IBM Flex System Enterprise Chassis by using the command-line interface.

The IBM Flex System Chassis Management Module (CMM) command-line interface (CLI) provides direct access to IBM Flex System management functions as an alternative to using the web-based user interface.

Using the CLI, you can issue commands to control the power and configuration of the CMM and other components that are in an IBM Flex System Enterprise Chassis. The command-line interface also provides access to the text-console command prompt on each compute node through a Serial over LAN (SOL) connection.

You can access the CMM CLI through the following connections:

- ▶ An Ethernet connection to the CMM
- ▶ A Telnet connection to the IP address of the CMM
- ▶ A Secure Shell (SSH) connection to the CMM

You can initiate connections from the client system by using standard remote communication software. No special programs are required. You must authenticate with the CMM before you issue commands.

Use telnet or ssh program, enter the IP address, and select the protocol. You can then log in to the Chassis Management Module CLI interface, as shown in the Figure 4-68.

```
login as: USERID
Using keyboard-interactive authentication.
password:

Hostname:          MM5CF3FC25E3B7
Static IP address: 9.27.20.56
Burned-in MAC address: 5C:F3:FC:25:E3:B7
DHCP:              Disabled - Use static IP configuration.
Last login: Wednesday June 27 2012 12:07 from 9.27.20.38 (CIM)

system> █
```

Figure 4-68 Chassis Management Module login

Command-line interface guidelines

All commands have the basic structure as shown in Example 4-1.

Example 4-1 Command usage

```
command -option parameter
```

Selecting the command target

You can use the command-line interface to target commands to the CMM or to other devices in the IBM Flex System Enterprise Chassis. The command-line prompt indicates the persistent command environment, which is where commands are directed unless another target is specified. You can specify a command target by using the full target name or by using a target name that is relative to the persistent command environment.

Command targets are specified hierarchically, as shown in Figure 4-69.




Figure 4-69 Command target hierarchic view

Example 4-2 show how to use CLI commands.

Example 4-2 Command target usage

Use the `-T system:mm[1]` option to redirect a command to the CMM in bay 1.
 Use the `-T system:switch[1]` option to redirect a command to the I/O module in I/O bay 1.

Using the command-line interface

This section addresses how the CLI works with Chassis Management Module. Two CLI commands are run in the example environment:

1. `list` command

This command displays a list of devices present within the command target. It can be used to determine the physical configuration of the IBM Flex System Enterprise Chassis. This information includes how many CMMs are installed in the IBM Flex System Enterprise Chassis, and which CMM is set as primary.

To view all the components in the Chassis, run this command as shown in Figure 4-70.

```
system> list -l 2
system
    blade[1]  node01
    blade[2]  node02
    blade[3]  node03
    blade[4]  node04
    blade[5]  node05
    blade[6]  node06
    blade[8]  node08
    blade[9]  node09
    blade[10] node10
    blade[11] node11
    blower[1]
    blower[2]
    blower[3]
    blower[4]
    blower[5]
    blower[6]
    blower[7]
    blower[8]
    blower[9]
    blower[10]
    power[1]
    power[2]
    power[3]
    power[4]
    power[5]
    power[6]
    mm[1]     primary
    switch[1]
    switch[2]
    switch[3]
    mt[1]
    fanmux[1]
    fanmux[2]
system>
```

Figure 4-70 list command output

2. info command

This command displays information about IBM Flex System components and their configuration. To view the information about a compute node in bay 6, issue the **info** command as shown in Figure 4-71.

```
system> info -T blade[6]
Name: node06
UUID: 60C2 2B07 2C58 4C62 C18B 164D 19B0 60C0
Manufacturer: IBM (Not Available)
Manufacturer ID: 20301
Product ID: 305
Mach type/model: 789522X
Mach serial number: 101D88B
Manuf date: Not Available
Hardware rev: 0.0
Part no.: 00E0910
FRU no.: 00E0740
FRU serial no.: YL1011243000
```

Figure 4-71 info command output

For more information about the CLI, see the IBM Flex System Information Center:

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

IBM Flex System Manager

This part describes how to implement systems management of IBM PureFlex System using the IBM Flex System Manager.

This part includes the following chapters:

- ▶ Chapter 5, “Planning for IBM Flex System Manager management” on page 85
- ▶ Chapter 6, “IBM Flex System Manager initial configuration” on page 121
- ▶ Chapter 7, “Managing chassis components with IBM Flex System Manager” on page 259
- ▶ Chapter 8, “IBM Fabric Manager” on page 305
- ▶ Chapter 9, “Managing the KVM environment with IBM Flex System Manager” on page 319
- ▶ Chapter 10, “Managing the PowerVM environment with IBM Flex System Manager” on page 401
- ▶ Chapter 11, “Managing the VMware environment with IBM Flex System Manager” on page 495
- ▶ Chapter 12, “Managing the Hyper-V environment with IBM Flex System Manager” on page 555
- ▶ Chapter 13, “Mobile management” on page 571



Planning for IBM Flex System Manager management

This chapter describes general planning information about IBM Flex System Manager (FSM). It also addresses specific virtualization solution management prerequisites and considerations when you build certain virtual infrastructure that is managed through IBM Flex System Manager.

The following topics are covered:

- ▶ 5.1, “Planning for IBM Flex System Manager” on page 86
- ▶ 5.2, “Planning for the management of virtualized environments” on page 100

5.1 Planning for IBM Flex System Manager

This section describes general planning considerations for the implementation of FSM-based systems management.

The following subtopics are covered:

- ▶ Flex System Manager network integration architecture
- ▶ Planning for security
- ▶ Planning for Features on Demand
- ▶ Agents and tasks supported
- ▶ Planning for the management of networking infrastructure
- ▶ Planning for the management of storage infrastructure
- ▶ Planning for IBM Fabric Manager

For more information about FSM hardware and software, see 2.5, “IBM Flex System Manager” on page 20.

5.1.1 Flex System Manager network integration architecture

In an IBM Flex System Enterprise Chassis, you can configure separate management and data networks.

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

The management network is shown in Figure 5-1 on page 87 as the blue line. It connects the Chassis Management Module (CMM) to the compute nodes, the switches in the I/O bays, and the FSM. The FSM connection to the management network is through a special Broadcom 5718-based management network adapter (Eth0). The management networks in multiple chassis can be connected together through the external ports of the CMMs in each chassis by using a GbE top-of-rack switch.

The yellow line in the Figure 5-1 shows the production data network. The FSM also connects to the production network (Eth1) so that it can access the Internet for product updates and other related information.

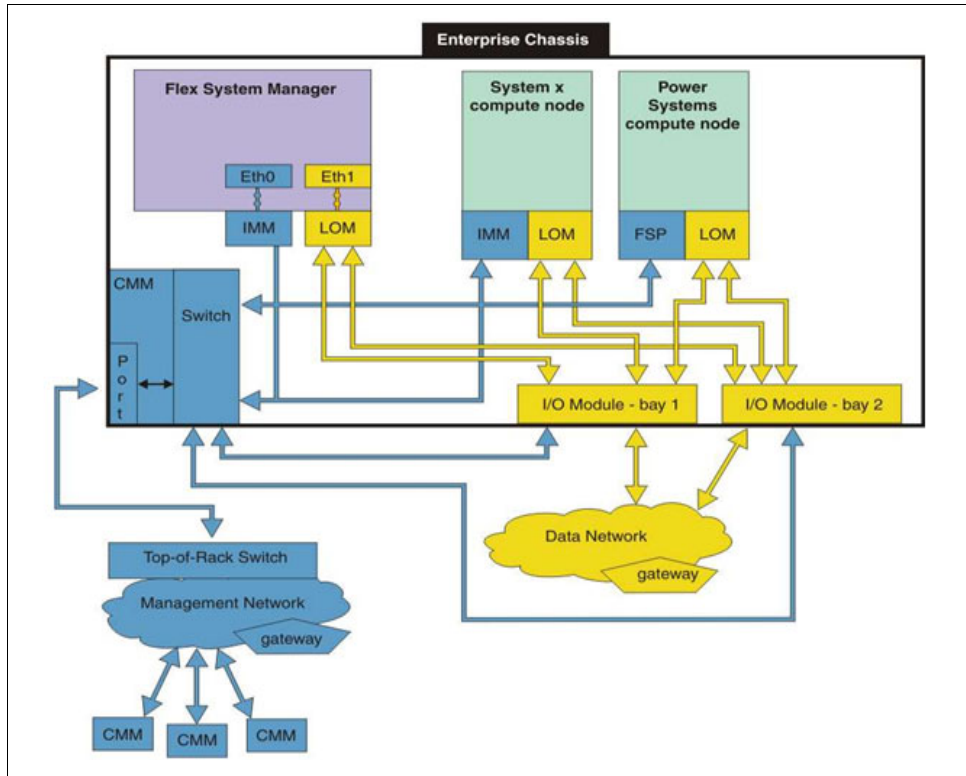


Figure 5-1 Management and production data network

One of the key functions that the data network supports is discovery of operating systems on the various network endpoints. Discovery of operating systems by the FSM is required to support software updates on an endpoint such as a compute node.

Management and production data networks are usually separate subnets. In such a case, the FSM management node uses both interfaces (Eth0 and Eth1). If data and management networks are combined into a single subnet, only Eth0 port must be configured with the IP address, and Eth1 must remain unconfigured.

5.1.2 Planning for security

The IBM Flex System products include features that can help you secure your environment. The following sections provide basic information about how some of these features work. You can use this information along with the documentation that comes with your IBM Flex System products to help you evaluate and implement the security plan for your environment.

As you evaluate the security requirements of your environment, remember that unsecured systems-management tools can damage hardware and software. It is important that you understand all security risks in your system environment and what you can do to minimize these risks.

An IBM Flex System Manager management software security policy is a set of security-related characteristics that define a particular level of protection from security exposures. Depending on its level, the security policy might include account-related policies, communication-protocol enablement, and event-tracking levels.

The management software enforces a chosen security policy only for the management node itself. The management software offers two types of security policy:

- ▶ *Legacy* Security Policy
- ▶ *Secure* Security Policy

Legacy Security Policy

The IBM Flex System Manager management software Legacy security policy is the least secure and most flexible setting that is available for your configuration.

The Legacy level of management software security policy provides flexibility in managing platform security, but this policy is the least secure. It allows the use of the following conditions:

- ▶ Weak password policies with minimal controls
- ▶ Manufacturing default passwords that do not have to be changed
- ▶ Unencrypted communication protocols such as Telnet, SNMPv1, TCP Command Mode, CIM-XML, FTP Server, and TFTP Server

Secure Security Policy

The IBM Flex System Manager management software Secure security policy is the most secure and least flexible setting that is available for your configuration.

The Secure security controls setting, or Secure policy, is the default security setting. It helps to ensure a secure chassis infrastructure and enforces the following conditions:

- ▶ Strong password policies with automatic validation and verification checks
- ▶ Updated passwords that replace the manufacturing default passwords after the initial setup
- ▶ Only secure communication protocols such as SSH and SSL
- ▶ Certificates to establish secure, trusted connections for applications that run on the management processors

Users, groups, and roles

Flex System Manager management software offers authentication and user administration options that enable you to specify user privileges for specific tasks and resources. User registry integration, integrity, confidentiality, and Secure Sockets Layer (SSL)-supported secure data transmission are other key elements of the management software security.

Management software user accounts are subjected to two interdependent processes: authentication and authorization. *Authentication* is used to determine the identity of the user and verify and validate that identity. *Authorization* checks the permissions of the authenticated user and controls access to resources according to the roles that are assigned to the user.

User authentication is the security mechanism by which a user's credentials that are used to access a system are verified. After authentication, a user can access the system. However, to access a specific resource or perform a specific task, the user must also have the appropriate authorization. Authentication prevents unauthorized management servers or rogue managed-system applications from accessing the resources.

To be authenticated, users are required to enter a user ID and password for the system that they want to access. The authentication process uses the configured user registry, which is from either the operating system, Lightweight Directory Access Protocol (LDAP), or the domain controller.

User authorization occurs when an authenticated user uses IBM Flex System Manager to perform a task on a resource. The authorization mechanism compares the user account, or the group to which the user belongs, to the role-based access control (RBAC) settings for that user or group. If a role exists that contains the authorizations necessary to complete that task on that specified resource, the task proceeds.

Users can access only the applications, tasks, and resources that their user accounts are authorized to access. The authorities that you grant to a user determine the console and resource information that the user can access, and the tasks that the user can perform on those resources.

The authorization process that IBM Flex System Manager performs when accessing a resource is independent of the authentication that is required to access that resource. For example, a user might be able to authenticate to and therefore access IBM Flex System Manager web interface or another resource by using IBM Flex System Manager web interface, but to perform a task on that resource, both the task and the resource must be authorized in the role settings that are assigned to that user or the authorization group to which the user belongs.

Centralized user management

A centralized management configuration uses a single user authentication repository for all of the Chassis Management Modules (CMMs) in a management domain. The user accounts that are created for the IBM Flex System Manager management software are used by all of the CMMs and compute node service processors in the chassis.

When you use the IBM Flex System Manager management software to place a chassis under centralized management, the Chassis Management Module (CMM) is configured to use the registry that is stored on the management node. The local user accounts in the CMM registry are disabled, and the new user account RECOVERY_ID is created for future authentication to the CMM (as long as it is configured to use the centralized user registry on the management node).

After the CMM detects the management node user registry, it uses the management node registry configuration to provision all of the managed resources in the chassis (except for network switches) so that they also use the central management node user registry. When you log in to an IMM or FSP on a compute node in a centrally managed chassis, you must use a user name and password that are stored in the IBM Flex System Manager user registry.

With centralized management, a single security policy is distributed and enforced on all of the Chassis Management Modules in a management domain. In addition, a single set of user accounts and a single password policy is in effect.

Chassis that are not centrally managed might have different security policies, different user accounts, and different passwords from those that are set for the management domain. If chassis are centrally managed, user accounts can be edited only by the management software.

For more information about IBM Flex System Manager security, see the Security topic in the IBM Flex System Information Center:

<http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8731.doc/security.html>

5.1.3 Planning for Features on Demand

IBM Features on Demand (FoD) is the capability to activate or “unlock” features integrated in IBM products. The feature is in the firmware or software, but is “locked” until the activation key is installed. There are several benefits from using FoD:

- ▶ You buy the features that you need now with the ability to grow your system later without costly rip and replace (Pay as You Grow).
- ▶ Allows for upgrades in the field.
- ▶ Feature activation can be done on server or chassis at the time of the server/chassis sale or later.
- ▶ FoD enables ease of installation, reduced inventory, and faster fulfillment of options.

Fulfillment process

You can activate Features on Demand by using these methods:

1. An FoD option ordered with server and installed during manufacturing (Figure 5-2).




Figure 5-2 Features on Demand field order option

2. An FoD option that you purchase separately or after the system sale (Figure 5-3 on page 91).




Figure 5-3 Features on Demand Post or respective order option

Features on Demand for IBM Flex System Manager

IBM Features on Demand (FoD) provides optional software that is available for IBM Flex System Manager Types 7955, 8731, and 8734, and IBM Flex System Manager management software. You can also use the management software to enable optional features for managed compute nodes.

FoD provides a convenient way to order and activate optional features from IBM through the management software web interface. You can also upload compute node Features on Demand keys to the management node and distribute the keys to managed compute nodes by using the management software.

Any Features on Demand software that you ordered with your IBM Flex System Manager Types 7955, 8731, and 8734 management node are preactivated. They do not require manual activation through the management software interface. If you did not order an FoD when you purchased your system, you can purchase it just like any other software and hardware option. You can redeem Features on Demand for the management software at this website:

<http://www.ibm.com/systems/x/fod/>

The IBM Flex System Manager management node ships with a preinstalled IBM Systems management stack. The part numbers to order Features on Demand (FoD) software entitlement licenses are shown in Table 5-1 on page 92 (for United States, Canada, Asia Pacific, and Japan) and Table 5-2 on page 92 (for Latin America and Europe/Middle East/Africa). The part numbers for the same features are different across geographies. Ask an IBM representative for specifics.

Table 5-1 FoD part numbers (United States, Canada, Asia Pacific, and Japan)

Description	Part number
Base feature set license	
IBM Flex System Manager Per Managed Chassis with 1 Year SW S&S	90Y4217
IBM Flex System Manager Per Managed Chassis with 3 Year SW S&S	90Y4222
Advanced feature set upgrade ^a	
IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 1 Year SW S&S	90Y4249
IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 3 Year SW S&S	00D7554
Fabric provisioning feature upgrade ^a	
IBM Flex System Manager Service Fabric Provisioning w/1 Yr S&S	90Y4221
IBM Flex System Manager Service Fabric Provisioning w/3 Yr S&S	90Y4226

a. The Advanced Upgrade and Fabric Provisioning licenses are applied on top of the IBM Flex System Manager base license.

Table 5-2 FoD part numbers (Latin America and Europe/Middle East/Africa)

Description	Part number
Base feature set license	
IBM Flex System Manager Per Managed Chassis with 1 Year SW S&S	95Y1174
IBM Flex System Manager Per Managed Chassis with 3 Year SW S&S	95Y1179
Advanced feature set upgrade ^a	
IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 1 Year SW S&S	94Y9219
IBM Flex System Manager, Advanced Upgrade, Per Managed Chassis with 3 Year SW S&S	94Y9220
Fabric provisioning feature upgrade ^a	
IBM Flex System Manager Service Fabric Provisioning w/1 Yr S&S	95Y1178
IBM Flex System Manager Service Fabric Provisioning w/3 Yr S&S	95Y1183

a. The Advanced Upgrade and Fabric Provisioning licenses are applied on top of the IBM Flex System Manager base license.

The IBM Flex System Manager base feature set offers the following functionality:

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

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

- ▶ Image management (VMControl Standard)
- ▶ Pool management (VMControl Enterprise)
- ▶ Advanced network monitoring and quality of service (QoS) configuration (Service Fabric Provisioning)

The Fabric Provisioning upgrade offers the following functionality:

- ▶ Advanced network monitoring and quality of service (QoS) configuration (Service Fabric Provisioning)

Fabric provisioning functionality is included in the advanced feature set. It is also available as a separate Fabric Provisioning feature upgrade for the base feature set. The Advanced Upgrade and the Fabric Provisioning feature upgrade are mutually exclusive; that is, either the Advance Upgrade or the Fabric Provisioning feature upgrade can be applied on top of the base feature set license, but not both.

Important: The Advanced Upgrade and Fabric Provisioning licenses are applied on top of the IBM Flex System Manager base license.

5.1.4 Features on Demand for components in the Chassis

There are many Features on Demand on the components. For example, there are a few options to activate to the I/O modules. Table 5-3 shows Features on Demand for I/O modules.

Table 5-3 Part numbers for ordering Feature on Demand entitlement licenses for I/O modules

Description	Part number
IBM Flex System EN2092 1Gb Ethernet Scalable Switch (10Gb Uplinks)	49Y4298
IBM Flex System EN2092 1Gb Ethernet Scalable Switch (Upgrade 1)	90Y3562
IBM Flex System EN4023 10Gb Scalable Switch (FoD 1)	94Y5158
IBM Flex System EN4023 10Gb Scalable Switch (FoD 2)	94Y5159
IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 1)	00D5845
IBM Flex System Fabric CN4093 Converged Scalable Switch (Upgrade 2)	00D5847
IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 1)	49Y4798
IBM Flex System Fabric EN4093 10Gb Scalable Switch (Upgrade 2)	88Y6037
IBM Flex System Fabric SI4093 System Interconnect Module (Upgrade 1)	95Y3318
IBM Flex System Fabric SI4093 System Interconnect Module (Upgrade 2)	95Y3320
IBM Flex System FC5022 16Gb Fabric Watch Upgrade	00Y3320
IBM Flex System FC5022 16Gb ISL/Trunking Upgrade	00Y3322
IBM Flex System FC5022 16Gb SAN Scalable Switch-Upgrade 1	88Y6382
IBM Flex System FC5022 16Gb SAN Scalable Switch-Upgrade 2	88Y6386
IBM Flex System IB6131 InfiniBand Switch (FDR Upgrade)	90Y3462

Table 5-4 on page 94 shows Features on Demand for compute nodes.

Table 5-4 Part numbers for ordering Feature on Demand entitlement licenses for compute nodes

Description	Part number
IBM Flex System CN4054 Virtual Fabric Adapter (SW Upgrade)	90Y3558
IBM Virtual Fabric Advanced Software Upgrade (LOM)	90Y9310
ServeRAID M5100 Series RAID 6 Upgrade for IBM Flex System	90Y4410
ServeRAID M5100 Series Performance Upgrade for IBM Flex System	90Y4412
ServeRAID M5100 Series SSD Caching Enabler for IBM Flex System	90Y4447

For more information about managing Features on Demand, see 6.10, “Manage Feature-on-Demand keys” on page 216.

5.1.5 Agents and tasks supported

IBM Flex System Manager provides four tiers of agents for managed systems. For each managed system, you need to select the tier that provides the amount and level of capabilities that you need for that managed system. Select the level of agent capabilities that you need for the type of managed system and the management tasks in your system.

IBM Flex System Manager has four agent tiers:

- ▶ Agentless in-band
 - Managed systems without any Flex System Manager client software installed. Flex System Manager communicates with the managed system through the operating system.
- ▶ Agentless out-of-band
 - Managed systems without any Flex System Manager client software installed. Flex System Manager communicates with the managed system through something other than the operating system, such as a service processor or a hardware management console.
- ▶ Platform Agent
 - Managed systems with Platform Agent installed. Flex System Manager communicates with the managed system through the Platform Agent.
- ▶ Common Agent
 - Managed systems with Common Agent installed. Flex System Manager communicates with the managed system through the Common Agent.

Table 5-5 on page 95 lists the agent tier support for the managed systems. Managed system types include x220, x222, x240, and x440 compute nodes supporting Windows, Linux, and VMware, p24L compute node supporting Linux, and p260, p270, and p460 compute nodes supporting IBM AIX, IBM i, and Linux.

Table 5-5 Agent tier support by management system type

Managed system type	Agent tier	Agentless in-band	Agentless out-of-band	Platform Agent	Common Agent
Compute nodes that run AIX		Yes	Yes	No	Yes
Compute nodes that run IBM i		Yes	Yes	Yes	Yes
Compute nodes that run Linux		No	Yes	Yes	Yes
Compute nodes that run Linux and supporting SSH		Yes	Yes	Yes	Yes
Compute nodes that run Windows		No	Yes	Yes	Yes
Compute nodes that run Windows and supporting SSH or DCOM		Yes	Yes	Yes	Yes
Compute nodes that run VMware		Yes	Yes	No	No
Other managed resources that support SSH or SNMP		Yes	Yes	No	No

Table 5-6 summarizes the management tasks that are supported by the compute nodes, depending on the agent tier.

Table 5-6 Compute node management tasks that are supported by the agent tier

Managed system type	Agent tier	Agentless in-band	Agentless out-of-band	Platform Agent	Common Agent
Command automation		No	No	No	Yes
Hardware alerts and status		No	Yes	Yes	Yes
Platform alerts		No	No	Yes	Yes
Health and status monitoring		No	No	Yes	Yes
File Transfer		No	No	No	Yes
Inventory (hardware)		No	Yes	Yes	Yes
Inventory (software)		Yes	No	Yes	Yes
Process Management		No	No	No	Yes
Power Management		No	Yes	No	Yes
Remote Control		No	Yes	No	No
Remote Command Line		Yes	No	Yes	Yes
Resource Monitors		Yes ^a	No	Yes	Yes
Update Manager (firmware)		Yes ^b	Yes	Yes	Yes
Update Manager (agent updates)		No	No	Yes	Yes

a. Supported for VMware and Hyper-V virtualized environments.

b. Supported for Windows environments

5.1.6 Planning for the management of networking infrastructure

IBM Flex System Manager management software Network Control provides advanced network management functions for IBM Flex System Enterprise Chassis network devices. Functions include discovery, inventory, network topology, health and status monitoring, and configuration of network devices. Network Control is a preinstalled plug-in that builds on the base management software capabilities. It integrates the launch of vendor-based device management tools, topology views of network connectivity, and subnet-based views of servers and network devices.

Network Control offers the following network-management capabilities:

- ▶ Discover network devices in your environment
- ▶ Review your network device inventory in tables or a network topology view
- ▶ Monitor the health and status of network devices
- ▶ Manage devices by groups: Ethernet switches, Fibre Channel over Ethernet, or Subnet
- ▶ View network device configuration settings, and apply templates to configure devices, including Converged Enhanced Ethernet quality of service (QoS), VLANs, and Link Layer Discovery Protocol (LLDP)
- ▶ View systems according to VLAN and subnet
- ▶ Run network diagnostic tools like ping and traceroute
- ▶ Create logical network profiles to quickly establish VLAN connectivity
- ▶ Simplified management of VM connections by configuring multiple characteristics of a network when virtual machines are part of a network system pool
- ▶ With management software VMControl, maintain network state (VLAN, ACLs) as a virtual machine is migrated (KVM)
- ▶ Management of virtual switches, including virtual Ethernet bridges
- ▶ Configuration of port profiles, a collection of network settings that is associated with a virtual system
- ▶ Automatic configuration of devices in network systems pools

IBM Flex System Manager Network Control provides facilities to discover, inventory, and monitor network devices, start vendor applications for configuration of network devices, and view groups of network devices. IBM Flex System Manager Network Control extends the network management functions of the IBM Flex System Manager product.

Table 5-7 on page 97 shows supported network switches and their management tasks.

Table 5-7 Supported I/O switches and management tasks

I/O module	EN2092 1 Gb Ethernet	EN4023 10 Gb Ethernet	EN4093 EN4093R 10 Gb Ethernet	CN4093 10 Gb Converged	SI4093 10 Gb Interconnect
Discovery	Yes	Yes	Yes	Yes	Yes
Inventory collection	Yes	Yes	Yes	Yes	Yes
Monitoring	Yes	Yes	Yes	Yes	Yes
Alerts	Yes	Yes	Yes	Yes	Yes
Protocol configuration	Yes	No	Yes	Yes	No
VLAN configuration	Yes	No	Yes	Yes	No
CEE configuration	No	No	Yes	Yes	No
EVB configuration	No	No	Yes	Yes	No
Stacked switch management	No	No	Yes	Yes	No

5.1.7 Planning for the management of storage infrastructure

Storage management with IBM Flex System Manager management software involves two software components: Storage Manager and Storage Control. Both components are included with the management software. Storage Manager is a standard management software capability that provides basic storage lifecycle management (Discovery, Inventory, Health, and Alerts). Storage Control is a preinstalled plug-in for the management software that expands storage support to mid-range and high-end storage devices. It is based on technology from IBM Tivoli® Storage Productivity Center.

IBM Flex System Enterprise Chassis and the management software offer these storage-management capabilities:

- ▶ Discovery of physical and virtual storage devices
- ▶ Support for virtual images on local storage across multiple chassis
- ▶ Inventory of physical storage configuration
- ▶ Health status and alerts
- ▶ Storage pool configuration
- ▶ Disk sparing and redundancy management
- ▶ Virtual volume management
- ▶ Support for virtual volume discovery, inventory, creation, modification, and deletion

Table 5-8 shows supported storage systems and their management tasks.

Table 5-8 Supported storage systems and management tasks

Storage system Management task	Flex System V7000	Storwize V3500 V3700 V7000	IBM SAN Volume Controller	IBM DS8000®	IBM XIV® Gen3
Discovery	Yes	Yes	Yes	Yes	Yes
Inventory collection	Yes	Yes	Yes	Yes	Yes
Monitoring	Yes	Yes	Yes	Yes	Yes
Physical topology	Yes	Yes	Yes	Yes	Yes
Logical topology	Yes	Yes	Yes	Yes	Yes
Server-to-storage mappings ^a	Yes	Yes	Yes	Yes	Yes
VMControl provisioning	Yes	Yes	Yes	Yes	Yes
NPIV support	Yes	Yes	Yes	No	Yes ^b
Monitor virtualized storage capacity ^c	Yes	Yes	Yes	No	No
Chassis Map	Yes	No	No	No	No
Software updates	Yes	No	No	No	No
IBM Electronic Service Agent™	Yes	No	No	No	No

a. Applies only to compute nodes running AIX, KVM, RHEL, and VMWare ESX.

b. VMControl provides limited support for IBM XIV storage that is connected through the N-Port ID Virtualization (NPIV) protocol to Power Systems compute nodes.

c. Ability to view and monitor capacity of IBM or non-IBM storage that is virtualized behind a V7000 array or the IBM SAN Volume Controller.

Table 5-9 shows supported Fibre Channel switches and their management tasks.

Table 5-9 Supported Fibre Channel I/O switches and management tasks

I/O module Management task	Flex System integrated switches		External switches
	FC3171 8 Gb FC	FC5022 16 Gb FC	IBM SAN24B-4 8 Gb FC
Discovery	Yes	Yes	Yes
Inventory collection	Yes	Yes	Yes
Monitoring	Yes	Yes	Yes
Alerting	Yes	Yes	Yes
Protocol configuration	Yes	No	No
Logical topology	Yes	Yes	Yes
NPIV support	Yes	Yes	Yes
Chassis Map	Yes	Yes	No
Software updates	Yes	Yes	No
Electronic Service Agent	Yes	Yes	No

With the Storage Control plug-in, you can manage an expanded set of storage subsystems and Fibre Channel switches. You can use Storage Control to discover and collect inventory, and monitor devices health.

Depending on the firmware levels of these devices, Storage Control supports native interfaces to the device, which simplifies configuration setup and improves device management reliability. These interfaces use Secure Shell (SSH) credentials. For information about configuring these credentials, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.director.cli.helps.doc/fqm0_r_cli_storage_cmds.html

After they are discovered, these devices are listed as being managed by Storage Manager in Flex System Manager.

Storage Control uses IBM Tivoli Storage Productivity Center technology. Therefore, several of the device support and operating environment conditions are related to IBM Tivoli Storage Productivity Center.

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

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

5.1.8 Planning for IBM Fabric Manager

IBM Fabric Manager (IFM) is a solution that you can use to quickly replace and recover compute nodes in your environment. It provides the following capabilities:

- ▶ I/O address assignment for initial compute node deployment and redeployment
- ▶ Slot-based I/O address assignment – Ethernet and FCoE MAC, FC WWNs, SAS WWNs, FC and SAS boot targets
- ▶ Pre-assignment allows LAN/SAN configuration before compute node installation
- ▶ Automatic reassignment on compute node swap (also called rip/replace)
- ▶ Failover Monitors or Event Automation Plans for automatic compute node failover
- ▶ Create standby compute node pools
- ▶ Configure boot target settings
- ▶ Provides I/O parameter and VLAN migration to standby compute nodes in case of hardware failure

IBM Fabric Manager is preinstalled on the FSM. It is also licensed as part of the FSM chassis license. If the FSM is not purchased, the stand-alone IBM Fabric Manager application can be licensed and installed in the environment to use against the Flex systems.

Because a management module failure results in a configuration loss, it is common to install a standby management module when using IBM Fabric Manager.

The IBM Fabric Manager configuration is not included in the management module configuration backup. The IBM Fabric Manager configuration is chassis based and does not transfer with the physical management module. When a management module is moved to a new chassis, it clears out its IBM Fabric Manager configuration, and the IBM Fabric Manager configuration must be reapplied on the new management module.

If the primary management module fails, the standby management module contains the IBM Fabric Manager configuration and takes over.

Boot from SAN: To take full advantage of the IBM Fabric Manager solution, consider setting up your server environment to boot from SAN.

For IBM Fabric Manager support on specific Flex System components, see the IBM Flex System Interoperability Guide:

<http://www.redbooks.ibm.com/fsig>

For more information about using IBM Fabric Manager, see the IBM Flex System Information Center:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.iofm.doc/dw1li_product_page.html

5.2 Planning for the management of virtualized environments

This section describes the planning steps for setting up the management of the specific virtual infrastructure that is deployed on IBM Flex System or IBM PureFlex System by using the Flex System Manager appliance. The following subtopics are covered:

- ▶ 5.2.1, “Virtualization and task supported” on page 100
- ▶ 5.2.2, “Planning for Linux KVM virtualization” on page 102
- ▶ 5.2.3, “Planning for PowerVM virtualization” on page 111
- ▶ 5.2.4, “Planning for VMware virtualization” on page 115
- ▶ 5.2.5, “Planning for Hyper-V virtualization” on page 118

5.2.1 Virtualization and task supported

System virtualization-management products and components in the IBM Flex System Enterprise Chassis integrate and interact to support the management of virtualized server, storage, and network resources.

Virtualization management is the software that enables the use and management of virtual server, storage, network, and image resources. With virtualization management, you can use your compute resources fully, deploy new workloads rapidly, monitor resource consumption, and maintain the availability of workloads. IBM Flex System Manager management software automates this complex set of tasks through administrator-defined policies. IBM VMControl is the single point of control for managing virtualized resources in one or more IBM Flex System Enterprise Chassis.

VMControl virtualization capabilities can help you simplify the management of virtual resources (server, storage, network, virtual appliance images) and pools of virtual resources. This simplification is achieved through the integrated provisioning of server, storage, and network resources when new workloads are deployed to the systems and system pools.

With VMControl, you can complete the following tasks:

- ▶ Discover existing image repositories in your environment and import external, standards-based images into your repositories as virtual appliances.
- ▶ Capture a running virtual server that is configured just the way you want, complete with guest operating system, running applications, and virtual server definition. When you capture the virtual server, a virtual appliance is created in one of your image repositories with the same definitions. This appliance can be deployed multiple times in your environment.

- ▶ Import virtual appliance packages that are in Open Virtualization Format (OVF) from the Internet or other external sources. After the virtual appliance packages are imported, you can deploy them in your data center.
- ▶ Deploy virtual appliances quickly to create new virtual servers that meet the demands of your changing business needs.
- ▶ Create, capture, and manage workloads.
- ▶ Create server system pools, which can be used to consolidate your resources and workloads into distinct and manageable groups.
- ▶ Deploy virtual appliances into server system pools.
- ▶ Manage server system pools, including adding hosts or more storage space and monitoring the health of the resources and the status of the workloads in them.
- ▶ Group storage systems together by using storage system pools to increase resource utilization and automation.
- ▶ Manage storage system pools by adding storage, editing the storage system pool policy, and monitoring the health of the storage resources.

There are three editions of VMControl:

- ▶ VMControl Express Edition manages virtual machines
- ▶ VMControl Standard Edition adds the ability to manage complete libraries of virtual images
- ▶ VMControl Enterprise Edition creates and enables the management of system pools. These pools are dynamic collections of computing resources that are used to support multiple virtual images that run concurrently

When you activate VMControl through the IBM Flex System Manager management software, a 90-day trial of VMControl Standard and Enterprise Editions begins. After the trial period ends, VMControl Express Edition remains, but VMControl Standard and Enterprise Edition are disabled. VMControl Standard and Enterprise Editions are available as an optional Features on Demand (Advanced Upgrade) in the IBM Flex System Manager management software.

VMControl discovery and inventory of virtual resources is supported for these hypervisor platforms:

- ▶ KVM RHEL 6.x
- ▶ VMware vCenter and VMware ESX
- ▶ Microsoft Hyper-V Server
- ▶ PowerVM

Table 5-10 on page 102 shows supported virtualization environments and their management tasks.

Table 5-10 Supported virtualization environments and management tasks

Virtualization environment	AIX and Linux ^a	IBM i	VMware ESXi with vCenter	Microsoft Hyper-V	Linux KVM
Management task					
Deploys virtual servers	Yes	Yes	Yes	Yes	Yes
Deploys virtual farms	No	No	Yes	No	Yes
Relocates virtual servers	Yes	Yes ^b	Yes	No	Yes
Maintenance mode	Yes	No	Yes	No	Yes
Imports virtual appliance packages	Yes	Yes	No	No	Yes
Captures virtual servers	Yes	Yes	No	No	Yes
Captures workloads	Yes	Yes	No	No	Yes
Deploys virtual appliances	Yes	Yes	No	No	Yes
Deploys workloads	Yes	Yes	No	No	Yes
Deploys server system pools	Yes	No	No	No	Yes
Deploys storage system pools	Yes	No	No	No	No

a. Linux on IBM Power Systems compute nodes.

b. Supported only for virtual servers that are running IBM i v7.1, TR4 PTF group SF99707 level 4, or later.

5.2.2 Planning for Linux KVM virtualization

The IBM FSM appliance can provide a set of capabilities to easily manage a KVM virtual infrastructure. It includes features such as high availability, virtual server relocation, capture, deployment, import appliance, and network multitenancy.

This section describes the requirements and support for the Linux kernel-based virtual machine (KVM) virtualization environment on IBM Flex System Manager VMControl.

There are two implementation models to manage a KVM virtual infrastructure:

- ▶ SAN storage-based model: A supported storage system that acts as the shared storage device.
- ▶ NFS storage-based model: The NFS server that acts as the shared storage device.

Network File System storage-based model

The Network File System (NFS) storage-based model has the following requirements:

- ▶ IBM Flex System Manager VMControl is activated
- ▶ An NFS x86_64 Red Hat Enterprise Linux (RHEL) server is set up and configured. Figure 5-4 shows the KVM virtualization environment with NFS storage.

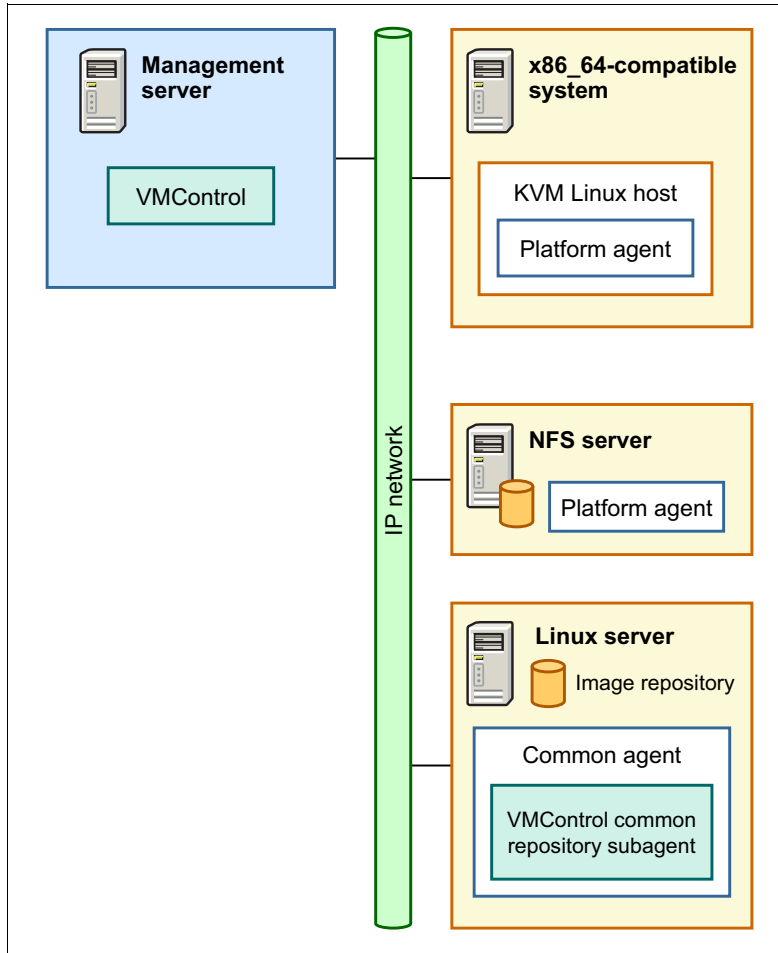


Figure 5-4 KVM virtual environment with NFS storage

- ▶ At least one NFS export on the NFS server is defined:
 - For image and disk inventory to work, the export path must end with /images, for example, /share/kvm/images.
 - If you are not setting up additional security in your environment, use the `no_root_squash` export option. For example, `cat /etc/exports` as shown Figure 5-5.

```
[root@kvm15 ~]# cat /etc/exports
/nfs/kvm/images 192.168.0.0/255.255.0.0 (rw,no_root_squash,sync,no_subtree_check)
[root@kvm15 ~]# █
```

Figure 5-5 `no_root_squash` export option

If you cannot change your NFS export setup, you can have image files that are inventoried from an export path ending in something other than /images. To do so, complete the following steps:

1. In the file /opt/ibm/director/lwi/conf/overrides/USMIKernel.properties, add a line for the following property: director.services.extendeddiscovery.nfs.suffix. For example, director.services.extendeddiscovery.nfs.suffix=/img-kvm.

This addition results in inventorying the image files within NFS export paths that end in /img-kvm instead of the default, /images.

2. Restart the IBM Flex System Manager after adding or changing the USMIKernel.properties file.

Remember: For consistency, image and disk files that are stored on NFS must have a .dsk, .img, or .raw extension.

3. Ensure that the NFS services are started. For example, you can run the command `service nfs start` as shown Figure 5-6.

```
[root@kvm15 ~]# service nfs restart
Shutting down NFS mountd:           [ OK ]
Shutting down NFS daemon:          [ OK ]
Shutting down NFS quotas:          [ OK ]
Shutting down NFS services:        [ OK ]
Starting NFS services:              [ OK ]
Starting NFS quotas:               [ OK ]
Starting NFS daemon:               [ OK ]
Starting NFS mountd:               [ OK ]
[root@kvm15 ~]# █
```

Figure 5-6 NFS service restart

4. The administrator must perform the following prerequisite tasks:
 - a. KVM Platform Agent is downloaded and installed as addressed in 9.2, “KVM platform agent installation” on page 320.
 - b. The NFS server is discovered, accessed, and inventoried by IBM Flex System Manager.
 - c. The image repository is set up as explained in 9.3, “Image repository for KVM” on page 328:
 - d. IBM Flex System Manager Common Agent is installed as explained in 9.3, “Image repository for KVM” on page 328.”
 - e. VMControl Common Repository subagent is installed as explained in 9.3, “Image repository for KVM” on page 328

- f. The shared NFS exported storage is mounted on the Image Repository server as shown in Figure 5-7.

```
[root@kvm2 ~]# mount
/dev/mapper/vg_kvm2-lv_root on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw,rootcontext="system_u:object_r:tmpfs_t:s0")
/dev/sda1 on /boot type ext4 (rw)
/dev/mapper/vg_kvm2-lv_home on /home type ext4 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
192.168.70.15:/nfs/kvm/images on /nfs/export type nfs (rw,vers=4,addr=192.168.70.15,clientaddr=192.168.70.2)
/nfs/export/RHEL6.1-20110510.1-Server-x86_64-DVD1.iso on /mnt/iso/1 type iso9660 (rw,loop=/dev/loop0)
[root@kvm2 ~]# █
```

Figure 5-7 NFS export that is mounted on the image repository server

- g. The image repository server is discovered and inventory is collected as shown Figure 5-8. For more information, see 9.3, “Image repository for KVM” on page 328.




Figure 5-8 Images Repositories window

The image repository is created from VMControl. For instructions to create an image repository, as shown chapter 9.3, “Image repository for KVM” on page 328

- h. One or more RHEL KVM hosts are set up and available:
- The KVM Platform Agent is downloaded and installed on the KVM hosts. For more information, see 9.2, “KVM platform agent installation” on page 320.
 - The KVM hosts are discovered, accessed, and inventoried from your IBM Flex System Manager. For more information, see 9.4, “Creating KVM storage system pools” on page 350.

- Set up storage by right-clicking a KVM host, selecting **System Configuration**, and selecting **Edit Host**, as shown in Figure 5-9.




Figure 5-9 Edit KVM host

- Click **Storage Pools**.
- Select the pool to create your virtual server disk as shown in Figure 5-10.




Figure 5-10 Create a KVM disk in a storage pool

Remember: When you configure KVM hosts, specify the fully qualified name as the host name, for example, `hostname.company.com`. Use the `hostname` command on the host to determine the system name. If the host is not configured with its fully qualified host name, the IBM Key Exchange providers might fail to exchange SSH keys during relocation. Also, ensure that the host name and IP address for the target system are recorded correctly in the Domain Name System (DNS) records.

SAN storage-based model

The SAN storage configuration looks more complex than the NFS solution but the block storage-based model offers better performance and more flexibility.

The picture in Figure 5-11 shows a KVM virtualization environment with SAN storage.




Figure 5-11 KVM virtualization environment with SAN storage

Implementing a SAN storage-based model has the following requirements:

- ▶ IBM Flex System Manager VMControl is activated.
- ▶ The Fibre Channel storage network is correctly cabled and configured with the appropriate Fibre Channel switches. KVM virtualization with VMControl supports only SAN storage over Fibre Channel. Typically, one of the fabric switches is configured with the zoning information. Additionally, VMControl requires that the Fibre Channel network has hard zoning enabled.

- ▶ One or more RHEL KVM hosts are set up and available:
 - Ensure that the RHEL KVM host is connected to the Fibre Channel network with a supported adapter.
 - The KVM Platform Agent is downloaded and installed.
 - KVM hosts are discovered, accessed, and inventoried from your IBM Flex System Manager.
- ▶ The SAN storage controllers (storage subsystems) are configured and storage pools are set up with the wanted storage space and RAID levels for virtual disk images. Neither VMControl or Storage Control will provision these RAID storage pools for you.

Requirement: Host definitions must include all worldwide port names (WWPNs) for the host (or hosts) they represent. The WWPNs are needed even if some ports are not physically connected or active. This process avoids the potential problem of mapping a single volume under different LUN IDs to the same host.

For example, assume that a KVM host has a Fibre Channel card with host ports WWPN1 and WWPN2. An IBM Storwize V7000 storage subsystem defines host definition KVM_Host1 for that host. Then, the host definition must contain both WWPN1 and WWPN2.

- ▶ A Fibre Channel switch provider is configured in environments where Brocade switches are used. This role can be handled by the Brocade SMI-S Agent or the Brocade Network Advisor.
- ▶ Storage subsystems, storage pools, and the Fibre Channel switch fabric are discovered and inventoried by Flex System Manager for shared access from endpoints in the KVM environment. These endpoints include KVM hosts and image repository servers as shown in Figure 5-11 on page 107.
 - Encryption keys are needed for the IBM Storwize V7000. The encryption keys are used for discovery enablement and to enable IBM FlashCopy®. If necessary, generate an encryption key file in OpenSSH format for your SAN device and store this file on your Flex System Manager server. For more information about generating an encryption key file for your storage, see your SAN storage device’s documentation.
 - Use the `manage7000` command to define your storage data source. This command pushes your pub key and enables SAN storage discovery and inventory collection through Storage Control. For more information, see 6.12, “Discover and manage external Storwize V7000” on page 234.

Tip: If you have many switches, zones, or zone groups that are defined on a fabric switch, the inventory collection task might show an error after the default Flex System Manager timeout period expires. However, zone inventory collection continues to run in the background.

- ▶ The image repository is set up and meets all of the following requirements. The image repository is used for storing and deploying virtual appliances.

The image repository server is connected to the Fibre Channel network with a supported Fibre Channel HBA. For more information about adapters, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.sdmm.adv.helps.doc/fnc0_r_network_ctrl_planning.html

- Common Agent is installed on your image repository server. For more information, see 9.3, “Image repository for KVM” on page 328.
 - VMControl Common Repository subagent is installed on the image repository server. For more information, see 9.2, “KVM platform agent installation” on page 320.
 - The image repository server is discovered and inventory is collected on it.
 - The image repository is created from VMControl. For more information about creating an image repository, see 9.3, “Image repository for KVM” on page 328.
- ▶ Verify that Flex System Manager and VMControl can manage the environment.
 - Run **dumpstcfg** to see the storage configuration information.

Example output:

```
Host Accessible Containers
-----
NAME: STORAGE SUBSYSTEM/POOL
IBM Host01: Storwize V7000-2076/RAID5_Pool_KVM
Storwize V7000-2076/RAID0_Pool_800GB
```

IBM Host01 is a KVM host, Storwize V7000-2076 is the storage subsystem, and the KVM host can access both the RAID5_Pool_KVM and RAID0_Pool_800GB storage pools. This output indicates that inventory collection has correctly modeled connectivity from the host to the storage.

Additionally, verify that the image repository server can access the SAN storage containers in the same way.

- Run **testluncreate** to verify that the SAN storage configuration is complete. The command tries to allocate a volume on a subsystem and storage pool, then attach it to a host. This host can be your image repository server.
- If **dumpstcfg** or **testluncreate** shows problems, there might be a configuration problem. Correct the problem and collect inventory again on each endpoint, farm, storage, and switch resource.

Supported hosts, Linux versions, and firmware versions

KVM virtualized environments must run on X-Architecture compute nodes. Hosts require Red Hat Enterprise Linux version 6.2, 6.3, or 6.4 with KVM installed.

Supported networks

VMControl supports the following network configurations for the KVM hypervisor:

- ▶ Virtual Ethernet Bridging (VEB)
- ▶ Virtual Ethernet Port Aggregator (VEPA) network (requires IBM Flex System Manager Network Control and that the host is in a network system pool)
- ▶ Limited support for KVM hypervisor networks

Tip: Use paravirtualized (virtio) drivers for enhanced performance. Use Virtio and e1000 model configurations for virtual network server adapters.

Supported storage

The model includes the following image repository and virtual disk storage options:

- ▶ NFS version 3 server running on RHEL version 6.2 and 6.3.
- ▶ NFS version 3 server running on RHEL version 6.2, 6.3, and 6.4 with KVM installed.
- ▶ Supported SAN devices. For more information about storage products support, see this website:
http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.common.nav.doc/network_integration_planning.html

Supported tasks

In the KVM virtualization environment, you can perform these tasks:

- ▶ Create and delete NFS storage pools on a host
- ▶ Create and delete NFS or SAN virtual disks
- ▶ Suspend or resume virtual servers and workloads (without release of resources)
- ▶ Create, edit, and delete virtual servers
- ▶ Power operations for virtual servers
- ▶ Relocate virtual servers
- ▶ Turn maintenance mode on and off for hosts that are in server system pools
- ▶ Import a virtual appliance package that contains one or more raw disk images
- ▶ Capture a workload or virtual server into a virtual appliance
- ▶ Deploy a virtual appliance package to a new virtual server with hardware and product customizations
- ▶ Deploy a virtual appliance package to an existing virtual server with adequate resources
- ▶ Start, stop, and edit a workload
- ▶ Create, edit, and delete server system pools
- ▶ Create, edit, and delete network system pools (if you are using IBM Flex System Manager Network Control with VMControl)
- ▶ Adjust the virtualization monitor polling interval for KVM by using the `KvmPlatformPollingInterval` parameter

Tip: To enable remote control access on your KVM, follow the instructions at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_t_access_kvm_remotely.html

KVM requirements

In addition to the packages required by the KVM platform agent, the `genisoimage.x86_64` package must also be installed for VMControl support. For more information, see 9.2, “KVM platform agent installation” on page 320.

Remember: These packages might be available from your installation software.

The following commands open required ports in the IPv4 firewall on the KVM host:

- ▶ `iptables -A INPUT -p tcp --dport 427 -j ACCEPT`
- ▶ `iptables -A INPUT -p udp --dport 427 -j ACCEPT`
- ▶ `iptables -A INPUT -p tcp --dport 22 -j ACCEPT`

- ▶ `iptables -A INPUT -p tcp --dport 15988 -j ACCEPT`
- ▶ `iptables -A INPUT -p tcp --dport 15989 -j ACCEPT`
- ▶ `service iptables save`

Considerations:

- ▶ The SSH service must be configured and running on the KVM host so that an SSH remote service access point for port 22 gets created for each host. These access points are in addition to the CIM RSAP on ports 15988 and 15989.
- ▶ When a SAN storage solution is being used, you must have at least several megabytes of free file system space under `/var/opt/ibm` and `/var/lib/libvirt` on the KVM host. The user that is employed to request access to the host from ISD must have authority to write to these directories.

Restrictions

For more information, see the Restrictions section at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_kvm.html

5.2.3 Planning for PowerVM virtualization

This section describes the requirements and support for AIX and Linux virtual appliances, virtual servers, and workloads in the Power Systems virtualization environment.

There are two types of architectures to implementing PowerVM base virtualization through Flex System Manager:

- ▶ Requirements and support for AIX using Network Installation Manager (NIM)
- ▶ Requirements and support for AIX, IBM i, and Linux using Storage Copy Services (SCS)

Requirements and support for AIX using Network Installation Manager

This section describes the requirements and support for AIX virtual appliances, virtual servers, and workloads in a Power Systems virtualization environment that relies on AIX NIM.

Requirements for AIX using NIM

The following diagram shows an example Power Systems virtualization environment for AIX virtual appliances, virtual servers, and workloads that rely on NIM.

Figure 5-12 shows an example Power Systems virtualization environment for AIX virtual appliances, virtual servers, and workloads that rely on NIM.

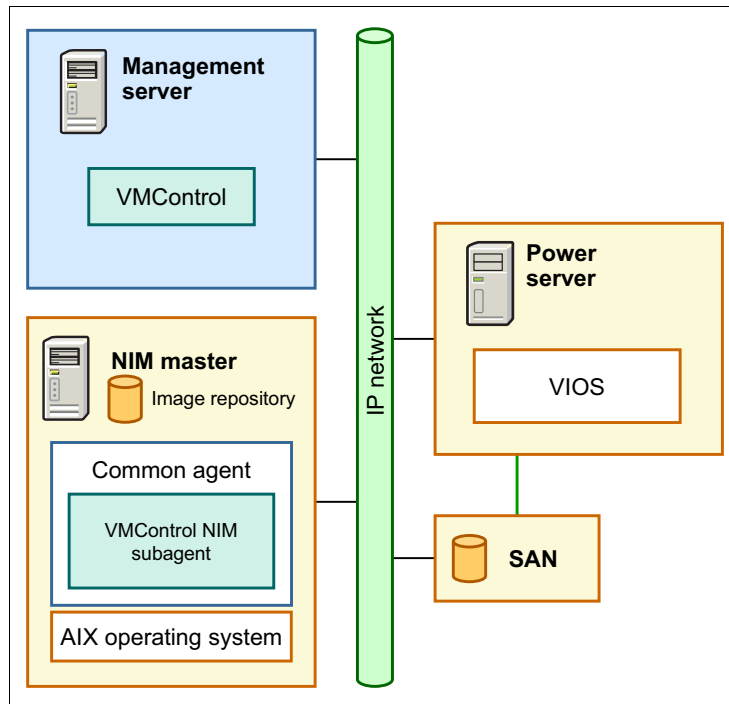


Figure 5-12 AIX using NIM system diagram

Implementation of a virtualization environment based on PowerVM has the following requirements:

- ▶ IBM Flex System Manager is installed on a supported server.
- ▶ IBM Flex System Manager VMControl Standard Edition or IBM Flex System Manager VMControl Enterprise Edition is activated.
- ▶ At least one NIM master is available.
- ▶ IBM Flex System Manager Common Agent and the VMControl NIM subagent are installed on the NIM master.

IBM Flex System Manager recognizes this NIM master as a VMControl image repository. The `/export/nim` filesystem in which the virtual appliances are stored must not be NFS mounted to the NIM master. The NIM master exports this file system itself, and NFS does not support export of a mounted file system.

Remember: The image repository is shown as a stand-alone server in the diagram. However, the image repository can also be on the same Power Systems server that hosts the AIX virtual servers that you can capture from and deploy to using VMControl.

- ▶ At least one IBM POWER7 compute node exists to host virtual servers that you can capture from and deploy to using VMControl.

- ▶ The IBM Power server is typically attached to a SAN as shown in the diagram. The SAN is used for the virtual disks of the virtual servers that are hosted by the IBM Power server. If you expect to use VMControl Enterprise Edition server system pools or do virtual server relocation on your own, a SAN is required. If not, disks that are local to the IBM Power server and virtualized by the Virtual I/O Server (VIOS) can be used as an alternative.
- ▶ Though not shown in the diagram, multiple VIOS virtual servers and multipath I/O (MPIO) are supported.

Supported operating systems and firmware versions

You must use the following AIX and firmware versions in this environment:

- ▶ NIM master: The NIM master must be AIX 6.1 TL03 or newer.

Requirement: The level of AIX on the NIM master must be the same or higher than the level of AIX on the virtual servers that you capture or the virtual appliances that you deploy.

- ▶ Virtual I/O Server (VIOS): For POWER7, use a minimum of VIOS 2.2.1.0 and all available updates.
- ▶ Virtual appliances: You can capture any AIX Version 5.3, AIX Version 6.1, or AIX Version 7.1 virtual server or workload as a virtual appliance. You can import or deploy any AIX Version 5.3, AIX Version 6.1, or AIX Version 7.1 virtual appliance.

Supported tasks

In this environment, you can perform these tasks:

- ▶ Create, edit, and delete virtual servers
- ▶ Relocate virtual servers
- ▶ Import a virtual appliance package that contains an AIX mksysb image
- ▶ Capture an AIX workload or virtual server, an AIX mksysb image file or NIM resource, or an IX lpp_source directory or NIM resource
- ▶ Deploy an AIX mksysb or lpp_source virtual appliance
- ▶ Group virtual servers to create a workload
- ▶ Start, stop, and edit a workload
- ▶ Create, edit, and delete system pools

Requirements for AIX, IBM i, and Linux using storage copy services

This section describes the requirements and support for AIX, IBM i, and Linux virtual appliances, virtual servers, and workloads in a Power Systems virtualization environment that relies on storage copy services (SCS).

Figure 5-13 shows an example Power Systems virtualization environment for AIX, IBM i, and Linux virtual appliances, virtual servers, and workloads that rely on SCS.

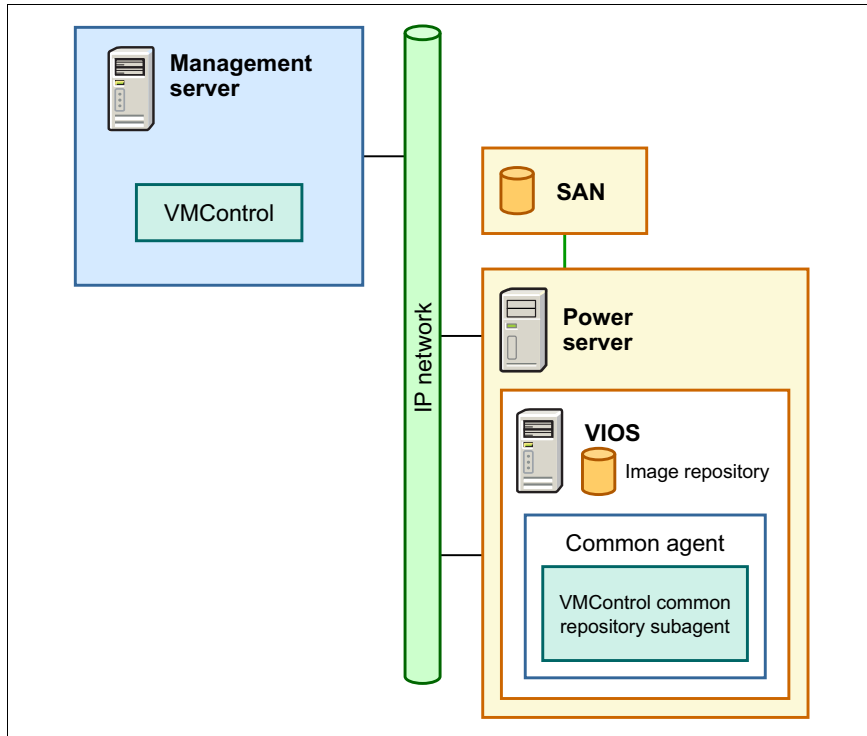


Figure 5-13 AIX, IBM i, and Linux using storage copy services system diagram

These environments have the following requirements:

- ▶ IBM Flex System Manager is installed on a supported server.
- ▶ IBM Flex System Manager VMControl Standard Edition or IBM Flex System Manager VMControl Enterprise Edition is activated.
- ▶ A Virtual I/O Server (VIOS) virtual server exists on an IBM Power server to host the image repository. The repository is used to store the raw disk images that are associated with your AIX, IBM i, and Linux virtual appliances.

Tip: You can have multiple repositories. However, repositories that are on separate IBM Power servers require special configuration. The image repository virtual servers must have access through a VIOS to the same shared SAN as the AIX, IBM i, and Linux virtual servers that they will capture and deploy.

- ▶ The IBM Flex System Manager Common Agent and VMControl Common Repository subagent are installed on the VIOS that you want to use as an image repository.
- ▶ At least one IBM POWER7 compute node exists to host virtual servers that you can capture from and deploy to using VMControl.

Consideration: If manual or automated virtual server relocation capabilities are needed, multiple IBM Power 6 or 7 servers are required.

- ▶ All AIX, IBM i, and Linux virtual servers to be captured from or deployed to using VMControl have their storage allocated from the SAN. They also must be provided through one or more VIOS virtual servers. These virtual servers must use virtual Ethernet connections that are provided through one or more VIOS virtual servers. These virtual servers must not have any physical devices that are allocated from the IBM Power server.
- ▶ For Virtual I/O Server Version 2.2, any virtual servers that you capture and any virtual appliances you deploy use the same storage pool as the image repository in which you store the virtual appliances.

Supported operating systems and firmware versions

You must use the following operating systems and firmware versions in this SCS-based Power Systems virtualization environment:

- ▶ IBM Flex System Manager: You can use any IBM Flex System Manager with VMControl Standard Edition or VMControl Enterprise Edition activated.
- ▶ Virtual I/O Server (VIOS): For POWER7, use a minimum of VIOS 2.2.1.0 and all available updates.
- ▶ IBM Power firmware: For POWER7 processor-based servers, use a minimum of FW7.2 and all available updates.

Supported tasks

In a Power Systems virtualization environment for AIX, IBM i, and Linux that relies on SCS, you can perform the following tasks:

- ▶ Create, edit, and delete virtual servers
- ▶ Import virtual appliance packages that contain an AIX, IBM i, or Linux raw disk image
- ▶ Capture an AIX, IBM i, or Linux workload or virtual server (logical partition)
- ▶ Deploy an AIX, IBM i, or Linux raw disk image virtual appliance
- ▶ Group virtual servers to create a workload
- ▶ Start, stop, and edit a workload

In a Power Systems virtualization environment for AIX and Linux that relies on SCS, you can perform the following additional tasks:

- ▶ Relocate virtual servers
- ▶ Create, edit, and delete system pools

5.2.4 Planning for VMware virtualization

This section describes the requirements and support for the VMware virtualization environment on IBM Flex System Manager VMControl.

VMware ESX and VMware ESXi hosts managed by VMware vCenter

Figure 5-14 shows a virtualization environment with VMware vCenter managing VMware ESX and VMware ESXi hosts.

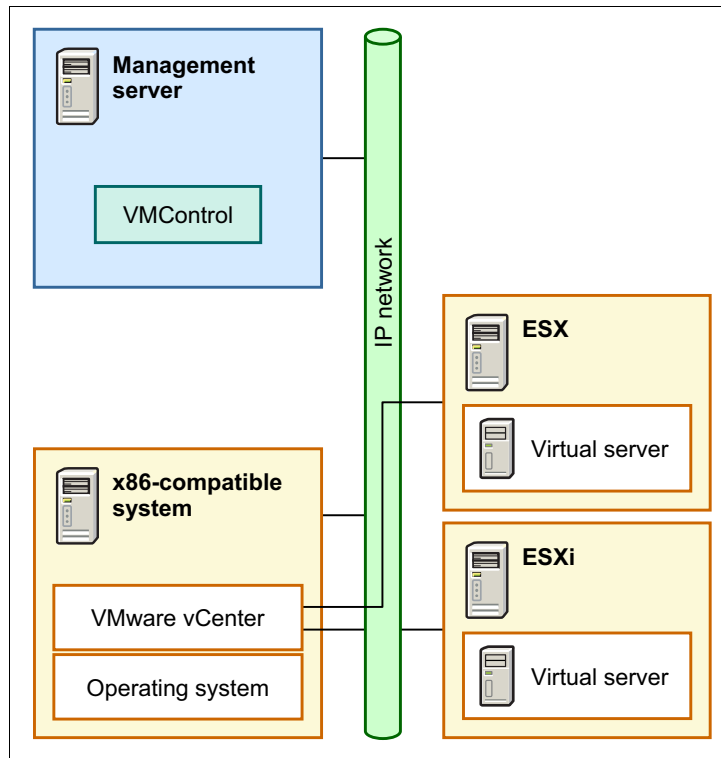


Figure 5-14 Diagram of VMware vCenter virtualization environment

This scenario has the following requirements:

- ▶ IBM Flex System Manager is installed on a supported server.
- ▶ IBM Flex System Manager VMControl is activated.

Remember: To start the VMware Infrastructure Client or the VMware vSphere Client from IBM Flex System Manager VMControl, the client must be installed on the IBM Flex System Manager system. It must also be on any system that you use to log in to the IBM Flex System Manager web interface.

- ▶ VMware vCenter is installed on an x86-compatible system. IBM Flex System Manager and VMControl require that the operating system (OS) that VMware vCenter is running on is an x86-compatible system with an OS based on Microsoft Windows.
- ▶ VMware ESXi exists to host virtual servers that you can manage by using VMControl. VMware ESXi is managed by VMware vCenter.
- ▶ VMware vCenter system is discovered and the request access task is complete. After the request access task completes, the Configure Access task shows the vCenter protocol in OK state.

Requirement: If you installed VMware vCenter with a non-default port number, you must create a VMware vCenter Server Discovery profile by using the Discovery Profile wizard. Specify the unique port number in the profile that you create. Then, use the profile to discover the VMware vCenter system. For more information, see “Managing discovery profiles” at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.director.discovery.helps.doc/fqm0_t_managing_discovery_profiles.html

Supported versions

VMControl supports the following virtualization software:

- ▶ VMware vCenter 4.x (capable of managing the following supported hosts): VMware ESX 4.x and VMware ESXi 4.x
- ▶ VMware vCenter 5.x (capable of managing the following supported hosts): VMware ESX 4.x, VMware ESXi 4.x, and VMware ESXi 5
- ▶ VMware ESX 4.x stand-alone software

Supported tasks

In the VMware vCenter virtualization environment, you can perform the following tasks:

- ▶ Create, edit, and delete virtual servers
- ▶ Create a DataCenter or Cluster by using the Create Virtual Farm wizard
- ▶ Add a host to a DataCenter or Cluster by using the Add host to farm function
- ▶ Remove a host from a DataCenter or Cluster by using the Remove host from farm function
- ▶ Relocate virtual servers
- ▶ Put a host into maintenance mode
- ▶ Remove a host from maintenance mode

VMware ESX stand-alone software

Figure 5-15 shows the VMware ESX virtualization environment.




Figure 5-15 Diagram of the VMware ESX virtualization environment

This scenario has the following requirements:

- ▶ IBM Flex System Manager is installed on a supported server.
- ▶ IBM Flex System Manager VMControl is installed on the IBM Flex System Manager.
- ▶ VMware ESX exists to host virtual servers that you can manage by using VMControl.

- ▶ IBM Flex System Manager Common Agent and the VMware ESX subagent are installed on the VMware ESX system.
- ▶ Using IBM Flex System Manager Network Control network system pools with VMware to provision new network configurations and provide automated network relocation requires special configuration. SNMP must be enabled on all the VMware hosts to be included in the network system pool. To enable SNMP, log on as root and issue the following command:


```
service snmpd start
```

Supported versions

VMControl supports the VMware ESX 4.x stand-alone virtualization software.

Supported tasks

In a VMware ESX virtualization environment, you can perform the following tasks:

- ▶ Create, edit, and delete virtual servers
- ▶ Create, edit, and delete virtual farms
- ▶ Relocate virtual servers

5.2.5 Planning for Hyper-V virtualization

Flex System Manager is able to manage basic tasks for the Microsoft Hyper-V hypervisor. You can start, stop, restart, suspend, create, and delete your virtual servers that are running on Microsoft hypervisors with the same tool that manages the other hypervisors on the market.

This section describes the requirements and support for the Windows Server 2008 and Windows Server 2012 Enterprise, Standard, and Datacenter x64 Editions with Hyper-V role enabled virtualization environment on IBM Flex System Manager VMControl.

Figure 5-16 shows the Windows Server 2008 and Windows Server 2012 with Hyper-V role enabled virtualization environment that is managed by IBM Flex System Manager.




Figure 5-16 Hyper-V architecture that is managed by FSM

This scenario has the following requirements:

- ▶ IBM Flex System Manager is available in your PureFlex chassis.
- ▶ IBM Flex System Manager VMControl is activated.

- ▶ Windows Server 2008 Standard, Enterprise, and Datacenter x64 Editions with Hyper-V role enabled is installed on an x86-compatible system.
- ▶ Windows Server 2012 Standard and Datacenter x64 Editions with Hyper-V role enabled is installed on an x86-compatible system.
- ▶ IBM Flex System Manager Common Agent or Platform Agent is installed on the host. For more information, see 12.1, “Initial setup tasks for a Hyper-V node” on page 556.

Supported versions

VMControl supports the following software virtualization versions:

- ▶ Windows Server 2008 Standard, Enterprise, and Datacenter x64 Editions with Hyper-V role enabled, Release 2 virtualization software.
- ▶ Windows Server 2012 Standard and Datacenter x64 Editions with Hyper-V role enabled, Release 3 virtualization software.

Supported tasks

In this environment, you can create, edit, and delete virtual servers, then start, stop, restart, and suspend your virtual servers.



IBM Flex System Manager initial configuration

This chapter describes the initial setup steps that are required for IBM Flex System Manager (FSM)-based systems management. This process includes the following tasks:

- ▶ Initial configuration of the FSM
- ▶ Discovery and inventory collection:
 - Chassis components
 - Operating systems
 - External storage devices
- ▶ Firmware updates:
 - The FSM
 - Chassis Management Module (CMM)
 - Compute nodes
 - I/O modules
- ▶ Operating system deployment

This chapter discusses the following topics:

- ▶ 6.1, “IBM Flex System Manager Setup Wizard” on page 123
- ▶ 6.2, “Updating Flex System Manager” on page 137
- ▶ 6.3, “Selecting chassis to manage” on page 138
- ▶ 6.4, “Configuring centralized user management” on page 142
- ▶ 6.5, “Configuring chassis components” on page 144
- ▶ 6.6, “Configuring compute nodes using Configuration Patterns” on page 150
- ▶ 6.7, “Deploying compute node images” on page 168
- ▶ 6.8, “System discovery, access, and inventory collection” on page 176
- ▶ 6.9, “Updating chassis components” on page 189

- ▶ 6.10, “Manage Feature-on-Demand keys” on page 216
- ▶ 6.11, “Flex System V7000 Storage Node initial configuration” on page 223
- ▶ 6.12, “Discover and manage external Storwize V7000” on page 234
- ▶ 6.13, “Overview of Flex System V7000 and Storwize V7000 systems management (Storage Control)” on page 241
- ▶ 6.14, “External Fibre Channel SAN switch discovery” on page 247
- ▶ 6.15, “Configuring network parameters (Network Control)” on page 256

6.1 IBM Flex System Manager Setup Wizard

IBM Flex System Manager (FSM) is an appliance that comes with all required software preinstalled. When this software stack is started for the first time, a startup wizard is initiated. This wizard guides you through the required configuration process, such as licensing agreements and Transmission Control Protocol/Internet Protocol (TCP/IP) configuration for the appliance.

When configuration is complete, the FSM is ready to manage the chassis in which it is installed and other chassis, up to four. After the chassis is managed, individual components, such as compute nodes and switches, can also be managed.

Requirement: At the time of writing, IBM Flex System Manager is required for any configuration that contains a Power Systems compute node.

It is also anticipated that IBM Flex System Manager is preconfigured to manage the initial chassis. In that event, the steps in this section are not required unless IBM Flex System Manager is being reinstalled.

FSM is based on an x86 compute node, and it has the same options for obtaining an initial console. You can use the Integrated Management Module II (IMM2) remote console. Or, use the supplied dongle and front port on the FSM node to connect directly to a keyboard, display, and mouse or a console manager unit.

To monitor the FSM startup process, connect a console to the FSM management node:

1. From the IBM Chassis Management Module web interface, right-click the **Flex System Manager node** in the Chassis Map, then select **Launch Compute Node Console**, as shown in Figure 6-1.




Figure 6-1 Launching Flex System Manager Console

2. The Launch Node Console window opens. Select **HTTPS** in the Protocol field and click **Launch**, as shown in Figure 6-2.

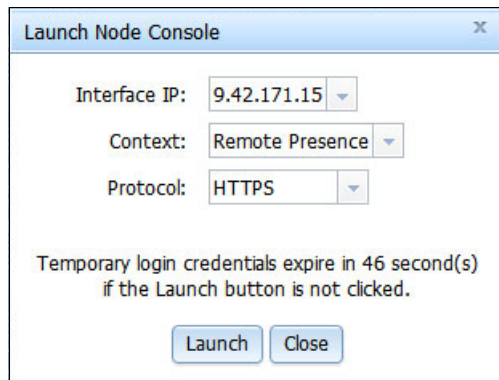


Figure 6-2 Launch Node Console window

Accept all security certificate exceptions if any. FSM's Remote Control window opens.

3. In the Remote Control window, click **Start remote control in single-user mode**, as shown in Figure 6-3. Clicking this button starts a Java applet on the local desktop that will be a console session to the FSM.




Figure 6-3 Starting remote console from IMM2

Figure 6-4 shows the Java console window opened to the FSM appliance before powering on.



Figure 6-4 FSM console in power off state

4. The FSM can be powered on from several locations, including the physical power button on the FSM, or from the CMM. For this example, the **Tools** → **Power** → **On** option from the remote console menu, as shown in Figure 6-5, is used.




Figure 6-5 Powering on the FSM from the remote console session

While the FSM powers up and boots, the process can be monitored. No input is accepted until the License Agreement window, which is shown in Figure 6-6, is displayed.



Figure 6-6 FSM license agreement

5. Click **I agree** to continue. The startup wizard's Welcome window opens as shown in Figure 6-7. Click **Next**.

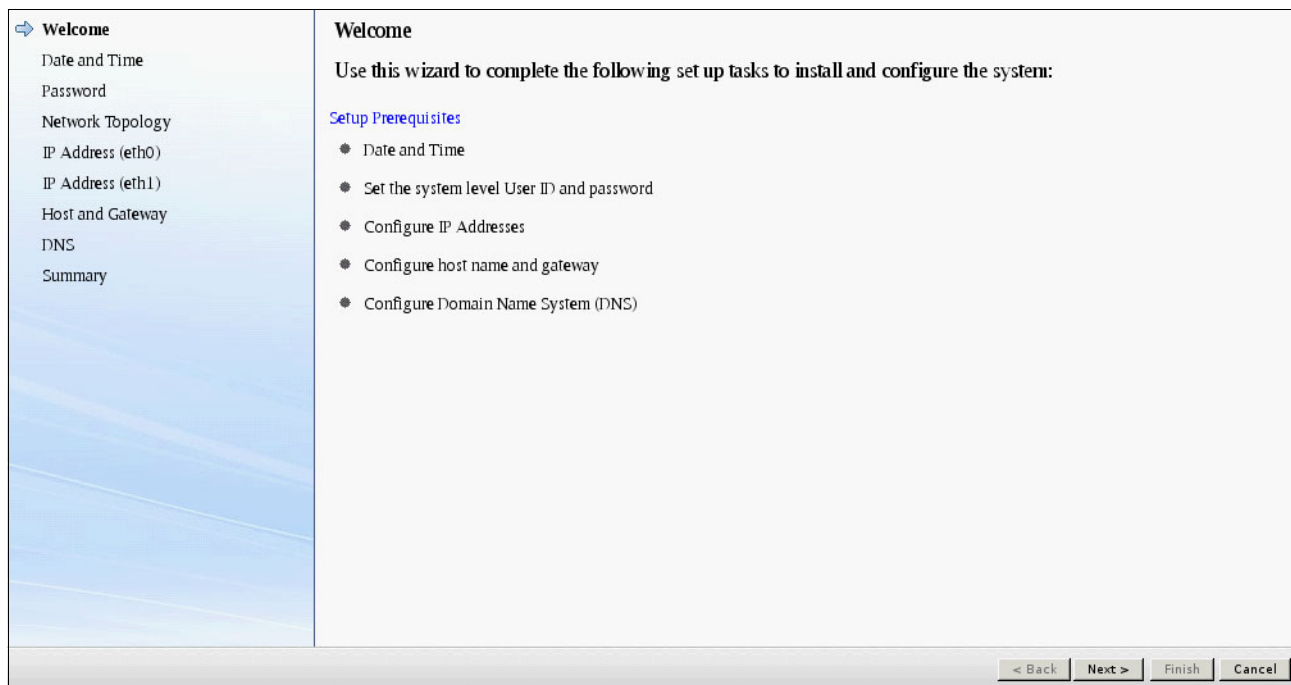


Figure 6-7 FSM Welcome window

- From the Date and Time window that is shown in Figure 6-8, set the time, date, time zone, and Network Time Protocol server, as needed. Click **Next**.

Figure 6-8 Setting the FSM date and time

- Create a user ID and password for accessing the GUI and CLI, as shown in Figure 6-9. User ID and password maintenance, including creating more user IDs, is available in IBM Flex System Manager after the startup wizard completes. Click **Next**.

Figure 6-9 FSM system level user ID and password step

- Network topology options include separate networks for management and data, or a single network for both data and management traffic from the chassis. Generally, have separate management and data networks. To simplify this example, a combined network is configured, as shown in Figure 6-10.

Welcome

Date and Time

Password

➔ **Network Topology**

IP Address (eth0)

IP Address (eth1)

Host and Gateway

DNS

Summary

Network Topology

There are two possible network topologies that can be configured for Flex System Manager (refer to network diagrams below). You can either separate your management and data networks by configuring eth0 for management and eth1 for data, or configure only eth0 to use the same network for both data and management.

[Learn more about planning and configuring your network topology.](#)

Select your current network topology configuration:

One network for both data and management traffic.

Separate networks for data and management traffic.

Single Management and Data Network

Enable IPv6 Stateless Auto Configuration

[Learn more](#)

Figure 6-10 FSM network topology options

- Click **Next** to continue to the IP network configuration.

The LAN adapter (eth0) is from the FSM management network that allows FSM to communicate on the chassis management network. Traffic from this adapter flows through the Chassis Management Module and uses the CMM physical connection to the network.

The LAN adapter for data network (eth1) is not available because we selected **One network for both data and management traffic** in the previous step. Eth1 represents the integrated Ethernet ports or LAN on motherboard (LOM) on the FSM management node. Traffic from this adapter flows through the Ethernet switch in the first I/O switch bay of the chassis.

The IP configuration for the eth0 interface is shown in Figure 6-11. This window allows the selection of Dynamic Host Configuration Protocol (DHCP) or static IP options for IPv4 and IPv6 addressing. Select the options that you want, enter the information as required, then click **Next**.

The screenshot shows a configuration window titled "Configure IP Address for eth0". On the left is a sidebar with the following menu items: Welcome, Date and Time, Password, Network Topology, **IP Address (eth0)** (highlighted with a blue arrow), Host and Gateway, DNS, and Summary. The main area contains the following text and controls:

Configure IP Address for eth0
Configure the IP addresses for LAN adapter (eth0). If the adapter is configured for DHCP and is unable to get an IP address, the management server will not start.

LAN interface address: 5C:F3:FC:5F:5E:8C eth0 (Management/Data Network)
Configured IPv6 Link-Local address: fe80:0:0:5ef3:fcff:fe5f:5e8c

IPv4 address:

- Obtain an IP address automatically
- Use the following IPv4 address:
 - Static IP address:
 - Network mask:

IPv6 address:

- Use DHCPv6 to configure IP settings
- Use the following IPv6 address:
 - Specify new static IPv6 address information and click Add:
 - IPv6 address:
 - Prefix length:

At the bottom right of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Figure 6-11 FSM LAN adapter configuration

10. After IP address assignment, the host name and gateway are configured as shown in Figure 6-12. Enter the host name, domain name, and default gateway address. Ensure that the IP address and the default gateway adapter are correct. Click **Next** to continue.

Tip: The host name of the FSM must be available on the domain name server.

Welcome	Configure Host and Gateway
Date and Time	Specify host name, domain name and the default gateway address.
Password	
Network Topology	*Host name: fsm1
IP Address (eth0)	*Domain name: its0.ral.ibm.com
⇒ Host and Gateway	*Default Gateway address: 9.42.170.1
DNS	
Summary	
	< Back Next > Finish Cancel

Figure 6-12 FSM host name and gateway configuration

11. You can enable the use of a Domain Name System (DNS) service and add the address of one or more servers and a domain suffix search order. Enter the information as shown in Figure 6-13 and click **Next** to continue.

Welcome

Date and Time

Password

Network Topology

IP Address (eth0)

Host and Gateway

⇒ DNS

Summary

Configure Domain Name System (DNS)

Enable DNS services and configure the search order for DNS servers and domain suffixes.

[See a list of services that require a working DNS.](#)

Enable DNS services

DNS server:

Add

List of DNS servers:

Up
Down
Remove

Domain suffix:

Add

List of domain suffixes:

Up
Down
Remove

Note: If you enable DNS services and are not using DHCP, make sure you update your DNS server with your specified host name, or network validation will fail.

< Back Next > Finish Cancel

Figure 6-13 FSM DNS services configuration

12. The summary window of all configured options is displayed as shown in Figure 6-14. To change a selection, click **Back**. If no changes are needed, click **Finish**.





Figure 6-14 FSM startup wizard summary window

Important: Do not check the **Perform network validation** check box if the Domain Name System (DNS) is not available or not configured.

Next, the final configuration and setup proceeds automatically without the need for more input. Figure 6-15 shows the processing status display.

System Setup Processing

This page shows processing information for the Setup wizard. After the setup tasks are completed, click Continue to proceed.

 The system setup is in progress.....

Setup task status and progress details:





Setup task	Status	Start time	Stop time
Date and Time	 Success	10/4/13 4:55:44 PM	10/4/13 4:55:44 PM
Setting password	 In Progress	10/4/13 4:55:44 PM	--
Host and Gateway	 Not Started	--	--

Figure 6-15 FSM system setup processing status

Figure 6-16 shows the message when the processing is complete.

System Setup Processing

This page shows processing information for the Setup wizard. After the setup tasks are completed, click Continue to proceed.

 Congratulations. All setup tasks completed.

Setup task status and progress details:





Setup task	Status	Start time	Stop time
Date and Time	 Success	10/4/13 4:55:44 PM	10/4/13 4:55:44 PM
Setting password	 Success	10/4/13 4:55:44 PM	10/4/13 4:56:24 PM
Host and Gateway	 Success	10/4/13 4:56:24 PM	10/4/13 4:56:24 PM

Figure 6-16 FSM system setup processing completed

Figure 6-17 shows the message when the server is being started.

 **Attention:** The web server is being restarted as part of the setup process. Network setup and validation can take up to 15 minutes, after which the setup process will continue for approximately 30 minutes. If there are network errors, you will receive notification within 15 minutes, after which the setup process can continue unattended. Do not close this page or refresh your browser window.


 Please wait while the network settings are being applied

Figure 6-17 FSM startup

Figure 6-18 shows the startup process display.

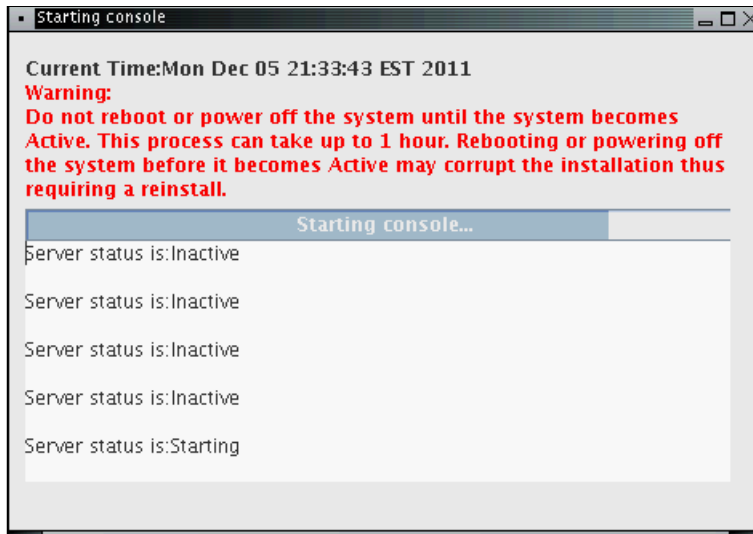


Figure 6-18 FSM startup status

13. When the startup is completed, the local browser on the FSM also starts. Accept any security certificate exceptions. With the security exceptions cleared, the login window of the IBM Flex System Manager GUI is displayed. Enter the credentials that you entered in the startup wizard, and click **Log in** as shown in Figure 6-19.



Figure 6-19 FSM login window

The Home window with the initial setup tasks that must be completed to configure the FSM for the first time opens as shown in Figure 6-20.




Figure 6-20 FSM Home window with the initial setup tasks

The startup wizard and initial login are complete. The FSM is ready for further configuration and use. The example used a console from the remote console function of the IMM2 initiated through the CMM. A secure browser session can now be started directly to the FSM management IP interface (eth0).

6.2 Updating Flex System Manager

When you first log in to the FSM web console, the Initial Setup window opens as shown in Figure 6-21. The goal of this window is to provide the logical steps to follow to update Flex System components.

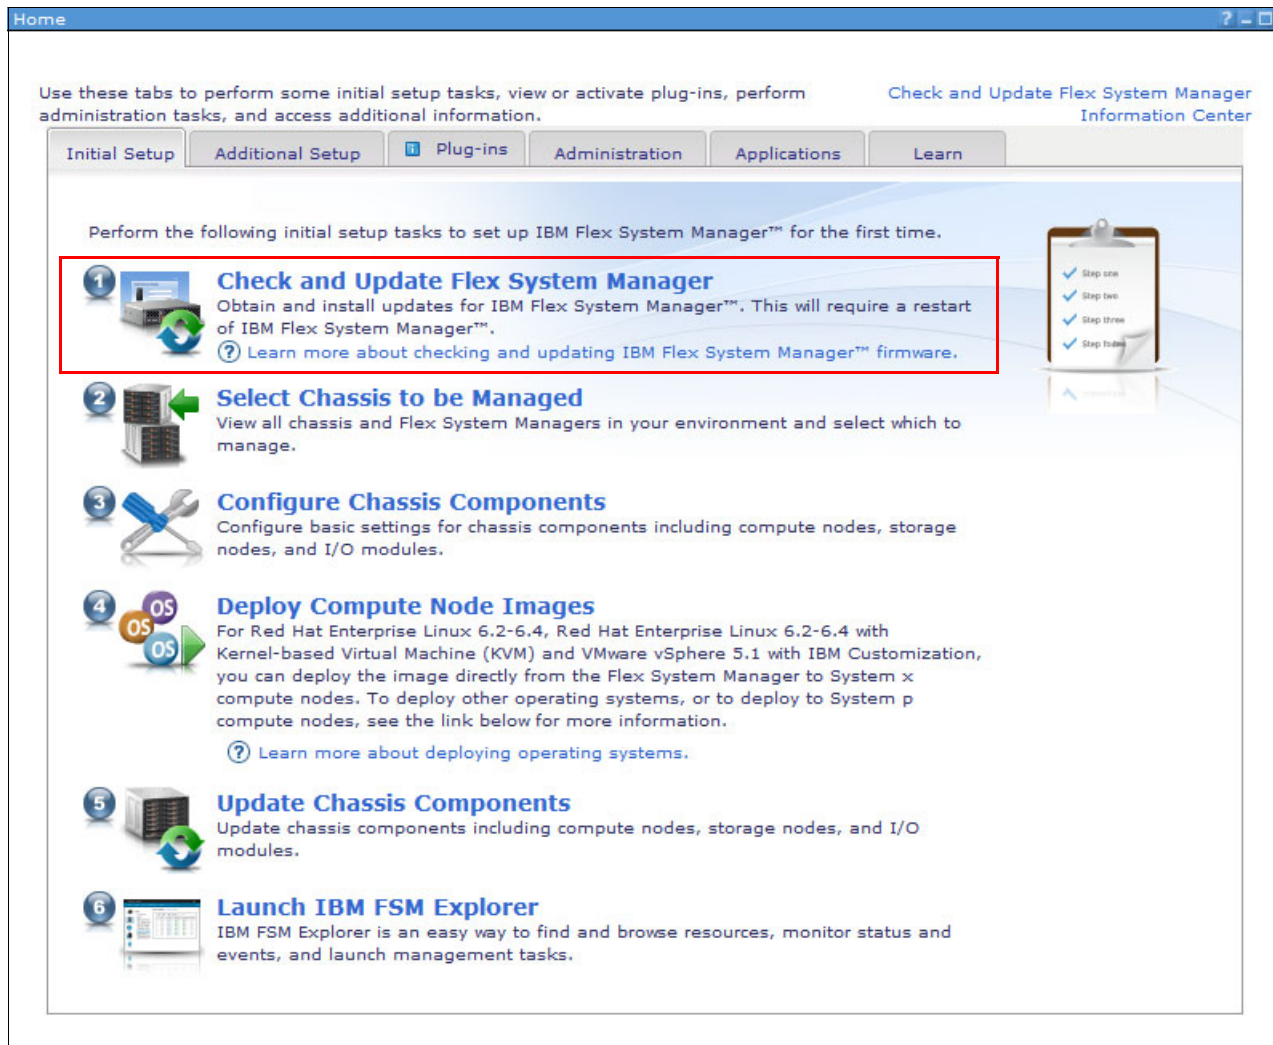


Figure 6-21 FSM initial setup

Step 1 is to check for updates for the FSM. When you select this option, the FSM attempts to contact the IBM Fix Central site to download updates. If a connection to the Internet is not available, a message is displayed that prompts for a local directory on the FSM from which to import the updates. The updates can be downloaded manually from the IBM Fix Central website and then manually copied to the FSM. For more information, see 6.9.1, “Acquiring updates for chassis components” on page 191. FSM updates include both software and hardware stack updates. If the firmware is updated through an FSM update, the FSM needs to be rebooted to activate the installed updates.

6.3 Selecting chassis to manage

Most tasks in IBM Flex System Manager can be performed with more than one method when you are using the GUI. In this example, the most common method is shown.

After FSM is set up initially, it discovers any available chassis. Selections can then be made as to which chassis are managed by the current FSM. To select chassis, perform these steps:

1. From the Initial Setup tab in the Home window, click the **Select Chassis to be Managed** link, as shown in Figure 6-22.




Figure 6-22 FSM Initial Setup tab

A list of available chassis is displayed as shown in Figure 6-23.

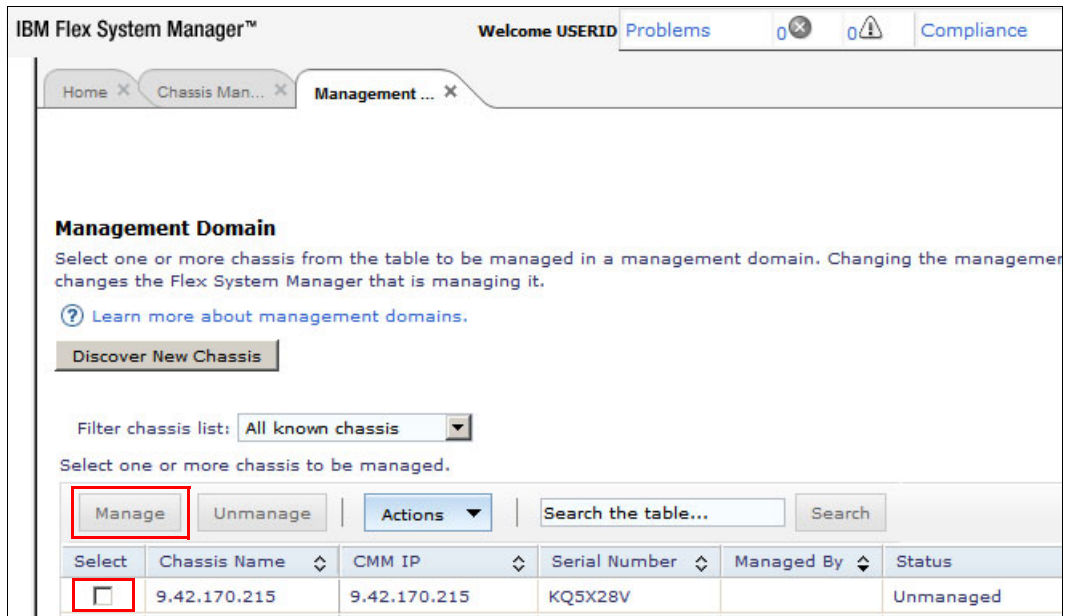


Figure 6-23 FSM chassis selection for management

2. Select the check box for the chassis that you want to manage and click **Manage**. The Manage Chassis window opens, which lists the selected chassis as shown in Figure 6-24. An option is available to choose Centralized User Management (see “Centralized user management” on page 89). Another option allows you to automatically assign IPv6 Unique Local Addresses (ULAs) to the chassis components. Click **Manage**.




Figure 6-24 FSM Manage Chassis options

3. If you choose “Authenticate using the Flex System Manager user registry”, you need to specify current CMM credentials to authenticate to the selected chassis and CMM recovery credentials to allow access to the chassis in case the FSM management node becomes unavailable, as shown in Figure 6-25. Optionally, you can also specify the FSM administrative account that will be associated with the managed chassis. Click **OK**.



The image shows a dialog box titled "Management Credentials" with a blue header bar. It is divided into three sections: "CMM Credentials", "CMM Recovery Credentials", and "IBM FSM Management Credentials (optional)".

CMM Credentials
Please enter a User ID and Password with administrator or supervisor authority to authenticate to the selected chassis: 9.42.170.215

*CMM User ID:

*CMM Password:

CMM Recovery Credentials
When a chassis is centrally managed, local credentials are disabled and the RECOVERY_ID account is created to allow local access to the CMM in the event that the IBM Flex System Manager™ is unavailable. Provide a password for initial use of the RECOVERY_ID account, and ensure that this password is kept current as part of your recovery plan. You will be required to change this password the first time you use it.

Recovery User ID:

*Recovery Password:

*Verify Password:

IBM FSM Management Credentials (optional)
To change the administrator account that will be associated with the CMM in audit and error logs, change the account credentials below.

Figure 6-25 CMM Management Credentials: Centralized user management

4. FSM begins to establish management relationships with the selected chassis as shown in Figure 6-26. In the Message column, FSM displays the current phase of the process.




Figure 6-26 FSM manage chassis process

5. After the successful completion of the manage chassis process, click **Done** as shown in Figure 6-27.

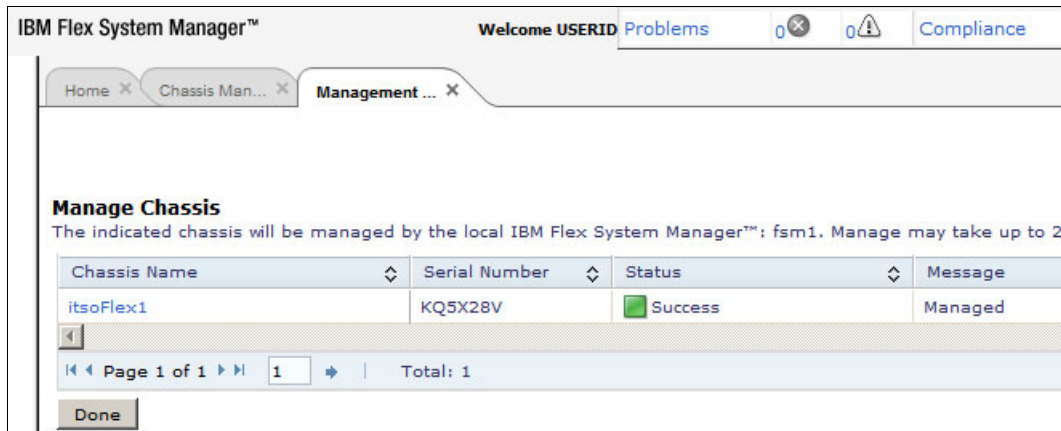


Figure 6-27 FSM manage chassis steps completed

The original IBM Flex System Manager Management Domain window opens with the target chassis as the chassis managed by IBM Flex System Manager (Figure 6-28).




Figure 6-28 FSM management domain with the managed chassis

The Enterprise Chassis is now managed by the IBM Flex System Manager.

6.4 Configuring centralized user management

Use the Flex System Manager management software to change centralized management settings for a chassis.

The option to manage a chassis with the centralized management node user registry is available when you first select a chassis for management on the Management Domain page in the Flex System Manager management software web interface (see 6.3, “Selecting chassis to manage” on page 138).

When you use the management software to place a chassis under centralized management, the Chassis Management Module (CMM) is configured to use the registry that is stored on the management node. The local user accounts in the CMM registry are disabled, and the new user account RECOVERY_ID is created on the CMM for future authentication to the CMM (as long as it is configured to use the centralized user registry on the management node).

If you make changes to the disabled local CMM accounts (for example, if you change a password), the changes have no effect on the RECOVERY_ID account. In centralized user management mode, the RECOVERY_ID account is the only CMM account that is activated and operational.

After the CMM detects the management node user registry, it uses the FSM management node registry configuration to provision all of the managed resources in the chassis (except for network switches) so that they also use the central management node user registry.

After a chassis is managed in centralized user management mode, the management node becomes the account manager for the chassis; you can log in to the CMM using accounts from the management node user registry. If a chassis is in centralized management mode, and the management node fails, you can use the RECOVERY_ID account to log in to the CMM to take recovery actions to restore account-management functions on the CMM until the management node is restored or replaced.

Changing from decentralized to centralized user management

A centralized management configuration uses a single user authentication repository for all of the Chassis Management Modules (CMMs) in a management domain.

The command-line interface (CLI) is used to update a managed chassis from decentralized to centralized user management mode.

Note: You cannot change a chassis from decentralized to centralized user management mode in the management software web interface; you must use the CLI. The web interface enables you to unmanage a chassis, and re-manage the chassis in centralized user management mode. However, unmanaging a chassis deletes all of the chassis settings, and is more complicated than using the `manageChassis` command and its options to change the chassis user management mode to centralized.

To update the chassis from decentralized to centralized user management mode in the management software CLI, run the following command:

```
smcli manageChassis --Uc -c <userid:password@x.x.x.x> --Cu <centralized user ID>
--Cp <centralized password> --Rp <RECOVERY_ID password>
```

The following variables in the command are defined:

- ▶ `<userid:password@x.x.x.x>` represents the administrator credentials and IP address for the target chassis.
- ▶ `<centralized user ID>` is an administrator user ID with supervisor authority on the management node. This account is used to request access to the CMM on behalf of the management node and managed nodes after the CMM is centrally managed.
- ▶ `<centralized password>` is the password for the centralized user ID.
- ▶ `<RECOVERY_ID password>` is the password for the CMM recovery account, which has the user ID RECOVERY_ID.

Changing from centralized to decentralized user management

User management is decentralized when a CMM uses its own user registry (and not that of the management node) or uses an external user registry, such as an external Lightweight Directory Access Protocol (LDAP) server.

Note: You cannot change a chassis from centralized to decentralized user management mode in the management software web interface; you must use the CLI. The web interface enables you to unmanage a chassis, and re-manage the chassis in decentralized user management mode. However, unmanaging a chassis deletes all of the chassis settings, and is more complicated than using the `rmCentrallyManagedChassis` command and its options to change the chassis user management mode to decentralized.

To change the chassis to decentralized user management mode from the management software CLI, run the `rmCentrallyManagedChassis` command, as shown in Example 6-1.

Example 6-1 CLI commands to decentralize a chassis

```
USERID@fsm1:~> smcli lsCentrallyManagedChassis
```

List of centrally managed chassis:

```
Chassis 1:  
  Hostname: 9.42.170.215  
  UUID: 2C684A86292E3D288C23725C87D0E7C7  
  OID: 23,680
```

```
USERID@fsm1:~>smcli rmCentrallyManagedChassis -u 2C684A86292E3D288C23725C87D0E7C7  
Chassis unmanaged successfully
```

Note: When the `rmCentrallyManagedChassis` command completes, the chassis is still managed. The chassis no longer uses the management node user registry. You must now request access to the chassis again using the chassis credentials.

6.5 Configuring chassis components

The next step in the initial setup tasks is to configure chassis components. Before the component can be configured, it must be discovered by the FSM, access must be granted to the component object in the FSM, and inventory must be collected on it.

The following tasks are typically associated with the configuration of chassis components:

- ▶ Request access to the compute nodes, I/O modules, and storage nodes
- ▶ Collect inventory on the chassis components
- ▶ Configure compute nodes using Configuration Patterns
- ▶ Configure the Chassis Management Module using Configuration Patterns
- ▶ Configure I/O modules using Configuration Templates

During the Manage Chassis configuration task (see 6.3, “Selecting chassis to manage” on page 138), the FSM discovers the components inside the managed chassis, requests access to them, and collects inventory on the components on which access was granted successfully.

If you are unable to grant access to a specific component, you can manually initiate an access request and then inventory collection. You can use the Initial Setup tab to verify access to the components and request access and collect inventory, if needed, by performing the following steps:

1. From the Initial Setup tab in the Home window, click **Configure Chassis Components**, as shown in Figure 6-29.

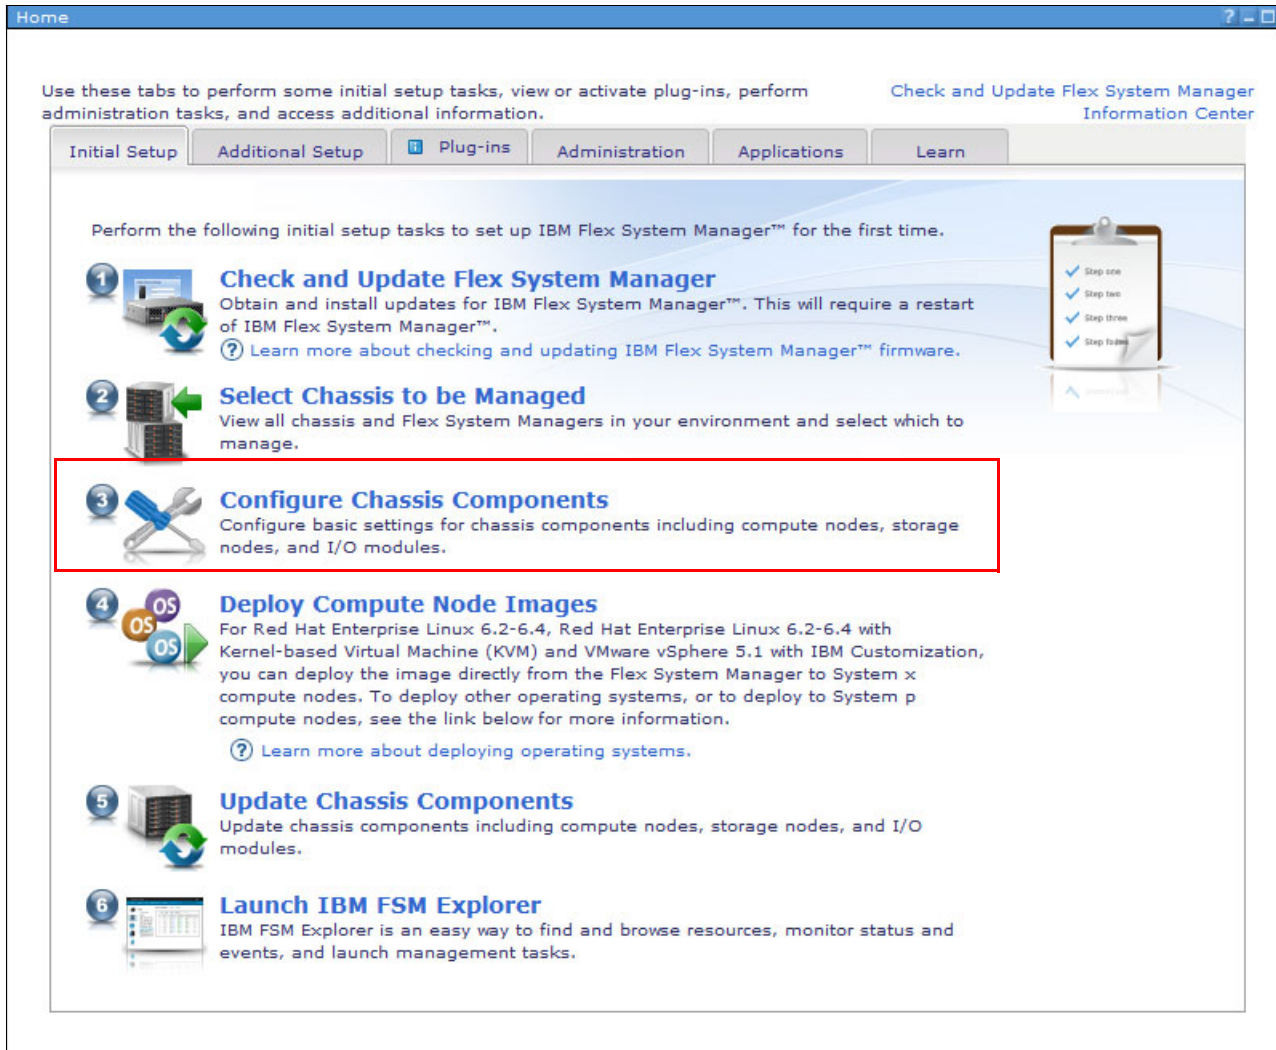


Figure 6-29 FSM Initial Setup tab in the Home window

- In the Configure Chassis Components window, you can verify how many compute nodes have been discovered, and how many nodes have full access, as shown in Figure 6-30. If you do not have access to some compute nodes, click **Request Access to Compute Nodes** to request full access to them.

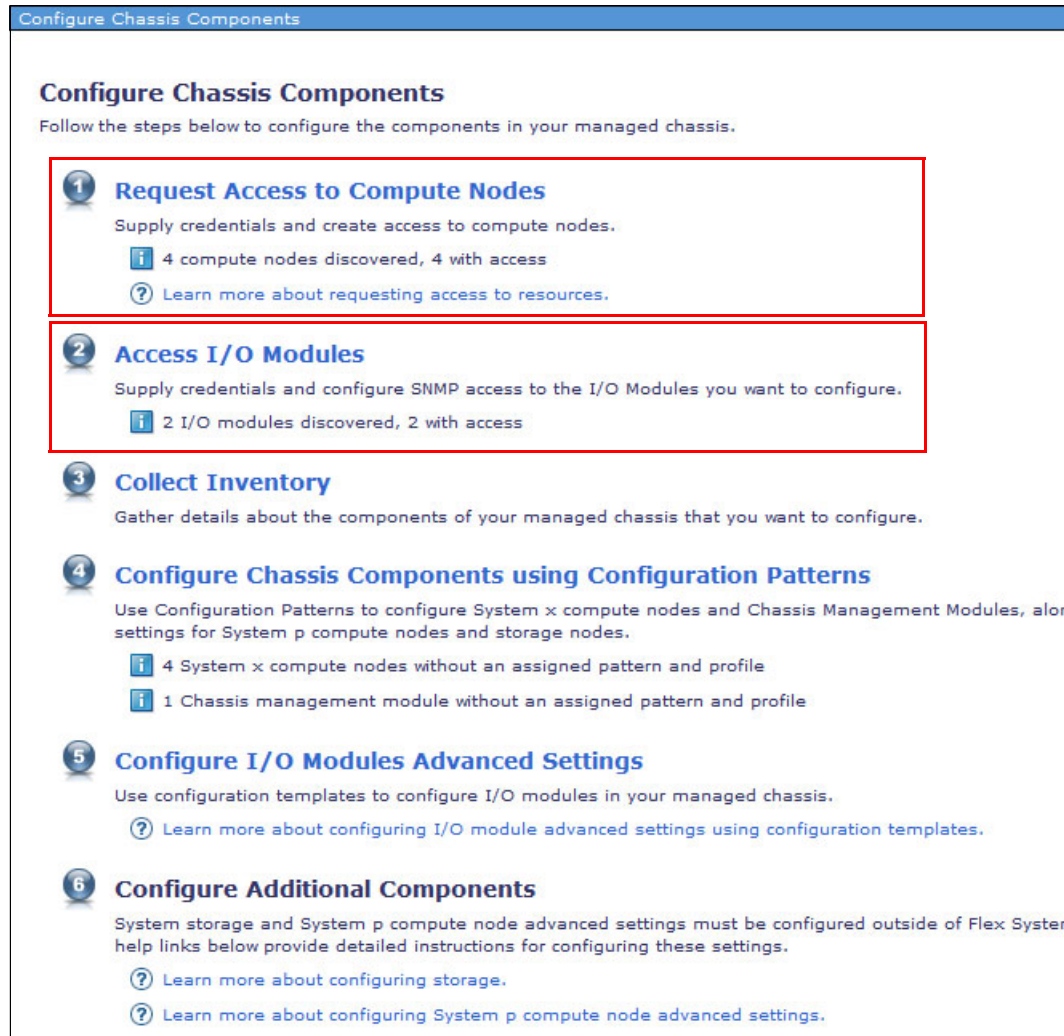


Figure 6-30 Configure Chassis Components window

In the same Configure Chassis Components window, you can also verify how many I/O modules have been discovered, and how many of them have full access, as shown in Figure 6-30.

If some I/O modules do not have access or have partial access, click **Access I/O Modules** to request access to them.

Select an I/O module and click **Request Access**, as shown in Figure 6-31.

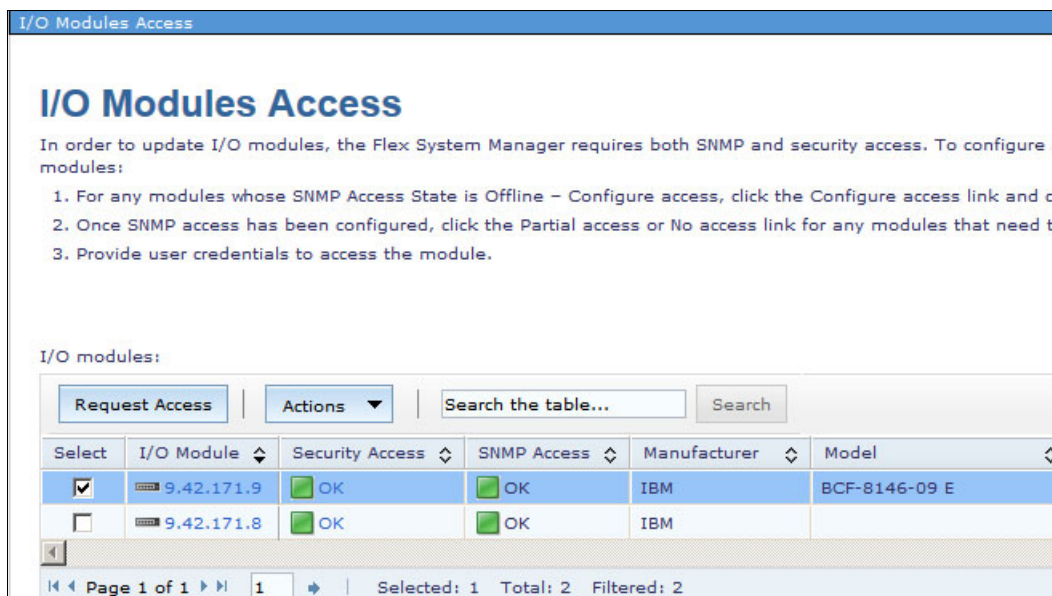


Figure 6-31 I/O Modules Access window

From the Configure Access window, you can request access, verify that management protocols are enabled, and verify their access status, as shown in Figure 6-32.

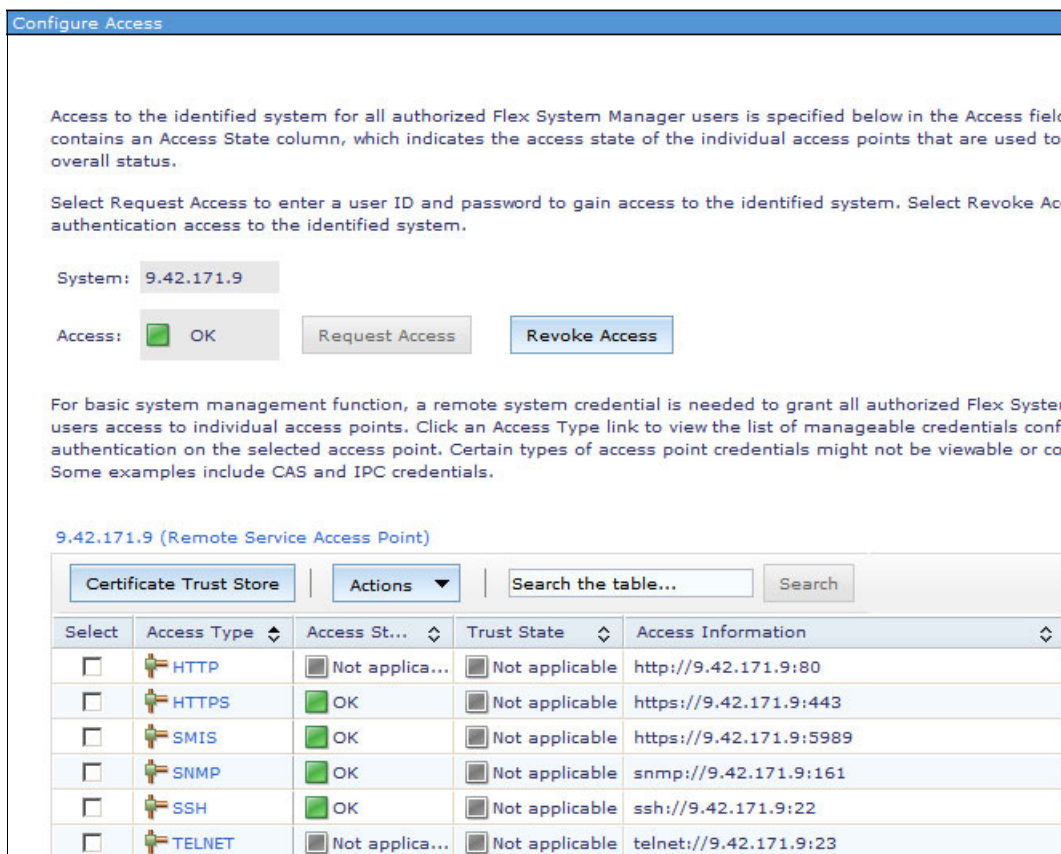


Figure 6-32 Configure access for I/O modules

Requesting access: If there is no access to the I/O module or access is partial (No access or Partial access is listed in the Access Status column), the Request Access button will be unlocked, and you can request access to the I/O module by clicking **Request Access** and supplying I/O module credentials.

3. Collect inventory on all chassis components by clicking **Collect Inventory** in the Configure Chassis Components window, as shown in Figure 6-33.

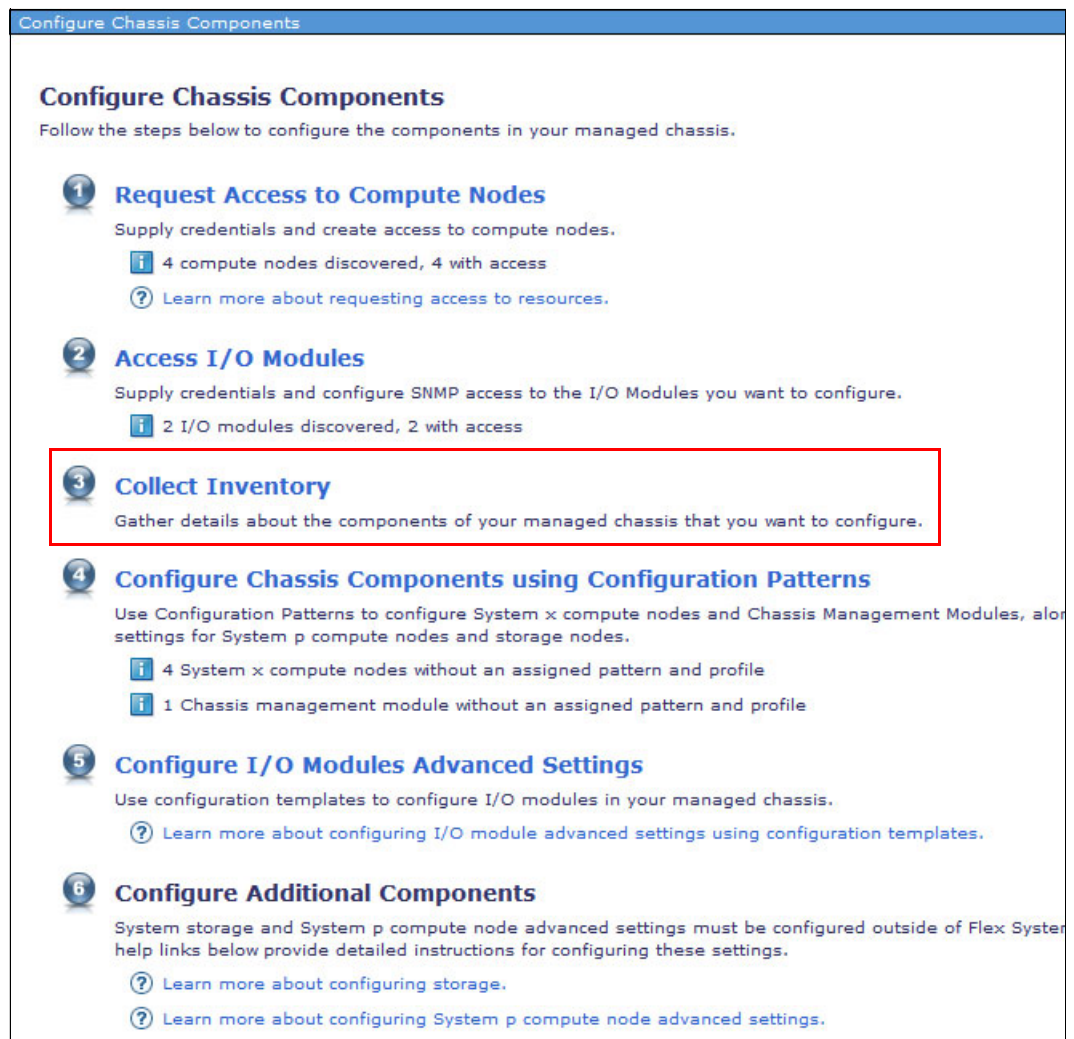


Figure 6-33 Configure Chassis Components window

Click **OK** in the Launch Job window to start the task, as shown in Figure 6-34.

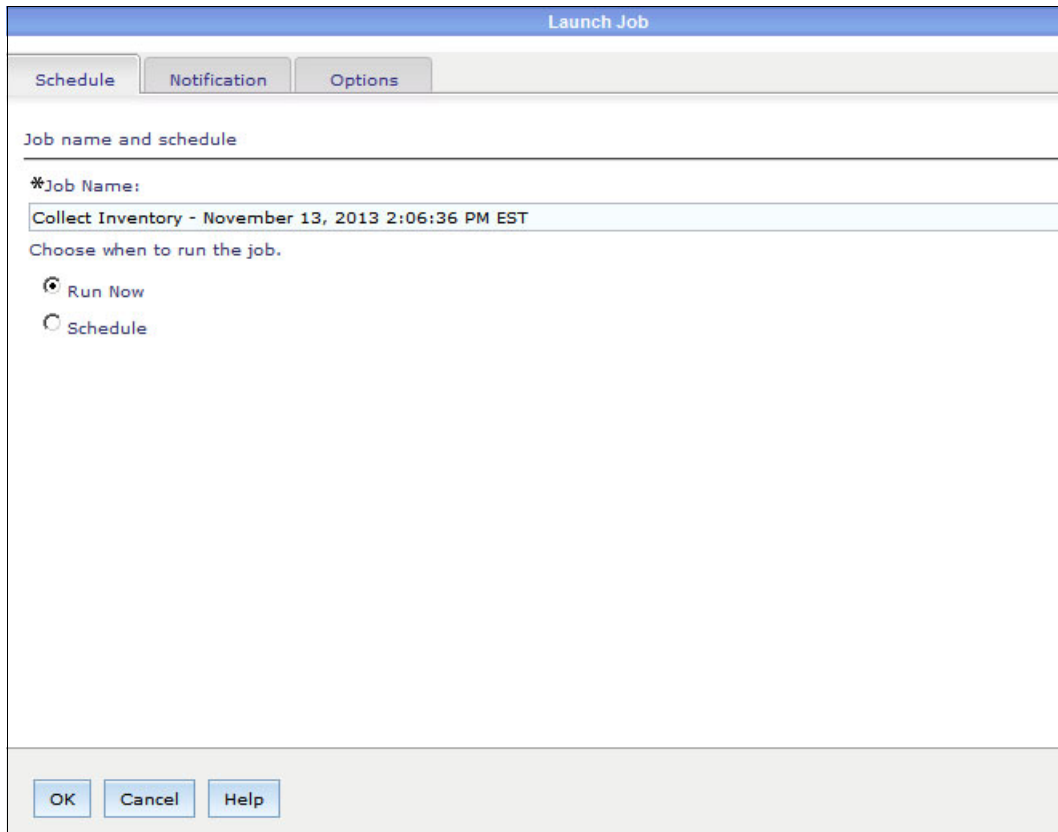


Figure 6-34 Launch the Collect Inventory job

Check that the job started successfully, as shown in Figure 6-35. You can monitor job progress by clicking **Display Properties**.

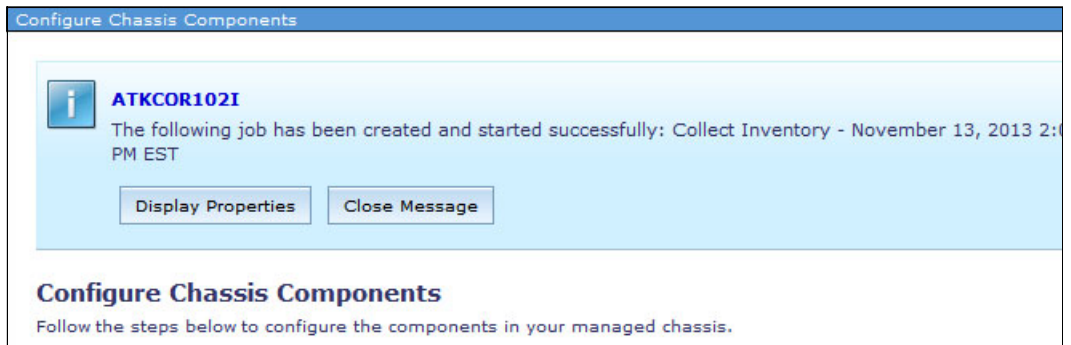


Figure 6-35 Job launch informational message

Check the job status and progress as shown in Figure 6-36.

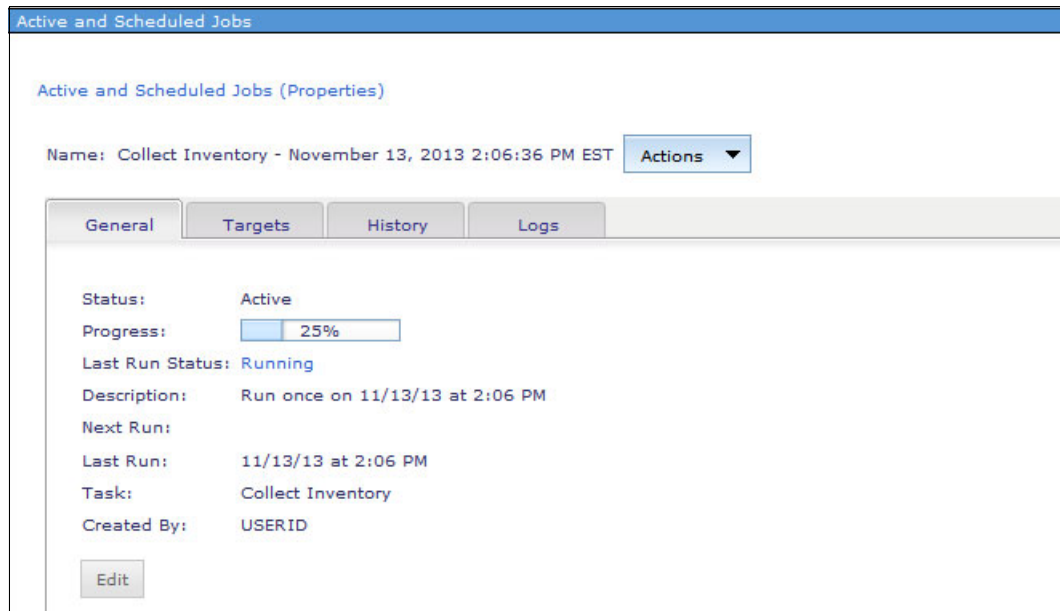


Figure 6-36 Job properties

Wait until the job is completed. You are now ready to perform configuration tasks on the compute nodes and I/O modules.

6.6 Configuring compute nodes using Configuration Patterns

This section describes the following topics:

- ▶ 6.6.1, “Overview of Configuration Patterns” on page 150
- ▶ 6.6.2, “Creating and applying compute node Configuration Patterns” on page 153
- ▶ 6.6.3, “Automating compute node failover with Configuration Patterns” on page 161

6.6.1 Overview of Configuration Patterns

You can use Configuration Patterns to provision or pre-provision X-Architecture compute nodes using a common Configuration Pattern that can be deployed to multiple compute nodes. Configuration Patterns enable you to configure local storage, network adapters, boot order, and Integrated Management Module (IMM) and Unified Extensible Firmware Interface (UEFI) settings.

The diagram in Figure 6-37 illustrates the concept of Configuration Patterns.




Figure 6-37 Configuration Patterns concept

Use Configuration Patterns to define and manage the server patterns, server profiles, policies, and standby server pools. Before you create a new server pattern on the Configuration Patterns page in the IBM Flex System Manager web interface, consider the following suggestions:

- A *server pattern* represents a compute node configuration that is deployed before an operating system is installed. It includes local storage configuration, network adapter configuration, boot settings, and other IMM and UEFI firmware settings.
- *Server profiles* are generated automatically when a server pattern is deployed. One profile is created for each target compute node. Each server profile represents the specific configuration of a single compute node and contains system-unique information (for example, assigned IP addresses and Media Access Control (MAC) addresses).
- Identify compute nodes for configuration that have a common hardware configuration. A server pattern is used to apply the same configuration settings to compute nodes with the same hardware.
- If you want to create a server pattern from an existing compute node, make sure that the compute node is discovered and unlocked.
- Identify the aspects of configuration that you want to customize for the server pattern (for example, local storage, network adapters, boot settings, IMM settings, and UEFI settings).

Server patterns

A server pattern represents a compute node configuration that is deployed before an operating system is installed. It includes local storage configuration, network adapter configuration, boot settings, and other IMM and UEFI firmware settings.

When you define a server pattern, select the category patterns and address pools that you need for the configuration that you want for a specific group of compute nodes. You can define multiple server patterns to represent different configurations in your data center. When a server pattern is deployed to multiple compute nodes, multiple server profiles are generated automatically (one profile for each compute node). Each profile inherits settings from the parent server pattern, which enables you to control a common Configuration Pattern from a single place.

Note: When you create a new server pattern from scratch, you are required to define the boot settings for compute nodes. If you deploy the new server pattern to compute nodes, the existing boot order on the compute nodes is overwritten with the default boot order settings in the new server pattern.

When you create server patterns, make sure that you create them for each compute node type. For example, create a server pattern for all IBM Flex System x240 compute nodes and a server pattern for all IBM Flex System x440 compute nodes. Do not apply a server pattern created for one compute node type to a different compute node type.

To ensure that your Configuration Patterns are not lost if the management node fails, back up the management software after you create or modify Configuration Patterns.

The category patterns within a server pattern correspond to the firmware settings for a compute node type. Most of the firmware settings that you might configure directly on the compute node IMM and UEFI can also be configured through Configuration Patterns in the management software web interface. However, some settings are not supported by Configuration Patterns, and other settings are not yet available.

Server profiles

Server profiles are generated automatically when a server pattern is deployed. One profile is created for each target compute node. Each server profile represents the specific configuration of a single compute node and contains system-unique information (for example, assigned IP addresses and MAC addresses).

When a pattern is deployed, an individual system profile is generated for each target system. You can edit a pattern and save changes, and any dependent system profiles are automatically updated and redeployed to their associated systems. You can move an existing profile from one system to another by unassigning the profile, and then redeploying the profile to another system.

Important: Systems retain their identification information (for example, host name, IP address, and virtual MAC address) when a profile is unassigned. To avoid name and address conflicts, any identification information about the original system must be cleared before the unassigned profile is deployed to a different system.

Each server profile represents the specific configuration of a single compute node and contains information that is unique to a compute node. The server profile is activated as part of the IMM startup process. After a server profile is activated for a compute node, any subsequent configuration changes are done by editing the appropriate server pattern or category pattern associated with the profile. This enables you to control a common Configuration Pattern from a single place.

If a compute node needs to be moved or repurposed, you can reassign a server profile from one compute node to another.

You can deploy a server pattern to a compute node or to an empty chassis bay. In either case, the profile is associated with the chassis bay. If you replace an existing compute node, you must redeploy the server profile associated with that bay to activate the profile on the new compute node. If you first deploy a server pattern to an empty bay, you must redeploy the server profile associated with that bay after a compute node is installed.

Note: To ensure that your Configuration Patterns are not lost if the management node fails, back up the management software after you create or modify Configuration Patterns.

6.6.2 Creating and applying compute node Configuration Patterns

Use the following procedure to create and apply a Configuration Pattern on the compute node:

1. From the Initial Setup tab in the Home windows, select **Configure Chassis Components**, then click **Configure Chassis Components using Configuration Patterns**, as shown in Figure 6-38.




Figure 6-38 Configure Chassis Components window

IBM Flex System Manager web interface (FSM Explorer) opens in a new window, as shown in Figure 6-39. Click **Create a new server pattern from an existing compute node** to continue.

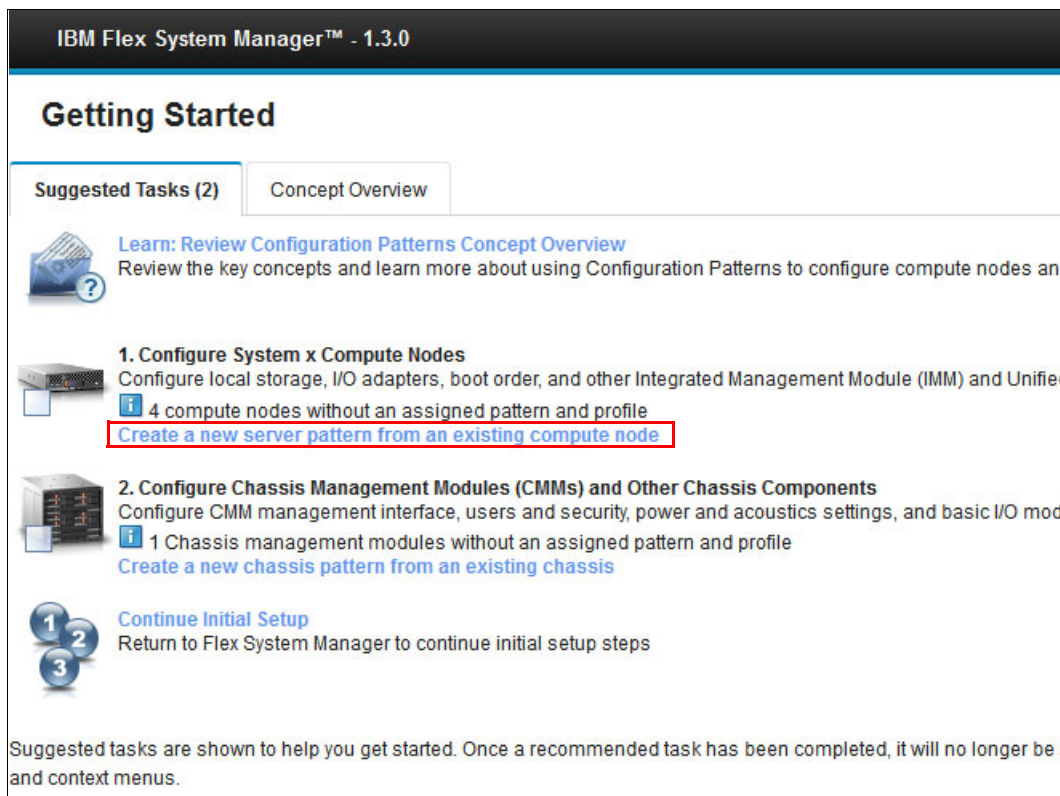


Figure 6-39 Configuration Patterns: Getting Started

Existing compute node: If you plan to create your new Configuration Pattern from the existing compute node, you need to configure node settings first using its UEFI user interface. Node settings include local storage, network adapters, boot order, and Integrated Management Module (IMM) and Unified Extensible Firmware Interface (UEFI) settings.

2. Choose the node on which the pattern will be based as shown in Figure 6-40. Click **Next**.

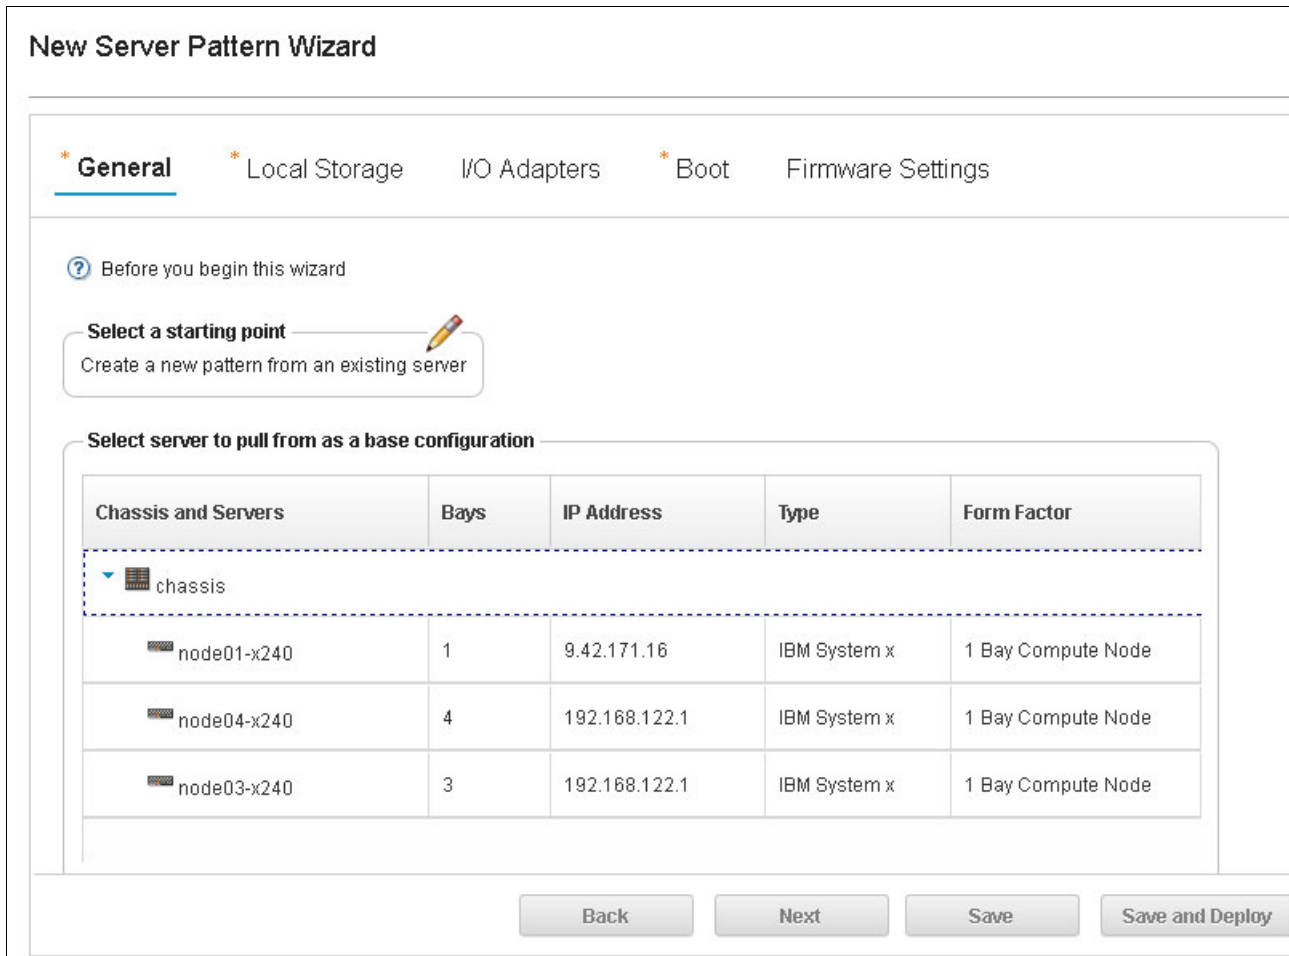


Figure 6-40 Select the node on which to base the pattern

New pattern from scratch: If you want to create a server pattern from scratch, click the pencil icon in the “Select a starting point” box to show an expanded starting point view, as shown in Figure 6-41.





Figure 6-41 Select a starting point expanded view


3. Choose a name for the new pattern as shown in Figure 6-42. Click **Next**.

New Server Pattern Wizard

* **General** * Local Storage I/O Adapters * Boot Firmware Settings

? Before you begin this wizard

Select a starting point — 
Create a new pattern from an existing server

Select server to pull from as a base configuration — 
chassis — node01-x240

Specify pattern name and description

* Name:

Description (limit of 500 characters)

Back **Next** Save Save and Deploy

Figure 6-42 Name the new server pattern

- The next window provides options for local storage configuration. You can specify a new storage configuration or keep the existing storage configuration on the target, or disable local disks as shown in Figure 6-43. Choose the required option and click **Next**.




Figure 6-43 New Server Pattern Wizard: Local Storage

- The next window provides options for IO adapters as shown in Figure 6-44. The I/O adapters based on the original model are the same for the new pattern. Click **Next**.

New Server Pattern Wizard

* General * Local Storage I/O Adapters * Boot Firmware Settings

? If desired you can modify adapter addressing and define additional adapters to match the hardware you expect to configure with this pattern

Graphic view [↗](#) I/O adapter addressing: ? **Burned in**

Advanced Settings | | More ▾

Location	Type	I/O Bay	Configuration Pattern	I/O Addressing
▼ Compute Node				
▼ LOM Fabric Connector	Virtual Fabric	1-2	Learned-Adapter-4.1	
▼ Port 1	Virtual Fabric	1	Learned-Port-4.1	Burned in Addresses
Function 1 — NIC	Ethernet	1		Burned in Addresses
Function 3 — NIC	Ethernet	1		Burned in Addresses
Function 5 — NIC	Ethernet	1		Burned in Addresses

Back **Next** **Save** **Save and Deploy**

Figure 6-44 New Server Pattern Wizard: I/O Adapters

6. The next window allows changes to the boot mode or allows you to keep the same as the original node as shown in Figure 6-45. Click **Next**.

The screenshot shows the 'New Server Pattern Wizard' interface, specifically the 'Boot' tab. The wizard has five tabs: 'General', 'Local Storage', 'I/O Adapters', 'Boot', and 'Firmware Settings'. The 'Boot' tab is active and underlined. Below the tabs, there is a descriptive text: 'This pattern can be used to configure boot order for Legacy Only boot environments, and SAN boot targets for UEFI or Legacy environments.' Underneath this text, the 'System boot mode' is set to 'Keep existing boot mode', which is selected with a radio button. Other options include 'UEFI Only Boot', 'UEFI First, Then Legacy', and 'Legacy Only Boot'. Below the radio buttons, there are three tabs for boot order configuration: 'Primary Boot Order', 'Wake on LAN (WoL) Boot Order', and 'SAN Boot'. The 'Primary Boot Order' tab is currently selected. A blue information box with an 'i' icon contains the message: 'Boot order can only be configured if Legacy Only Boot is selected as the system boot mode.' At the bottom of the wizard, there are four buttons: 'Back', 'Next', 'Save', and 'Save and Deploy'.

Figure 6-45 New Server Pattern Wizard: Boot

7. The next window provides the ability to use firmware settings from the original node or to update those settings as shown in Figure 6-46. Click **Next**.

New Server Pattern Wizard

* General * Local Storage I/O Adapters * Boot **Firmware Settings**

⚠ Some learned System Information pattern information was set to the default and might need to be modified.

⚠ Some learned Management Interface pattern information was set to the default and might need to be modified.

i Firmware settings have been learned from "chassis - node01-x240" and learned patterns have been auto-generated. Any l...

Integrated Management Module (IMM) and Server Firmware Settings (UEFI)

Select existing or create new category patterns as desired to include in this server pattern.

Category	Pattern			
System Information:	<input type="text" value="Learned-System_Info-4"/>	?		
Management Interface:	<input type="text" value="Learned-Management-4"/>	?		
Power Schedule And Capping:	<input type="text" value="Learned-Power-4"/>	?		
Performance And Recovery:	<input type="text" value="Learned-Performance-4"/>	?		
Device And IO Ports:	<input type="text" value="Learned-Devices_IO-4"/>	?		

Back **Next** **Save** **Save and Deploy**

Figure 6-46 New Server Pattern Wizard: Firmware Settings

- Choose **Save and Deploy** to store the new server pattern and deploy it. The new server pattern can be deployed to one or more nodes as shown in Figure 6-47.

Deploy Server Pattern - node 1 pattern

Deploy the server pattern to one or more individual servers, or groups of servers (e.g. chassis). On deploy, one server profile is created for each

*Pattern To Deploy: (1 bay compute pattern)

*Profile Activation:

Available Servers Selected Servers


Add Placeholder Chassis

Name	Bay	Deploy Status
chassis		
node01-x240	1	<input checked="" type="checkbox"/> Ready
node03-x240	3	<input checked="" type="checkbox"/> Ready
Empty Bay	6	<input checked="" type="checkbox"/> Ready
Empty Bay	7	<input checked="" type="checkbox"/> Ready
Empty Bay	8	<input checked="" type="checkbox"/> Ready

Name	Bay	Deploy Status
chassis		
node04-x240	4	<input checked="" type="checkbox"/> Ready

Figure 6-47 One node is selected to which to deploy the new pattern

If the node that you are planning to deploy is powered on, you receive the following message as shown in Figure 6-48.

 Some of the servers you selected are online. To fully activate the profile, these servers will be restarted after deployment:
node04-x240

Do you want to deploy the pattern and restart the servers?

Figure 6-48 Choose to deploy the new pattern and restart the node

6.6.3 Automating compute node failover with Configuration Patterns

Configuration Patterns can be set up to monitor for a hardware fault and fail the node over to another one that is in the standby server pool.

Note: Each of the nodes must be configured to use SAN storage only and must be compatible with virtual addressing for its I/O adapters.

Follow these steps to set up a node in the standby server pool:

1. Open the FSM Explorer web interface by clicking **Launch IBM FSM Explorer** on the Initial Setup tab in the Home window. The main FSM Explorer window opens, as shown in Figure 6-49.

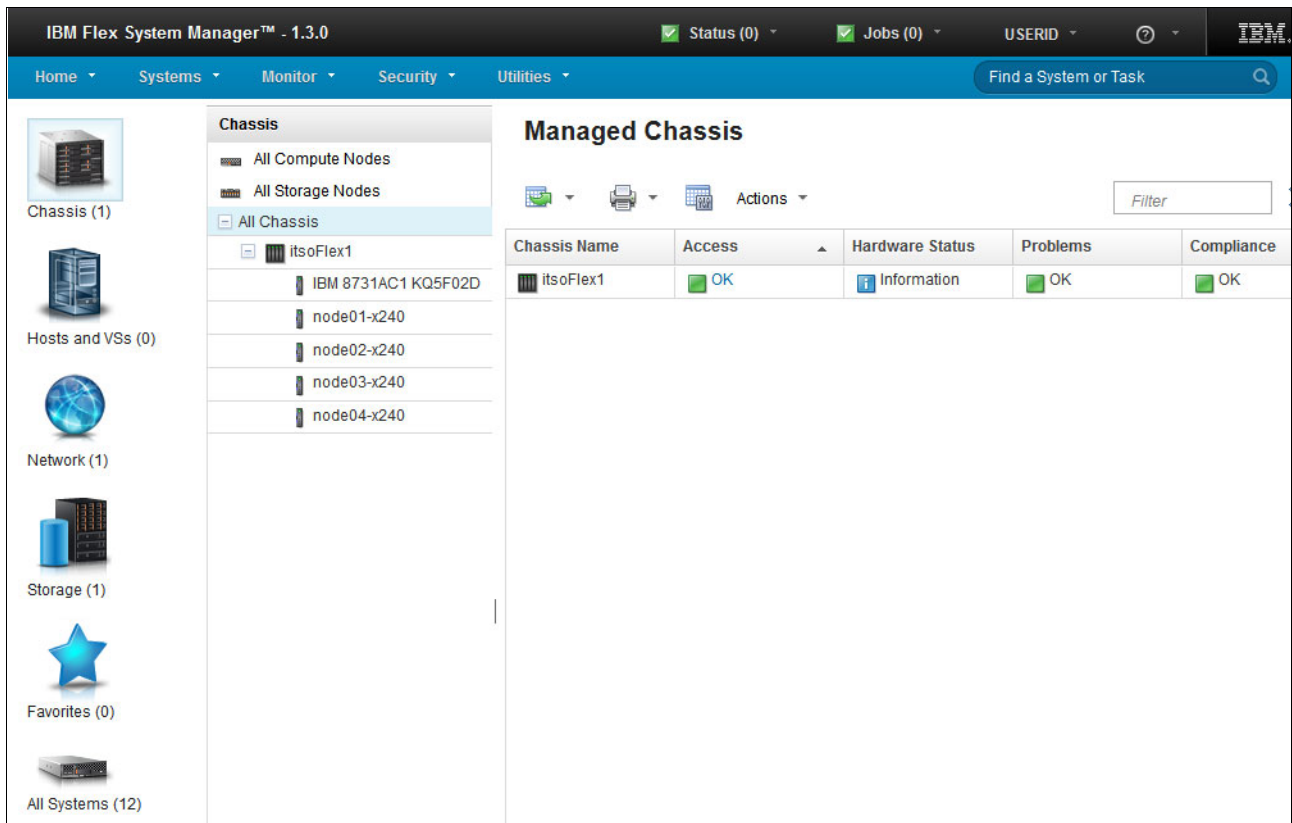


Figure 6-49 FSM Managed Chassis window

2. Select **Systems** → **Configuration Patterns** to open the Configuration Patterns window. Within the Configuration Patterns window, on the left column under Servers heading, choose **Policies** as shown in Figure 6-50.

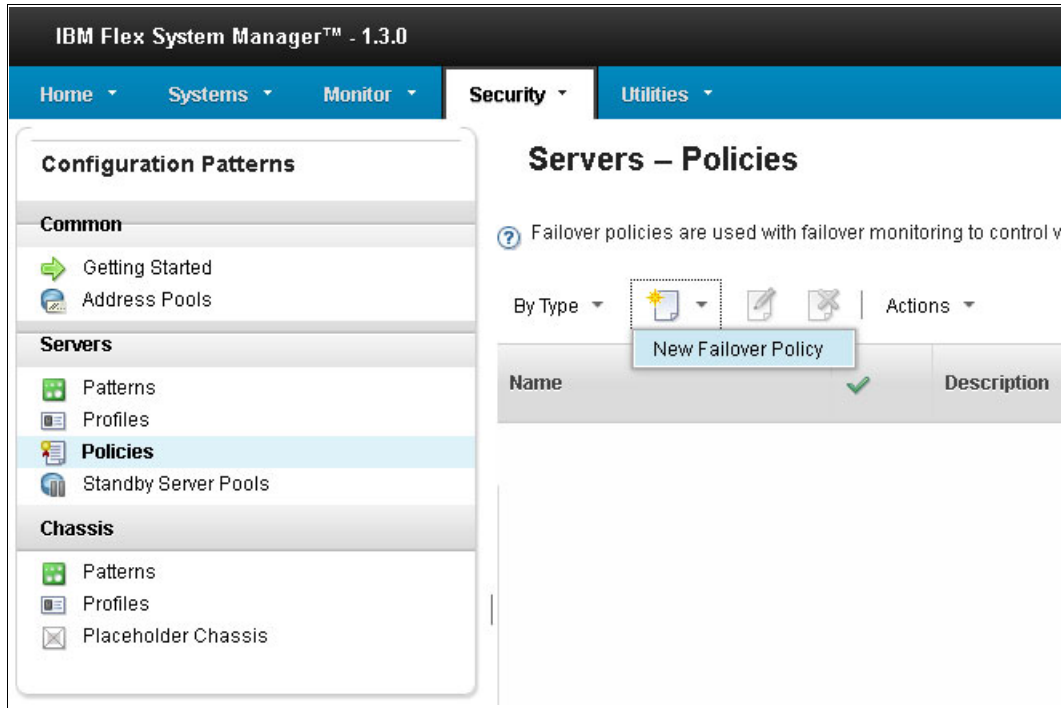


Figure 6-50 Create new failover server policy

3. Choose the criteria for when a failover will be triggered. Both a CPU failure and Memory failure are set up as shown in Figure 6-51.

New Failover Policy

New Failover Policy

Specify name and description

*Name:

Description (limit of 500 characters):

Monitor for the following conditions to initiate failover

- Power off
- CPU failure
- Blade communication errors
- Blade removal
- HDD failure
- Denied power
- Memory failure
- Voltage warnings
- Predictive failure analysis(PFA) events

Choose failover options

- Only failover to standby servers that are powered off [?](#)
- Reset the IMM on the failed server to factory defaults [?](#)

Note: VLAN settings on attached switch are not automatically copied and on failover will need to be migrated manually.

Figure 6-51 Choose criteria for when a failover will be triggered

Select **Create** when finished.

4. Choose servers to be in the standby pool by selecting **Standby Server Pools** under **Servers** in the left column as shown in Figure 6-52.




Figure 6-52 Creating a new standby server pool

- In the New Standby Server Pool window, name the new pool and choose the servers to add into the pool as shown in Figure 6-53. Click **Create**.

New Standby Server Pool

Specify name and description

* Name:

Description (limit of 500 characters):

Choose one or more servers to include in this standby pool to use as failover targets

Available Servers

Name	Bay	Form Factor
▼ chassis	...	
node03-x240	3	1 Bay Compute No
node04-x240	4	1 Bay Compute No

➤

➤➤

➤➤

➤➤

✂

Selected Servers

Name	Bay	For
▼ chassis		
node01-x240	1	1 B

Figure 6-53 Select Create after naming the new pool and adding nodes

- From the main Configurations Pattern window, choose **Start Failover Monitoring** as shown in Figure 6-54.






Figure 6-54 Select to start failover monitoring



- From the drop-down menus, choose the failover policy and the target standby server pool. Then, choose the server to monitor from the list. We selected node 3 as shown in Figure 6-55. Click **Start**.

Start Failover Monitoring

i If a monitored server fails and the server profile is migrated to a standby server, the failed server will be powered off and ... 1:3

Select a failover policy and a target standby server pool

Compute Node hardware failure  

X-Architecture servers  

Choose one or more servers to monitor for failover based on above policy

Available Servers

Name	Bay	Access	Form Factor
chassis	...	<input checked="" type="checkbox"/> OK	
node02-x240	2	<input checked="" type="checkbox"/> OK	1 Bay Compute N
node04-x240	4	<input checked="" type="checkbox"/> OK	1 Bay Compute N

Selected Servers

Name	Bay	Access	For
chassis			
node03-x240	3	<input checked="" type="checkbox"/> OK	1 B

Start

Figure 6-55 Start Failover Monitoring

The failover configuration is now completed. To verify failover operations, within the Chassis Manager tab in the FSM, select the node and right-click for the option to fail over to the standby server pool as shown in Figure 6-56.

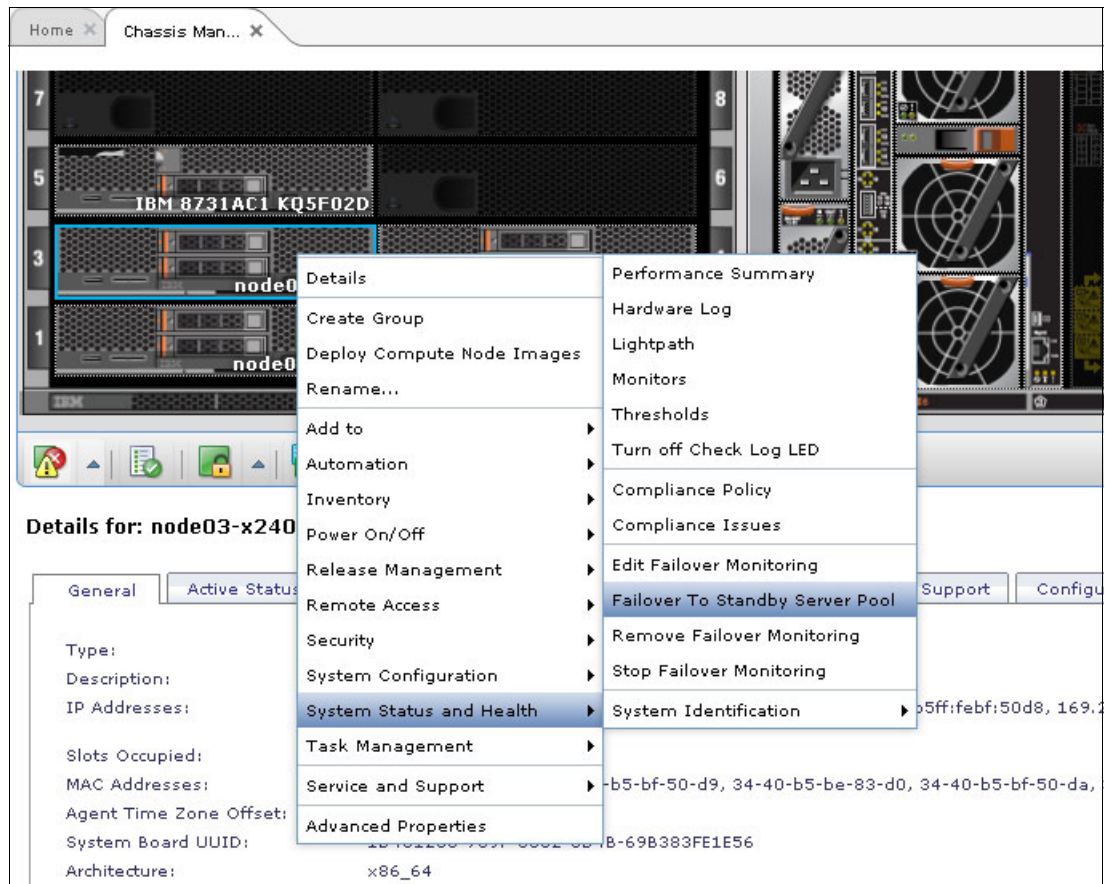


Figure 6-56 Trigger a failover to the standby server pool

Return to the Initial Setup tasks window.

6.7 Deploying compute node images

The Flex System Manager provides the ability to deploy operating system images to one or more X-Architecture compute nodes. The Deploy compute node image task is capable of mass operating-system deployment. The management software supports up to 56 X-Architecture compute nodes for concurrent deployment.

Considerations:

- ▶ You can use the Deploy compute node image task to install operating systems on X-Architecture compute nodes only.
- ▶ At this time, the Deploy compute node image task only provisions to local (internal) disks. SAN disks are not supported at this time.
- ▶ If you deploy an image to a compute node that already has an operating system installed, the existing operating system will be overwritten.

The following operating systems are supported:

- ▶ VMware vSphere Hypervisor (ESXi) 5.1 with IBM customization. A version of the IBM-customized VMware vSphere Hypervisor is preloaded on the IBM Flex System Manager management node.
- ▶ Red Hat Enterprise Linux 6.2, 6.3, and 6.4. When you import the Red Hat Enterprise Linux ISO image, it will generate three different OS image profiles: Minimal, Basic, and Virtualization. You can deploy any of these images to X-Architecture compute nodes.

When you deploy the Virtualization OS image profile, the Kernel-based Virtual Machine (KVM) Platform Agent is automatically installed and configured on the compute node. If the VMControl plug-in is active on the IBM Flex System Manager management node, the deployed operating system will be added automatically as one of the virtual servers or hosts that VMControl can manage.

If the VMControl plug-in is not activated before you deploy the image and you want to manage the virtual server or host through VMControl, you must activate the VMControl plug-in and then collect inventory on the virtual server or host manually for the deployed operating system.

6.7.1 Importing operating system images

The IBM Flex System Manager management node supports a maximum of two operating system images in local storage. A version of the IBM-customized VMware vSphere Hypervisor is preloaded on the IBM Flex System Manager management node. Therefore, you can import one additional operating system image on the IBM Flex System Manager management node and then deploy that image to X-Architecture compute nodes.

Multiple images: If you already have two images loaded on the IBM Flex System Manager management node, you will need to first delete one of those images before attempting to import another image. For example, you can delete the IBM customized VMware vSphere image if needed to allow for two versions of Red Hat Enterprise Linux 6.2 and 6.3.

Complete the following steps to delete an operating system image from the IBM Flex System Manager management node:

1. Access the management node using Secure Shell (SSH).
2. Log in to the IBM Flex System Manager CLI user interface using a user account with administrator privileges, such as USERID.
3. Use the **smcli lsosimages** command to list all stored images, as shown in Example 6-2.

Example 6-2 List all stored images on the IBM Flex System Manager management node

```
USERID@fsm1:~> smcli lsosimages

OS Name: esxi5.1
OS Profiles:
esxi5.1-x86_64-install-Virtualization

OS Name: rhels6.3
OS Profiles:
rhels6.3-x86_64-install-Basic
rhels6.3-x86_64-install-Minimal
rhels6.3-x86_64-install-Virtualization
USERID@fsm1:~>
```

4. Use the `smcli deleteosimage -o <os image name>` command to delete an image, as shown in Example 6-3.

Example 6-3 Delete an OS image from the IBM Flex System management node

```
USERID@fsm1:~> smcli deleteosimage -o rhels6.3
```

Complete the following steps to import an operating system to the IBM Flex System Manager management node:

1. Obtain a licensed ISO image of the operating system.
2. Use one of the following methods to copy the ISO image to an accessible directory on the management node, such as `/home/USERID`:

- a. Use Secure Copy Protocol (SCP)

Use an SCP tool on a notebook or workstation attached to the management or data network to send the ISO image to the management node, as shown in Example 6-4. Send the ISO image from the IBM Flex System Manager management node where you are logged in using a user account with administrator privileges, such as USERID.

Example 6-4 SCP command from a remote workstation

```
scp * USERID@<management_node_host_name>:/home/USERID/*
```

- b. Copy the ISO image from a USB storage device:

- i. Insert a USB storage device into the USB port on the front of the management node.
- ii. Access the management node using SSH.
- iii. Log in using a user account with administrator privileges, such as USERID.
- iv. Use the command `lsmediadev` to list the storage media devices that are available for use on the system and identify the USB storage device. The resulting output will be similar to Example 6-5, where `vd1` is the USB storage device.

Example 6-5 List media devices and mount USB storage device

```
USERID@fsm1:~>lsmediadev
device=/dev/vd1,mount_point=/media/vd1,type=3,description=USB flash memory
device
USERID@fsm1:~>mount /dev/vd1
```

- v. Copy image file from the USB storage device to the FSM node as shown in Example 6-6.

Example 6-6 Copy image file from USB storage device to the Flex System management node

```
USERID@fsm1:~>cp /media/vd1/RHEL6.3-2012.0-Server-x86_64.iso /home/USERID/.
```

3. Import the ISO image using the command `smcli importosimage`.

Example 6-7 Importing as OS image into the Flex System management node

```
USERID@fsm1:~>smcli importosimage /home/USERID/RHEL6.3-2012.0-Server-x86_64.iso
```

After importing the image, it will be displayed in the Image to Deploy column (within the Flex System Manager GUI).

6.7.2 Deploying a new image

You can launch the “Deploy Compute Node Images” task from the Initial Setup tab on the Home page, the Chassis Map view, or the IBM FSM Explorer console in the management software web interface. All three ways to launch the “Deploy Compute Node Images” task are shown:

- ▶ On the Initial Setup tab, click **Deploy Compute Node Images**, as shown in Figure 6-57.




Figure 6-57 FSM Initial Setup: Deploy Compute Node Images

- ▶ In the Chassis Map view, right-click the compute node and select **Deploy Compute Node Images**, as shown in Figure 6-58.

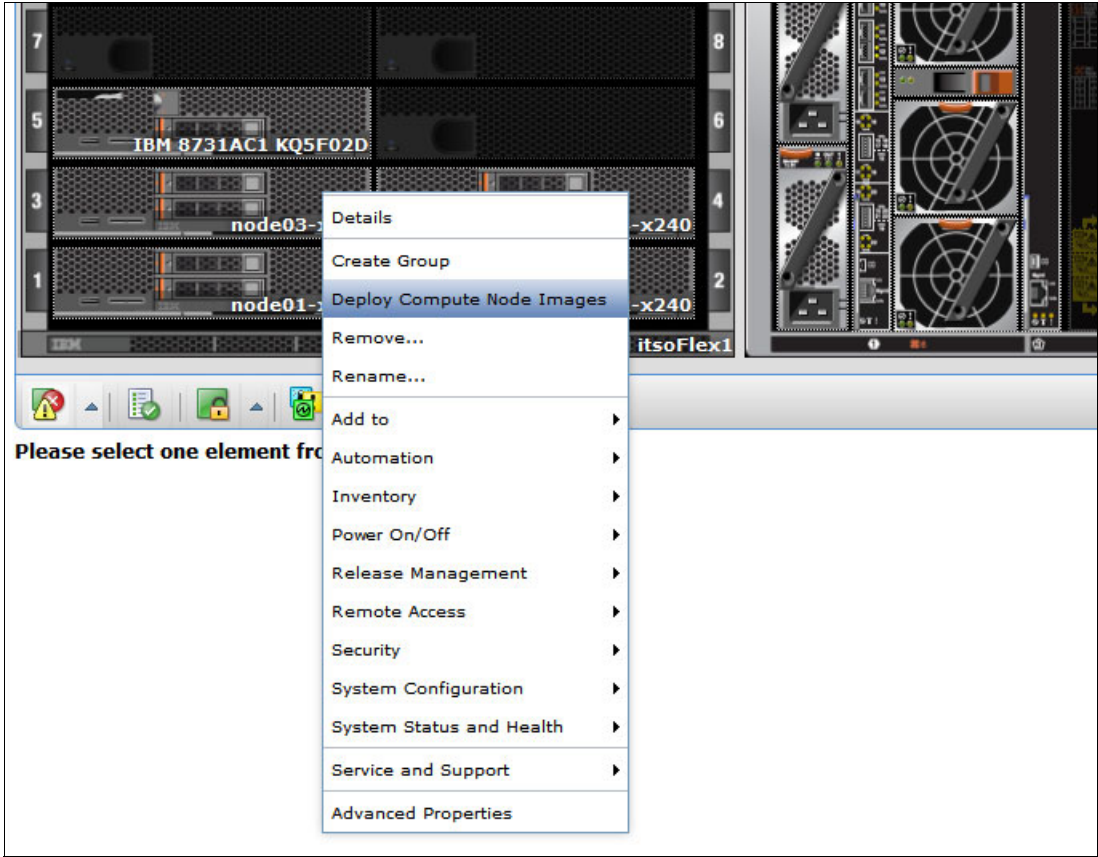


Figure 6-58 Chassis Map: Deploy Compute Node Images

- ▶ In the FSM Explorer, click **Systems** → **Deploy Compute Node Images**, as shown in Figure 6-59.

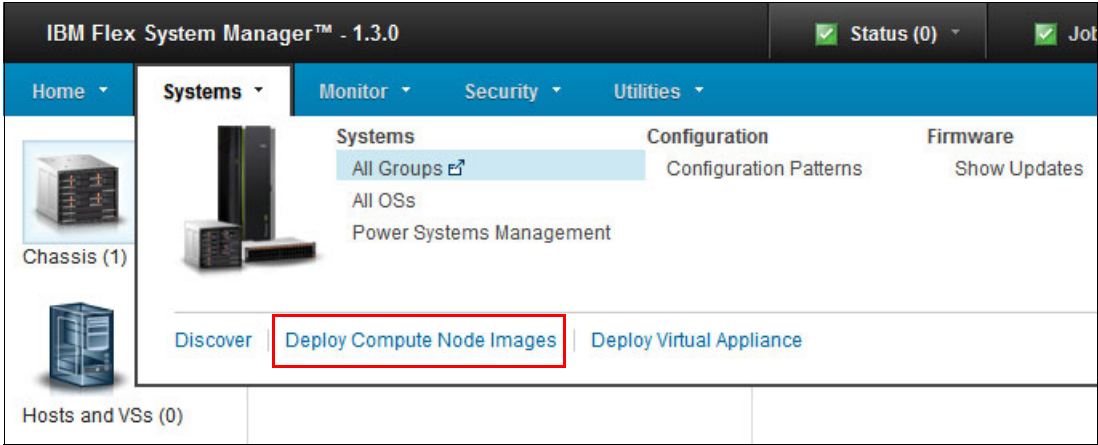


Figure 6-59 FSM Explorer: Deploy Compute Node Images

Perform the following steps to deploy a new operating system image (the numbers correspond to the numbers in Figure 6-60):

1. From the Deploy Compute Node Images window that is shown in Figure 6-60, select the compute nodes that you want to deploy. Multiple nodes can be selected.

The screenshot shows the 'Deploy Compute Node Images' window in IBM Flex System Manager. At the top, there are navigation tabs: Home, Systems, Monitor, Security, and Utilities. Below the tabs, there's a search bar 'Find a System'. The main heading is 'Deploy Compute Node Images'. A note states: 'Note: Before you begin, validate the IBM Flex System Manager network port being used to attach to the data network is configured to be on the same as the network ports on the compute node.' Below the note, there's a checkbox 'Apply this image to all rows:' followed by a dropdown menu showing 'esxi5.1-x86_64-install-Virtualization'. To the right of this is a link 'Learn more about importing additional images'. Below this is a row of tabs: 'Deploy Images', 'Global Settings', 'Remote Control', and 'Actions'. The 'Deploy Images' tab is selected. Below the tabs is a table with the following columns: Chassis and Node, Bay, Access State, Discovered Operating System, Deploy Status, MAC Address, and Image to Deploy. The table contains four rows of nodes. The first two rows have their 'Image to Deploy' dropdown menus selected, and their 'Chassis and Node' checkboxes are checked. Red boxes and arrows point to these elements with labels: '1. Select compute nodes to deploy' (pointing to the checked checkboxes), '2. Choose the image to deploy' (pointing to the selected dropdown menus), '3. Set OS credentials' (pointing to the 'Global Settings' tab), and '4. Deploy OS' (pointing to the 'Deploy Images' tab).

Chassis and Node	Bay	Access State	Discovered Operating System	Deploy Status	MAC Address	Image to Deploy
itsoFlex1						
node01-x240	1	OK	Not Discovered	Ready	34:40:B5:BE:7D:00	esxi5.1-x86_64-install-Virtualization
node02-x240	2	OK	Not Discovered	Ready	34:40:B5:BE:8E:90	esxi5.1-x86_64-install-Virtualization
node03-x240	3	OK	Not Discovered	Ready	34:40:B5:BE:83:D0	esxi5.1-x86_64-install-Virtualization
node04-x240	4	OK	Not Discovered	Ready	34:40:B5:BE:9D:58	esxi5.1-x86_64-install-Virtualization

Figure 6-60 Deploy Compute Node Images window

2. Choose the image to deploy from the Image to Deploy column. Alternatively, you can specify one image for multiple nodes by selecting the “Apply this image to all rows” check box.
3. Click **Global Settings** to set operating system administrative credentials (password for user root). This password will be used across all OS deployments through the FSM.
4. Click **Deploy Images** to begin the deployment process. In the Warning window that opens, click **Deploy**.

5. In the Deploy Compute Node Images job window (see Figure 6-61), you can choose to deploy the image immediately (Run Now) or at a scheduled time. You can also enter options to be notified when the job completes. Click **Submit**.

* **Schedule** Notifications Options

* Job Name:
Deploy Compute Node Images - Thursday, September 12, 2013 - 11:52:15 AM Mountain Daylight Time

When to run:
 Run Now
 Schedule

Back Next Submit Cancel

Figure 6-61 Deploy Compute Node Images job window

6. The status of the job can be viewed in the Jobs menu in the IBM Flex System Manager Explorer view as shown in Figure 6-62.

The screenshot displays the IBM Flex System Manager interface. At the top, it shows 'IBM Flex System Manager™ - 1.3.0' and a 'Jobs (159)' dropdown. A 'Deployment in progress' job is highlighted, with details including its name, status (Active), progress (25%), and last run status (Running). A sidebar on the right shows a list of jobs, with the current job selected. Below the job details, there is a section titled 'Deploy Compute Node Images' with a table of compute nodes.

Name and Node	Bay	Access State	Discovered Operating System	Deploy Status	MAC Address	Image to Deploy
Flex1						
node01-x240	1	OK	Not Discovered	Compute Node Restar	34:40:B5:BE:7D:00	rhels6.3-x86_64-install-Basic
node02-x240	2	OK	9.42.171.22	Ready	34:40:B5:BE:8E:90	esxi5.1-x86_64-install-Virtualization
node03-x240	3	OK	Not Discovered	Ready	34:40:B5:BE:83:D0	esxi5.1-x86_64-install-Virtualization

Figure 6-62 Deploy Compute Node Images job status

The Deploy Status column is updated with the installation steps as the job runs, including the following steps:

- ▶ Node created
- ▶ Bootable ISO mounted
- ▶ Boot order updated
- ▶ Compute node restarted
- ▶ Installing OS compute node
- ▶ OS successfully installed
- ▶ Node ready

Note: The deployment method currently supports the allocation of IPv6 addresses or the use of Dynamic Host Configuration Protocol (DHCP) for IPv4 (if a DHCP server is available on the same network as the Flex System management node).

After deployment, if you are not using the listed network address allocations or the OS is not reachable, follow these steps to configure network settings:

1. Right-click the target node from the Deploy Compute Node Images page.
2. Choose **Remote Control**.
3. The process uses the IMM to access the node console and then logs in to the OS using the password entered through the Global deployment settings.
4. Configure the networking as required to communicate with the Flex System management node (either VMWare or Red Hat network configuration steps).

6.8 System discovery, access, and inventory collection

To manage a resource within an environment or view inventory data about it, that resource must first be discovered. After access is granted, an inventory must be collected. The resource is recognized and added to the comprehensive list of native resources and native attributes for the system. Discovery and inventory collection are the two primary tasks that are used to connect to supported network resources and collect information about them.

The following topics are covered:

- ▶ 6.8.1, “Discovery basics” on page 176
- ▶ 6.8.2, “Operating system discovery” on page 179
- ▶ 6.8.3, “Requesting access to the discovered operating system” on page 182
- ▶ 6.8.4, “Collecting operating system inventory” on page 184

6.8.1 Discovery basics

Discovery is the process by which IBM Flex System Manager identifies and establishes connections with network-level resources that IBM Flex System Manager can manage. These resources include compute nodes, operating systems, switches, and external storage devices. Use system discovery to identify resources within your environment, collect data about those resources, and establish connections with them.

A *discovery protocol* is any network communication protocol that IBM Flex System Manager uses during the discovery process to discover a resource. The default discovery profile uses a predetermined list of protocols. When you specify a single IP address, a single host name, or a single range of IP addresses, system discovery uses one or more protocols. These protocols are based on the selected target resource type. With a discovery profile, you can refine the target resource type and configure specific protocols that you want to use.

The communication protocols that IBM Flex System Manager uses during discovery depend on the protocols that are used by the target resource type. You need to decide about the different protocols only when you create or edit a discovery profile. The Discovery Profile wizard helps you select and configure the correct protocol for the type of resource that you want to discover.

When you are discovering many resources, network traffic that is associated with the discovery process might cause timeouts. These timeouts might result in some discoverable resources remaining undiscovered. To help prevent this problem, use one or more discovery profiles. With a discovery profile, you can target specific resources and limit the number of communication protocols that are used during discovery.

By default, IBM Flex System Manager supports the following discovery protocols:

- ▶ Agent manager discovery

Agent manager discovery specifically targets the discovery of Tivoli Common Agents. In the Tivoli paradigm, Service Location Protocol (SLP) is not supported. Management nodes must contact an agent manager that knows about the agents in their environment. You can select the agent managers that you want to use in discovery.
- ▶ Common Agent Services discovery

This discovery uses SLP discovery, with which clients can locate servers and other services in the network.

- ▶ Common Information Model (CIM) discovery
 CIM discovery uses the Service Location Protocol (SLP) for discovery. With CIM discovery, clients can locate servers and other services in the network.
- ▶ Interprocess communication (IPC) discovery
 IPC is the process by which programs send messages to each other. Sockets, semaphores, signals, and internal message queues are common methods of interprocess communication. IPC is also a mechanism of an operating system that enables processes to communicate with each other within the same computer or over a network. IPC uses services that IBM Flex System Manager provides that components use to communicate with each other. By using these services, a server task can communicate with an agent task that is running on a target.
- ▶ Secure Shell (SSH) discovery
Secure Shell is a command interface and protocol that is based on UNIX for securely accessing a remote computer. With SSH discovery, you can specify either a single IP address or a range of IP addresses upon which to run discovery.
- ▶ Simple Network Management Protocol (SNMP) discovery
 SNMP is a network management standard that is widely used in Internet Protocol networks. SNMP runs management services by using a distributed architecture of management systems and agents. SNMP provides a method of managing network hosts, such as workstation and server computers, routers, bridges, and hubs from a centrally located computer that runs the network-management software.
- ▶ Storage Management Initiative Specification (SMI-S) discovery
 With SMI-S discovery, clients can locate servers and other services in the network. This design specification was developed by the Storage Networking Industry Association (SNIA). It specifies a secure and reliable interface with which storage management systems can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools that are used to manage them.
- ▶ Windows distributed component object model (DCOM) discovery
 Use Windows DCOM (an extension of the Microsoft Component Object Model (COM)) to support objects that are distributed across a network configuration. Use DCOM to specify either a single IP address or a range of IP addresses on which to run discovery.

The system discovery task can be started in one of the following ways:

- ▶ From the FSM Explorer web interface using the search capabilities: Type system discovery in the search box in the upper-right corner of the window, then click **System Discovery** in the Tasks panel, as shown in Figure 6-63.




Figure 6-63 FSM Explorer: System Discovery

- ▶ From the Plug-ins tab in the Home window: Click **System Discovery** under Discovery Manager, as shown in Figure 6-64.




Figure 6-64 FSM Home window: System Discovery

6.8.2 Operating system discovery

One of the most common uses of the discovery tasks is the operating system discovery and inventory collection. This task must be completed every time that a new operating system, hypervisor, or guest virtual machine is deployed when it needs to be managed by the Flex System Manager.

Perform the following steps to discover an operating system or hypervisor:

1. Open the System Discovery window by using one of the methods described in 6.8.1, “Discovery basics” on page 176.
2. In the System Discovery window, select a required discovery option, enter an IP address or a range of IP addresses, and select a resource type to discover (or leave **All** to discover all resource types), as shown in Figure 6-65. Click **Discover Now**.

System Discovery

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a single IP address or host name, discover resources of the same type for a range of IP addresses, or use a discovery profile. Discovery profiles enable you to customize discoveries, including importing IP addresses, and requesting access to and collecting discovered resources.

[Learn more about using discovery](#)

Select a discovery option:
Single IPv4 address

IP address:
9 . 42 . 171 . 21

Select the resource type to discover:
Operating System

Discover Now

Schedule...

Advanced Tasks
[Create new profile](#)
[Manage discovery profiles](#)
[Discovery jobs](#)

Figure 6-65 Enter IP address to discover

3. A blue informational message is displayed that indicates that the job is started, and the Processing discovery protocols status is displayed, as shown in Figure 6-66. Click **Display Properties** to check the job status, if needed.

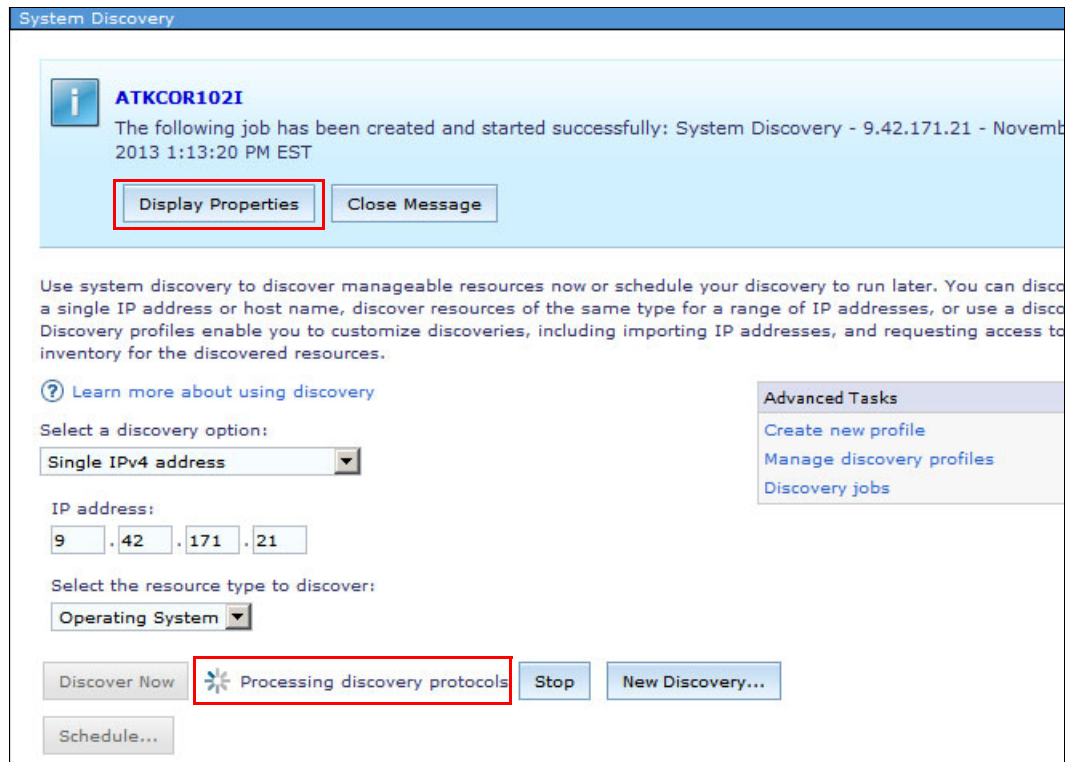


Figure 6-66 Discovery job information

Wait until the progress bar reaches 100%, which indicates that the discovery is complete, as shown in Figure 6-67.

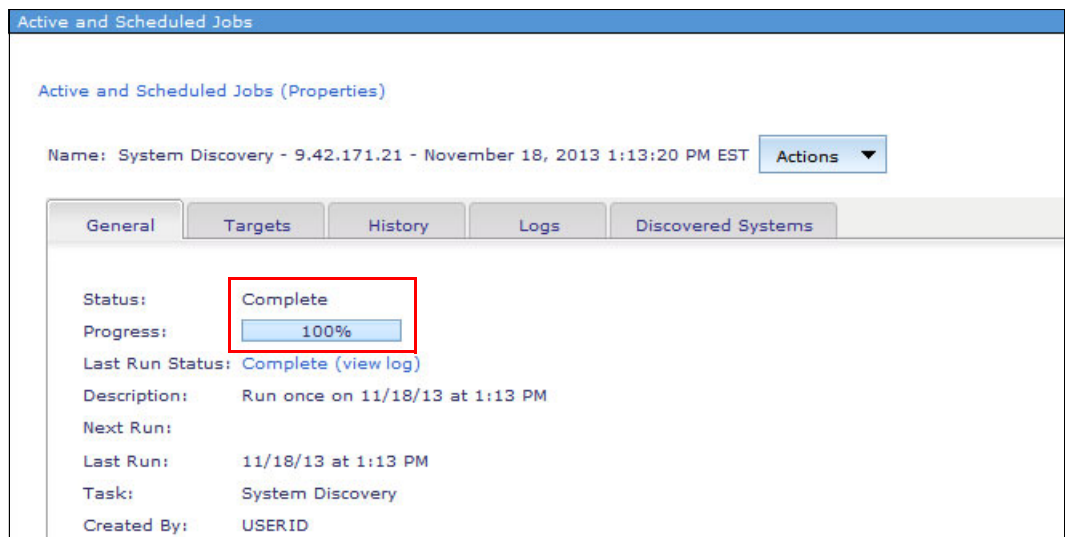


Figure 6-67 Discovery completed

- The list of discovered systems will be displayed under Discovered Manageable Systems, as shown in Figure 6-68.

System Discovery

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a single IP address or host name, discover resources of the same type for a range of IP addresses, or use a discovery profile. Discovery profiles enable you to customize discoveries, including importing IP addresses, and requesting access to inventory for the discovered resources.

[Learn more about using discovery](#)

Select a discovery option:
 Single IPv4 address

IP address:
 9 . 42 . 171 . 21

Select the resource type to discover:
 Operating System

Discover Now

Schedule...

Advanced Tasks

- Create new profile
- Manage discovery profiles
- Discovery jobs

Discovered Manageable Systems:

Actions | Search the table... Search




Select	Name	Discovered	Type	Access	Problems
<input type="checkbox"/>	 9.42.171.21	New	Operating Sys...	 No access	 OK

Figure 6-68 Discovered Manageable Systems

6.8.3 Requesting access to the discovered operating system

The discovered manageable operating systems (see 6.8.2, “Operating system discovery” on page 179) are displayed in the System Discovery window as shown in Figure 6-68 on page 181. Perform the following steps to request access to the newly discovered object:

1. Select one or multiple discovered systems with no access (No access is listed in the Access column). Then, click **Actions** → **Security** → **Request Access**, as shown in Figure 6-69.

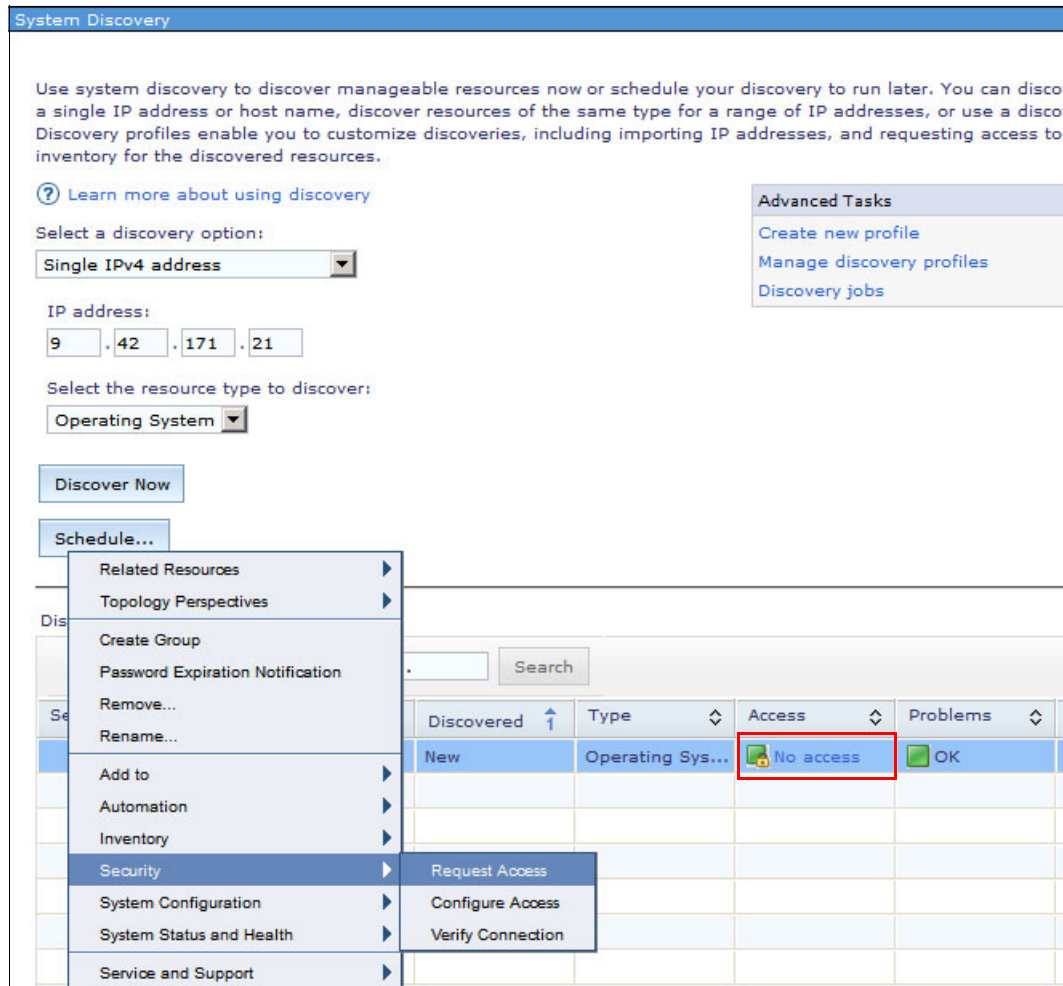


Figure 6-69 Newly discovered systems

2. In the Request Access window, enter operating system credentials and click **Request Access**, as shown in Figure 6-70.

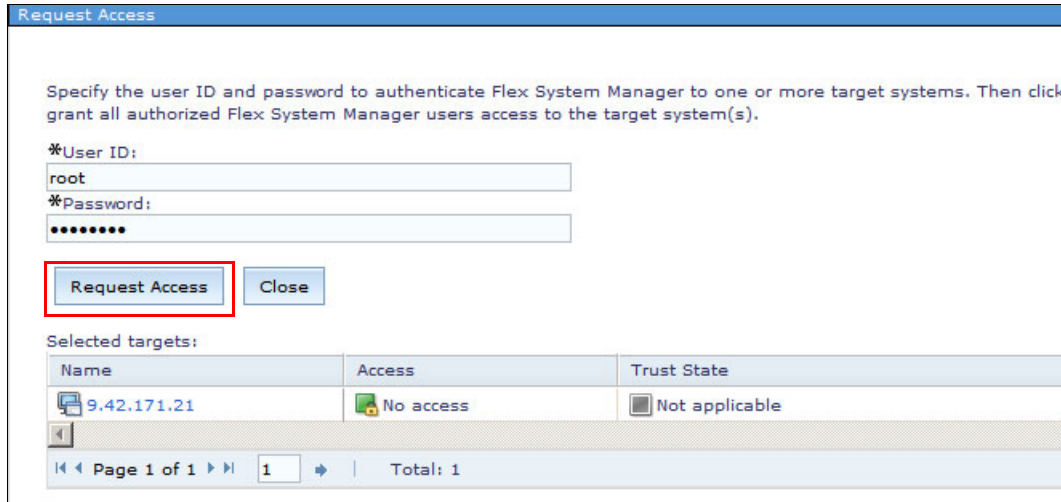


Figure 6-70 Granting access to the object

3. After you request access to the object, ensure that access is granted (OK is displayed in the Access column), as shown in Figure 6-71. Click **Close**.

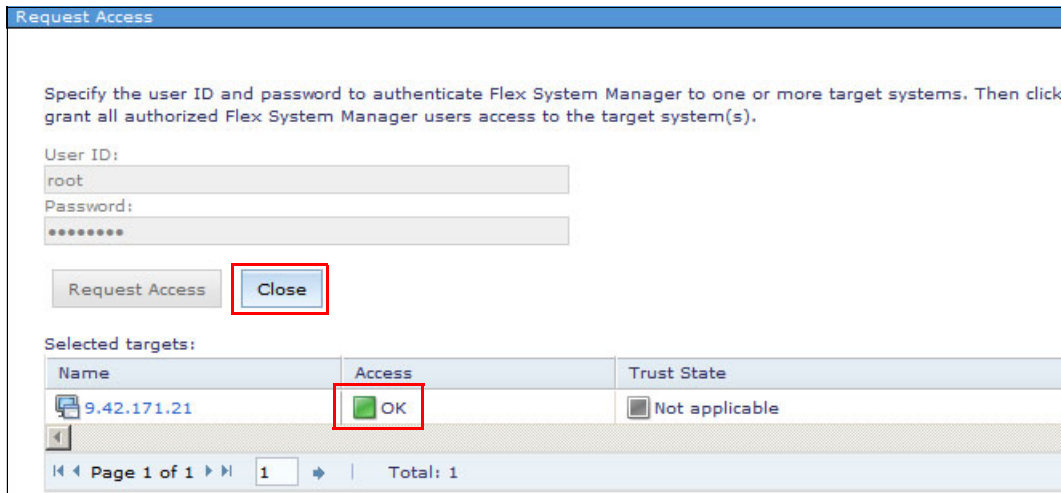


Figure 6-71 Access is granted

Object discovery: Additional objects that are associated with the discovered OS might also be discovered, such virtual switches in the hypervisors.

6.8.4 Collecting operating system inventory

The manageable operating systems with access granted (see 6.8.2, “Operating system discovery” on page 179 and 6.8.3, “Requesting access to the discovered operating system” on page 182) are displayed in the System Discovery window, as shown in Figure 6-72.

The screenshot shows the 'System Discovery' window. At the top, there is a title bar and a descriptive paragraph. Below this, there are configuration options: a 'Learn more about using discovery' link, a 'Select a discovery option:' dropdown menu set to 'Single IPv4 address', an 'IP address:' field with input boxes for '9', '.42', '.171', and '.21', and a 'Select the resource type to discover:' dropdown menu set to 'Operating System'. There are 'Discover Now' and 'Schedule...' buttons. On the right side, there is an 'Advanced Tasks' panel with links for 'Create new profile', 'Manage discovery profiles', and 'Discovery jobs'.

Below the configuration section is a section titled 'Discovered Manageable Systems:' containing a table with columns for 'Select', 'Name', 'Discovered', 'Type', 'Access', 'Problems', and 'Co'. The table lists four items: '9.42.171.21' (Operating Sys...), 'vSwitch0-9.42.171.21' (Switch), 'vSwitchUSB0-9.42.171.21' (Switch), and 'node01-x240' (Server). All items have 'OK' in the 'Access' and 'Problems' columns.

Select	Name	Discovered	Type	Access	Problems	Co
<input checked="" type="checkbox"/>	9.42.171.21	New	Operating Sys...	OK	OK	
<input type="checkbox"/>	vSwitch0-9.42.171.21	New	Switch	OK	OK	
<input type="checkbox"/>	vSwitchUSB0-9.42.171.21	New	Switch	OK	OK	
<input type="checkbox"/>	node01-x240	Previous	Server	OK	OK	

Figure 6-72 Discovered objects with access granted

Perform the following steps to collect inventory on the newly discovered object:

1. In the System Discovery window, select the discovered objects and click **Actions** → **Inventory** → **Collect Inventory**, as shown in Figure 6-73.




Figure 6-73 Inventory collection

2. In the Launch Job window, to begin the inventory collection, select **Run Now** and click **OK**, as shown in Figure 6-74.




Figure 6-74 Run collect inventory

3. A blue informational message is displayed that indicates that the job is started, as shown in Figure 6-75. You can click **Display Properties** to check the job status.

The screenshot shows the 'System Discovery' window. At the top, a blue informational message box titled 'ATKCOR102I' states: 'The following job has been created and started successfully: Collect Inventory - November 18, 2013 2:00 PM EST'. Below the message are two buttons: 'Display Properties' (highlighted with a red box) and 'Close Message'. Below the message box, there is a section for configuring discovery options. It includes a dropdown menu for 'Select a discovery option:' set to 'Single IPv4 address', an 'IP address:' field with input boxes containing '9', '.42', '.171', and '.21', and another dropdown menu for 'Select the resource type to discover:' set to 'Operating System'. There are 'Discover Now' and 'Schedule...' buttons. To the right of the configuration section is an 'Advanced Tasks' panel with links for 'Create new profile', 'Manage discovery profiles', and 'Discovery jobs'. At the bottom, a table titled 'Discovered Manageable Systems:' displays a list of discovered resources.

Select	Name	Discovered	Type	Access	Problems
<input checked="" type="checkbox"/>	9.42.171.21	New	Operating Sys...	OK	OK
<input type="checkbox"/>	vSwitch0-9.42.171.21	New	Switch	OK	OK
<input type="checkbox"/>	vSwitchUSB0-9.42.171.21	New	Switch	OK	OK
<input type="checkbox"/>	node01-x240	Previous	Server	OK	OK

Figure 6-75 Collect inventory information

4. Wait until the job is completed.

The inventory is now collected. You can view collected inventory in the View and Collect Inventory window (shown in Figure 6-76) by selecting one or more systems and clicking **Actions** → **Inventory** → **View and Collect Inventory** in the System Discovery window.

Collected Items

- Summary
- Hardware Devices
 - Ethernet Port
 - Fibre Channel Port
 - Memory
 - Port Controller
 - Processor
- Network Configuration
- Physical Hardware
 - Card
 - Chassis
 - Chip
 - Physical Connector
 - Physical Memory
 - Physical Port
 - Slot
- Related Systems
- System Internals
- System Software
- Virtual Configuration

System Summary

System name: 9.42.171.21
 Type: Operating System
 Access State: Full Access/Communication OK
 Last Collected: November 18, 2013 4:57 PM
 Protocols: CIM

Software Summary

Software Type: Hypervisor
 Software Version: 5.1.0
 Build Number: 799733
 System BIOS: B2E1/22D/1.21

Network Summary

Hostname:
 IP Addresses: fe80:0:0:0:3640:b5ff:febf:IPV4_1_0_, 9.42.171.16, fe80:0:0:0:3640:b5ff:febf:169.254.95.101, 9.42.171.21, 169.254.95.118, IPV4_0_0_
 MAC Addresses: 3440B5BE7D00, 3440B5BF4D70, 3440B5BF4D72, 3440B5BE7D04, 3640B5BF4D73, 3440B5BF4D71

Asset Summary

Manufacturer: IBM
 Model: AC1
 Machine Type: 8737
 Serial Number: KQ9M03F
 Architecture: x86_64
 UUID: 35f25f27-d8fd-0f31-9b61-3384af289b01

Utilization Summary

Number of Processors: 2
 Number of Cores: 16
 Max Processor Speed: 4000 MHz
 Processor Family: Intel® Xeon®

Related Systems

Actions | Search the table... Search

Select	Name	System n...	Type	Software ...	So
<input type="checkbox"/>	itsoFlex1	9.42.171.21	System Chassis		
<input type="checkbox"/>	node01-x240	9.42.171.21	Management ...		
<input type="checkbox"/>	vSwitch0-9.42.171.21	9.42.171.21	Switch		
<input type="checkbox"/>	vSwitchUSB0-9.42.171.21	9.42.171.21	Switch		

Page 1 of 1 | 1 | Selected: 0 Total: 4 Filtered: 4

Installed Firmware

Actions | Search the table... Search

Figure 6-76 View inventory collected

Tip: You can also invoke Collect Inventory and View and Collect Inventory tasks for any managed object from the IBM Flex System Manager Explorer user interface, as shown in Figure 6-77.

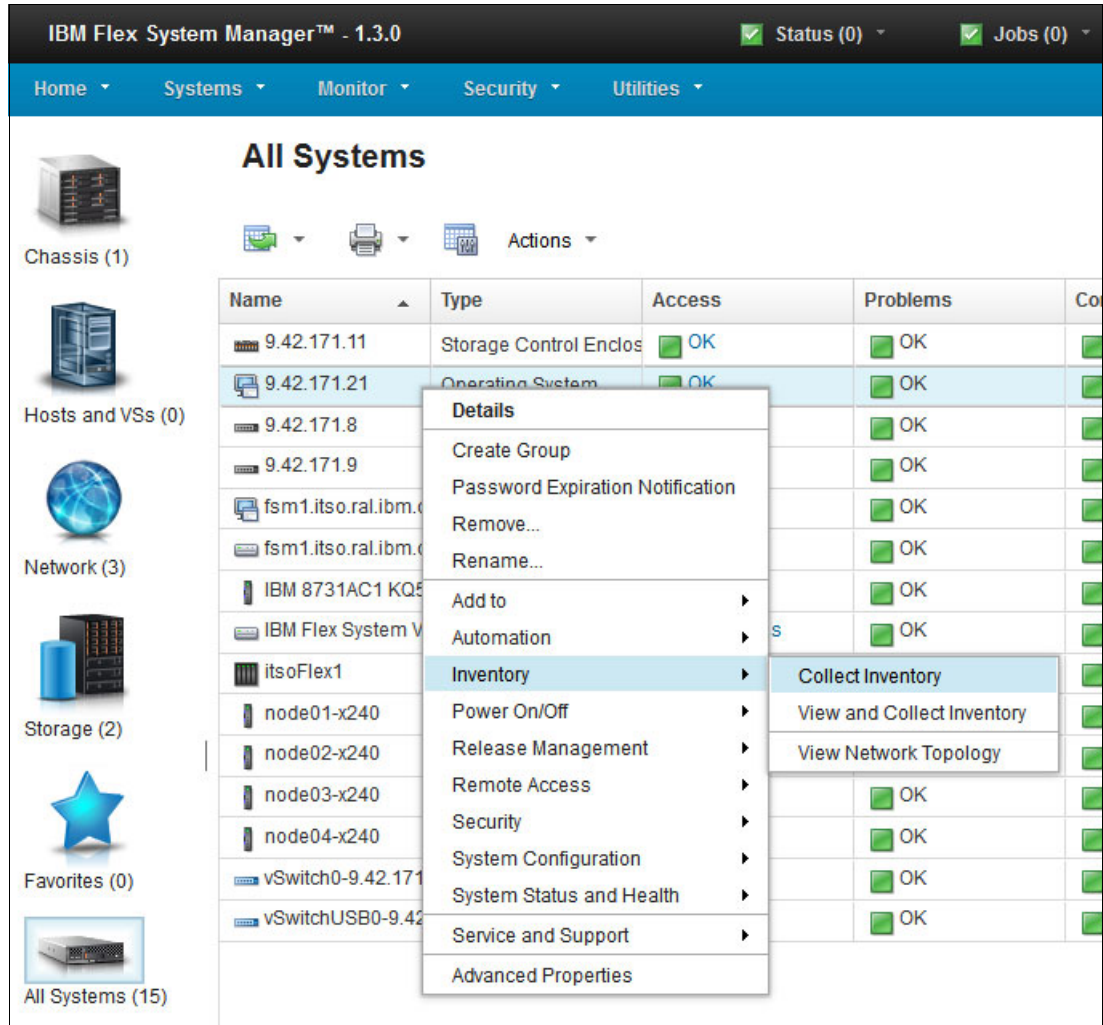


Figure 6-77 Flex System Manager Explorer: View and Collect Inventory

6.9 Updating chassis components

The next step in the FSM initial configuration tasks is to update firmware or software on the chassis components including compute nodes, storage nodes, and I/O modules. The following topics are covered:

- ▶ 6.9.1, “Acquiring updates for chassis components” on page 191
- ▶ 6.9.2, “Updating the CMM firmware” on page 199
- ▶ 6.9.3, “Updating compute node firmware” on page 201
- ▶ 6.9.4, “Updating I/O module firmware” on page 205
- ▶ 6.9.5, “Compliance policies” on page 211

Chassis components can be updated from the Update Chassis Components window (see Figure 6-79 on page 191) which can be opened by clicking Update Chassis Components on the Initial Setup tab in the Home window, as shown in Figure 6-78.



Figure 6-78 Initial Setup: Update Chassis Components

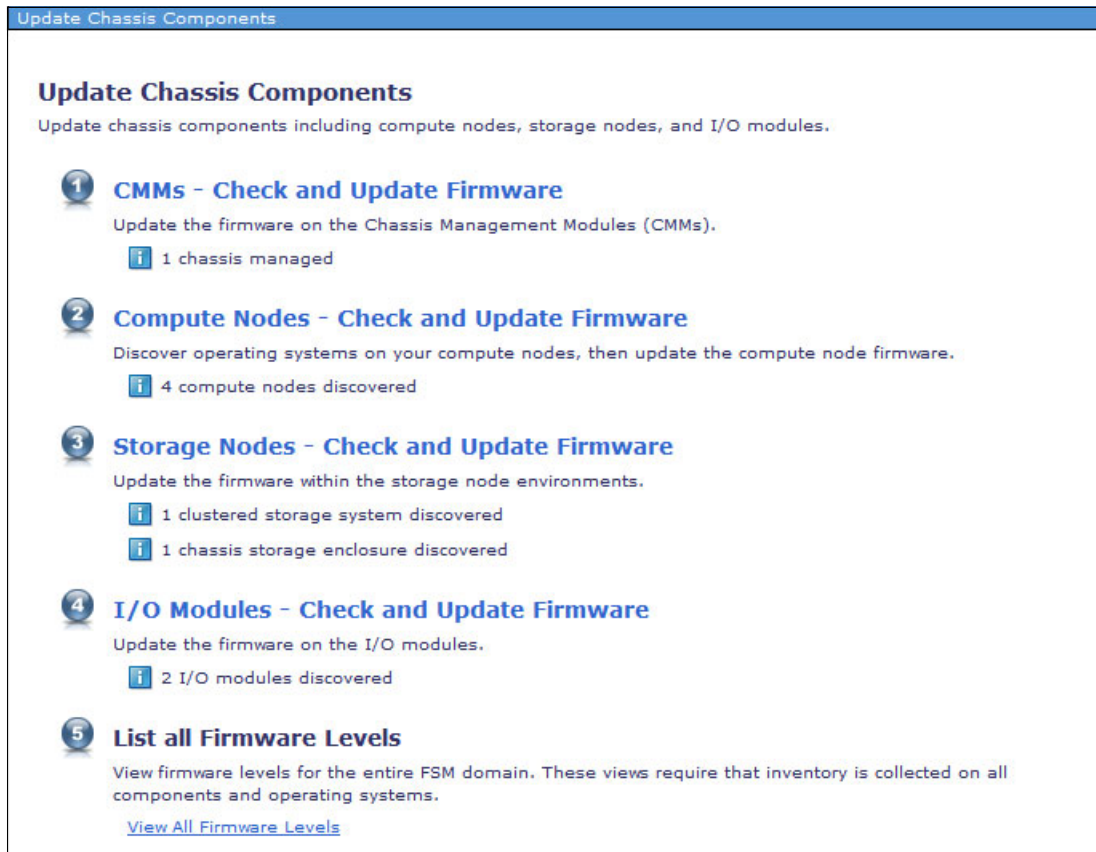


Figure 6-79 Update Chassis Components window

6.9.1 Acquiring updates for chassis components

The IBM Flex System Manager Update Manager is responsible for obtaining and applying chassis, switch, system firmware, and certain operating system updates from IBM. In addition, the Update Manager is used to update the FSM itself. The updates can be obtained by Internet connection from the FSM. They can also be downloaded manually from IBM to another workstation, then copied to the FSM by FTP or Secure Copy Protocol (SCP) connection. After the files are copied to the FSM, they must be imported into the Update Manager. First, you need to set up the Internet connection.

Direct Internet connection

To set up and test the Internet connection, perform these steps:

1. Starting from the Home page, click the **Plug-ins** tab. The Plug-ins window lists all of the managers that are available on the FSM, as shown in Figure 6-80.




Figure 6-80 FSM list of manager plug-ins

- From the list of manager plug-ins, click **Update Manager** to display the window that is shown in Figure 6-81.

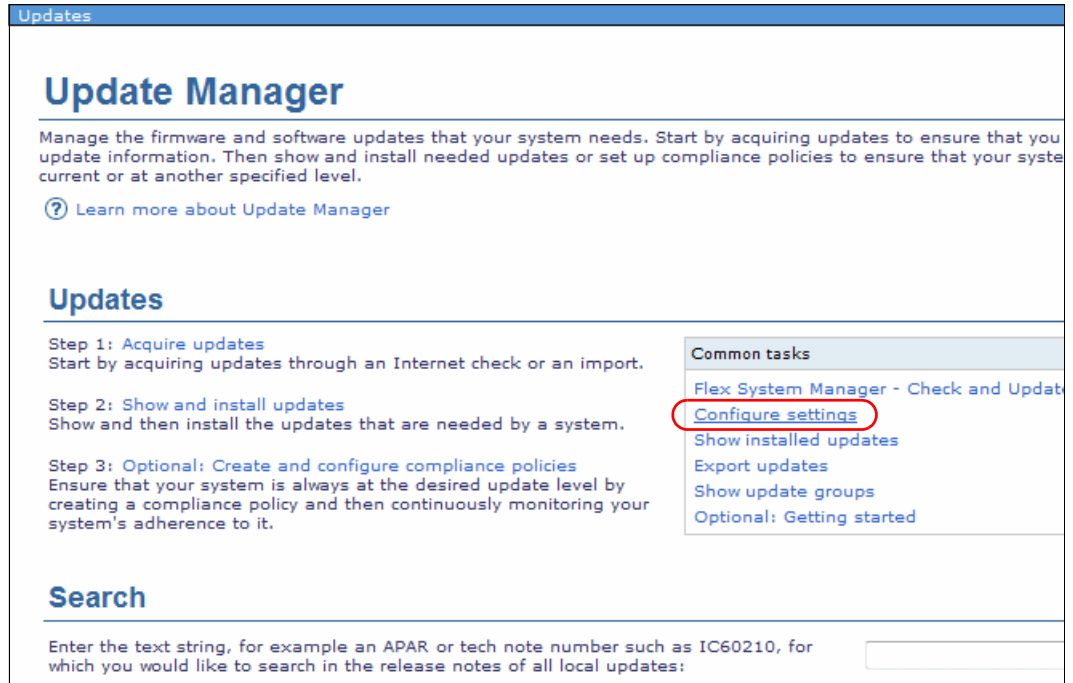


Figure 6-81 FSM Update Manager window

3. In the Common task box, click the link for **Configure settings** to open the window that is shown in Figure 6-82.

This window allows for the configuration of a direct Internet connection, or the configuration settings to use an existing proxy server.

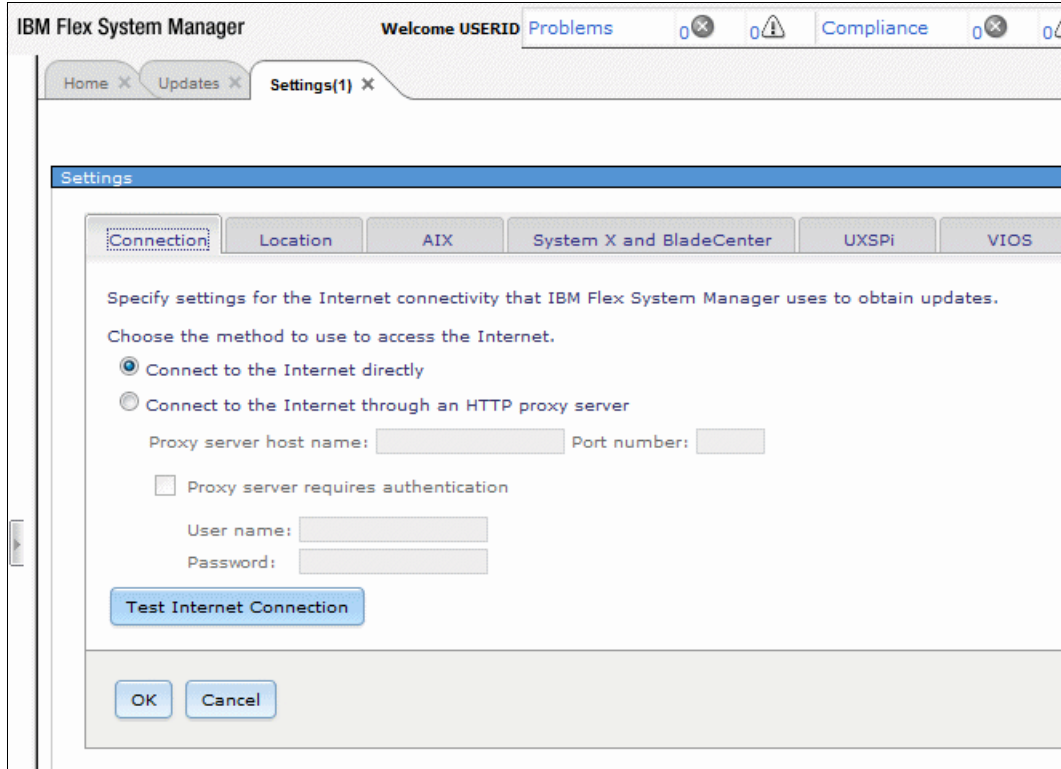


Figure 6-82 FSM Update Manager Internet connection settings

4. With the settings complete, click **Test Internet Connection** to verify the connection.
The test attempts to make a connection to a target IBM server. During the test, a progress indicator displays, as shown in Figure 6-83.

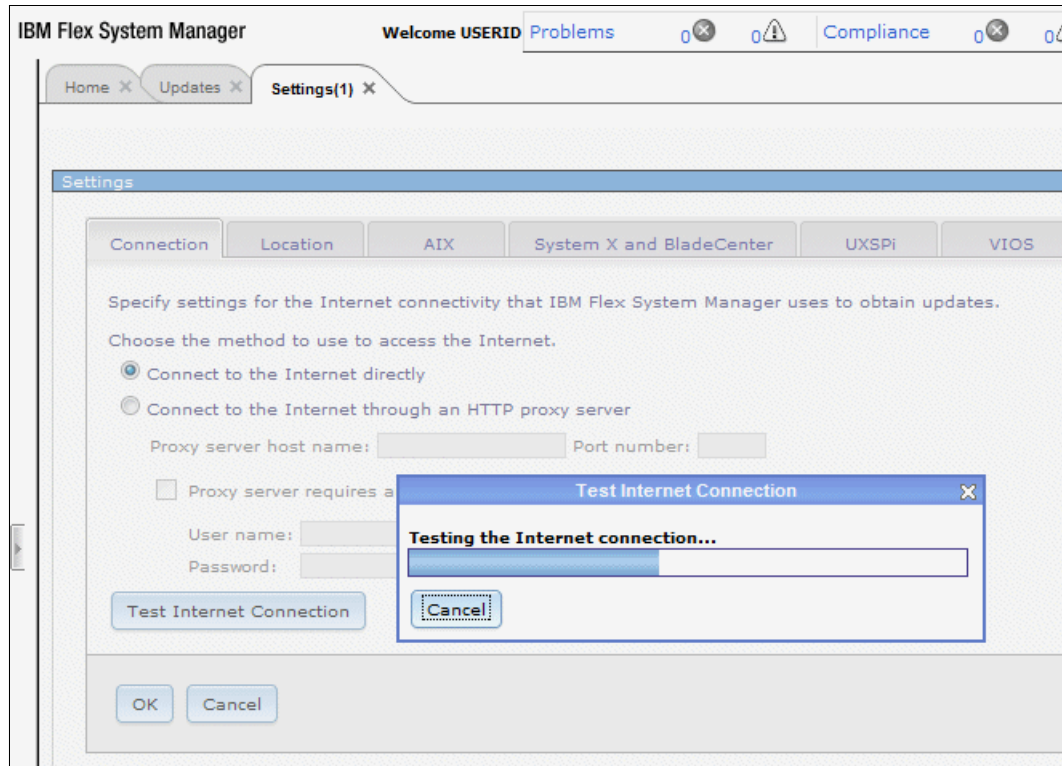


Figure 6-83 FSM testing Internet connection for Update Manager

A message is displayed upon successful completion, as shown in Figure 6-84.

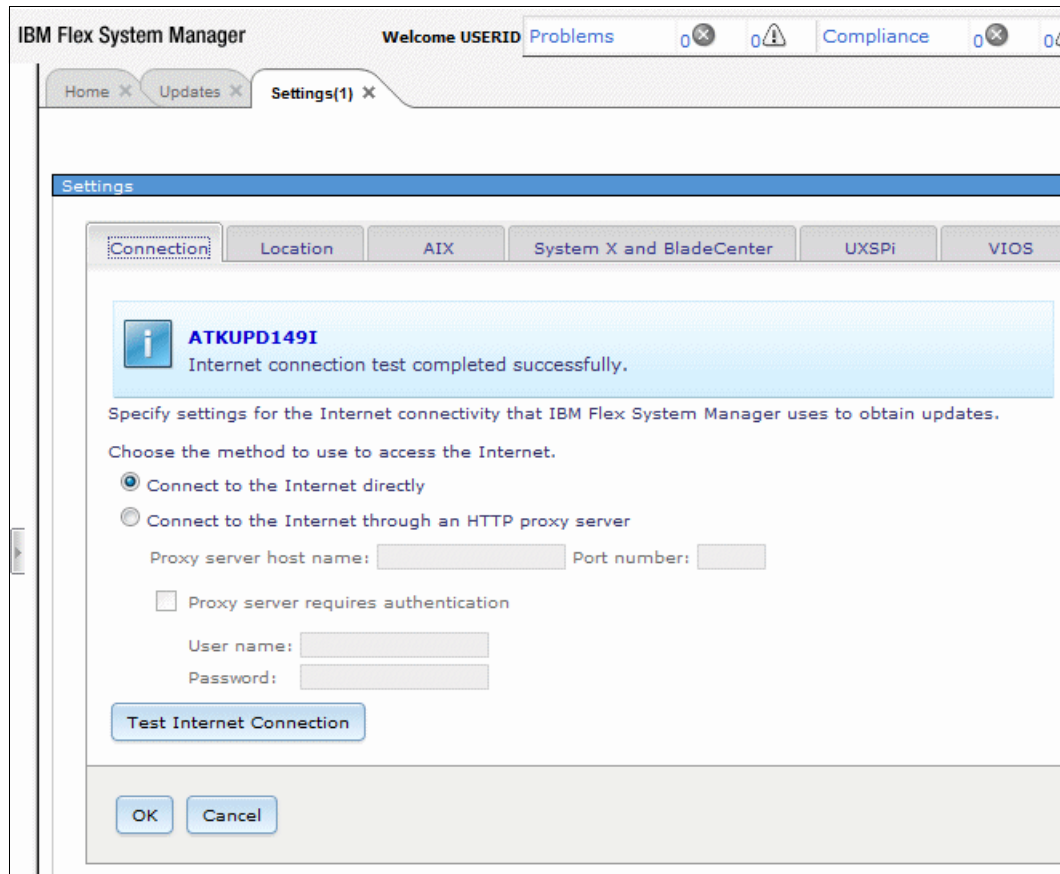


Figure 6-84 Successful Internet connect test for Update Manager

With the test successful, the Update Manager can obtain update packages directly from IBM.

If a direct Internet connection is not allowed for the FSM, the following steps show how to import updates files into Update Manager.

Importing update files

If the FSM does not have Internet connection, all updates can be downloaded from IBM Fix Central and then imported into the FSM manually. The process includes downloading firmware packages to the local administrator's workstation, then copying them to the FSM node using scp protocol, and then importing them into the FSM repository through the CLI.

Perform the following steps:

1. Go to IBM Fix Central (<http://www.ibm.com/support/fixcentral>), log in using your IBM ID and password (IBM ID is required to download firmware from IBM Fix Central). If you do not have an IBM ID, you can register at this website: <http://www.ibm.com/account>.

2. Select required PureFlex System components, as shown in Figure 6-85. Click **Continue**.

Select product **Find product**

Select the product below.

When using the keyboard to navigate the page, use the **Alt** and **down arrow** keys to navigate the selection lists.

Product Group
PureSystems

Select from PureSystems
PureFlex System

Select from PureFlex System
Chassis

Select from Chassis
Enterprise Chassis

Select from Enterprise Chassis
8721

Operating system
All

Continue

Figure 6-85 Fix Central: Component selection

- From the Select Fixes page select the firmware packages that you want to download as shown in Figure 6-86.




Figure 6-86 Select firmware packages

- Enter you chassis machine type and serial number and click **Continue**. Agree with Terms and Conditions.
- Download required files from the Fix Central.
- Copy downloaded files from your local workstation to the /home/<username> directory on the FSM node using the scp protocol.

7. Import the updates from the `/home/<username>` directory into the FSM repository using `smcli importupd` command, as shown in Example 6-8.

Example 6-8

```
USERID@fsm1:~> smcli importupd -v /home/USERID
ATKUSC206I Generating SDDs for path: "/home/USERID".
ATKUPD293I Update "ibm_fw_imm2_1aoo42y-2.60_anyos_noarch" was successfully
imported to the library.
ATKUPD293I Update "ibm_fw_uefi_b2e126e-1.31_anyos_32-64" was successfully
imported to the library.
ATKUPD573I Running compliance for all new updates that were found.
ATKUPD286I The import updates task has completed successfully.
USERID@fsm1:~>
```

The updates are now ready to be applied to the chassis components.

6.9.2 Updating the CMM firmware

The FSM can push firmware updates to the Chassis Management Module (CMM). When a Flex Chassis is set to managed, inventory is collected automatically. For more information about setting the Flex Chassis to be managed, see 6.3, “Selecting chassis to manage” on page 138. From this point, when updates are imported (either manually or from the IBM site) a compliance check is run against the CMM. This process compares any new CMM firmware with the currently installed CMM firmware. If a newer firmware is found, a message is displayed in the Update Chassis Components window as shown in Figure 6-87.




Figure 6-87 Update Chassis Components: CMMs

To update the firmware, perform these steps:

1. Click **Install updates** (see Figure 6-87 on page 199). A list of newer available firmware is displayed that can be deployed to the CMM, as shown in Figure 6-88. Select the update, and click **Install**.

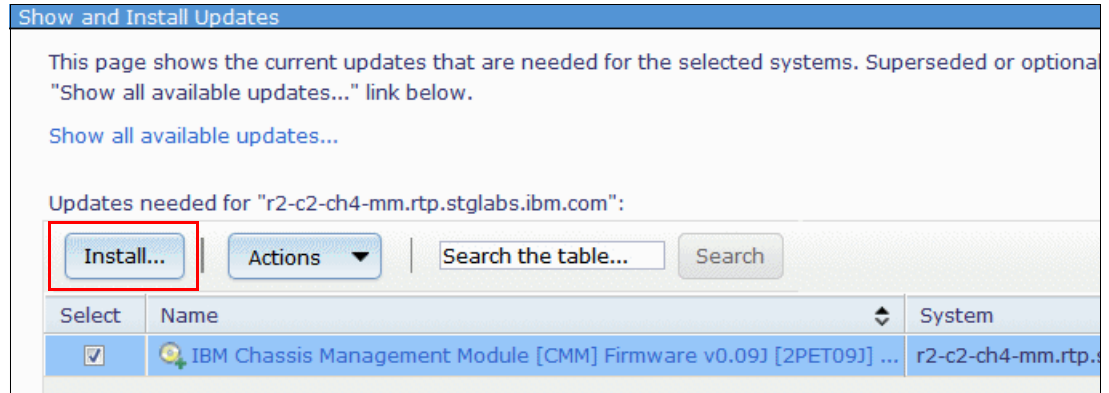


Figure 6-88 CMM Show and Install Updates window

2. An installation wizard is displayed. Click **Next** through the initial Welcome window. The Restarts section shows whether restart of the device is required, as shown in Figure 6-89. Click **Next**.

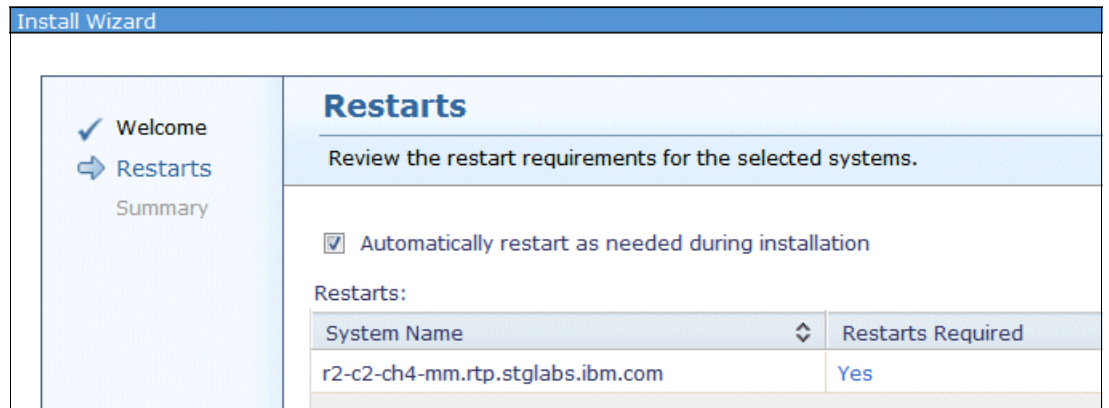


Figure 6-89 Install Wizards Restarts window

3. The Summary window opens as shown in Figure 6-90. Click **Finish**.




Figure 6-90 Install Wizard Summary window

4. Select to either **Run Now** or **Schedule** as shown in Figure 6-91, and click **OK**.

The screenshot shows a window titled "Launch Job" with three tabs: "Schedule", "Notification", and "Options". The "Schedule" tab is selected. Below the tabs, the text "Job name and schedule" is displayed. A label "*Job Name:" is followed by a text input field containing "Install Updates - June 26, 2012 11:58:20 PM EDT". Below the input field, the text "Choose when to run the job." is displayed. There are two radio buttons: "Run Now" (selected) and "Schedule".

Figure 6-91 Install Wizard Schedule Job window

The update is deployed to the CMM, which is then restarted as part of the job. After the CMM is restarted, inventory will be run against it, and a compliance check to compare the newly installed firmware against the FSM update repository. This process ensures that the firmware is up to date, and the CMM has a compliant version.

6.9.3 Updating compute node firmware

Firmware updates for the compute nodes are applied in-band, i.e. through the operating system (OS). For this reason, you must first discover and inventory the OS or hypervisor that is running on the compute node (see 6.8, "System discovery, access, and inventory collection" on page 176).

After the OS is discovered, accessed and inventoried, firmware updates can be pushed. After updates are imported as addressed in 6.9.1, "Acquiring updates for chassis components" on page 191, a compliance check will occur automatically against all discovered systems. With discovered operating systems, the FSM compares its repository of updates against the firmware inventory available.

To determine whether a compute node needs updates and apply these updates if needed, perform these steps:

1. In the Update Chassis Components window, click **Compute Nodes - Check and Update Firmware**, as shown in Figure 6-92.

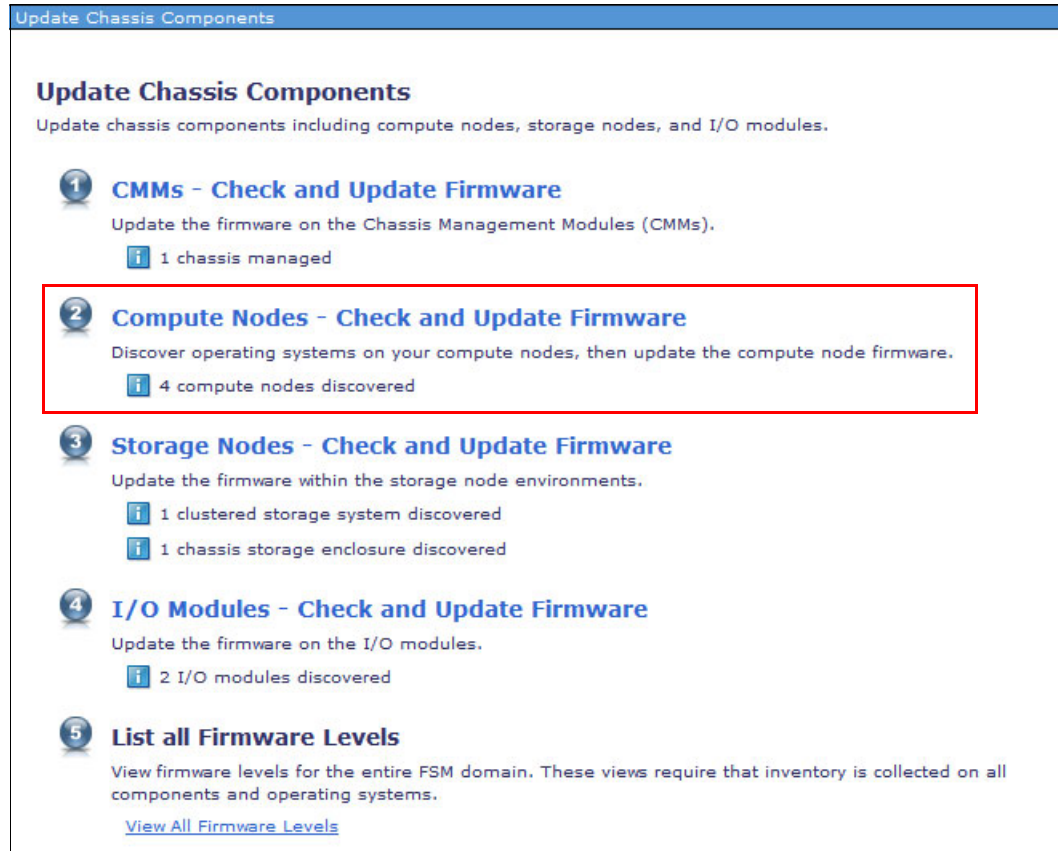


Figure 6-92 Update Chassis Components window: Check for compute node updates

2. In the Compute Nodes - Check and Update Firmware window, click **Check for Updates**, as shown in Figure 6-93.

Discovery, access, and inventory: We already discovered operating systems, granted access to them, and collected their inventory, as described in 6.8, “System discovery, access, and inventory collection” on page 176. Alternatively, you can discover, request access, and collect inventory for the operating systems from the Compute Nodes - Check and Update Firmware window.




Figure 6-93 Compute Nodes - Check and Update Firmware window: Check for Updates

3. In the Acquire Updates window (see Figure 6-94), select the method of acquiring updates and click **OK** to start the Acquire Updates job.

Previously imported updates: If you already imported updates (see “Importing update files” on page 196) you still need to specify a valid path in the Acquire Updates window. If no updates will be found, an error message appears indicating that no updates were found, but it can be safely ignored.




Figure 6-94 Acquire Updates window

4. After the Acquire Updates job is completed, click **Show and Install Updates**, as shown in Figure 6-95.




Figure 6-95 Acquire Updates window: Show and Install Updates

- Review the list of available updates as shown in Figure 6-96. The updates can be selected individually or all at one time, and deployed by clicking **Install**.




Figure 6-96 Show and Install Updates list

- Use the Install Wizard that is shown in 6.9.2, “Updating the CMM firmware” on page 199 to run the updates immediately or schedule them.

6.9.4 Updating I/O module firmware

Downloading updates for I/O modules (for example, Ethernet and Fibre Channel switches) is similar to downloading updates for compute nodes. This example shows updates for an IBM Flex System Fabric EN4093 10Gb Scalable Switch.

Compute node firmware updates are pushed through the Operating System (OS) by using specific OS protocols (DCOM for Windows, SSH for Linux). Updates for I/O modules are pushed out over SFTP. For this reason, the I/O module must support SFTP, otherwise an external FTP or TFTP server might need to be provided.

Ethernet I/O modules:

- ▶ EN2092, EN4093/EN4093R, and CN4093 switches require a TFTP server to host updates before applying them. The FSM node can be configured as a TFTP server for these purposes.
- ▶ EN2092, EN4093/EN4093R, and CN4093 switches must be configured to use the menu-based CLI (ibmnos-cli).

Perform the following steps to prepare the EN2092, EN4093/EN4093R, and CN4093 switches for the software updates through the FSM management node:

- Enable `ibmnos-cli` on the switch by issuing the commands listed in Example 6-9.

Example 6-9 Enabling `ibmnos-cli` on the Ethernet switch

```
EN4093>enable
EN4093#config t
EN4093(config)#boot cli-mode ibmnos-cli
EN4093(config)#copy running-config startup-config
EN4093(config)#reload
```

2. Enable TFTP server on the FSM management node. Open Update Manager settings (see “Direct Internet connection” on page 192), then click **System x and BladeCenter** tab, as shown in Figure 6-97. Click **OK**.

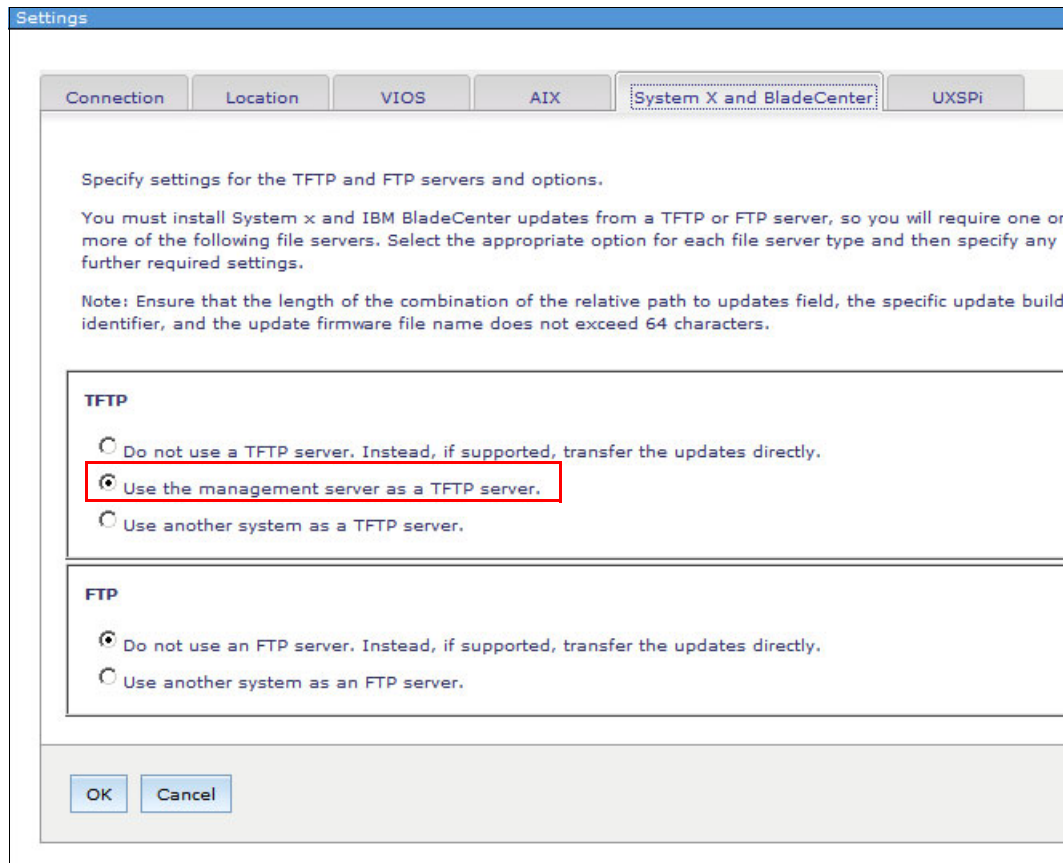


Figure 6-97 Update Manager settings: System x and BladeCenter

External TFTP server: You can configure FSM to use an external TFTP server to host software updates for the EN2092, EN4093/EN4093R, and CN4093 switches.

To update the I/O module firmware, perform these steps:

1. In the Update Chassis Components window, click **I/O Modules - Check and Update Firmware**, as shown in Figure 6-98.

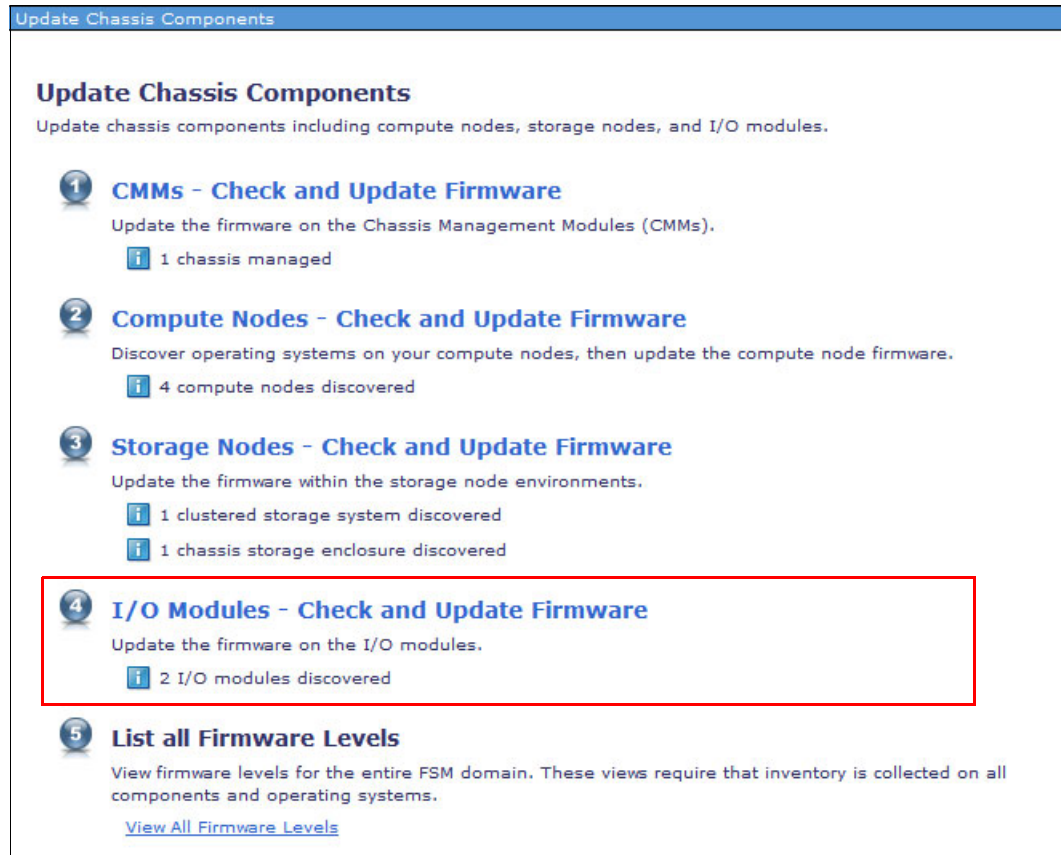


Figure 6-98 Update Chassis Components window: Check for I/O module updates

2. In the I/O Modules - Check and Update Firmware window, click **Check for Updates**, as shown in Figure 6-99.

Discovery, access, and inventory: We already discovered I/O modules, granted access to them, and collected their inventory, as described in 6.5, “Configuring chassis components” on page 144. Alternatively, you can discover, request access, and collect inventory for the I/O modules from the I/O Modules - Check and Update Firmware window.




Figure 6-99 I/O Modules - Check and Update Firmware window: Check for Updates

3. In the Acquire Updates window (see Figure 6-100), select the method of acquiring updates and click **OK** to start the Acquire Updates job.

Previously imported updates: If you already imported updates (see “Importing update files” on page 196) you still need to specify a valid path in the Acquire Updates window. If no updates will be found, an error message appears indicating that no updates were found, but it can be safely ignored.

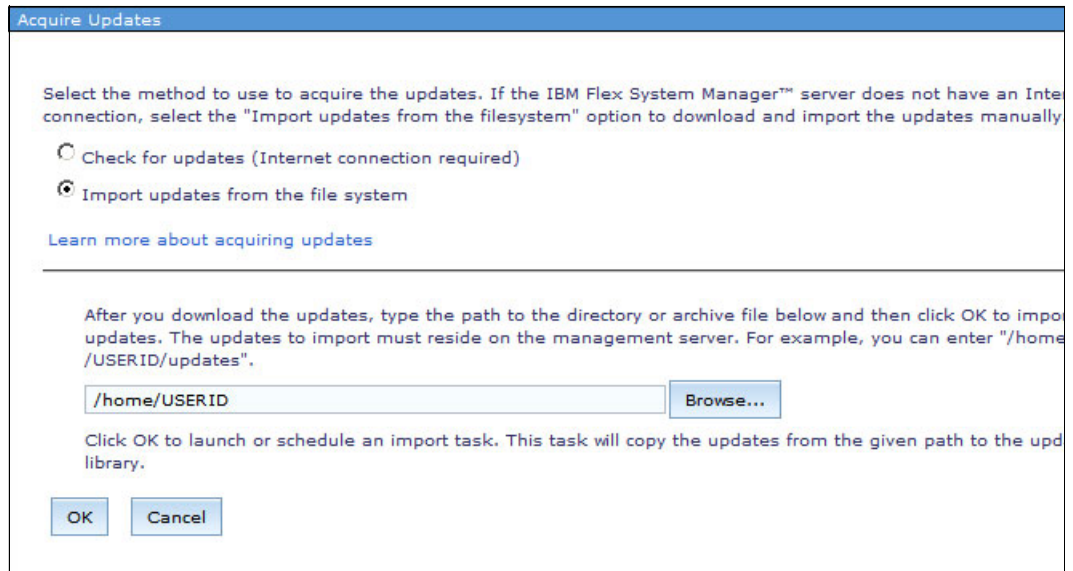


Figure 6-100 Acquire Updates window

4. After the Acquire Updates job is completed, click **Show and Install Updates**, as shown in Figure 6-95 on page 204.




Figure 6-101 Acquire Updates window: Show and Install Updates

- Review the list of available updates as shown in Figure 6-96 on page 205. The updates can be selected individually or all at one time, and deployed by clicking **Install**.

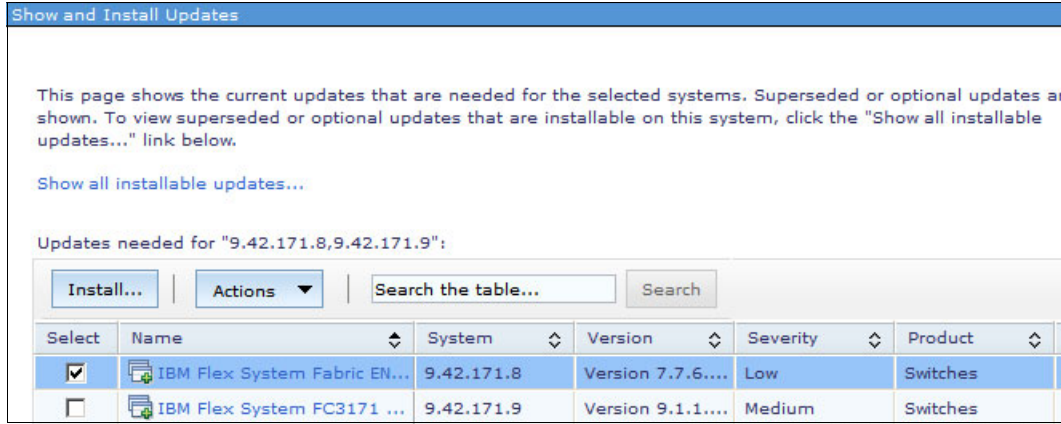


Figure 6-102 Show and Install Updates list: I/O modules

Use the Install Wizard that is shown in 6.9.2, “Updating the CMM firmware” on page 199 to run the updates immediately or schedule them.

6.9.5 Compliance policies

You can also track and determine needed updates for systems by configuring a compliance policy. To create a compliance policy, perform these steps:

1. Click the **Update Manager** link from the Plug-ins tab on the Home page, as shown in Figure 6-103.




Figure 6-103 Plug-ins window: Update Manager

2. The main Update Manager window opens. This window provides a jumping point for many options such as acquiring updates, creating compliance policies and viewing currently downloaded updates. Click **Optional: Create and configure compliance policies** to create a compliance policy, as shown in Figure 6-104.




Figure 6-104 Show Update Groups

3. Select a system or group of systems to apply the updates to. Keep in mind that for compute nodes, the OS is required to push updates. Generally, create compliance policies against specific devices, and not the All Systems group. For compute nodes, create a compliance policy against the All Operating Systems group or just against a group of Chassis. In this example, navigate to the All Operating Systems group and click **Show Compliance Policies**. This process displays any previously created compliance policies. Figure 6-105 shows that a policy does not yet exist, so add one by clicking **Add**.

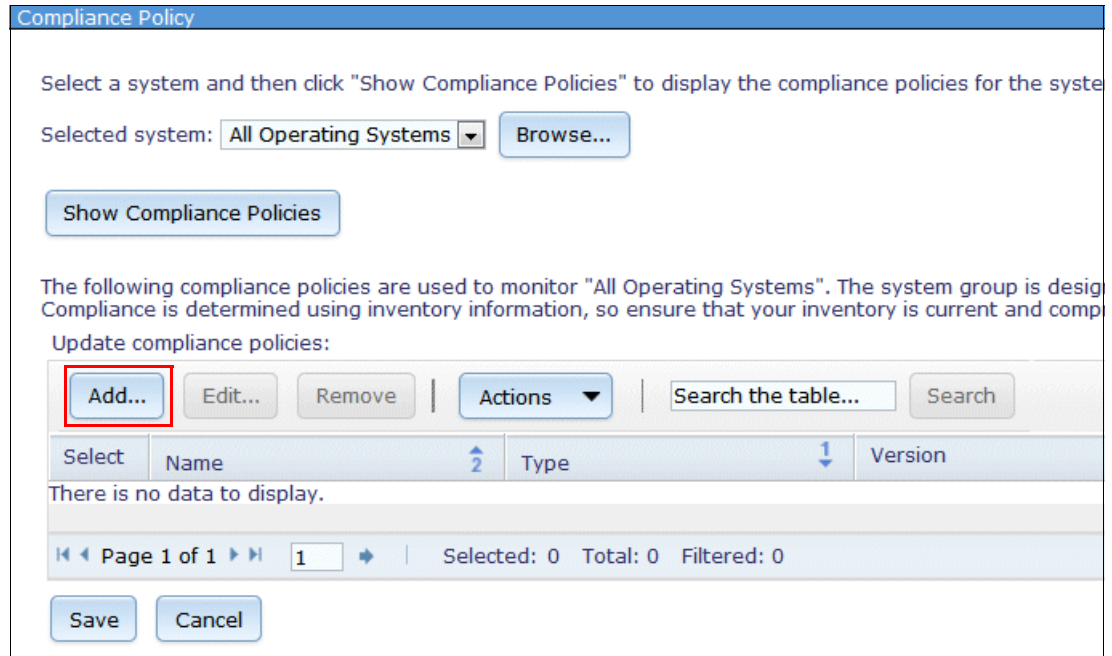


Figure 6-105 Window with no compliance policy

- In this example, select **All IBM System x and BladeCenter updates** and click **Add** as shown in Figure 6-106.




Figure 6-106 Policy Group

- The selection is moved to the right column as shown in Figure 6-107. Click **OK**.

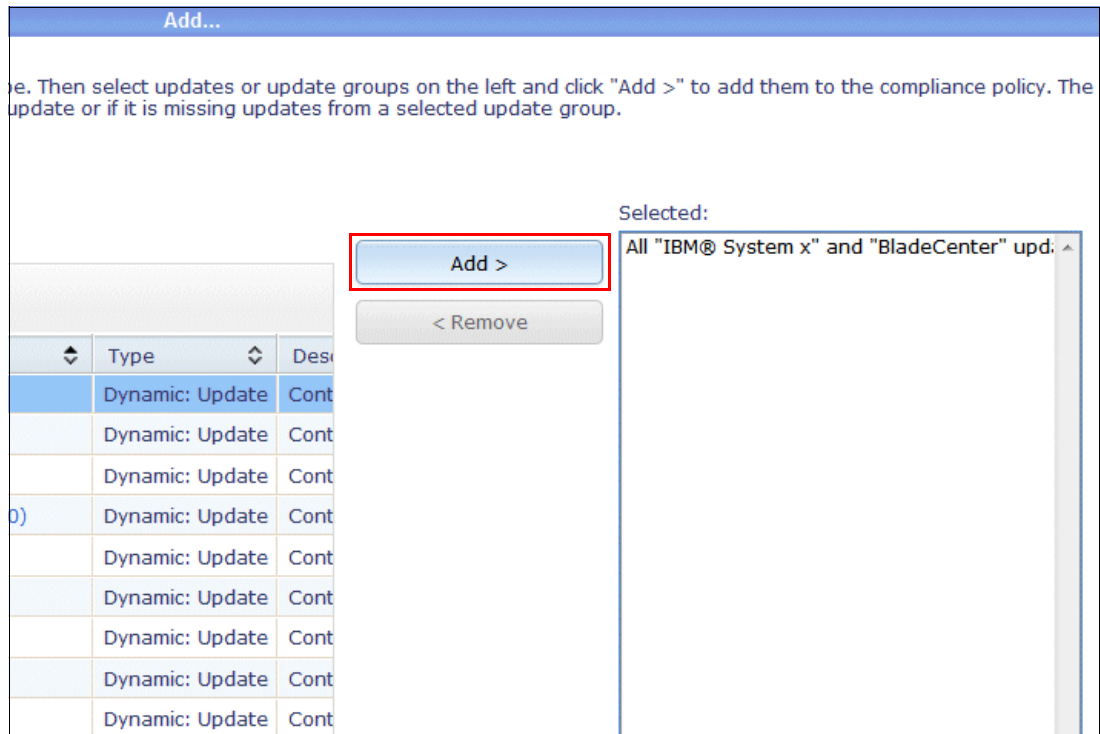


Figure 6-107 Policy Group moved to the Selected box

6. In the next window, select **Save for the Compliance** to take effect.

The main Update Manager window opens with a new compliance graph as shown in Figure 6-108.




Figure 6-108 Update Compliance window

- You can see that six systems need minor updates. Click the **6 systems** hyperlink and review which systems need updates. From here, all systems or certain systems can be selected by clicking **Actions** → **Select All**. Then, click **Actions** → **Release Management** → **Show and Install Updates**, as shown in Figure 6-109, to deploy the updates.




Figure 6-109 Show and Install Updates

6.10 Manage Feature-on-Demand keys

With Features on Demand keys, you can view and install optional features for IBM Flex System Manager management software and managed IBM Flex System resources.

You can only enable Features on Demand keys on compute nodes through the management software if the compute node is configured to boot a UEFI-compatible operating system. If the compute node is configured in the Setup Utility to boot in legacy mode, you cannot activate Features on Demand keys.

Install Features on Demand keys by using the management software web interface. If you use the IMM2 interface or any other method to import and install FoD keys, you must restart the management software to activate the software.

The Manage Features on Demand Keys window has a view list from which you can select either IBM FSM keys or All keys. You can use the FSM keys view to add and remove only the

keys that apply to the management software. Use the All keys view to add and remove keys from the management software and managed compute nodes.

For more information about how to obtain FoD keys, see 5.1.3, “Planning for Features on Demand” on page 90.

To manage the Features on Demand keys that are installed, or to import a new key, perform the following steps:

1. On the Home page, click the **Administration** tab and scroll to the bottom. Under Features on Demand tasks, click **Manage Features on Demand Keys** as shown in Figure 6-110.

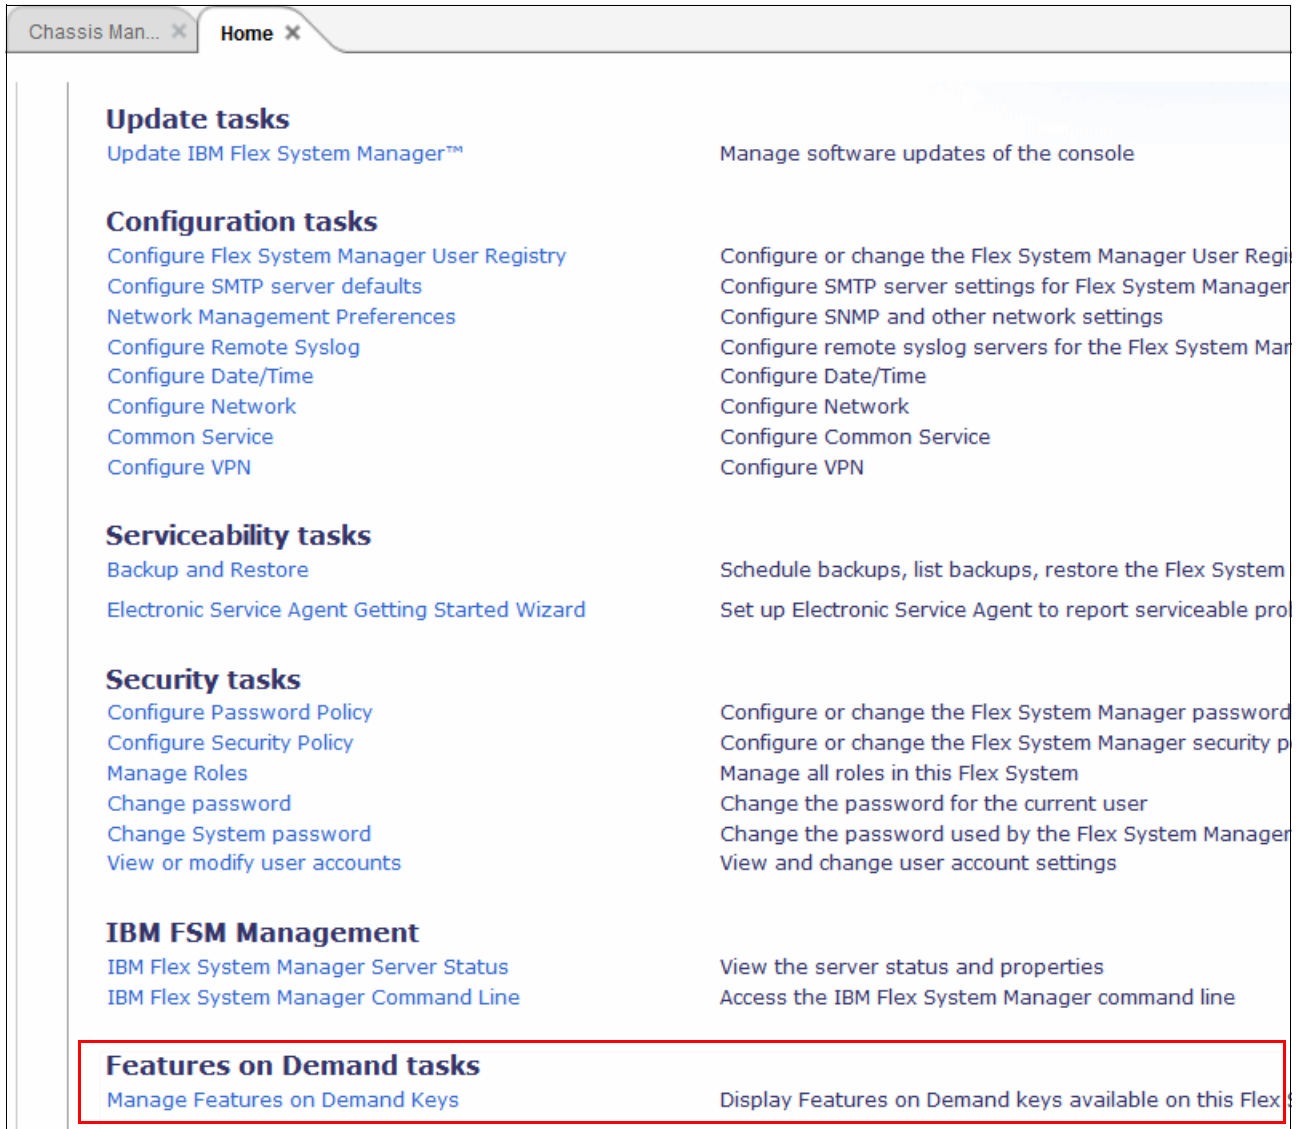


Figure 6-110 Features on Demand tasks

2. The Manage Features on Demand Keys window opens and the installed keys are displayed. Select the keys that you want to view.

To view management software keys, select **IBM FSM keys** from the **View** list as shown in Figure 6-111.

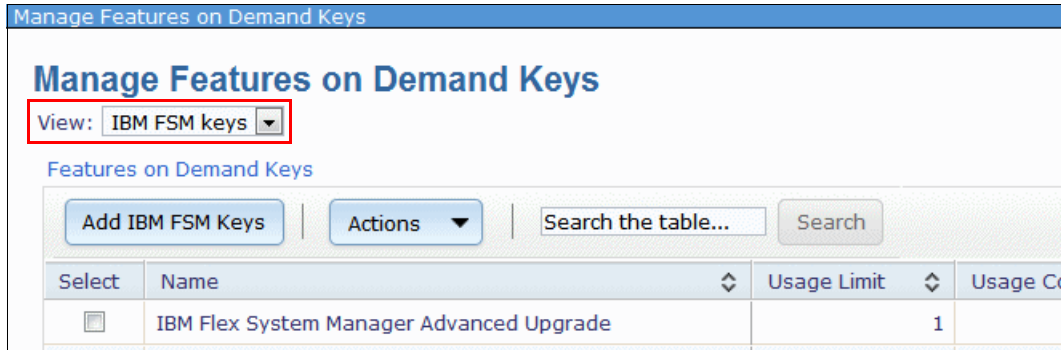


Figure 6-111 FSM keys

To view all IBM Flex System keys, including keys for managed resources, select **All keys** from the **View** list as shown in Figure 6-112.

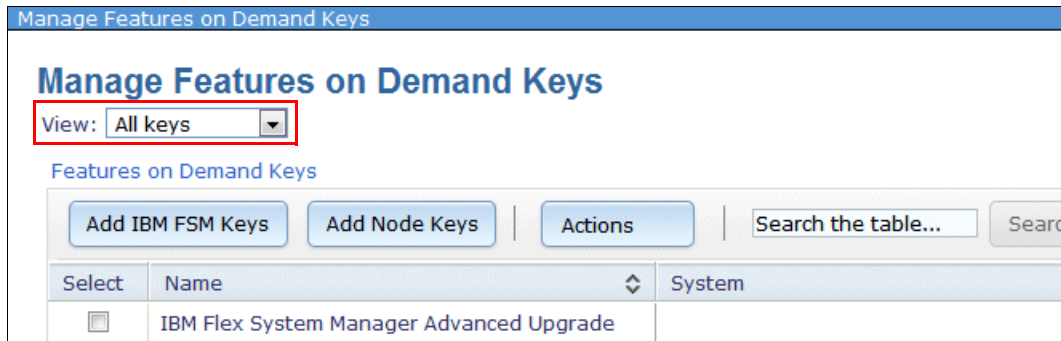


Figure 6-112 All keys

3. To import a key, click **Add IBM FSM Keys** or **Add Node Keys** (depending on the view).
If you click **Add IBM FSM Keys**, an Add IBM FSM Keys window opens that enables you to import a key.

If you click **Add Node Keys**, a window opens that enables you to import keys for compute nodes. When you click **Add Node Keys**, a list of configuration templates opens as shown in Figure 6-113. Click **Create**.

Manage Features on Demand Keys

Manage Features on Demand Keys

View: ▾

Configuration Templates

Use configuration templates to deploy settings on one or more systems.

Configuration Templates

▾

Select	Name	Deployed	Plan Count
<input type="checkbox"/>	8GbSANSwitchProtocolConfigurationTemplate	No	0
<input type="checkbox"/>	Boot Sequence Predefined Template	No	1
<input type="checkbox"/>	Ethernet1GbSwitchProtocolConfigurationTempl...	No	0
<input type="checkbox"/>	Ethernet1GbSwitchVLANConfigurationTemplate	No	0
<input type="checkbox"/>	IPv4AddressPoolConfigurationTemplate	No	0
<input type="checkbox"/>	IPv6AddressPoolConfigurationTemplate	No	0
<input type="checkbox"/>	OperatingSystemCreatei5AccountTemplate	No	0
<input type="checkbox"/>	OperatingSystemCreateLinuxAccountTemplate	No	1
<input type="checkbox"/>	OperatingSystemCreateWindowsAccountTempl...	No	1

Figure 6-113 Configuration Templates

- A configuration template must be created that includes FoD keys that need to be deployed. Click **Create**.
- Figure 6-114 is displayed. Under Template Type, select **Server (via CIM protocol)** to deploy FoD keys to an IMM2 of a compute node. For template, select **Feature Activation Manager Configuration** and provide a name. Then click **Continue**.

Create

Template type:

Server (via CIM protocol) ▼

Configuration to create a template:

Feature Activation Manager Configuration

Description:
Feature Activation Manager Configuration

*Configuration template name:

FOD Key

Configuration template description:

Automatically deploy this configuration template when notified of a matching resource
(This option is enabled only if automatic deploy is supported by the selected configuration.)

Continue Cancel

Figure 6-114 Configuration Template creation

- In the wizard, there are two choices as shown in Figure 6-115.

Feature Activation Configuration

✓ Welcome
➔ Key Redemption Method
Summary

Key Redemption Method

The key files to be used for key activation can be redeemed from Key Mar

Obtain activation keys from Key Management System (KMS)

Upload activation keys from a local system

Figure 6-115 Key Redemption Method window

If you select the Key Management System (KMS), the options that are shown in Figure 6-116 are available. To pull in keys from the KMS, a connection to the Internet is required. Additionally, an IBM ID must be configured appropriately and the appropriate authorization codes must be available as provided when the FoD keys were purchased. If a proxy is required for connection, that can be configured when you click **Next**.

The screenshot shows a web interface titled "Feature Activation Configuration". On the left is a navigation pane with the following items: "Welcome" (checked), "Key Redemption Method" (checked), "KMS Login" (selected with a right-pointing arrow), "Connection", "Automatic Reboot", and "Summary". The main content area is titled "KMS Login" and contains the following text: "The user ID, password, and authorization codes will be used to log in more authorization codes to the list." Below this text are four input fields: "*User ID:", "*Password:", "*Confirm password:", and "Feature authorization code:". Each of the first three fields has an asterisk indicating it is a required field. Below the "Feature authorization code:" field is a button labeled "Add to List". Below that is a list box labeled "*Feature authorization code list:" with a "Remove" button to its right. At the bottom left of the main content area, there is a legend: "* Required field".

Figure 6-116 KMS Login

If you select to upload keys from the local system, Figure 6-117 is displayed. Select the FoD keys that need to be imported and add them to the list, then click **Next**.




Figure 6-117 Upload Key Files window

7. The Automatic Reboot window opens as shown in Figure 6-118. You can choose to reboot the device if needed for the FoD key to take effect. Click **Next** and then **Finish**.



Figure 6-118 Automatic Reboot window

- The newly created template is displayed in the list as shown in Figure 6-119. Select the key, click **Deploy**, and point it to a system that you want the key deployed to.

Configuration Templates

Use configuration templates to deploy settings on one or more systems.

Configuration Templates

Deploy Create Create Like Edit Delete Actions Search the t

Select	Name	Deployed	Plan Count	T
<input type="checkbox"/>	8GbSANSwitchProtocolConfigurationTemplate	No	0	N
<input type="checkbox"/>	Boot Sequence Predefined Template	No	1	C
<input type="checkbox"/>	Ethernet1GbSwitchProtocolConfigurationTempl...	No	0	N
<input type="checkbox"/>	Ethernet1GbSwitchVLANConfigurationTemplate	No	0	N
<input checked="" type="checkbox"/>	FOD Key	No	0	S
<input type="checkbox"/>	IPv4AddressPoolConfigurationTemplate	No	0	S
<input type="checkbox"/>	IPv6AddressPoolConfigurationTemplate	No	0	S
<input type="checkbox"/>	OperatingSystemCreatei5AccountTemplate	No	0	O
<input type="checkbox"/>	OperatingSystemCreateLinuxAccountTemplate	No	1	O
<input type="checkbox"/>	OperatingSystemCreateWindowsAccountTempl...	No	1	O
<input type="checkbox"/>	OperatingSystemImmediatePowerOff	No	0	O
<input type="checkbox"/>	OperatingSystemImmediateRestart	No	0	O
<input type="checkbox"/>	OperatingSystemIPv4NetworkTemplate	No	1	O
<input type="checkbox"/>	OperatingSystemIPv6NetworkTemplate	No	1	O
<input type="checkbox"/>	ServerEnableSerialOverLAN	No	0	S

Page 1 of 2 1 Selected: 1 Total: 25 Filtered: 25

Figure 6-119 FoD key deployment

6.11 Flex System V7000 Storage Node initial configuration

This section discusses initial configuration tasks for the IBM Flex System V7000 Storage Node. The following topics are covered:

- ▶ 6.11.1, “Creating a new system on the V7000 Storage Node” on page 224
- ▶ 6.11.2, “Flex System V7000 Storage Node Setup wizard” on page 228

6.11.1 Creating a new system on the V7000 Storage Node

The following procedure guides you through the necessary steps to create a new system on the V7000 Storage Node when using the FSM web user interface:

1. Open a web browser and point it to the IP address of the FSM and log in. The menu panel shown in Figure 6-120 displays, giving you a number of selections.

Select **Launch IBM FSM Explorer** from the menu list.

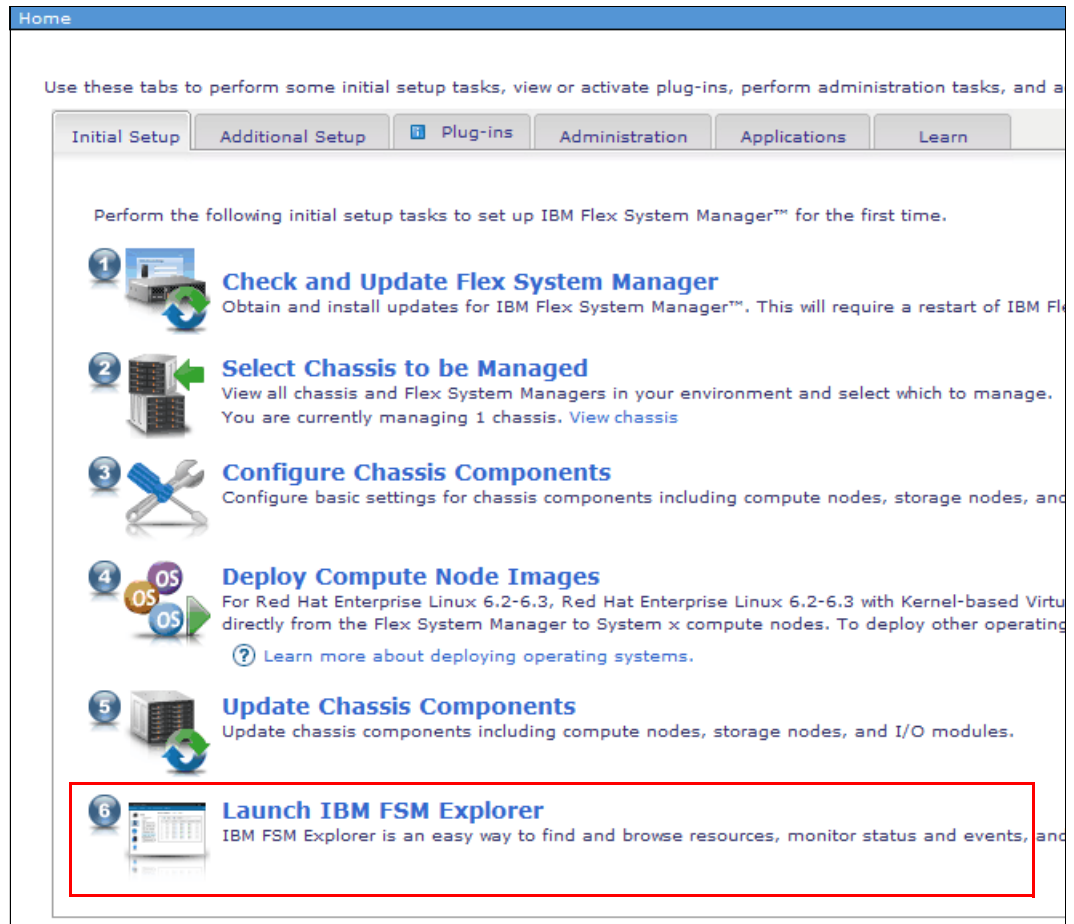


Figure 6-120 Launch IBM FSM Explorer

Notice that a new browser tab is opened, which allows you to select the applicable enclosure from the Chassis Map as shown in Figure 6-121.




Figure 6-121 Select and launch the chassis in the Chassis Manager

2. In the Chassis Manager, select the applicable chassis that will launch the chassis map for that chassis, as shown in Figure 6-122.

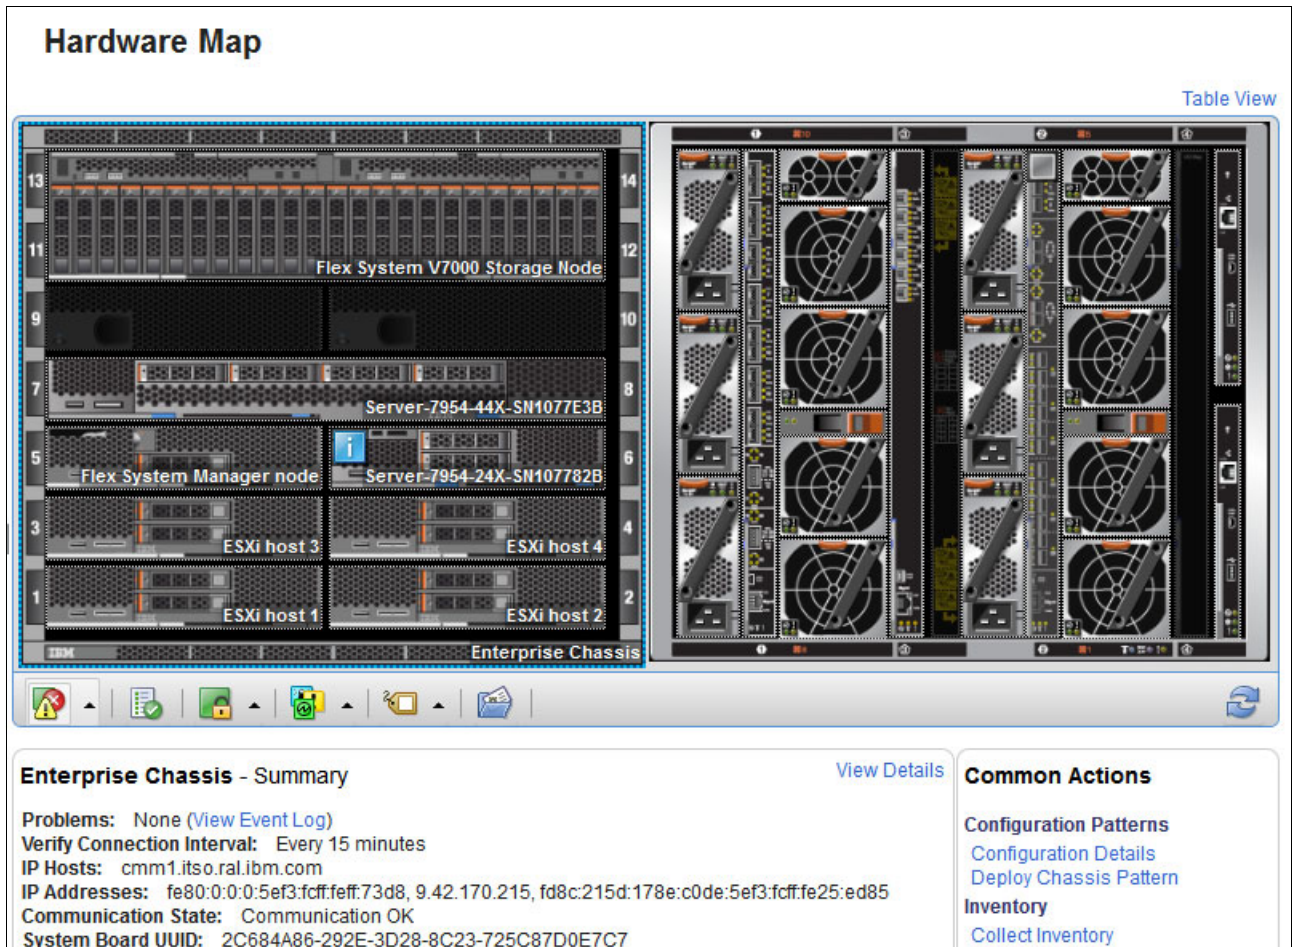


Figure 6-122 IBM Flex System Manager: Hardware Map

3. Right-click the V7000 Storage Node in the chassis map, then select **Remote Access** and click **Launch IBM Flex System V7000** as shown in Figure 6-123 to start the Initial Setup wizard.

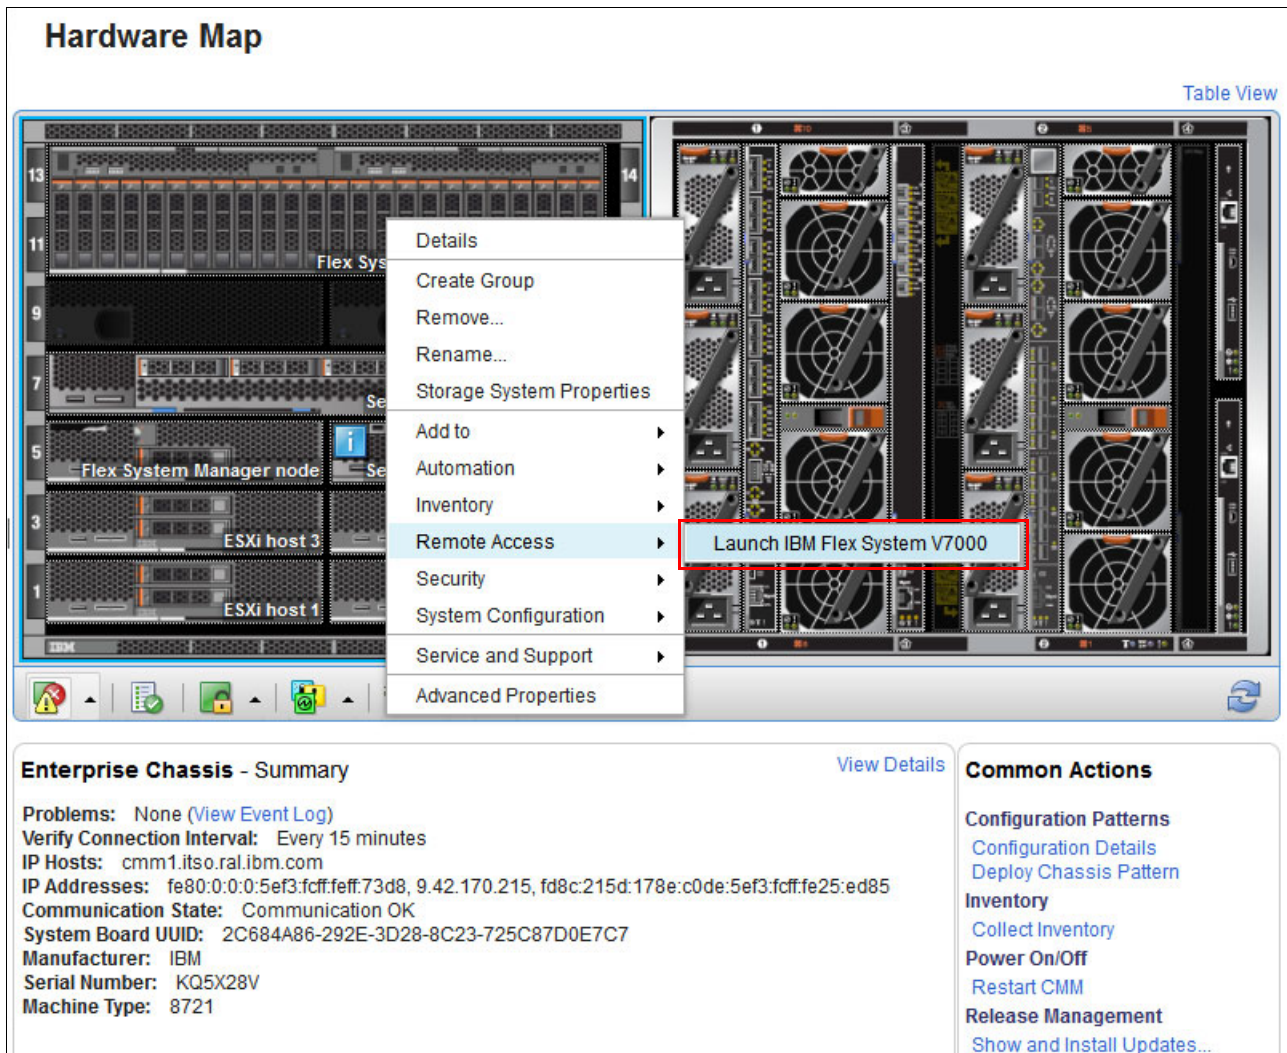


Figure 6-123 Launch Storage Manager (V7000)

4. The next window is a welcome window from the IBM Flex System V7000 Storage Node interface, asking to either create a new system (cluster) or add to an existing system, as shown in Figure 6-124 on page 227.

In this example, we are creating a new system. Select **Create a new system** and then click **Next**.



Figure 6-124 IBM Flex System V7000 Storage Node first-time setup welcome window

5. In the window shown in Figure 6-125, select whether you are using an IPv4 or IPv6 management IP address and type the IP address (you can use either DHCP or the static address that was assigned). The subnet mask and gateway will already list defaults, which you can edit.

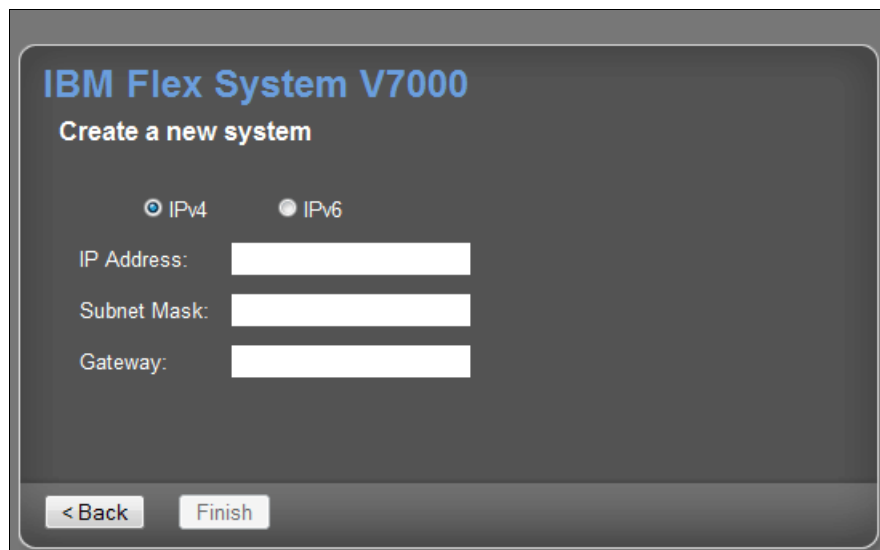


Figure 6-125 Create a new storage cluster

6. Click **Finish** to set the management IP address for the system. System initialization begins and might take several minutes to complete.

When system initialization is complete, the V7000 login window opens. Log in using the default credentials (superuser and passw0rd). Then, you are prompted to change the default password. System Setup is launched automatically. The setup wizard takes you through the steps to configure basic system settings, such as time and date, system name, and hardware detection and verification.

6.11.2 Flex System V7000 Storage Node Setup wizard

After the initial configuration described in 6.11, “Flex System V7000 Storage Node initial configuration” on page 223 is complete, the IBM Flex System V7000 Storage Node Welcome window opens (Figure 6-126).



Figure 6-126 IBM Flex System V7000 Storage Node Welcome window

Tip: During the initial setup of the Flex System V7000, the installation wizard asks for various information that you need to have available during the installation process. If you do not have this information or choose not to configure some of the items at this time, you can configure them later through the GUI.

Click **Next**, and perform the following steps:

1. Read and accept the license agreement, as shown in Figure 6-127. Click **Next** after accepting the license agreement.

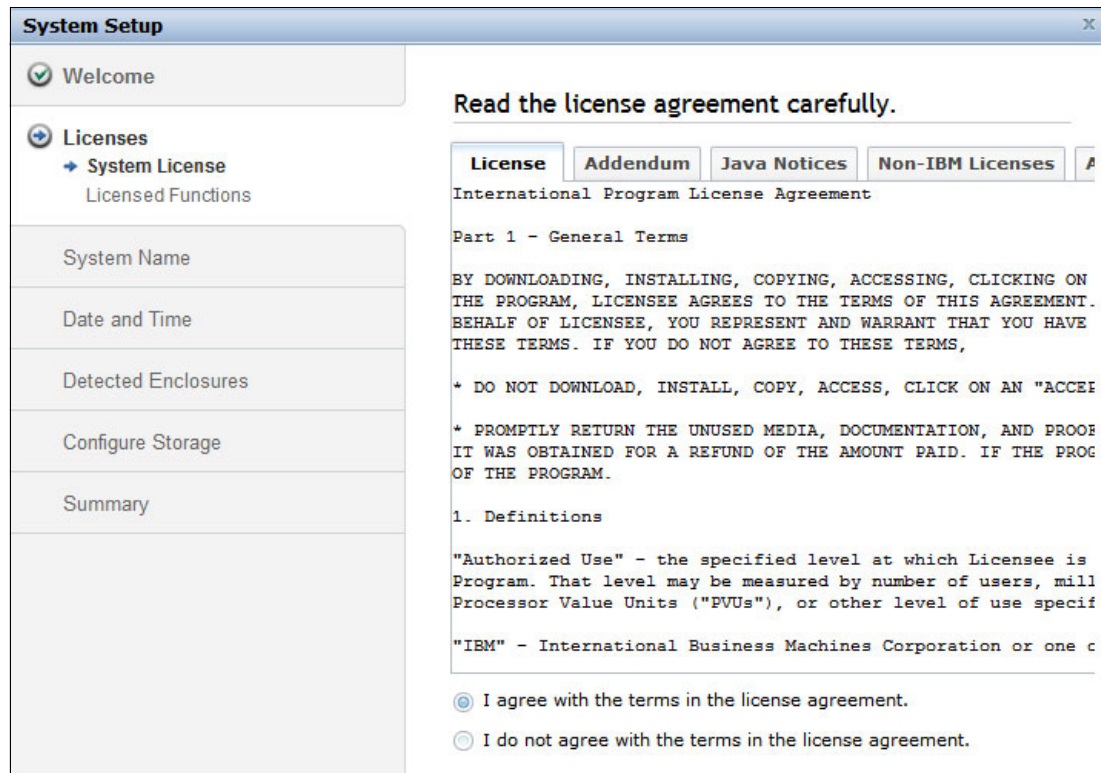


Figure 6-127 Setup wizard: License agreement

2. Optional: Specify licensed functions, as shown in Figure 6-128. The System Licenses include the External Virtualization limit, Remote Copy limit, and IBM Real-time Compression™ limit. Click **Apply** and select **Next**.

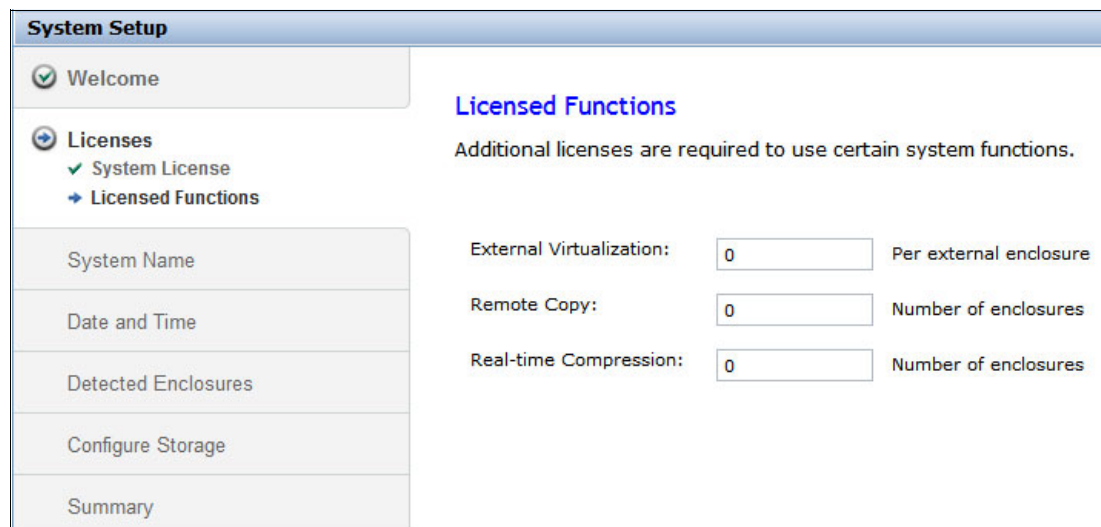


Figure 6-128 Setup wizard: Licensed Functions

3. Specify a system name, as shown in Figure 6-128 on page 229. Click **Apply** and select **Next**.

The screenshot shows the 'System Setup' wizard interface. On the left, a vertical list of steps includes 'Welcome', 'Licenses', 'System Name', 'Date and Time', 'Detected Enclosures', 'Configure Storage', and 'Summary'. The 'System Name' step is currently selected and highlighted. The main content area on the right is titled 'System Name' and contains the instruction 'Enter a name for the system.' Below this, there is a label 'System Name:' followed by a text input field containing the value 'ITSO_V7000_Cluster_9.42.171.20'.

Figure 6-129 Setup wizard: Set system name

4. Set up the system date and time as shown in Figure 6-130. Click **Apply** and select **Next**.

The screenshot shows the 'System Setup' wizard interface. On the left, a vertical list of steps includes 'Welcome', 'Licenses', 'System Name', 'Date and Time', 'Detected Enclosures', 'Configure Storage', and 'Summary'. The 'Date and Time' step is currently selected and highlighted. The main content area on the right is titled 'Date and Time' and contains two sections. The first section, 'Current Date and Time', has a text input field displaying 'Nov 19, 2013 6:39:43 PM'. The second section, 'Time Zone', has a dropdown menu showing '(GMT-5:00) US Eastern Time' and a 'Use Browser Settings' button below it.

Figure 6-130 Setup wizard: Set date and time

- Verify that all hardware is detected by the system correctly as shown in Figure 6-131. Click **Apply** and select **Next**.

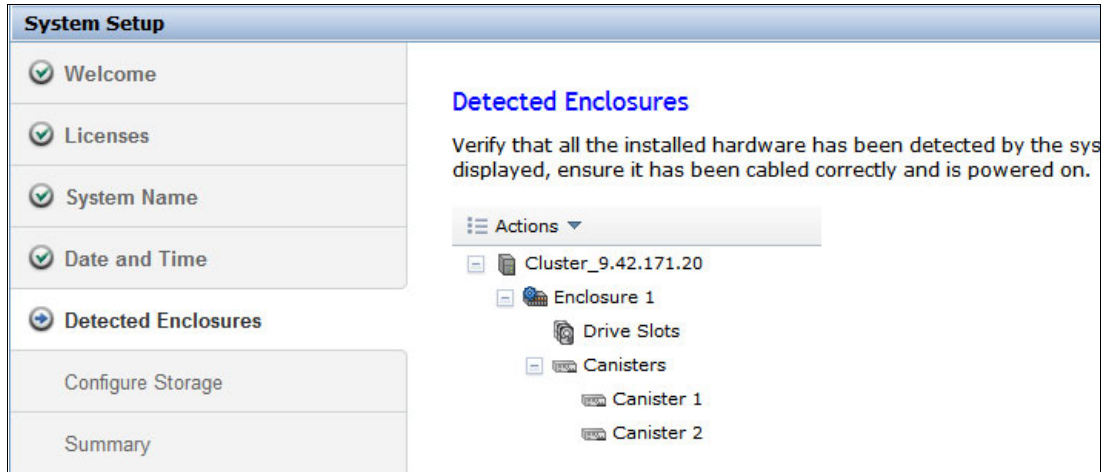


Figure 6-131 Verify hardware

- Do not select **Yes** to automatically configure internal storage if you are creating a customized storage layout. Click **Next**.
- Verify the configuration settings in the Summary window, as shown in Figure 6-132.

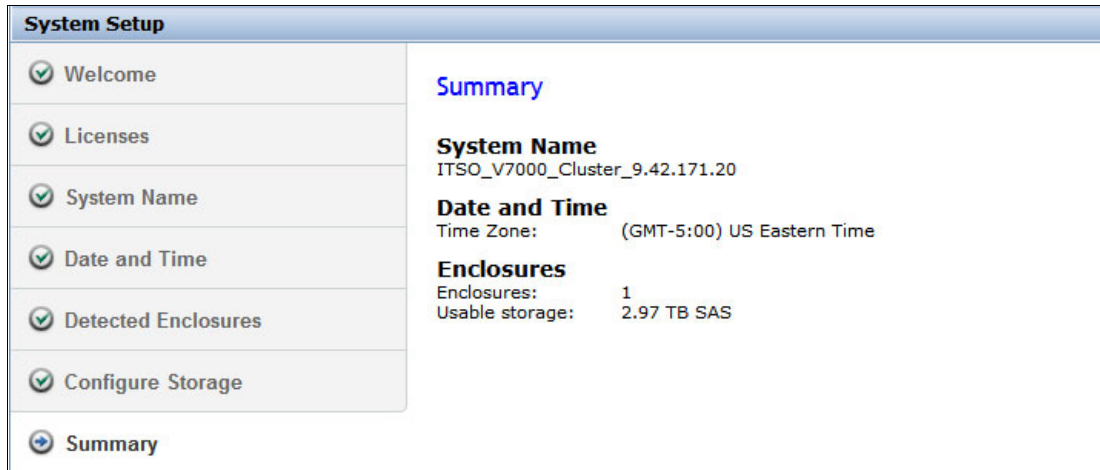


Figure 6-132 V7000 Setup wizard: Summary

- Click **Finish** to complete the Setup wizard task and log in to IBM Flex System V7000 Storage Node. You log in as a superuser with your newly defined password.

- After you successfully log in, the IBM Flex System V7000 Storage Node System Details window opens. IBM Flex System V7000 Storage Node initial configuration is complete and the cluster is up and running, as shown in Figure 6-133.

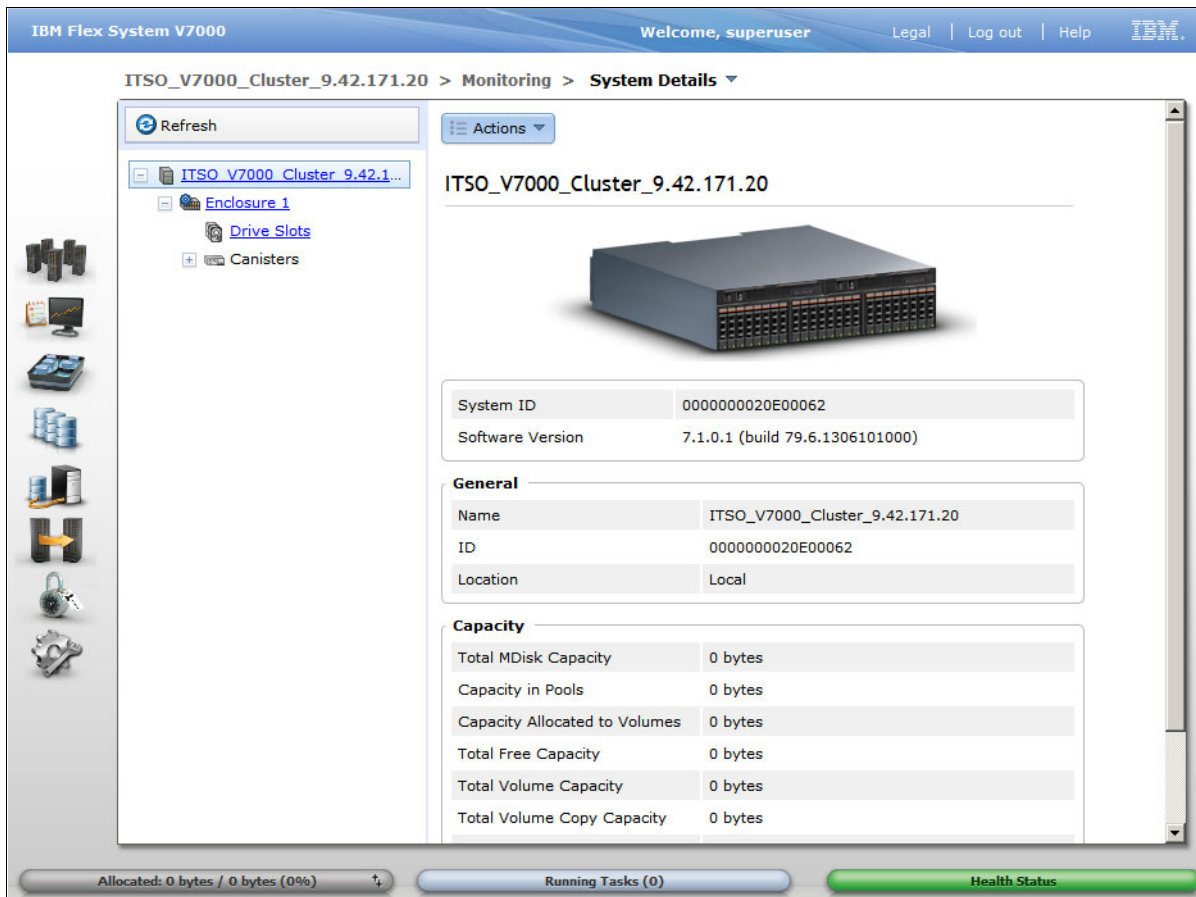


Figure 6-133 IBM Flex System V7000 Storage Node home overview window

After you create a new system on the V7000 Storage Node through the FSM, the FSM attempts to automatically discover it, request access to it, and collect its inventory by running the “Storage Auto Discovery with IP <storage system IP address>” job. If the job ends with the error “Could not connect to any specified IP address”, after the V7000 system comes online, you can rerun it from the Active and Scheduled Jobs window by selecting the job and clicking **Actions** → **Run Now**, as shown in Figure 6-134.




Figure 6-134 Active and Scheduled Jobs: Storage Auto Discovery

The V7000 cluster system appears in the list of storage systems, as shown in Figure 6-135.




Figure 6-135 Discovered V7000 cluster system

6.12 Discover and manage external Storwize V7000

As part of the comprehensive consolidated management approach, storage management is integrated into the FSM. After your storage appliance is discovered and access is requested by the FSM, you can start managing it as shown in Figure 6-136.




Figure 6-136 Flex System Manager Storage Management

6.12.1 Discover an IBM Storwize V7000

Discovery of the IBM Storwize V7000 is performed through the command-line interface (CLI).

To use the command **manageV7000**, the user name **superuser** must be active on the IBM Storwize V7000 and have a password. If the user name **superuser** is not active, you must transfer the key manually and use the **mkdatasource** command, as described in this information center:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.storage.helps.doc%2Fconfiguring_v7000_or_svc_storage.html

Requirement: Automatic discovery of the Storwize V7000 and the CLI commands requires V7000 software version 6.3.0.0 or later.

Check the firmware level on the V7000 as shown in Figure 6-137.




Figure 6-137 Firmware level on V7000

The new IBM Storwize V7000 storage subsystem comes with the pre-configured IPv4 address 192.168.70.151. After the installation of the Storwize V7000 hardware, the user ID superuser, with the password passw0rd, is created. The management software then assigns an IPv6 address to the V7000 storage subsystem. If the IPv6 address is accessible to the management software, the default IPv4 address is disabled and the management software manages the chassis through the new IPv6 address. If the IPv6 address is inaccessible, the default IPv4 address is used to manage the V7000.

Consideration: All new V7000 storage subsystems come with the same default IPv4 address. Therefore, you must install them one at a time until they are all managed.

If you need to change the IP address of the V7000 storage subsystem, log in to the V7000 web interface and change the IPv4 or IPv6 address as shown in Figure 6-138.

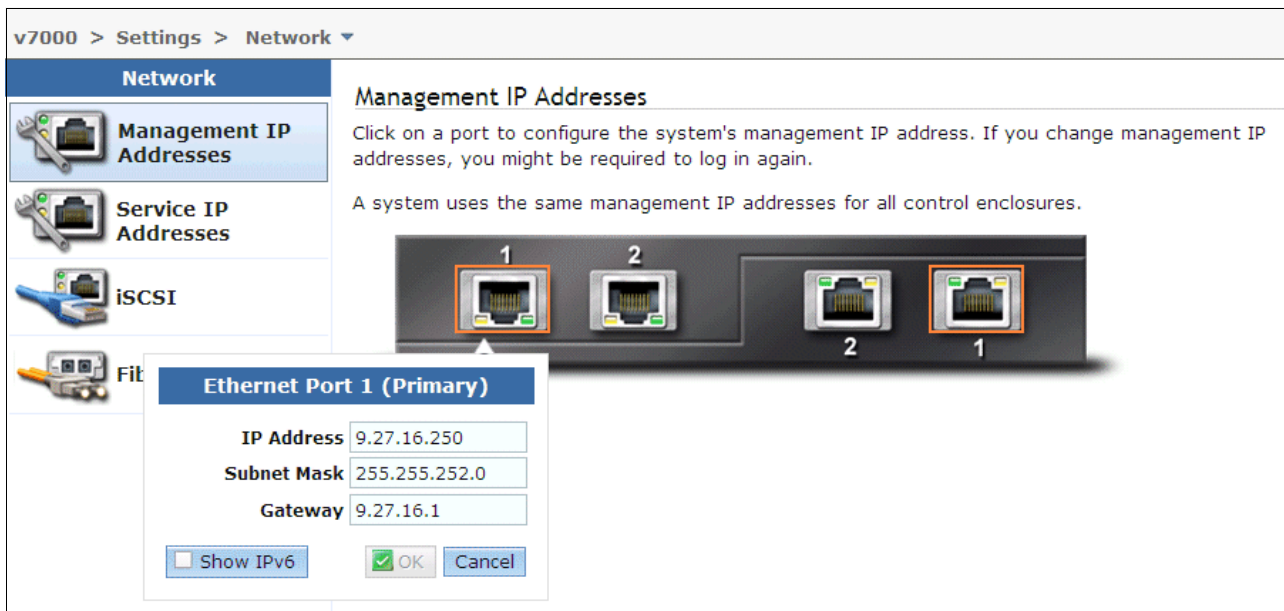


Figure 6-138 V7000 management IP address

For more information about using the storage subsystem web interface, see this website:

<http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp>

Adding a V7000

If your V7000 is not managed by any Flex System Management Node or other IBM System Management Software, enter this command from the management software CLI:

```
smcli manageV7000 -i V7000_IP_address -u admin_user_ID -p admin_password
```

In this command, *V7000_IP_address* is the IP address of the V7000, and *admin_user_ID* and *admin_password* are the administrator's user ID and password.

Managing an existing V7000 from several FSMs

Important: If you have more than one management node (FSM) on your network, you can use the **manageV7000** command only on the management node (FSM) that manages the V7000. If you run the **manageV7000** command on an additional management node, the V7000 key generated by the initial management node is replaced by a key generated by the subsequent management node.

If you need to manage a V7000 from another Flex System Manager, complete the following steps for each additional management node:

1. Use the management software CLI on the initial management node to copy the file `/home/USERID/.ssh/id_rsaV7000` from the primary management node to the `/home/USERID/.ssh` directory on the other FSM as shown in Figure 6-139.

```
USERID@FSM-5CF3FC5F54EF:~> scp /home/USERID/.  
./          ../          .bash_history  .fonts/      .mozilla/    .ssh/  
USERID@FSM-5CF3FC5F54EF:~> scp /home/USERID/.ssh/  
authorized_keys2  id_rsaV7000      id_rsaV7000.pub  known_hosts  
USERID@FSM-5CF3FC5F54EF:~> scp /home/USERID/.ssh/id_rsaV7000 IPaddress_of_the_other_FSM:/home/USERID/.ssh
```

Figure 6-139 Copy USERID key to another FSM node

2. From the CLI of the additional management node (the additional FSM), run the following command to manage the V7000:

```
smcli mkdatasource -c svc -f /home/USERID/.ssh/id_rsaV7000 -v v7000 -i  
<V7000_ip_address>
```

The `<V7000_ip_address>` is the IP address of the V7000.

Removing an existing V7000

To remove an existing V7000 from the list of managed devices, enter the command:

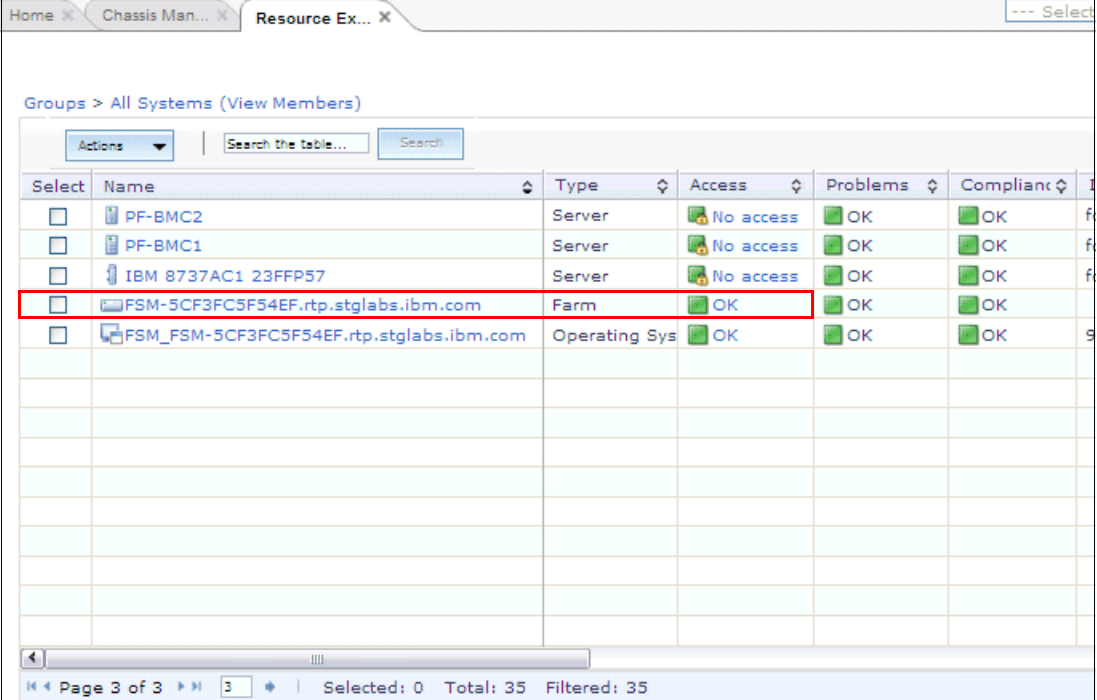
```
smcli unmanageV7000 -i V7000_IP_address
```

The *V7000_IP_address* is the IP address of the V7000.

6.12.2 Collect inventory on the discovered V7000

Discovery and inventory collection must be run before you can display storage systems in Storage Control. To do so, perform these steps:

1. After you run the **manageV7000** command, the V7000 is discovered by your Flex System management node. A new Farm object type with OK access appears in the Resource Explorer as shown in Figure 6-140.



The screenshot shows the Resource Explorer interface with a table of discovered systems. The table has the following columns: Select, Name, Type, Access, Problems, and Compliance. The data rows are as follows:

Select	Name	Type	Access	Problems	Compliance
<input type="checkbox"/>	PF-BMC2	Server	No access	OK	OK
<input type="checkbox"/>	PF-BMC1	Server	No access	OK	OK
<input type="checkbox"/>	IBM 8737AC1 23FFP57	Server	No access	OK	OK
<input type="checkbox"/>	FSM-5CF3FC5F54EF.rtp.stglabs.ibm.com	Farm	OK	OK	OK
<input type="checkbox"/>	FSM_FSM-5CF3FC5F54EF.rtp.stglabs.ibm.com	Operating Sys	OK	OK	OK

Figure 6-140 Storage Farm object discovered

2. Collect inventory on V7000 storage by clicking **Inventory** → **View and Collect Inventory** as shown in Figure 6-141.

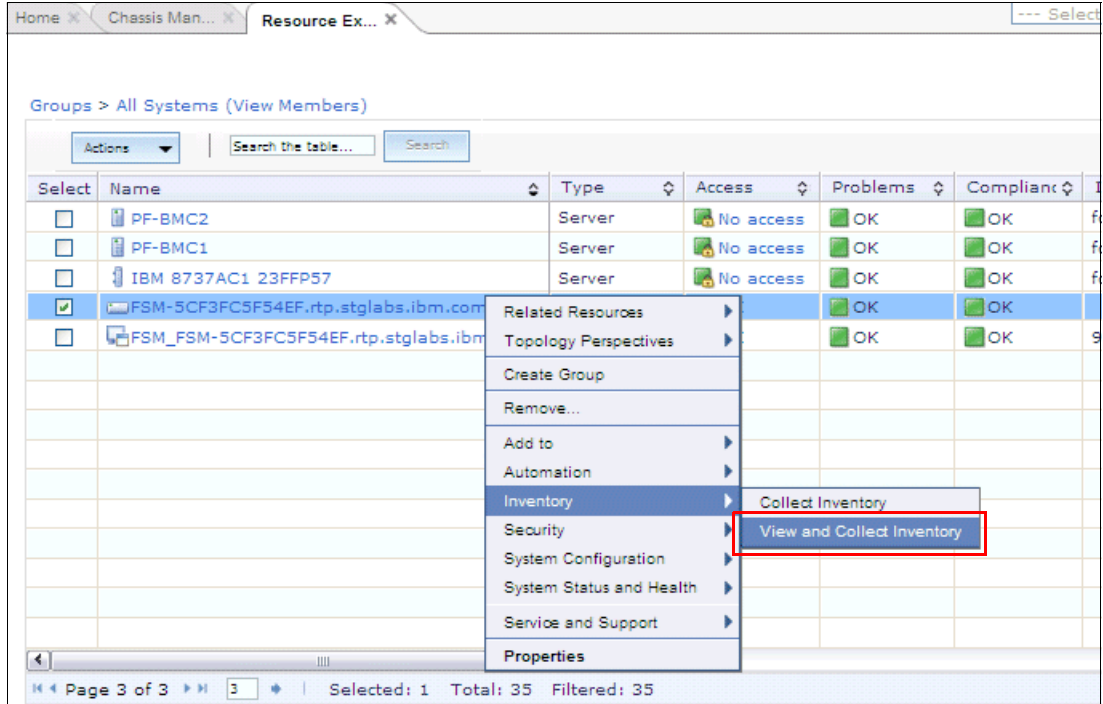


Figure 6-141 View and Collect Inventory on Storage Farm object

3. Click **Collect Inventory** to begin the inventory collection as shown in Figure 6-142.

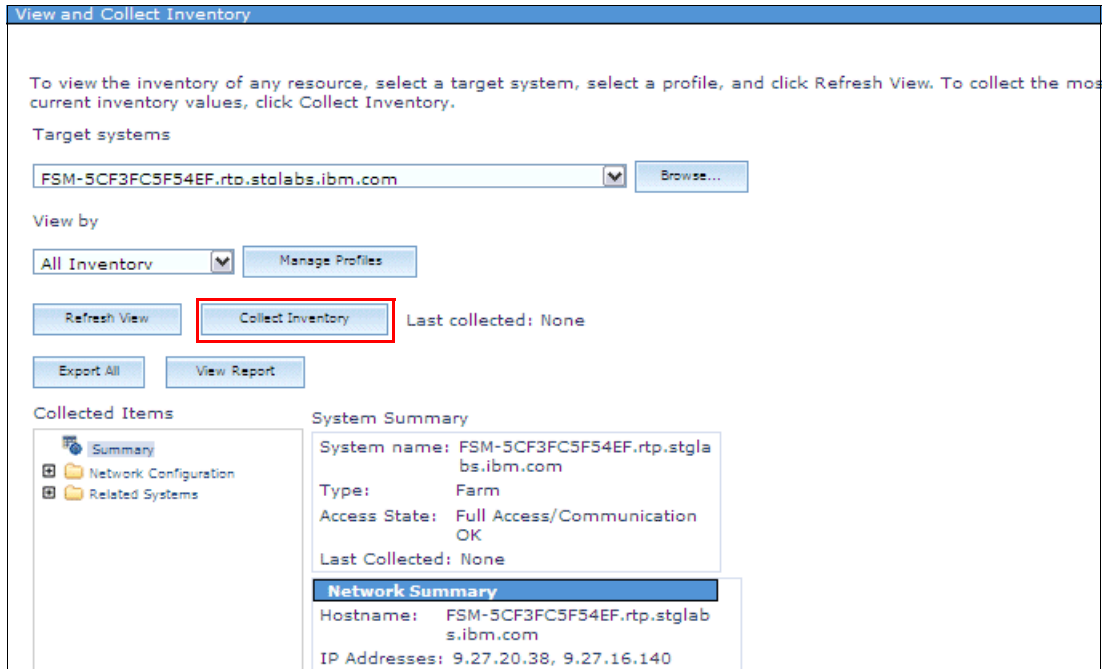


Figure 6-142 Collect Inventory on V7000 Farm object discovered

4. Click **OK** to run the collection task as shown in Figure 6-143.

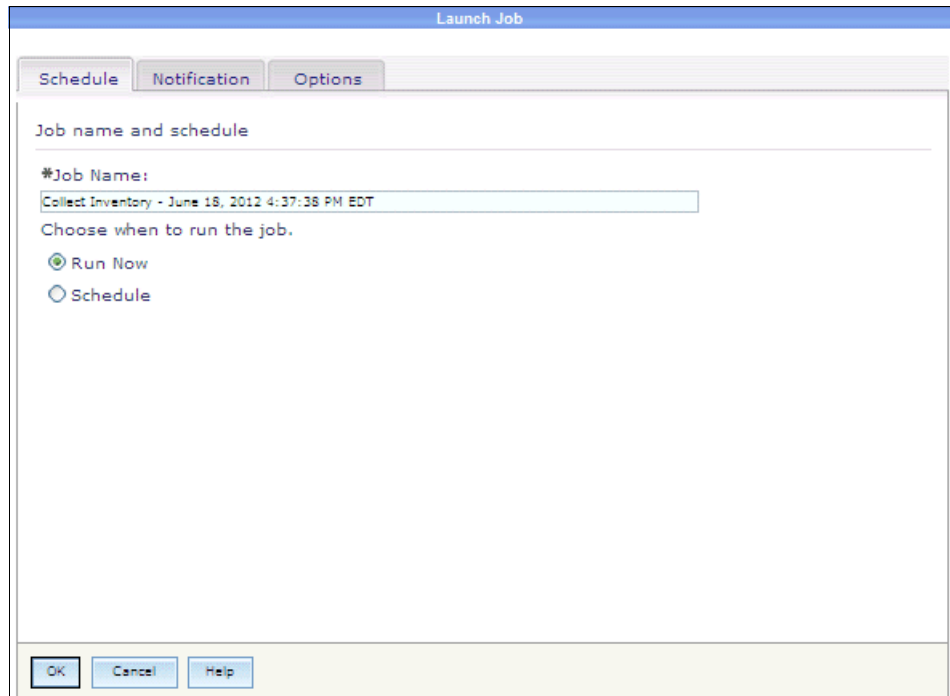


Figure 6-143 Collection task on Storage Farm object

A blue box message is displayed that indicates that the job has been started and created successfully, as shown in Figure 6-144.




Figure 6-144 Storage farm inventory collection started successfully

Wait until it is complete as shown in Figure 6-145.




Figure 6-145 Collect Inventory on a farm object completed

5. If some errors are displayed, click the **Logs** tab (Figure 6-145) to get more information, as shown in Figure 6-146.

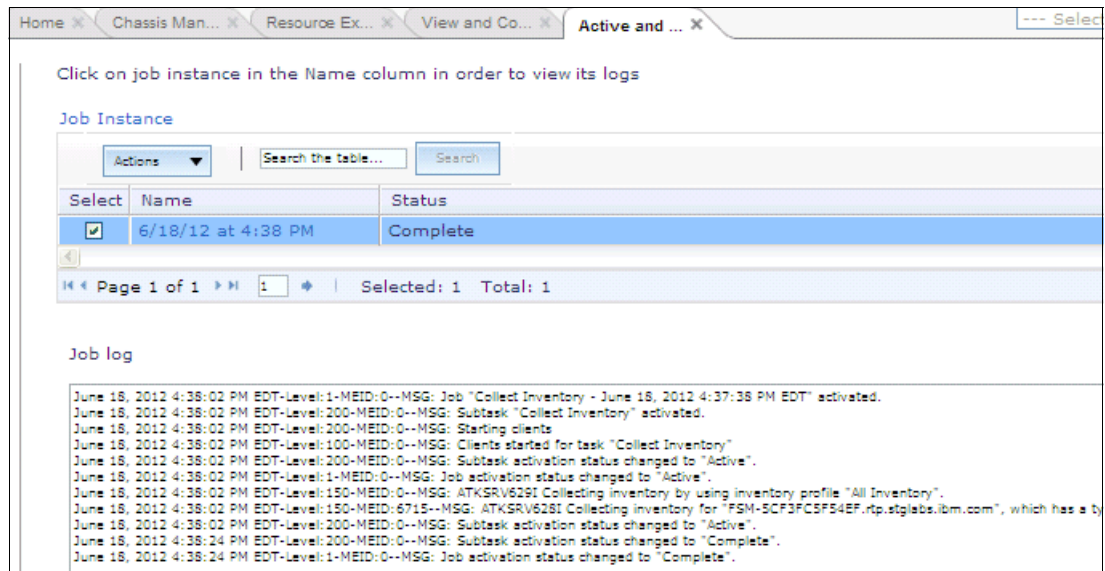


Figure 6-146 Storage farm inventory collect log

6.13 Overview of Flex System V7000 and Storwize V7000 systems management (Storage Control)

Your storage devices can be viewed and managed in one central location. The Storage Management Summary window provides an introduction to your storage systems. The Storage Management Summary is started from the Plug-ins tab of the Home page as shown in Figure 6-147.

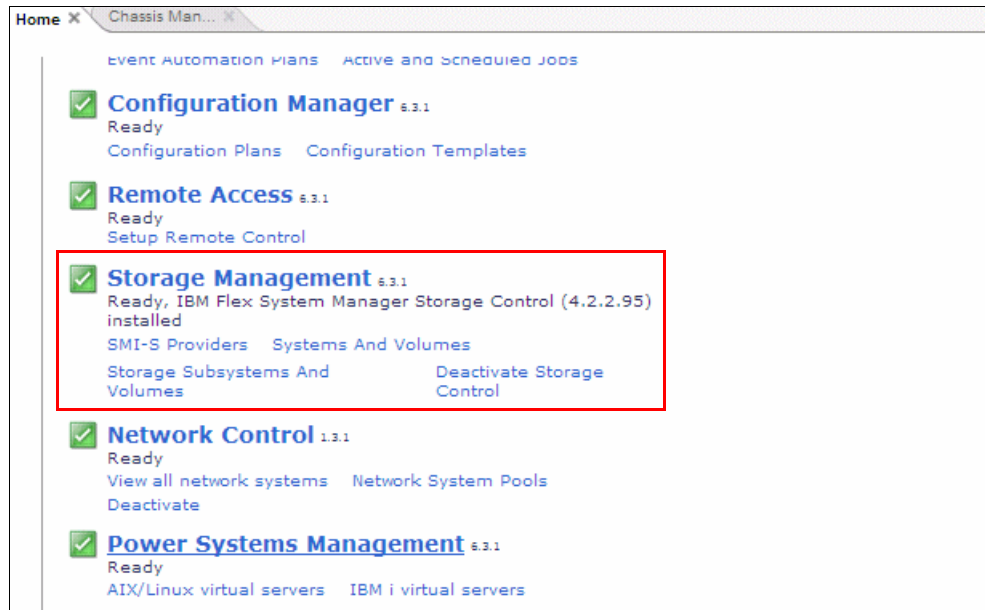


Figure 6-147 Go to Storage Control from Home page

This action opens the Storage Management window that is shown in Figure 6-148.




Figure 6-148 Storage Control main window

The Storage Management summary window is divided into these areas:

- ▶ Capacity Summary
 - A pie chart represents your disk capacity in each of these categories.
- ▶ Total configured capacity (to volumes)
 - Number of GB of the volumes that are assigned.
- ▶ Total available capacity (for volumes)
 - Number of GB of disk pool size unassigned to volumes, but available for creating future new volumes.

Total available capacity represents the remaining total storage array or storage pool space that can be used to create volumes. When you are creating a storage array or pool, the configured capacity is zero and the available capacity is the pool size. Creating more volumes decreases the amount of available capacity. Available capacity is a measurement of the current quantity of usable storage.

The *capacity measurement* is a snapshot that is created when inventory is collected on the storage arrays. Inventory collection can be configured to run on a schedule to periodically update the capacity information. This configuration is only possible if storage volumes and pools are created after the initial inventory collection.

For local storage, capacity information that is collected for attached storage devices is limited to the Total raw capacity. The Total configured capacity and Total available capacity are not included in the Capacity Summary for these storage devices.

The tasks that can be performed on an external V7000 storage system are summarized in the following sections.

Server to Storage Mapping View

The Server to Storage Mapping View displays a table that shows what storage resources are associated with your virtual servers. Select the servers to view, as shown in Figure 6-149.

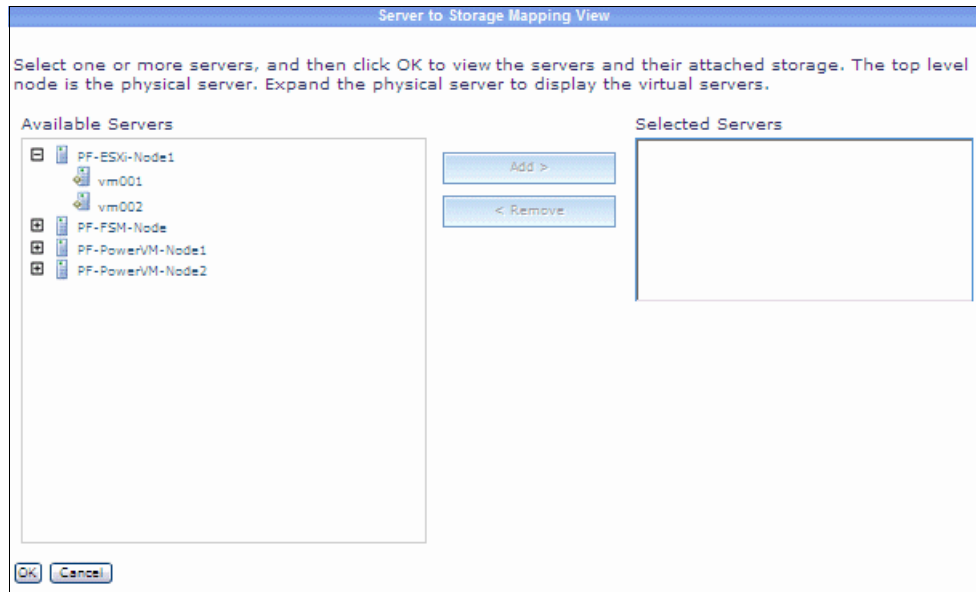


Figure 6-149 Server to Storage Mapping View

Storage to Server Mapping View

This task displays a table that shows the virtual server disks and the storage pools that contain them. Select the storage resources to view as shown in Figure 6-150.

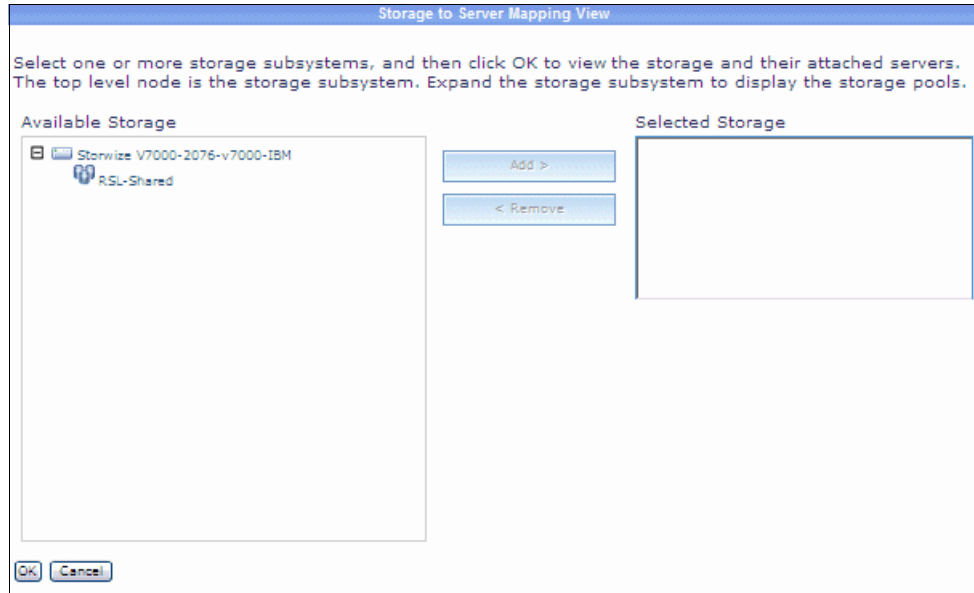


Figure 6-150 Storage to Server Mapping View

View and Manage Storage Volumes

This task works with your currently defined storage volumes to change, add, or delete them. Details about status and capacity are also available.

This task manages the assignment of network storage to individual systems, from the context of a specific host system. You can create a volume from network storage to be assigned to a selected host system. You can also delete a volume that is assigned to a selected host system. The create volumes function simplifies the allocation process by determining the best fit storage system and by creating any necessary RAID arrays automatically.

View and Apply Storage Templates

This task works with *storage templates*, which are predefined images for certain storage devices. Storage templates are used to ensure uniformity among common storage elements, and save time and effort on repetitive tasks.

Storage templates, as shown in Figure 6-151, can be used for these tasks:

- ▶ To clone the storage configuration of a system, or save a storage volume template from an existing server. Then, apply the saved template to another system. This process can be used for duplication (clustering, or virtual server hosts) or for saving the storage configuration for backup or disaster recovery purposes.
- ▶ As a starting point for creating more volumes, start with a base template and change it as needed.

Configuration Templates

Use configuration templates to deploy settings on one or more systems.
Configuration Templates

Deploy Create Create like Edit Delete Actions Search the table... Search

Select	Name	Deploy	Plan	Cour	Type	Subtype
<input type="checkbox"/>	8GbSANSwitchProtocolConfigurationTemplate	No	0		Network	I/O Module Fibre switch
<input type="checkbox"/>	Boot Sequence Predefined Template	No	1		Chassis	Processor
<input type="checkbox"/>	Ethernet1GbSwitchProtocolConfigurationTemplate	No	0		Network	I/O Module Ethernet Switch
<input type="checkbox"/>	Ethernet1GbSwitchVLANConfigurationTemplate	No	0		Network	I/O Module Ethernet Switch
<input type="checkbox"/>	IPv4AddressPoolConfigurationTemplate	No	0		Server	Server (via CIM protocol)
<input type="checkbox"/>	IPv6AddressPoolConfigurationTemplate	No	0		Server	Server (via CIM protocol)
<input type="checkbox"/>	OperatingSystemCreatei5AccountTemplate	No	0		Operating System	Operating System (via CIM protocol)
<input type="checkbox"/>	OperatingSystemCreateLinuxAccountTemplate	No	1		Operating System	Operating System (via CIM protocol)
<input type="checkbox"/>	OperatingSystemCreateWindowsAccountTemplate	No	1		Operating System	Operating System (via CIM protocol)
<input type="checkbox"/>	OperatingSystemImmediatePowerOff	No	0		Operating System	Operating System (via CIM protocol)
<input type="checkbox"/>	OperatingSystemImmediateRestart	No	0		Operating System	Operating System (via CIM protocol)
<input type="checkbox"/>	OperatingSystemIPv4NetworkTemplate	No	1		Operating System	Operating System (via CIM protocol)
<input type="checkbox"/>	OperatingSystemIPv6NetworkTemplate	No	1		Operating System	Operating System (via CIM protocol)
<input type="checkbox"/>	ServerEnableSerialOverLAN	No	0		Server	Server (via CIM protocol)
<input type="checkbox"/>	ServerEnableServiceProcessorRedirection	No	0		Server	Server (via CIM protocol)

Page 1 of 2 Selected: 0 Total: 24 Filtered: 24

Figure 6-151 Storage configuration templates

Capacity Details table

For each type of storage, this table indicates the available capacity, usable capacity, total capacity, number of systems, and number of disk drives. The entries in this list correspond to storage groups that have the same name as shown Figure 6-152. If you select an entry, a table of storage subsystems that are members of the selected group is displayed.

Capacity Details					
Location	Available GB	Usable GB	RAW GB	Systems	Disk Drives
Local Storage	185	185	0	4	2
BladeCenter Storage	0	0	0	0	0
Network Storage	75643	80498	80509	1	22
Total	75828	80683	80509	5	24

Storage system not being discovered? [Learn more](#)

Figure 6-152 Capacity Details

6.14 External Fibre Channel SAN switch discovery

Figure 6-155 shows the test environment. One SAN switch that is installed in the chassis has a pass-through capability and is connected to the IBM SAN B80 switch.

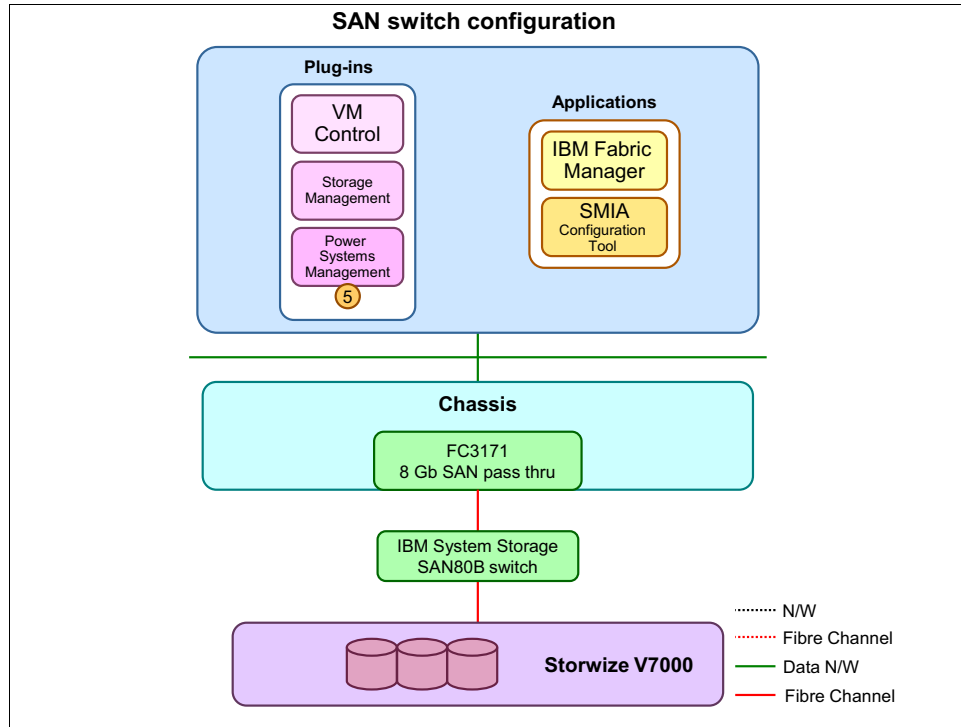


Figure 6-155 Test SAN environment

To add the Brocade SAN switch to FSM, perform these general steps:

1. Obtain the IP address from Brocade switch UI.
2. Ensure that the FSM is in the local Domain Name System (DNS) or added to the configured workstation `/etc/hosts` file.
3. Start the SMIA application.
4. Start the configuration tool.
5. Add the switch as a new fabric to the SMIA.
6. Use the `mkdatasource` command from FSM.
7. Collect inventory on the "Farm".
8. Collect inventory on the switch objects.

To add the Brocade SAN switch, perform these detailed steps:

1. Figure 6-156 shows the Applications tab main window. Click **Start** to run SMIA.




Figure 6-156 Applications tab in the FSM

2. Log in to your external switch (Figure 6-157 shows the SAN switch main window). Click **Switch Admin**.




Figure 6-157 SAN switch main window

3. Note the IP address in the **Network** tab as shown in Figure 6-158.

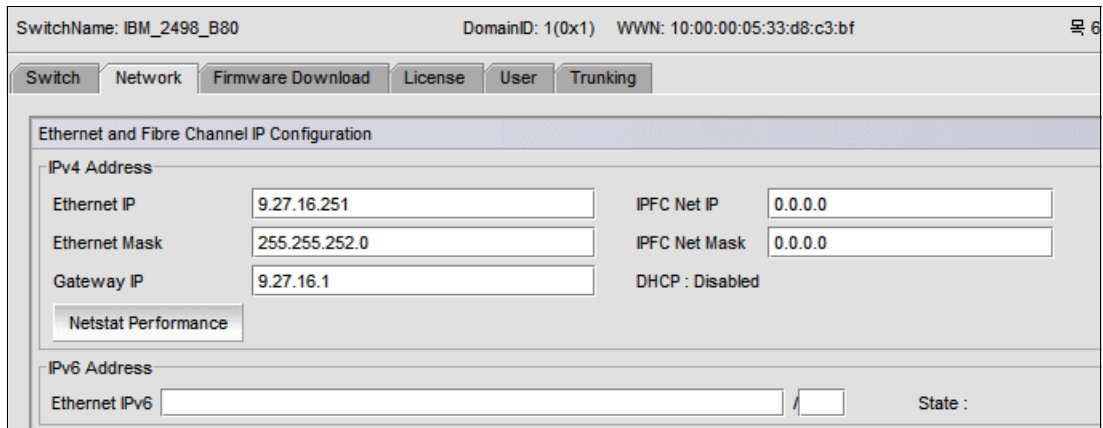


Figure 6-158 Ethernet IP address of the IBM SAN B80 switch

4. Check that SMIA is running as shown in Figure 6-159. Click **SMIA Configuration Tool**.




Figure 6-159 Check SMIA status

5. Enter the SAN switch credentials as shown in Figure 6-160.




Figure 6-160 SAN switch login window

6. Click the **Home** tab, and then click **Fabric Discovery** as shown in the Figure 6-161.

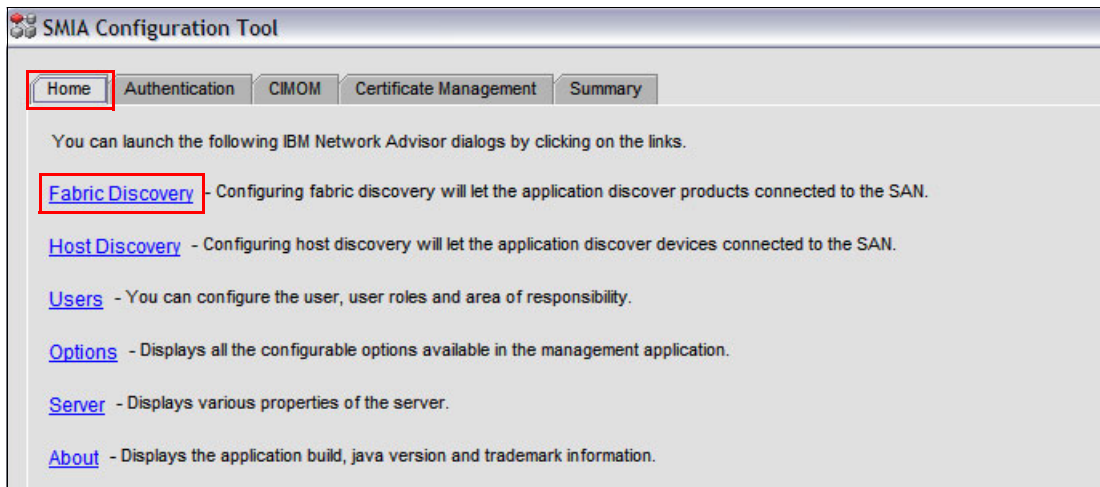


Figure 6-161 SMIA Configuration Tool window

7. Click **Add** as shown in Figure 6-162.




Figure 6-162 Discover Fabrics Main window

8. On the IP Address tab, enter a fabric name for the Top of Rack (TOR) switch, IP address, and login credentials (the defaults are admin/password), as shown in Figure 6-163.




Figure 6-163 Add Fabric Discovery window

9. Figure 6-164 shows the added SAN switch.




Figure 6-164 Discover Fabrics Main window

10. Click the **CIMOM** tab from the main interface as shown in Figure 6-165.

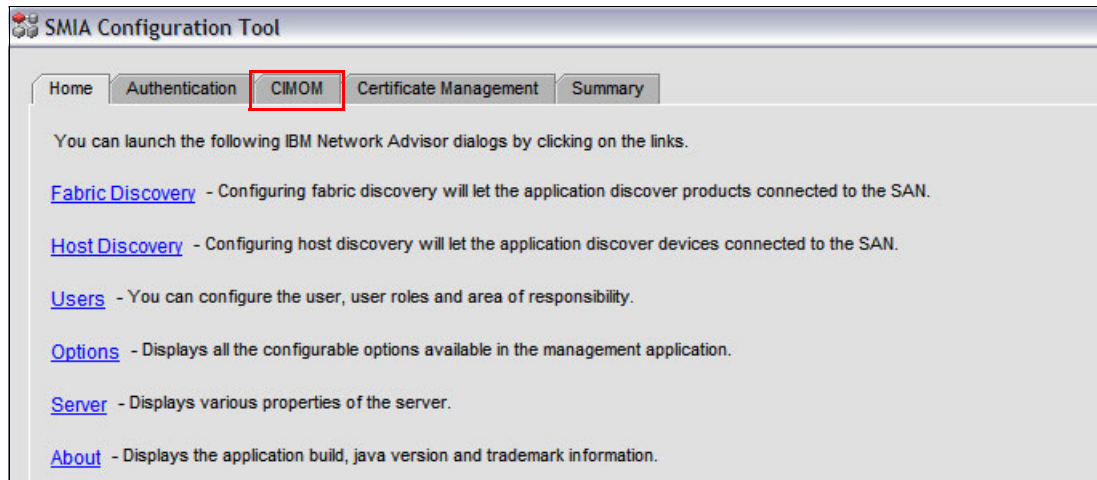


Figure 6-165 SMIA main window

11. Check **Enable SSL** to enable the SMI Agent port (25989), as shown in Figure 6-166.




Figure 6-166 CIMOM window

12. From the FSM CLI, run the `mkdatasource` command to add the data source as shown in Figure 6-167.

```
USERID@FSM-5CF3FC5F54EF:~> smcli mkdatasource -c fabric -t https -i 9.27.20.38 -p 25989
-u USERID -w Password -n "interop"
Adding the data source ...

The data source was added successfully.
```

Figure 6-167 FSM CLI window

13. Collect inventory on a farm as shown in Figure 6-168.

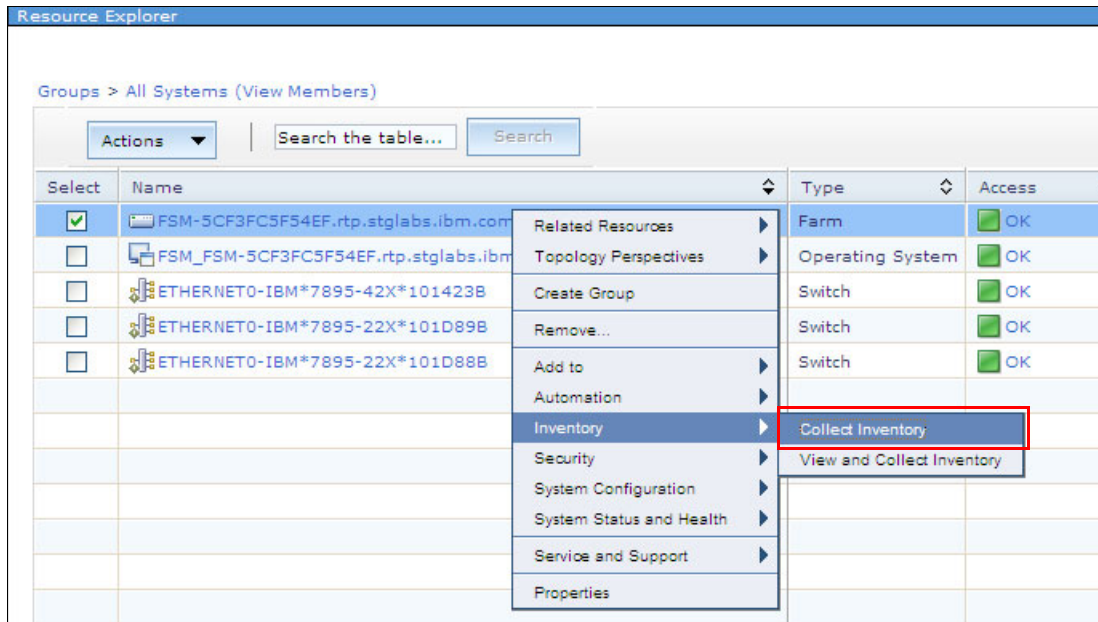


Figure 6-168 Run Collect Inventory

14. Check the logs as shown in Figure 6-169.




Figure 6-169 Check logs

Figure 6-170 shows that the new switch object is added in the Resource Explorer.




Figure 6-170 SAN switch added

6.15 Configuring network parameters (Network Control)

With IBM Flex System Manager, you can manage your entire network and network devices if the network devices are discovered and have full access. The Network Control window is shown in Figure 6-171.

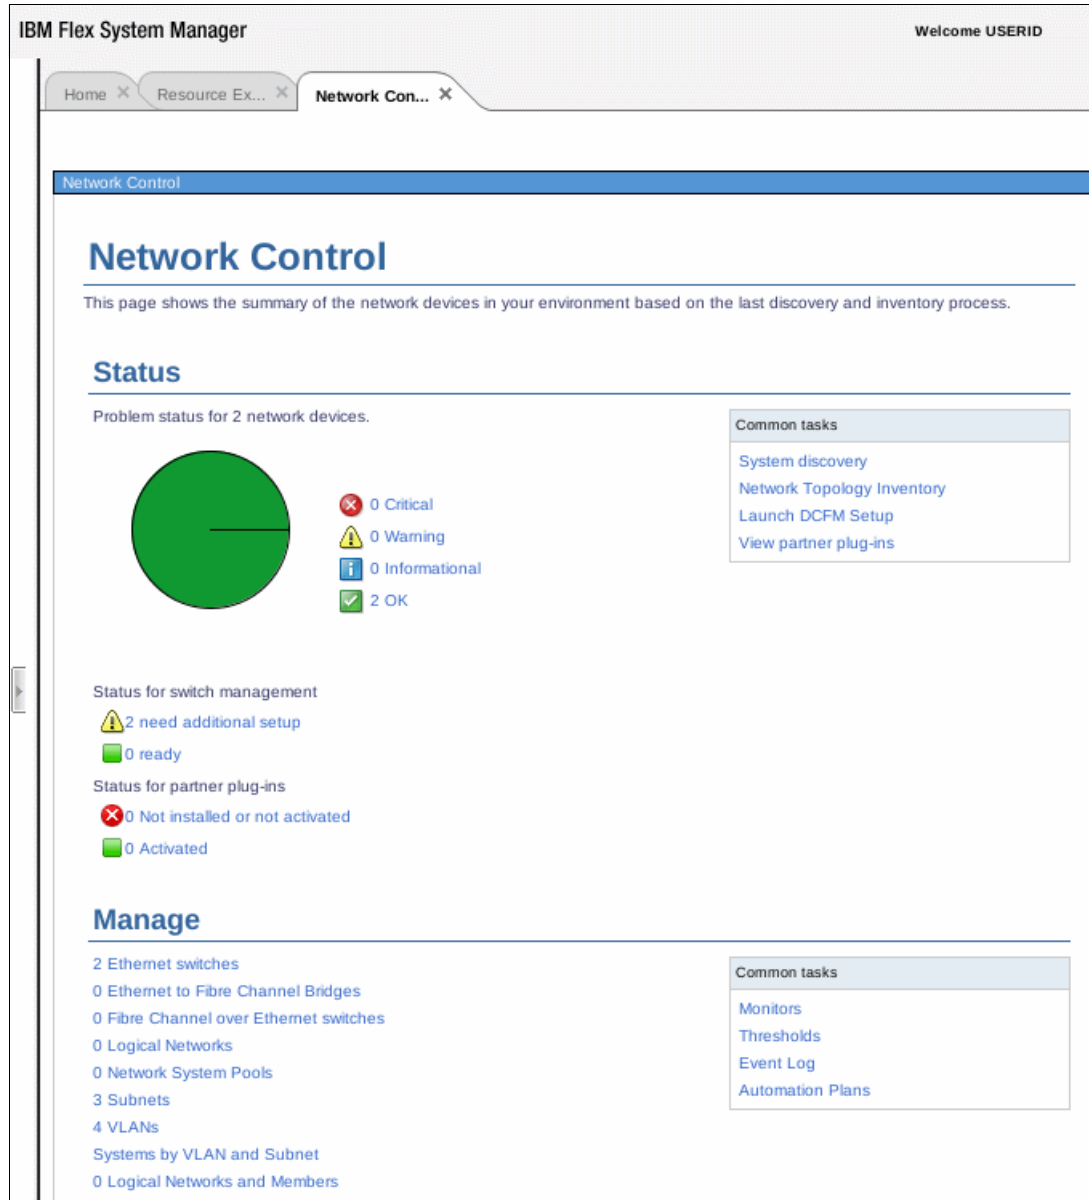


Figure 6-171 Managing network devices

Use IBM Flex System Manager Network Control to manage network devices in your managed systems environment by performing the following tasks:

- ▶ Discovering network systems

Use the Discovery task to collect an extended set of resources and relationships for network systems.

- ▶ Collecting and viewing inventory for network systems
Use the View and Collect Inventory task in IBM Flex System Manager Network Control to view and manage an extended set of resources and relationships for discovered network systems.
- ▶ Configuring network systems with configuration plans and templates
You can use the configuration manager to create, view, edit, delete, deploy, and schedule virtual local area network (VLAN) and protocol configuration templates to be deployed on supported network resources.
- ▶ Managing network system pools and logical networks
Use network system pools and logical networks to effectively manage your virtual and physical networks.
- ▶ Managing network systems health
IBM Flex System Manager provides facilities to monitor and troubleshoot network systems health.
- ▶ Working with network device groups
Use the Resource Explorer task to view and manage network systems in IBM Flex System Manager.
- ▶ Collecting and viewing Network Topology inventory
Use IBM Flex System Manager Network Control to work with network inventory in a topology view.

For information about the Network Pools concept and how to configure them to manage the virtualization environments, see 9.5, “Creating KVM network system pools” on page 352. For more information about Network Control, see the Network Control section in the IBM Flex System Information Center at this website:

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



Managing chassis components with IBM Flex System Manager

This chapter addresses IBM Flex System Manager (FSM) capabilities for chassis hardware component management. FSM offers a wide range of hardware management options. The examples use the FSM graphical user interface to illustrate common hardware management interfaces and tasks.

This chapter includes the following sections:

- ▶ 7.1, “Using FSM Explorer” on page 260
- ▶ 7.2, “Using the Chassis Map” on page 264
- ▶ 7.3, “Using the Event Log” on page 269
- ▶ 7.4, “Automating tasks with event automation plans” on page 271
- ▶ 7.5, “Handling problems with Service and Support Manager” on page 279
- ▶ 7.6, “Integrating Flex System Manager with an enterprise monitoring system” on page 288
- ▶ 7.7, “Monitoring system status and health” on page 288
- ▶ 7.8, “Remote management” on page 298

7.1 Using FSM Explorer

The IBM FSM Explorer console provides a next generation user interface that provides views of your resources and helps you to manage your systems-management environment. It provides a resource-based view of your environment with intuitive navigation of those resources.

You can view basic information about your resources just by hovering over them; you do not have to click to access information about them.

You can perform the following tasks in IBM FSM Explorer:

- ▶ Configuring local storage, network adapters, boot order, Integrated Management Module (IMM) settings, and Unified Extensible Firmware Interface (UEFI) settings for one or more compute nodes before you deploy operating-system or virtual images to them. (See 6.6, “Configuring compute nodes using Configuration Patterns” on page 150.)
- ▶ Installing operating system images on X-Architecture compute nodes. (See 6.7, “Deploying compute node images” on page 168.)
- ▶ Navigating resources, viewing the properties of resources, and performing basic management tasks, such as powering on and off, collecting inventory, and working with LEDs.
- ▶ Using the Chassis Map to edit compute node details, view server properties, and manage compute node actions
- ▶ Working with resource views, such as All Systems, Chassis and Members, Hosts, Virtual Servers, Network, and Storage
- ▶ Visual monitoring of status and events
- ▶ Visual monitoring of job status

For other tasks, you are launched from the IBM FSM Explorer into IBM Flex System Manager in a separate browser window. You can return to the IBM FSM Explorer window after you complete those tasks. As more tasks become available in IBM FSM Explorer, you will need to launch the IBM Flex System Manager less often.

The FSM Explorer can be started from the initial setup tab on the IBM Flex System Manager Home page. Click **Launch IBM FSM Explorer** as shown in Figure 7-1.



Figure 7-1 Initial Setup: Launch FSM Explorer

A separate browser window opens to present the IBM FSM Explorer as shown in Figure 7-2.

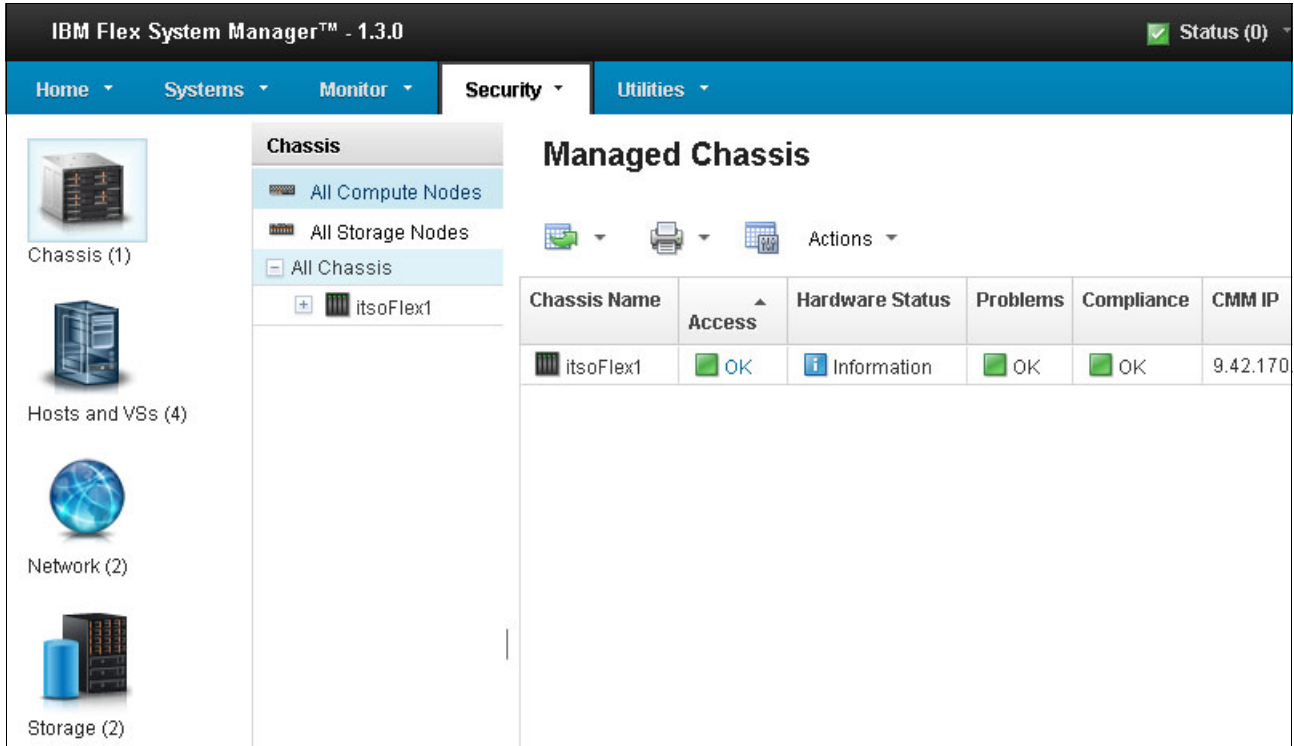


Figure 7-2 Main dashboard view of FSM Explorer

Jobs and status can be viewed by hovering your mouse over the status indicators at the upper right, as shown in Figure 7-3.

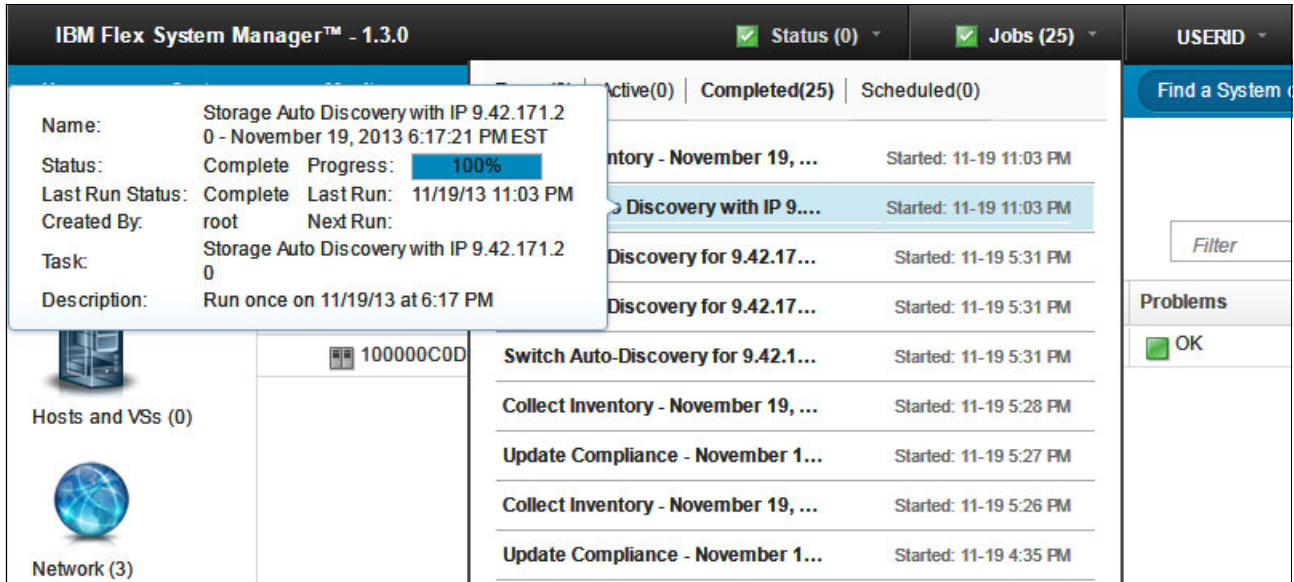


Figure 7-3 Pop-up view of job status

You can search for groups, resources, and tasks from a single location in the FSM Explorer. If a task is not available in IBM FSM Explorer, you are launched into IBM Flex System Manager to complete the task.

Figure 7-4 shows a search started for “chassis” and a dynamic list of items that match what you are typing.

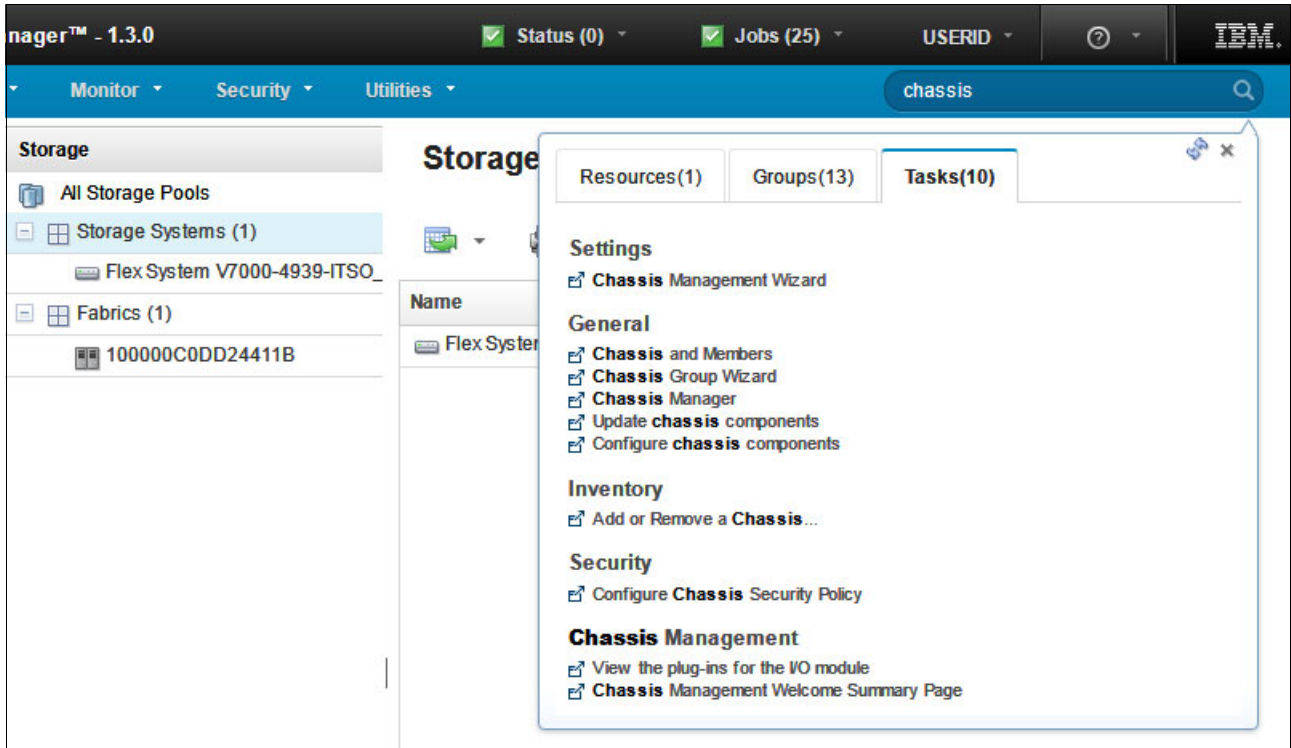


Figure 7-4 Dynamic list of items provided during a search

Clicking one of the search results performs the respective action - displays more information about the object for the Resources, displays group members for Groups, or start the task for Tasks.

7.2 Using the Chassis Map

You can view chassis properties and manage a chassis with the Chassis Map in the IBM FSM Explorer management software web interface. To do so, perform these steps:

1. In the FSM Explorer dashboard (the default view when you open FSM Explorer from the Home window), click **Chassis** group in the upper-left corner, and then click *<chassis name>* to open the Chassis Map for the selected chassis, as shown in Figure 7-5.



Figure 7-5 IBM FSM Explorer: Chassis Hardware Map

Move the mouse over **Home** (you do not need to do a mouse click) and click **FSM Explorer** to switch to the dashboard view from other FSM Explorer views, as shown in Figure 7-6.




Figure 7-6 FSM Explorer dashboard

The graphical Chassis Map is a visual representation of the front and back of the chassis and its components (see Figure 7-7). It shows you where your hardware components are located physically, and it is a central point of management from which you can get hardware configuration and status information. You can also perform various actions.

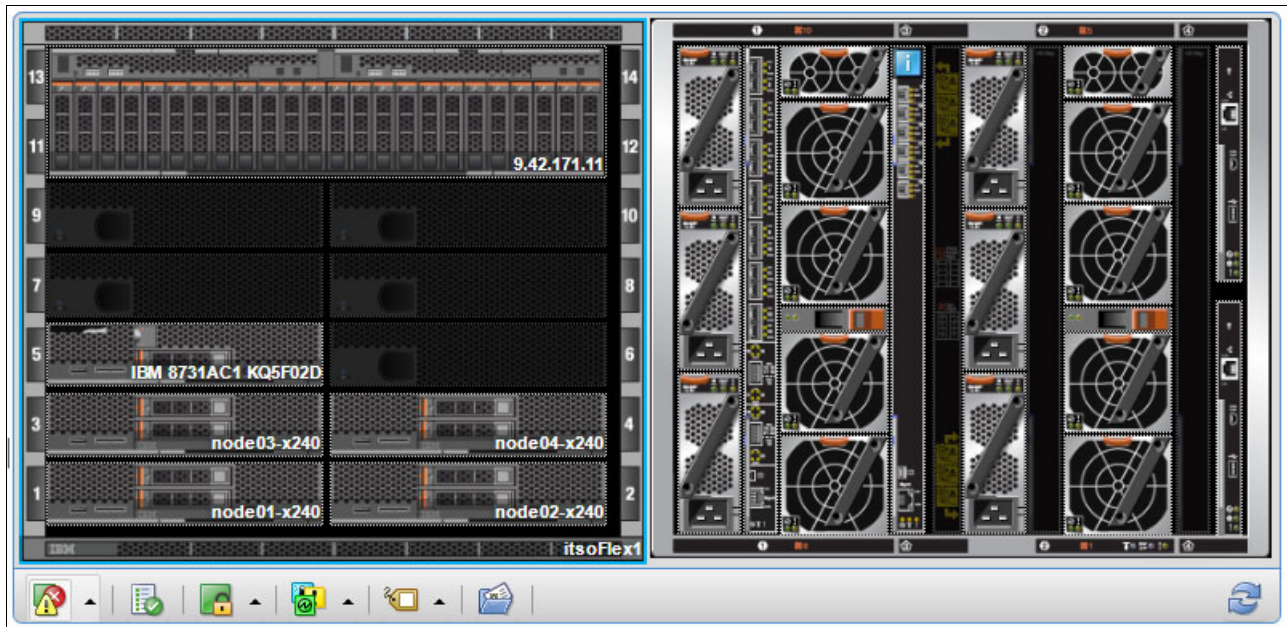


Figure 7-7 Chassis Map (graphical view)

- Click the compute node to get summary information, the option to view details, and relevant actions. The Common Actions area for the specific node is displayed under the Chassis Map or on the right side of it depending on the width of the web browser window, as shown in Figure 7-8.



Figure 7-8 Chassis Map with selected compute node and available Common Actions area

- The full list of actions and details that are relevant to the selected chassis component can be displayed by right-clicking it, as shown in Figure 7-9.




Figure 7-9 List of actions for the selected component (compute node)

- Use the Chassis Map to quickly identify problematic components of your chassis. For illustration purposes, we simulated a link failure on the switch module, and the Chassis Map showed an error with one of the I/O modules. Move the mouse pointer over the I/O module to see information about the switch and its problems, as shown in Figure 7-10. Click **View All Status** to open the Active Status window.

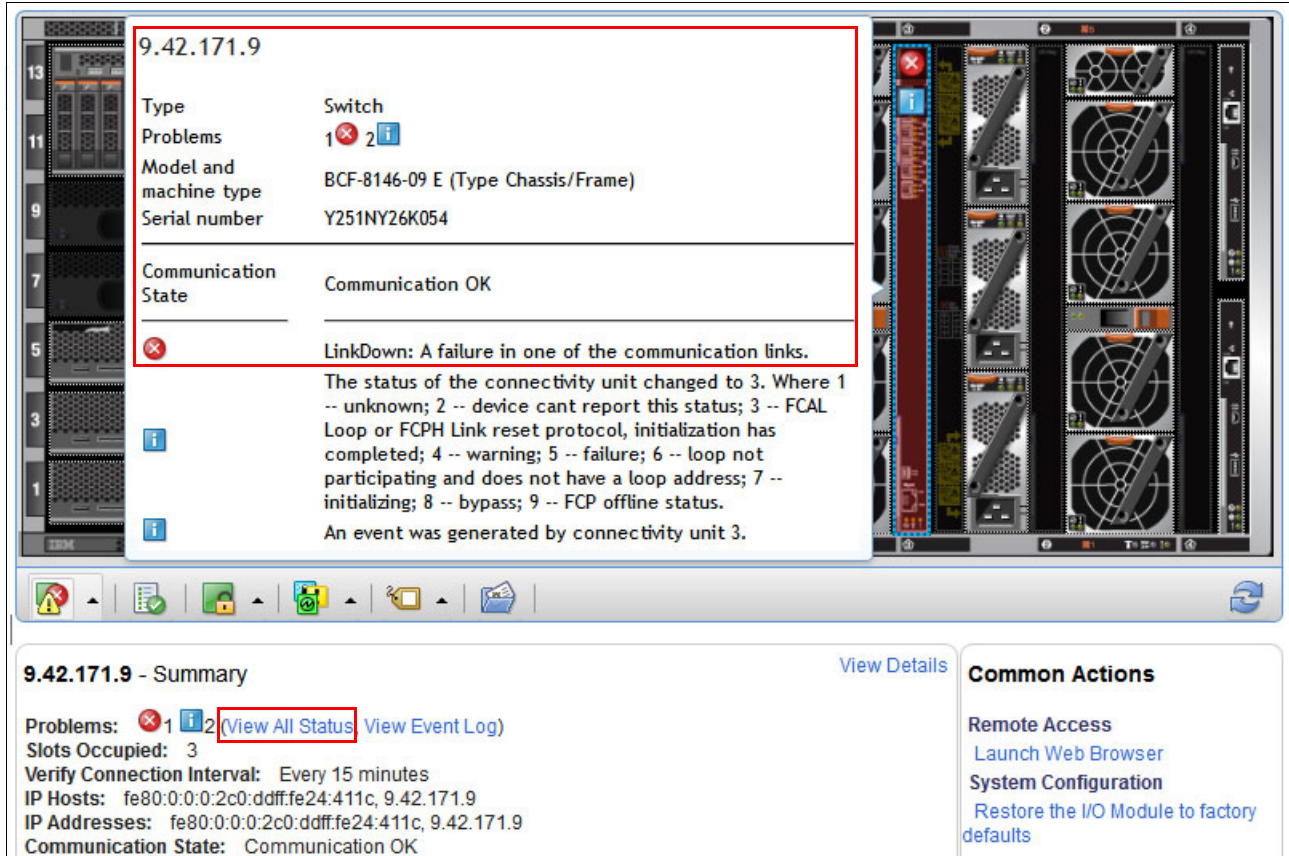


Figure 7-10 Caption with brief information about an I/O module in the Chassis Map

- The Active Status window provides a centralized interface that you can use to get a quick snapshot of the resources that trigger a status set entry. Currently, only entries that are related to the selected I/O module are listed, as shown on Figure 7-11. You also can ignore status-set entries to prevent them from displaying with an elevated status in the future.

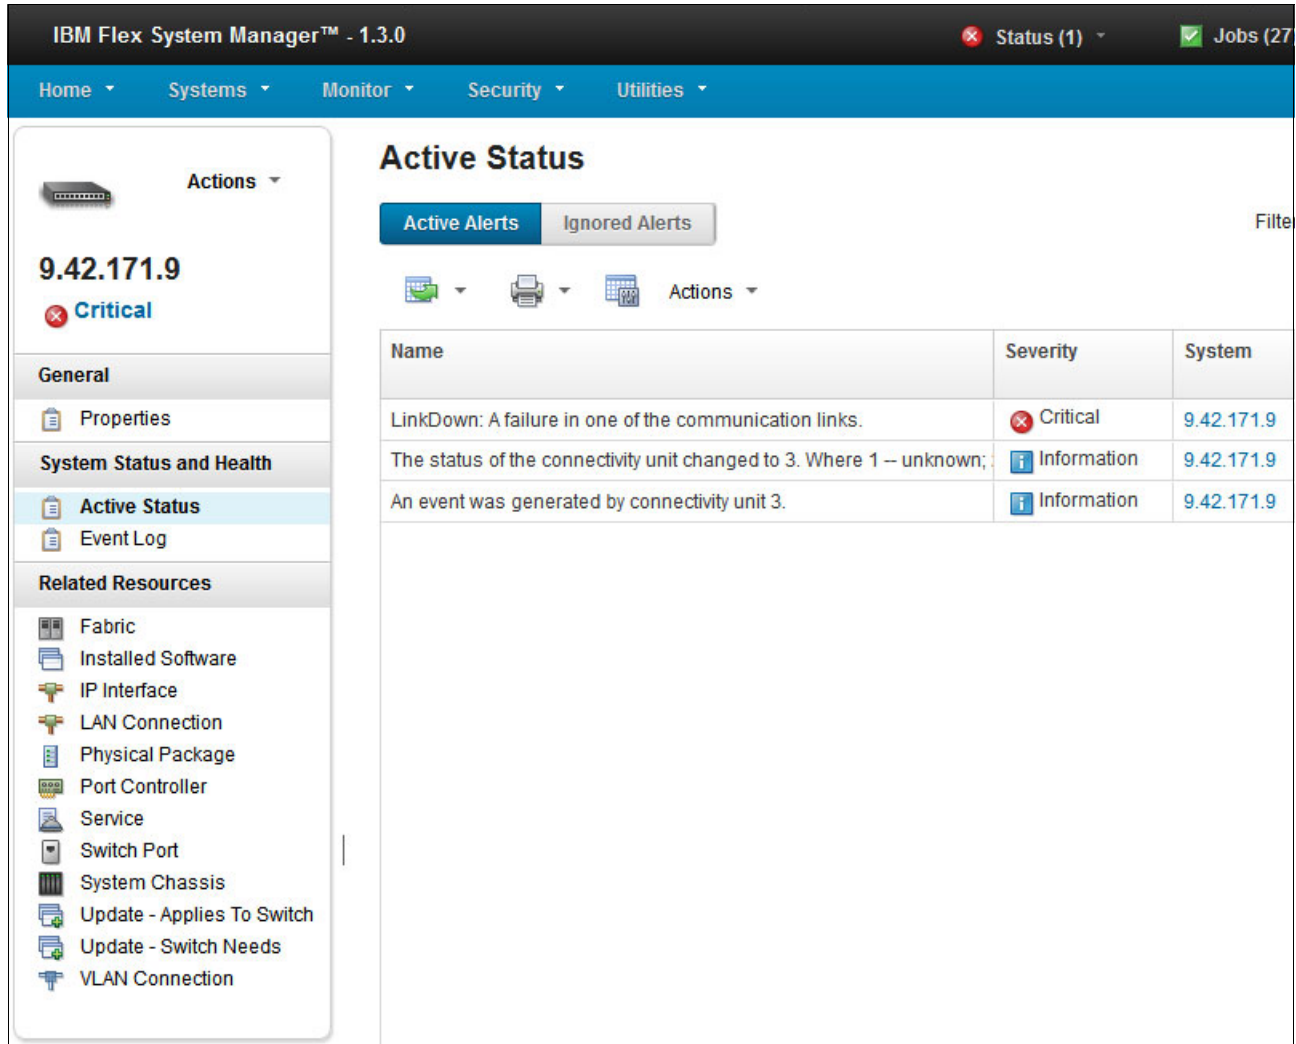


Figure 7-11 FSM Explorer: Active Status window

All status-set entries, including problems and compliance issues, are displayed. Double-clicking an entry displays further details about the entry.

7.3 Using the Event Log

An *event* is an occurrence of significance to a task or resource. Examples of events include the completion of an operation, the failure of a hardware component, or exceeding a processor threshold. The **event log** task displays all events that the FSM receives from any resource to which you have authority to view events.

To open the Event Log for events that are reported by a specific source, perform these steps:

1. Select a component from the Chassis Map.
2. Click **View Event Log** in the Summary area below the Chassis Map, as shown in Figure 7-12.




Figure 7-12 Opening the event log for the specific component

Move the mouse over **Monitor** in the FSM Explorer user interface and click **Event Log** to view all events for all chassis components, as shown in Figure 7-13.

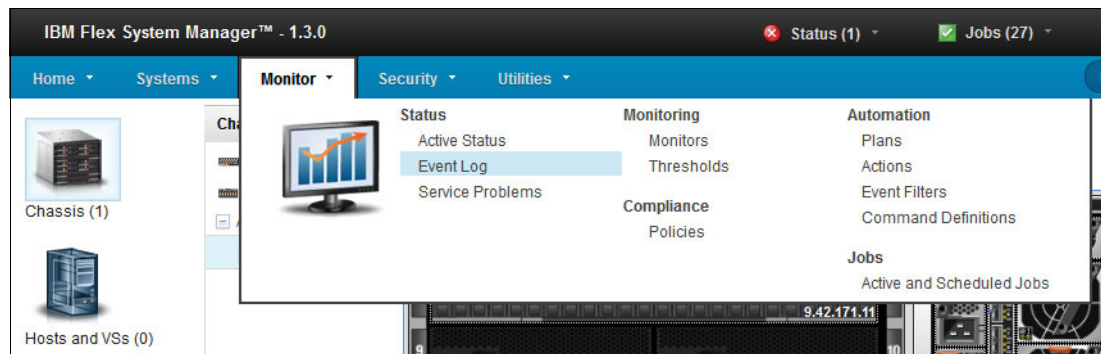


Figure 7-13 FSM Explorer: Event Log

In the Event Log window, click **Event Filter** and select the filter criteria that you want to use. The default filter is All Events. The event log displays the events that have been received by IBM Flex System Manager and match the filter criteria. See Figure 7-14 for possible choices.



Figure 7-14 Event filter drop-down list

Remember: The number of events that are displayed is limited by the event log preferences settings. By default, the event log displays the last 500 events that occurred over the last 24 hours. Use the Event Log Preferences window to change the defaults.

View the properties for the event in the table or click the event to view more properties and details. You can also use the Filter field to filter the event log based on specific keywords.

7.4 Automating tasks with event automation plans

Use *event automation plans* to automate tasks in your systems management environment.

Create event automation plans and apply them to specific systems to be notified by email, for example, when a specified threshold is reached or a specified event occurs. Or, you can configure an event automation plan to start a program on a system in response to the event. These plans are composed of event filters and event actions. The plans are triggered by events. Event automation plans are a powerful feature to automate a huge variety of manual tasks in your environment. These tasks can significantly reduce labor costs.

You have identified an event, created an event filter for it, and defined an event action. Now, automate triggering the action that is based on the filtered event. For this example, set up an automated email notification for an event that indicates a Predictive Failure Analysis (PFA) alert for memory dual inline memory modules (DIMMs).

To create the automation task, perform these steps:

1. Type event automation in the search field of the FSM Explorer and click **Event Automation Plans** in the Tasks pane, as shown in Figure 7-15.

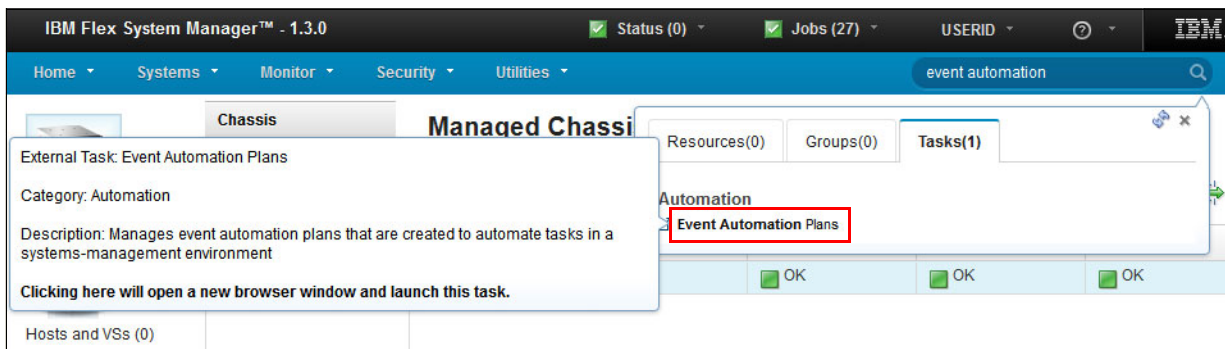


Figure 7-15 Launching Event Automation Plans task

2. In the opened Event Automation Plans window (Figure 7-16), click **Create**.

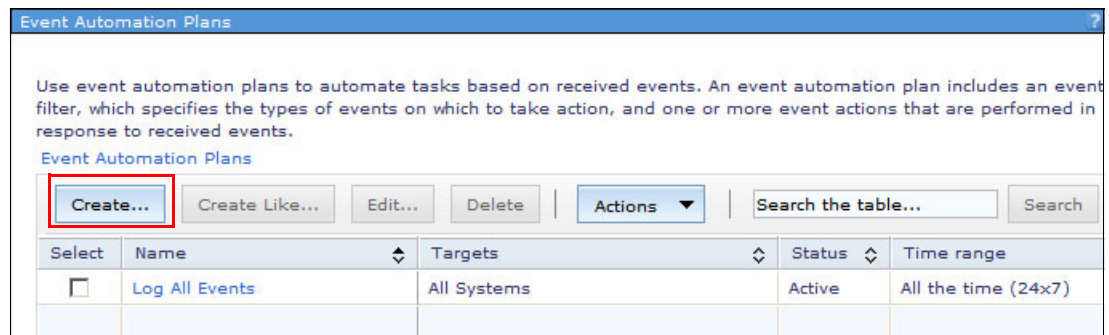


Figure 7-16 Event Automation Plans window

3. The Welcome window of the Event Automation Plan wizard is displayed (Figure 7-17). In the Welcome window, specify whether you want to show the Welcome window the next time you use the wizard and click **Next**.

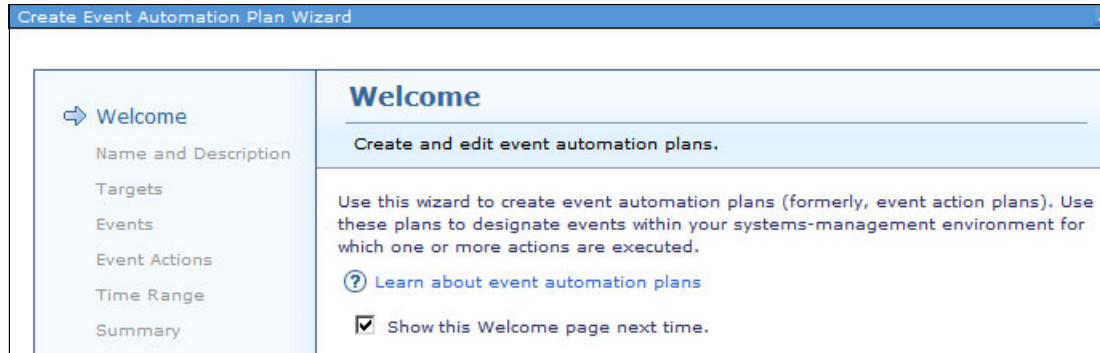


Figure 7-17 Create Event Automation Plan Wizard Welcome window

4. The Name and Description window opens (Figure 7-18). Enter a descriptive name for the event automation plan that you are creating.




Figure 7-18 Create Event Automation Plan Wizard Name and Description window

- In the Targets window, select the systems that the event automation plan will monitor for specific generated events. Select **All Systems** in the Available list and click **Add** to move them to the Selected list (Figure 7-19). Click **Next**.

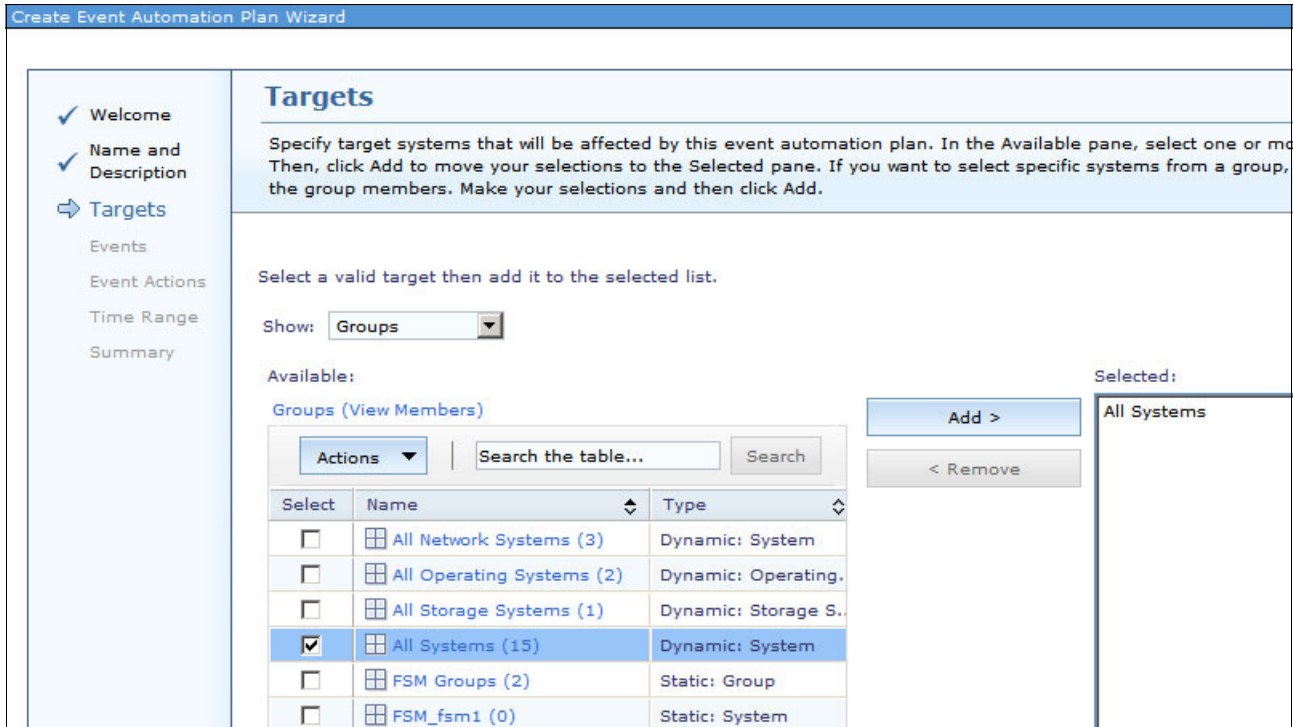


Figure 7-19 Create Event Automation Plan Wizard Targets window

Common event filters are predefined simple filters that monitor for events of common interest in systems management. For example, the Disks event filter is triggered by any hard disk events, and the Fans event filter is triggered by any fan events. The Event Automation Plan Wizard provides several common event filters so that you can create event automation plans quickly and easily.

- To monitor specific events that are not included in the common event filters, select **Advanced Event Filters** (Figure 7-20). Select the **Hardware Predictive Failure Alert events** filter and click **Next**. You can also create your own custom filter, if required.

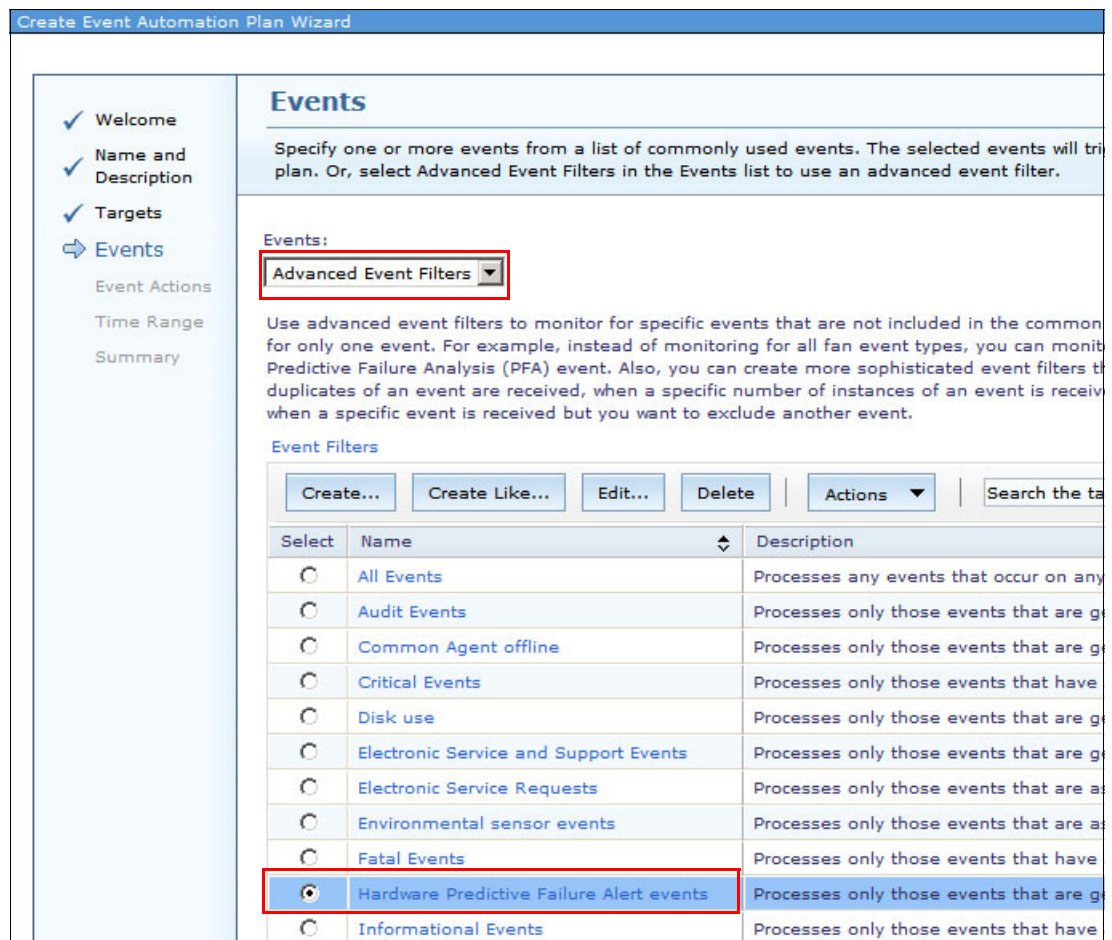


Figure 7-20 Create Event Automation Plan Wizard Events window

- The Event Actions window opens (Figure 7-21). Click **Create** to define a new event action.

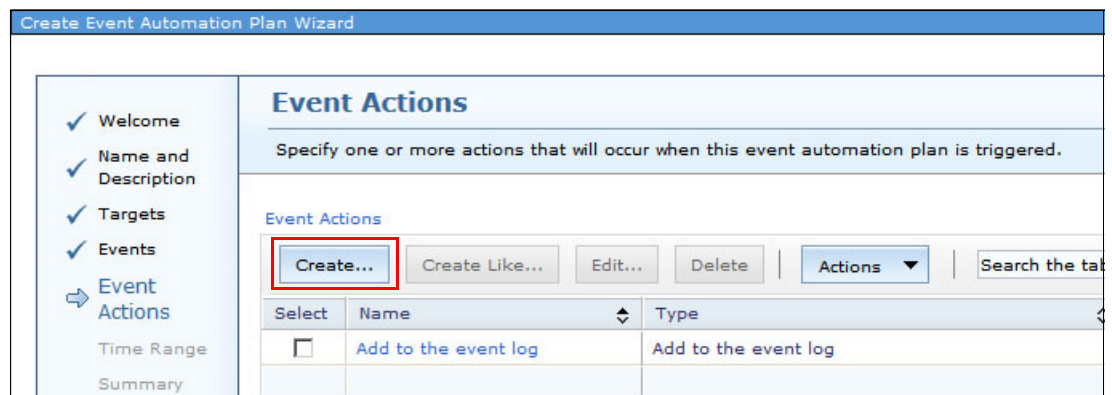


Figure 7-21 Create Event Automation Plan Wizard Event Actions window

- Select the required action as shown in Figure 7-22 and click **OK**. In our example, we selected **Send an e-mail (Internet SMTP)**.




Figure 7-22 Create Action window: Choose an action

9. Customize the selected “Send an e-mail” action. In addition to the action name and description, you need to enter the Send-to e-mail address, the Reply-to e-mail address, E-mail server, and E-mail port number. You can also specify the information to include in the email subject and body. Customize that by adding or removing predefined event variables from the Event variable list box. See Figure 7-23.

Create Action

E-mail [? Learn more](#)

*Action name:
Electronic Service Notification

Description:
[Empty text area]

*Send-to e-mail address:
hwsupport@mycompany.com

*Reply-to e-mail address:
fsmadmins@mycompany.com

*E-mail (SMTP) server (for example, smtp.mycompany.com):
smtp.mycompany.com

*E-mail (SMTP) port:
25

Subject of message:
&date &system

Body of message:
&text

Select an event variable and text field to insert the variable in the target text field. You can also specify the appropriate language and time zone.

[? Learn more about using event variables](#)

Event variable:
Date the event occurred (&date)

Target text field:
Subject of message:

Insert

Language:
English

Time zone:
America/New_York - Eastern Standard Time - EST

Test OK Cancel

Figure 7-23 Create Action window: The details for sending an email action

10. Click **Test** to ensure that the action is configured correctly. Click **OK**.

The newly created event action is displayed in the Event Actions list, as shown in Figure 7-24. Click **Next**.




Figure 7-24 Event Actions window

Remember: Event action history is not saved by default. Saving the history of an event action can provide useful information, such as when the event action ran and the event that triggered the action. Click **Actions** → **Start Saving History** to enable history for a selected event action.

11. Click **Next** in the Time Range window (Figure 7-25) because time settings cannot be modified for the built-in PFA filter.

For custom filters, you can select the time period over which you want to collect the events. You can select All the time (24x7) to enable the plan to be active all the time. Or, you can select Custom to choose specific days and a specific time for the plan to be active.




Figure 7-25 Create Event Automation Plan Wizard Time Range window

12. In the Summary window (Figure 7-26), verify the details of the event automation plan. If you need to make changes, click **Back**. Ensure that you specify whether you want to apply the event automation plan as soon as you finish creating it. Click **Finish**.

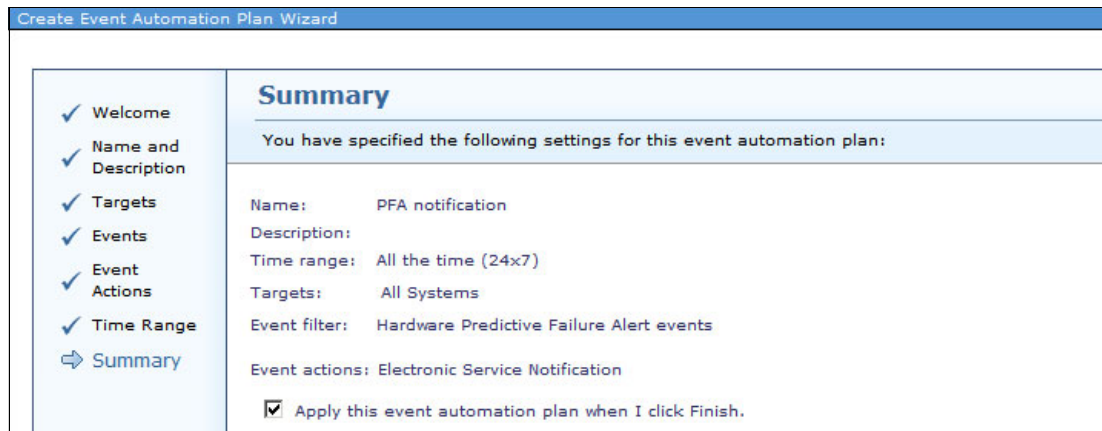


Figure 7-26 Create Event Automation Plan Wizard Summary window

The event automation plan is saved and displayed in the Event Automation Plans window, as shown in Figure 7-27.

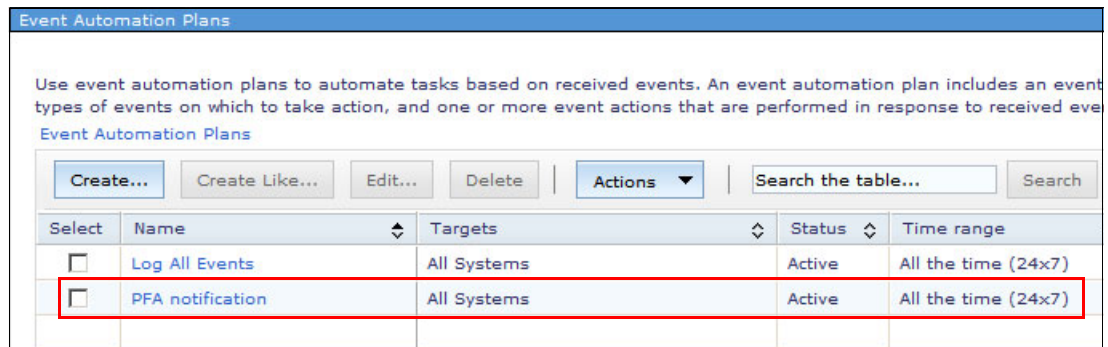


Figure 7-27 Event Automation Plans window

You can Edit, Delete, or Deactivate the automation plan from the Actions menu.

You can deactivate an event automation plan so that the specified events do not trigger the plan. When you want to enable the event automation plan again, you can activate the event automation plan. By default, event automation plans are activated.

7.5 Handling problems with Service and Support Manager

In certain cases, you need to contact IBM support about a hardware issue and submit supporting data for further analysis. This activity usually includes a number of time-consuming manual tasks. Such manual tasks can now be automated with Service and Support Manager.

Service and Support Manager is a plug-in for FSM. Service and Support Manager automatically detects serviceable hardware problems and collects supporting data for serviceable hardware problems that occur on your monitored endpoint systems. The IBM Electronic Service Agent tool is integrated with Service and Support Manager, and transmits serviceable hardware problems and associated support files to IBM support.

Service and Support Manager includes the following features:

- ▶ Automatically detects serviceable hardware problems to IBM support for all monitored systems.
- ▶ The integrated Electronic Service Agent tool securely transmits serviceable hardware problems, associated support files, and performance management data to IBM support.
- ▶ Collects and securely transmits scheduled system inventory and diagnostic support files to an IBM database. This inventory information is available to IBM support representatives when they are solving your problem.
- ▶ Communicates with IBM support through a secure connection that uses encryption and authentication.
- ▶ Includes the option to send email notifications when a serviceable problem is detected and a service request is opened.

Service and Support Manager begins automatically monitoring for serviceable hardware problems as soon as FSM is installed. However, activation by running the Getting Started Wizard is required to configure the Electronic Service Agent tool. This tool is integrated with Service and Support Manager, and securely transmits serviceable hardware problems and associated support files, inventory, and performance management data to IBM support.

To activate Electronic Service Agent, perform these steps:

1. Navigate to the FSM Home page. Click the **Plug-ins** tab, as shown in Figure 7-28.

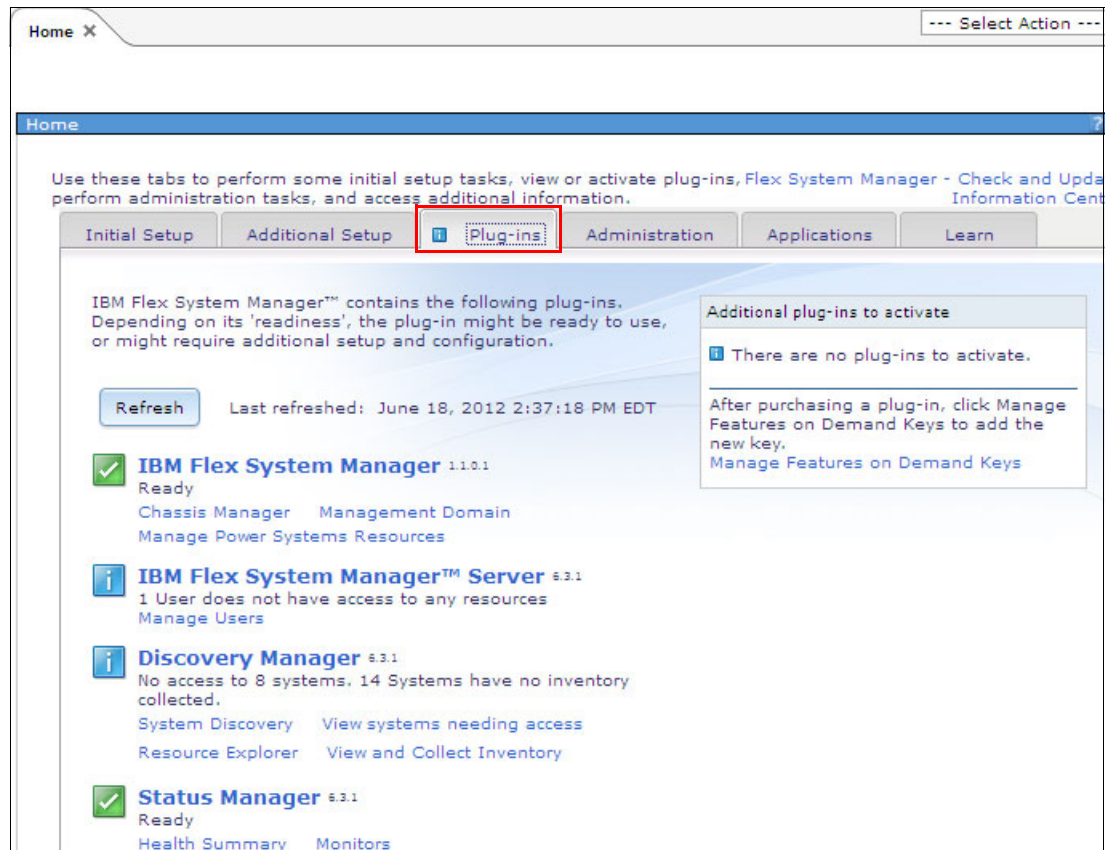


Figure 7-28 Plug-ins tab of FSM Home page

2. Locate the Service and Support Manager section, and click **Getting Started with Electronic Service Agent**, as shown in Figure 7-29.

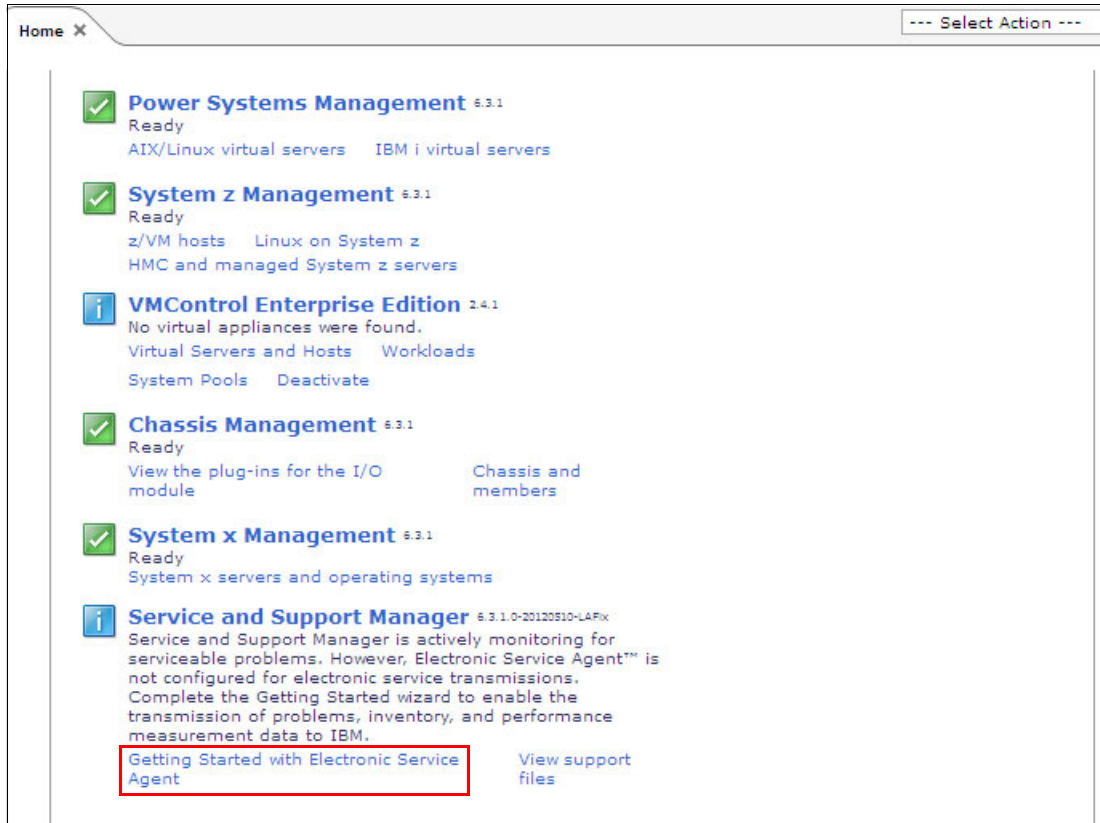


Figure 7-29 Service and Support Manager section on FSM Plug-ins window

3. Click **Next** in the Welcome window (Figure 7-30).

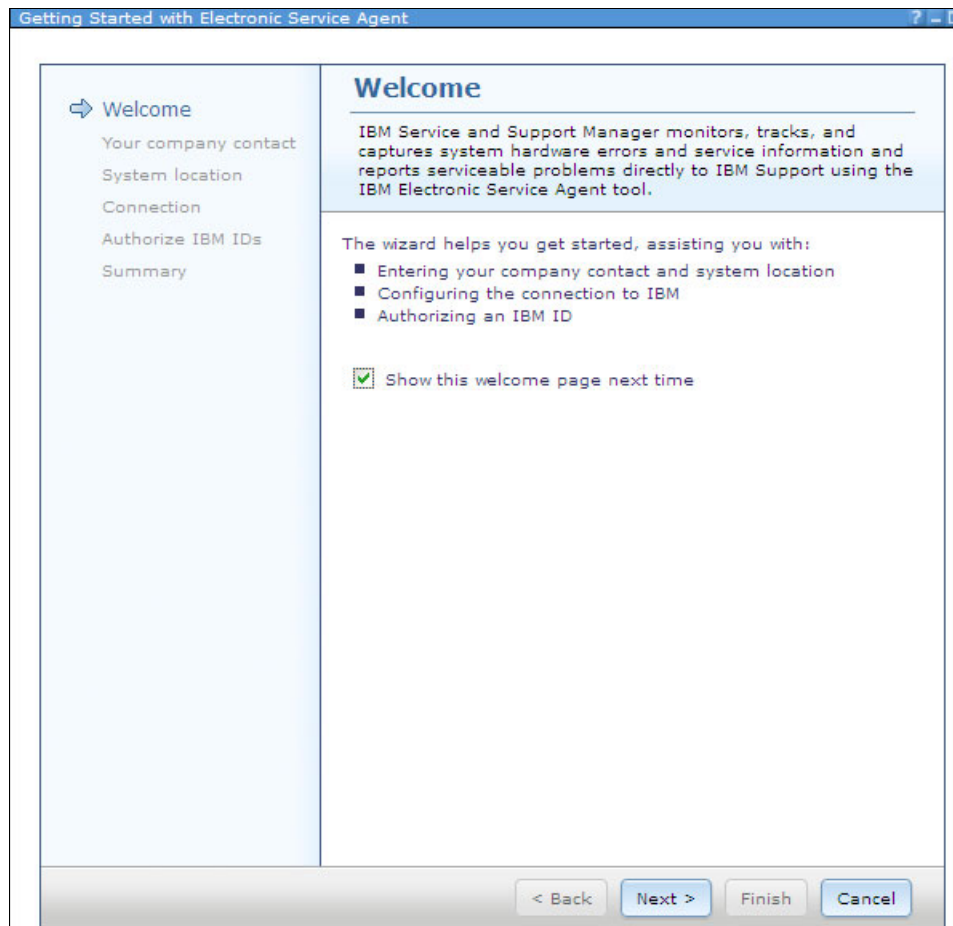


Figure 7-30 Getting Started with Electronic Service Agent Welcome window

4. Enter your company contact information and click **Next** (Figure 7-31).

✓ Welcome	<h3>Your company contact</h3> <p>Provide information about the person that IBM Support may contact about a problem reported by Electronic Service Agent.</p>	
⇒ Your company contact	*Contact name:	<input type="text"/>
System location	*Company name:	<input type="text"/>
Connection	*Telephone number:	<input type="text"/>
Authorize IBM IDs	Extension:	<input type="text"/>
Summary	Fax number:	<input type="text"/>
	Alternate fax number:	<input type="text"/>
	*E-mail:	<input type="text"/>
	Alternate e-mail:	<input type="text"/>
	Help desk number:	<input type="text"/>
	Pager number:	<input type="text"/>
	Street address Line 1:	<input type="text"/>
	Line 2:	<input type="text"/>
	Line 3:	<input type="text"/>
	City:	<input type="text"/>
	State or province:	<input type="text"/>
	*Country or region:	UNITED STATES <input type="button" value="v"/>
	Postal code:	<input type="text"/>

Figure 7-31 Getting Started with Electronic Service Agent: Your company contact window

5. Enter the system physical location information, which might differ from your company contact information (Figure 7-32). Click **Next**.

✓ Welcome	<h3>System location</h3> <p>Provide default information about the physical locations of your systems. Information can be overridden for specific systems by clicking Resource Explorer, selecting a system, and clicking Location under the Additional Properties heading.</p>	
✓ Your company contact	*Telephone number:	<input type="text"/>
⇒ System location	Extension:	<input type="text"/>
Connection	*Country or region:	<input type="text"/>
Authorize IBM IDs	*Street address:	<input type="text"/>
Summary	*City:	<input type="text"/>
	*State or province:	<input type="text"/>
	*Postal code:	<input type="text"/>
	*Building:	<input type="text"/>
	Floor:	<input type="text"/>
	Room number:	<input type="text"/>
	Row:	<input type="text"/>
	Aisle:	<input type="text"/>
	Displaced height (cm):	<input type="text"/>
	Altitude (meters):	<input type="text"/>
	Other information:	<input type="text"/>

Figure 7-32 Getting Started with Electronic Service Agent: System location window

- The Electronic Service Agent tool needs Internet access to securely transmit serviceable hardware problems, associated support files, and performance management data to IBM support. Enter the Internet proxy details if applicable, test the Internet connection, and click **Next** (Figure 7-33).

Figure 7-33 Getting Started with Electronic Service Agent: Connection window

- Optional: Enter any IBM IDs you might have to see the service information that is transmitted to IBM by Electronic Service Agent under your IBM account (Figure 7-34). Click **Next**.

Figure 7-34 Getting Started with Electronic Service Agent: Authorize IBM IDs window

8. Review the Summary window (Figure 7-35) and click **Finish**.



Figure 7-35 Getting Started with Electronic Service Agent: Summary window

9. Return to the Plug-ins tab of the FSM Home page, scroll down and click **Service and Support Manager**, as shown in Figure 7-36.

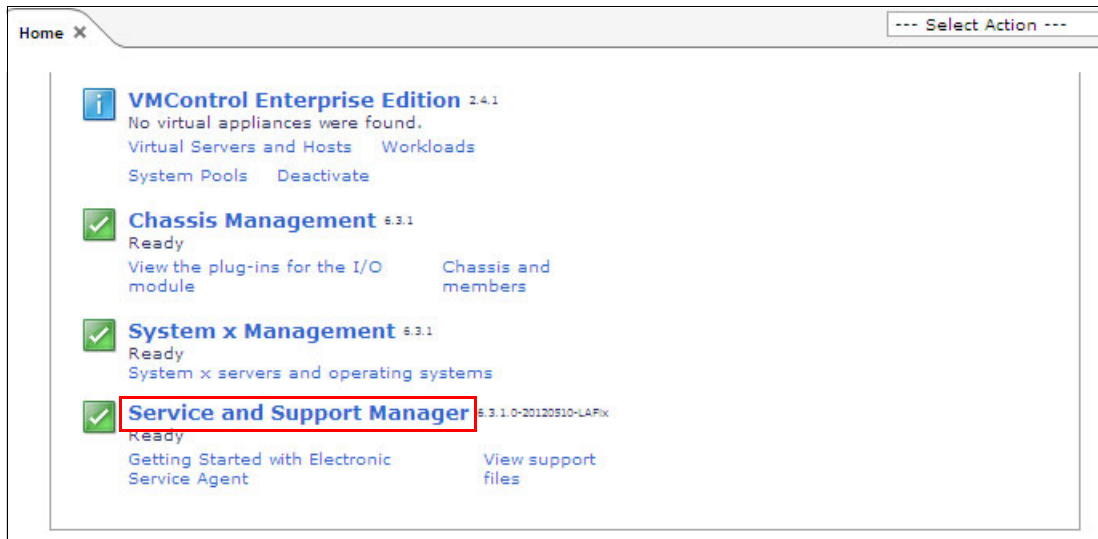


Figure 7-36 Service and Support Manager section in FSM Plug-ins window

10. The Service and Support Manager window opens (Figure 7-37). You can use this window to get an overview of the serviceable problems in the environment and a number of useful links. If you need to view, collect, and submit support files to IBM manually, click **Manage support files**. Click **Manage settings** to configure Service and Support Manager settings.

The screenshot shows the 'Service and Support Manager' interface. At the top, it says 'Service and Support Manager' and 'Manage serviceable problems on your systems.' Below this is a 'Problem Reporting' section with a pie chart titled 'Serviceable Problems for 37 Monitored Systems'. The pie chart shows 3 systems with serviceable problems (yellow slice) and 34 systems with no open serviceable problems (green slice). To the right of the pie chart are two status indicators: a yellow warning icon for '3 systems with serviceable problems' and a green checkmark icon for '34 systems with no open serviceable problems'. To the right of the pie chart is a 'Electronic Services Links' box with links for 'Serviceable Problems', 'All Problems', 'IBM Support Portal', and 'Open a service request'. Below the pie chart is a 'Recent Activity' section with four items: '3 serviceable problems require attention', '0 service requests being investigated by IBM', '0 requests have been updated in the last 24 hours', and '0 serviceable problems opened in the last 24 hours'. Below that is a 'Status' section with two items: 'Ready. Service and Support Manager is actively monitoring for serviceable problems and Electronic Service Agent™ is configured to automatically transmit problems, inventory, and performance measurement data to IBM.' and 'Dynamic System Analysis (DSA) status current. Service and Support Manager has verified that DSA collectors are current.' To the right of the status section is a 'Common Tasks' box with links for 'Manage support files', 'Send test problem...', 'Test connection to IBM...', and 'Verify DSA status'. At the bottom is a 'Setup and Configuration' section with links for 'Manage settings', 'Manage your system contacts', and 'Getting Started with Electronic Service Agent'. The 'Manage support files' and 'Manage settings' links are highlighted with red boxes.

Figure 7-37 Service and Support Manager window

11. Select **Automatically report problems for all systems** to allow Electronic Service Agent to report all problems to IBM support automatically (Figure 7-38).




Figure 7-38 Service and Support Manager settings

If you want to receive notifications about Service and Support Manager events, you can create an event automation plan as described in 7.4, “Automating tasks with event automation plans” on page 271. Use the **Electronic Service Requests** event filter in your automation plan to get notified every time that Service and Support Manager detects a serviceable hardware problem. The plan then opens an electronic service request with IBM support.

For more information about Service and Support Manager, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.esa.director.help/esa_kickoff.html

7.6 Integrating Flex System Manager with an enterprise monitoring system

If you have an enterprise monitoring system that is already implemented in your environment, you can create event automation plans to forward FSM events to it. FSM offers several predefined actions to help you. For more information, see 7.4, “Automating tasks with event automation plans” on page 271. Instead of selecting “Send an e-mail (Internet SMTP)”, you can choose one of the following advanced event actions:

- ▶ Send an IBM Tivoli Enterprise Console® event
After it is configured, this event action will forward FSM events to your Tivoli Enterprise Console server in the appropriate format.
- ▶ Send an SNMP trap reliably to an IBM Tivoli NetView® for IBM z/OS® host
After it is configured, this event action will forward FSM events to your Tivoli NetView host. For your Tivoli NetView for z/OS host to understand the Simple Network Management Protocol (SNMP) data that it receives from FSM, you need to load it with the FSM Management Information Base (MIB) files.
- ▶ Send an SNMP trap to an IP host
This event action can be used generally when you integrate FSM with a monitoring system that can receive SNMP traps. After it is configured, this action forwards FSM events to your enterprise monitoring system. For your monitoring system to understand the SNMP data, you must load it with the FSM MIB files.

Contact your IBM support representative to obtain the required FSM MIBs.

For more information about these and other event actions, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.director.automation.helps.doc/fqm0_c_ea_actions.html

7.7 Monitoring system status and health

The term *monitor* refers to a specific resource counter (for example, CPU Utilization) that you can watch (for real-time monitoring), record (for historical information), or set a threshold on (for alerting and automation).

The Monitors task provides the tools that you need to retrieve real-time status and quantitative data for specific properties and attributes of resources in your environment. You can also set thresholds for the monitors, graph the data that monitors retrieve, and drill down to quickly view the status of resources for each system. The specific monitors that are available vary based on the type of resource.

Explanation: For this example, a Microsoft Windows 2008 R2 server with Common Agent installed is used. For more information about agents, see 5.1.5, “Agents and tasks supported” on page 94.

To use a monitor, perform these steps:

1. Open the Monitors task from the FSM Explorer user interface by moving the mouse over the **Monitor** menu and clicking **Monitors**, as shown in Figure 7-39.

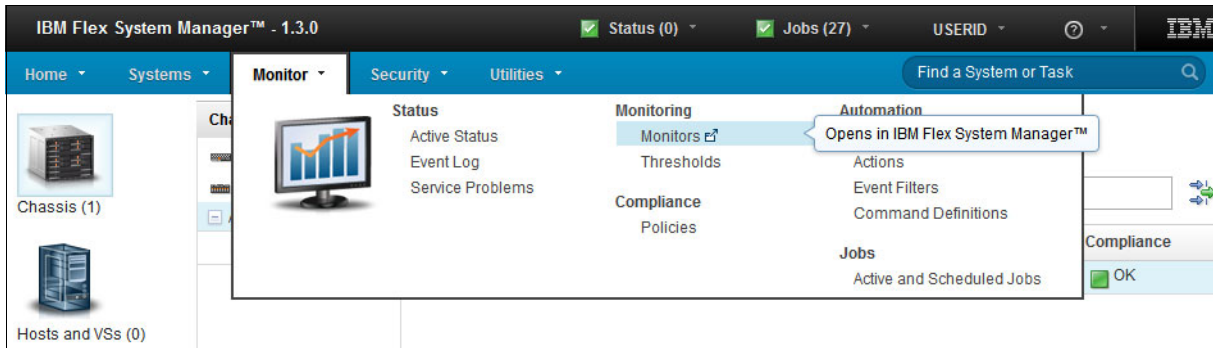


Figure 7-39 FSM Explorer: Monitors

2. Click **Browse** to select a system or group to view monitors, as shown in Figure 7-40.

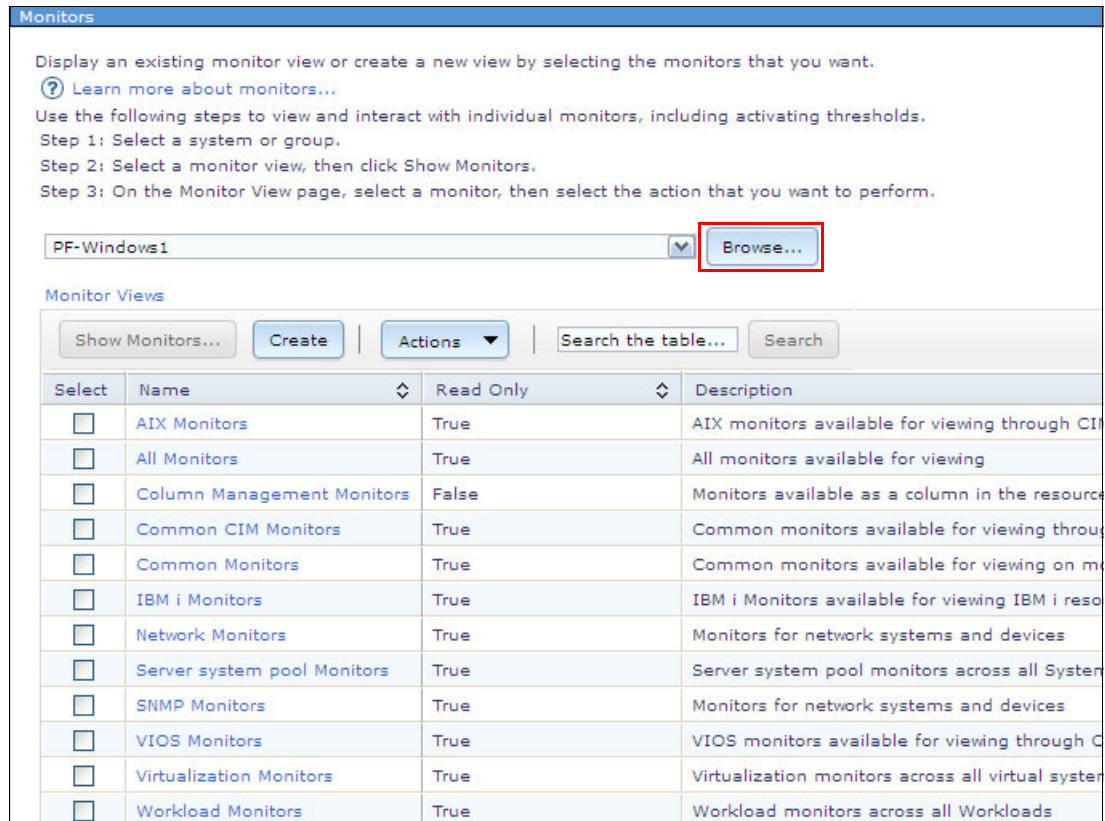


Figure 7-40 Monitors window

3. Select **Target Systems** from the Show menu to see individual systems. Select the system to monitor (in this example, **PF-Windows1**) and click **Add** to add it to the Selected list, as shown in Figure 7-41. If needed, you can choose Groups of systems, instead of individual systems. Click **OK** to proceed.

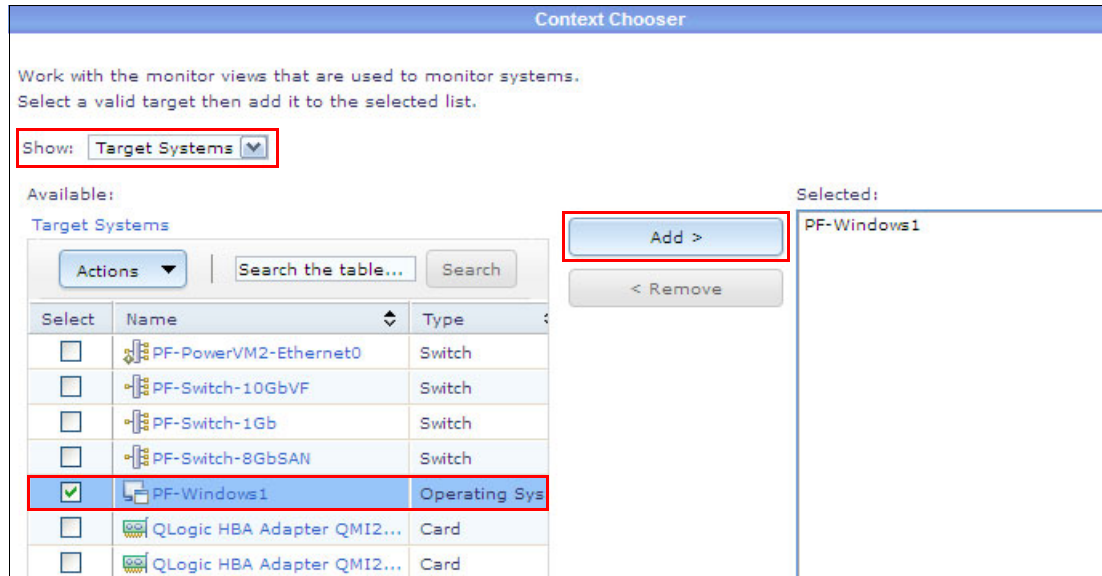


Figure 7-41 Monitors Context Chooser

4. FSM arranges available monitors in groups called *monitor views*. Each view represents a list of the most commonly available monitors in a category, for example, monitors that are supported by AIX. For this example, use the Common Monitors view.

The Common Monitors view contains some of the most common monitors for operating systems that are supported by FSM. When you create your own monitor view, more individual operating systems monitors might be available.

Click **Common Monitors** or select it, and click **Show Monitors**, as shown in Figure 7-42.

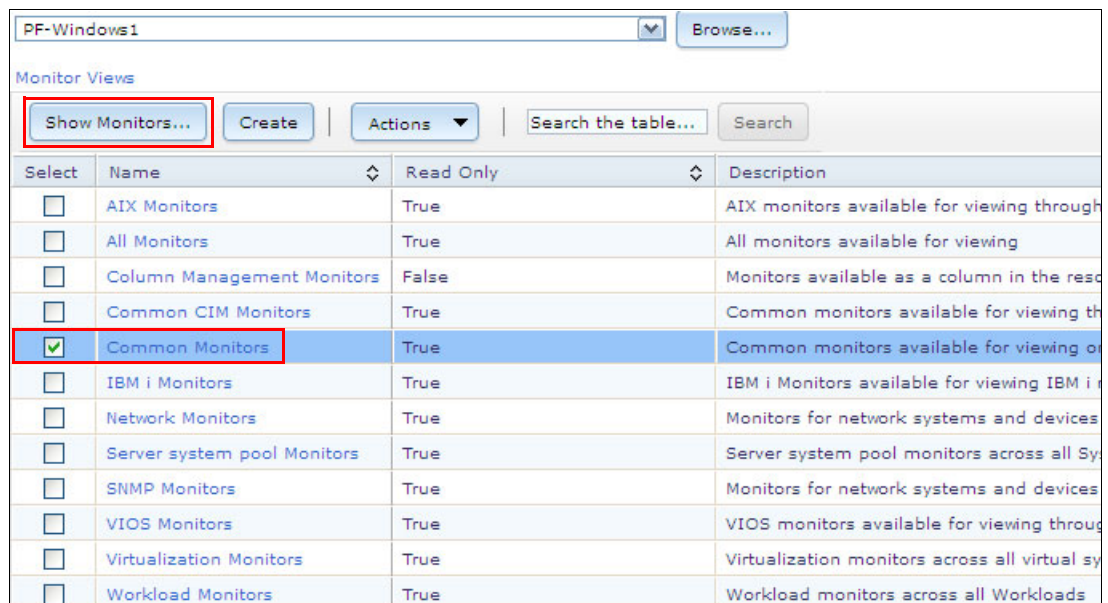


Figure 7-42 Monitor Views window with Common Monitors selected

- The Monitor View window shows all common monitors for the selected Windows target system. You can see the real-time values of the individual monitors and information about activated thresholds. There are no activated thresholds as shown in Figure 7-43.




Figure 7-43 Common Monitors Monitor View for the selected target system

- A *threshold* for a numeric monitor is a high or low limit that you do not want the monitored system resource to exceed. For both the high threshold and the low threshold, you have the option of specifying a warning value and a critical value. For example, a monitor that measures the percentage of used space on a disk drive might have a warning value of 80% and a critical value of 90%.

Select the **CPU % Utilization** monitor, and click **Actions** → **Activate Threshold** to configure a threshold (see Figure 7-44).




Figure 7-44 Clicking Activate Threshold in Actions menu for selected CPU % Utilization monitor

- When you activate a threshold on a numeric monitor, you get the window shown in Figure 7-45. Activating a threshold includes setting a number of options. Choose whether to generate an event when the threshold is exceeded and determine the amount of time that the threshold waits before it resends the information. Define Critical and Warning threshold values. When the monitored resource exceeds the specified value for any threshold limit, the monitor displays the appropriate icon for a warning or critical notification.

Configure the threshold to trigger an event in the case of high CPU Utilization on the selected Windows system. Select **Critical** and **Warning** under the “Monitor values that are too high” section and configure the values as shown in Figure 7-45. Set the Minimum duration to **20 seconds**. This setting ensures triggering an event if the CPU Utilization value exceeds the threshold for over 20 seconds. Click **OK**.




Figure 7-45 Threshold window

The activated Threshold is now visible on the CPU % Utilization monitor, as shown in Figure 7-46.

Select	Name	Monitor Name	Moni...	Thres...	Current	Warning	Cr
<input checked="" type="checkbox"/>	PF-Windows1	CPU % Utilization	Individual	Activated	1%	>= 85.0	>=
<input type="checkbox"/>	PF-Windows1	Disk 0 Workload	Individual		108119.8 bytes/sec		
<input type="checkbox"/>	PF-Windows1	Disk Space Rem...	Individual		26504.3 Megabytes Free		
<input type="checkbox"/>	PF-Windows1	Disk Space Used	Individual		14353.7 Megabytes Used		
<input type="checkbox"/>	PF-Windows1	IP Packets Rece...	Individual		0 Packets/sec		
<input type="checkbox"/>	PF-Windows1	IP Packets Rece...	Individual		9 Packets/sec		
<input type="checkbox"/>	PF-Windows1	IP Packets Sent...	Individual		2 Packets/sec		
<input type="checkbox"/>	PF-Windows1	Locked Memory	Individual		44.1 Megabytes		
<input type="checkbox"/>	PF-Windows1	Memory Usage	Individual		1360.8 Megabytes		
<input type="checkbox"/>	PF-Windows1	Primary file syst...	Individual		35%		
<input type="checkbox"/>	PF-Windows1	Process Count	Individual		66 Processes		
<input type="checkbox"/>	PF-Windows1	TCP Connections	Individual		3 TCP Connections		
<input type="checkbox"/>	PF-Windows1	UDP Datagrams...	Individual		3 Packets/sec		
<input type="checkbox"/>	PF-Windows1	UDP Datagrams...	Individual		1 Packets/sec		

Figure 7-46 Monitoring View window with activated threshold for CPU % Utilization monitor

In the example, the system's CPU is used at 100%. As shown in Figure 7-47, the Threshold Status turns to Critical. This change happens 20 seconds after the high CPU utilization started. This delay is because the Critical threshold value that we specified earlier is 95 and the minimum duration is 20 seconds.

Select	Name	Monitor Name	Moni...	Threshold Status	Current	Warning	Cr
<input checked="" type="checkbox"/>	PF-Windows1	CPU % Utilization	Individual	Critical	100%	>= 85.0	>=
<input type="checkbox"/>	PF-Windows1	Disk 0 Workload	Individual		108119.8 bytes/sec		
<input type="checkbox"/>	PF-Windows1	Disk Space Rem...	Individual		26504.2 Megabyt...		
<input type="checkbox"/>	PF-Windows1	Disk Space Used	Individual		14353.8 Megabyt...		
<input type="checkbox"/>	PF-Windows1	IP Packets Rece...	Individual		0 Packets/sec		
<input type="checkbox"/>	PF-Windows1	IP Packets Rece...	Individual		8 Packets/sec		
<input type="checkbox"/>	PF-Windows1	IP Packets Sent...	Individual		2 Packets/sec		
<input type="checkbox"/>	PF-Windows1	Locked Memory	Individual		44.1 Megabytes		
<input type="checkbox"/>	PF-Windows1	Memory Usage	Individual		1377.7 Megabytes		
<input type="checkbox"/>	PF-Windows1	Primary file syst...	Individual		35%		
<input type="checkbox"/>	PF-Windows1	Process Count	Individual		72 Processes		
<input type="checkbox"/>	PF-Windows1	TCP Connections	Individual		3 TCP Connections		
<input type="checkbox"/>	PF-Windows1	UDP Datagrams...	Individual		2 Packets/sec		
<input type="checkbox"/>	PF-Windows1	UDP Datagrams...	Individual		1 Packets/sec		

Figure 7-47 Monitor View window with Critical Threshold Status for CPU % Utilization monitor

The Critical problem is also displayed in the Active Status window.

In the example, you want to be notified by email in the event of high CPU utilization on your Windows system. Return to the Monitor View for the selected target system. Select the **CPU % Utilization** monitor, and click **Actions** → **Create Event Automation Plan**, as shown in Figure 7-48.




Figure 7-48 Use the Actions menu to create an Event Automation Plan for the selected monitor

8. Click **Next** in the Welcome window.
9. Enter a name and description for the event automation plan, as shown in Figure 7-49.




Figure 7-49 Create Event Automation Plan Wizard Name and Description window

10. Select threshold levels on which to filter. In this case, select **High - Critical** and **High - Warning**, as shown in Figure 7-50. If you want to receive an email when the Critical condition is resolved, select **Threshold resolved**, as well. Click **Next**.

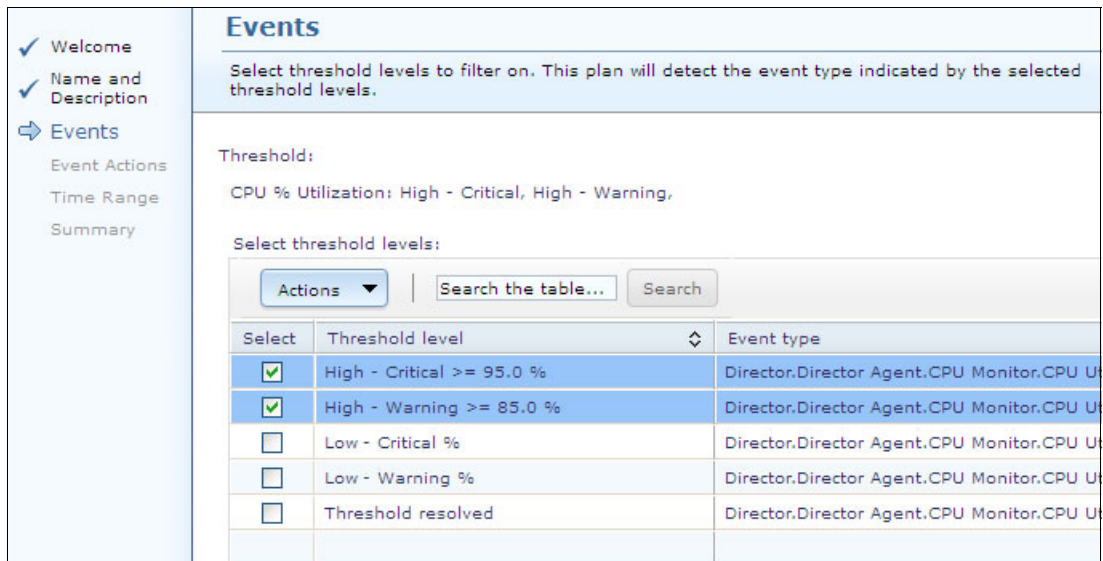


Figure 7-50 Create Event Automation Plan Wizard Events window

11. Select the **Add to the event log** and **Electronic Service Notification** event actions as shown in Figure 7-51. You created the Electronic Service Notification action in 7.4, "Automating tasks with event automation plans" on page 271. It sends an email notification with event information.

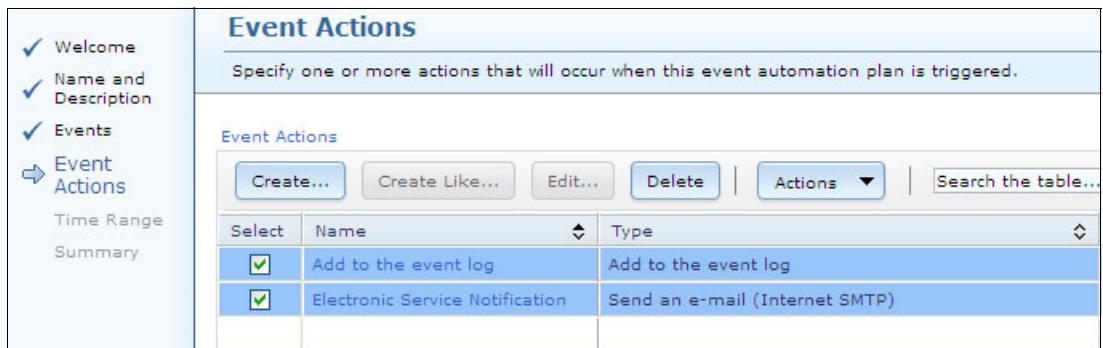


Figure 7-51 Create Event Automation Plan Wizard Event Actions window

12. In the example, the Windows server runs heavy CPU load operations during the weekend, so high CPU utilization is expected and considered normal. Therefore, you want to be notified for high CPU utilization only from Monday until Friday. In the Time Range window, select **Custom** and specify the time range constraints for the automation plan as shown in Figure 7-52. Click **Add** and then click **Next**.

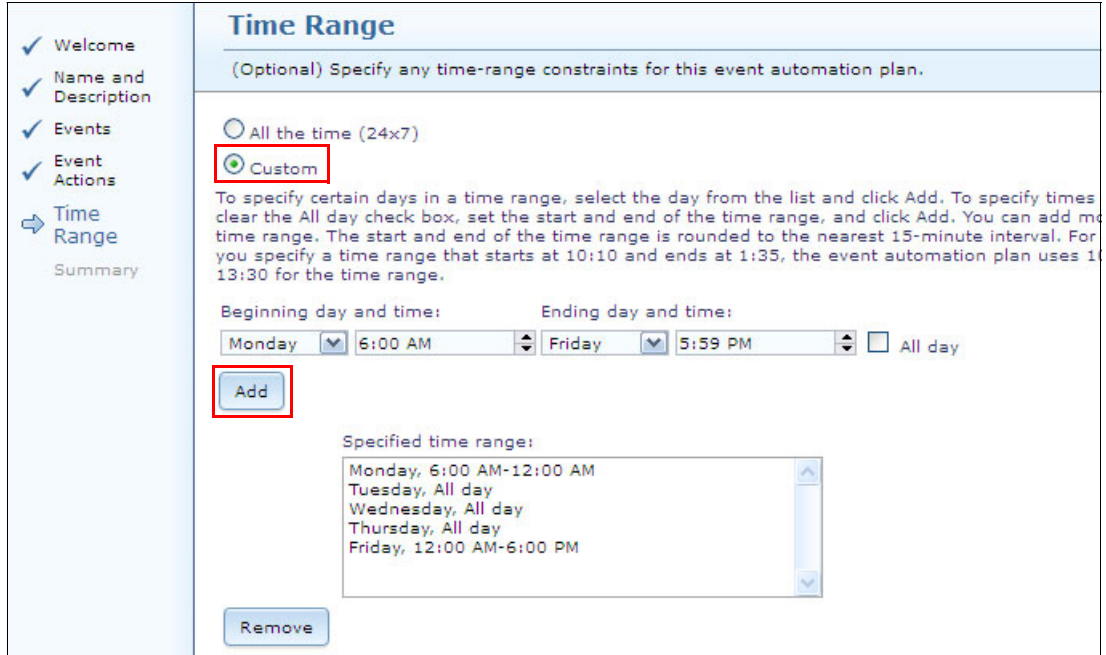


Figure 7-52 Create Event Automation Plan Wizard Time Range Custom window

13. Review the Summary window (Figure 7-53) and click **Finish**.

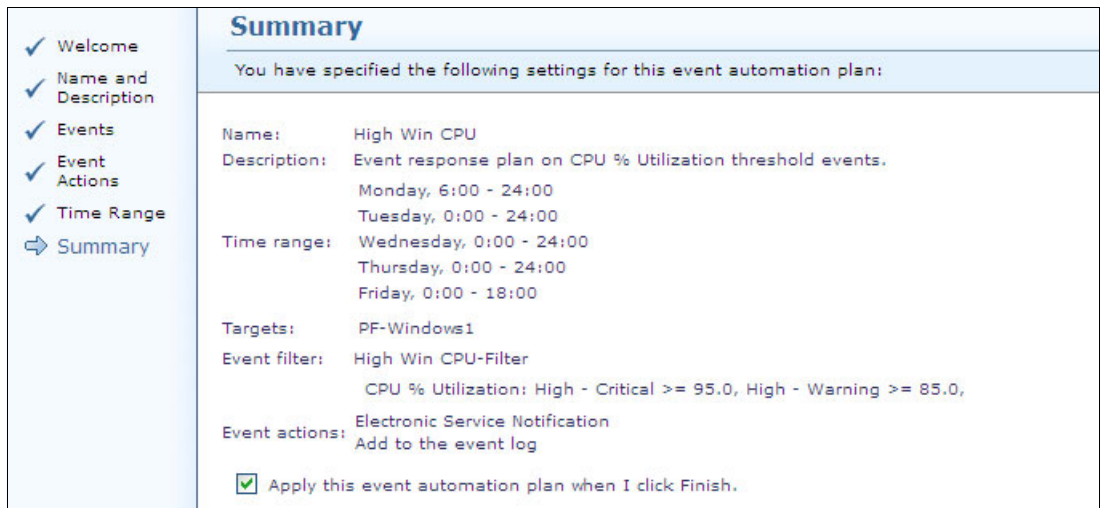


Figure 7-53 Create Event Automation Plan Wizard Summary window

14. In the FSM Explorer web interface, open the Event Automation Plans window (see 7.4, “Automating tasks with event automation plans” on page 271) to see the automation plan that you just created (Figure 7-54).

Select	Name	Targets	Status	Time range
<input type="checkbox"/>	High Win CPU	PF-Windows1	Active	Monday, 6:00-24:00, T
<input type="checkbox"/>	Hot air	All Systems	Active	All the time (24x7)

Figure 7-54 Event Automation Plans window

The automation plan is active. Any time that the CPU % Utilization threshold for your Windows system turns to the Critical or Warning state, you are notified. This notification is sent to the email address that is specified in the Electronic Service Notification event action settings.

15. Now, decrease the generated CPU load on your Windows system. Open the Event Log window. Enter PF-Windows1 in the Search field to filter only events that are related to your Windows system. Notice the Resolution events in Figure 7-55 that are displayed after you decrease the CPU load.

Select	Event Text	Source	Severity	Category	Date and Time
<input type="checkbox"/>	Monitor 'Individual threshold for CPU % ...	PF-Windows1	Information	Resolution	Jun 18, 20
<input type="checkbox"/>	Monitor 'Individual threshold for CPU % ...	PF-Windows1	Critical	Alert	Jun 18, 20
<input type="checkbox"/>	Monitor 'Individual threshold for CPU % ...	PF-Windows1	Information	Resolution	Jun 18, 20

Figure 7-55 Event Log window

16. Open the Thresholds window from the FSM Explorer (**Monitor** → **Thresholds**). Use this window (see Figure 7-56) as a quick way to view and manage thresholds that are set for the monitors on your resources.

Select	Name	Monitor N...	Monitor Ty...	Threshold...	Warning	Critical
<input type="checkbox"/>	PF-Windows1	CPU % Utilizat...	Individual	Activated	>= 85.0	>= 95.0

Figure 7-56 Thresholds window

7.8 Remote management

With the Remote Control application in IBM Flex System Manager management software, you can manage X-Architecture compute nodes as though you were at a local console. This configuration is useful in many cases, especially when you do not have any other way of accessing your system remotely.

Requirement: Remote Control is a Java Web Start application that requires the IBM or Oracle/Sun Java runtime environment (JRE) plug-in, Version 6.0, update 18 or later. You must obtain and install the JRE plug-in before you can use the Remote Control application.

To enable remote management, perform these steps:

1. Return to the Chassis Map and select a compute node. The Common Actions navigation bar gives you quick access to some of the most common tasks performed on the system, as shown in Figure 7-57.

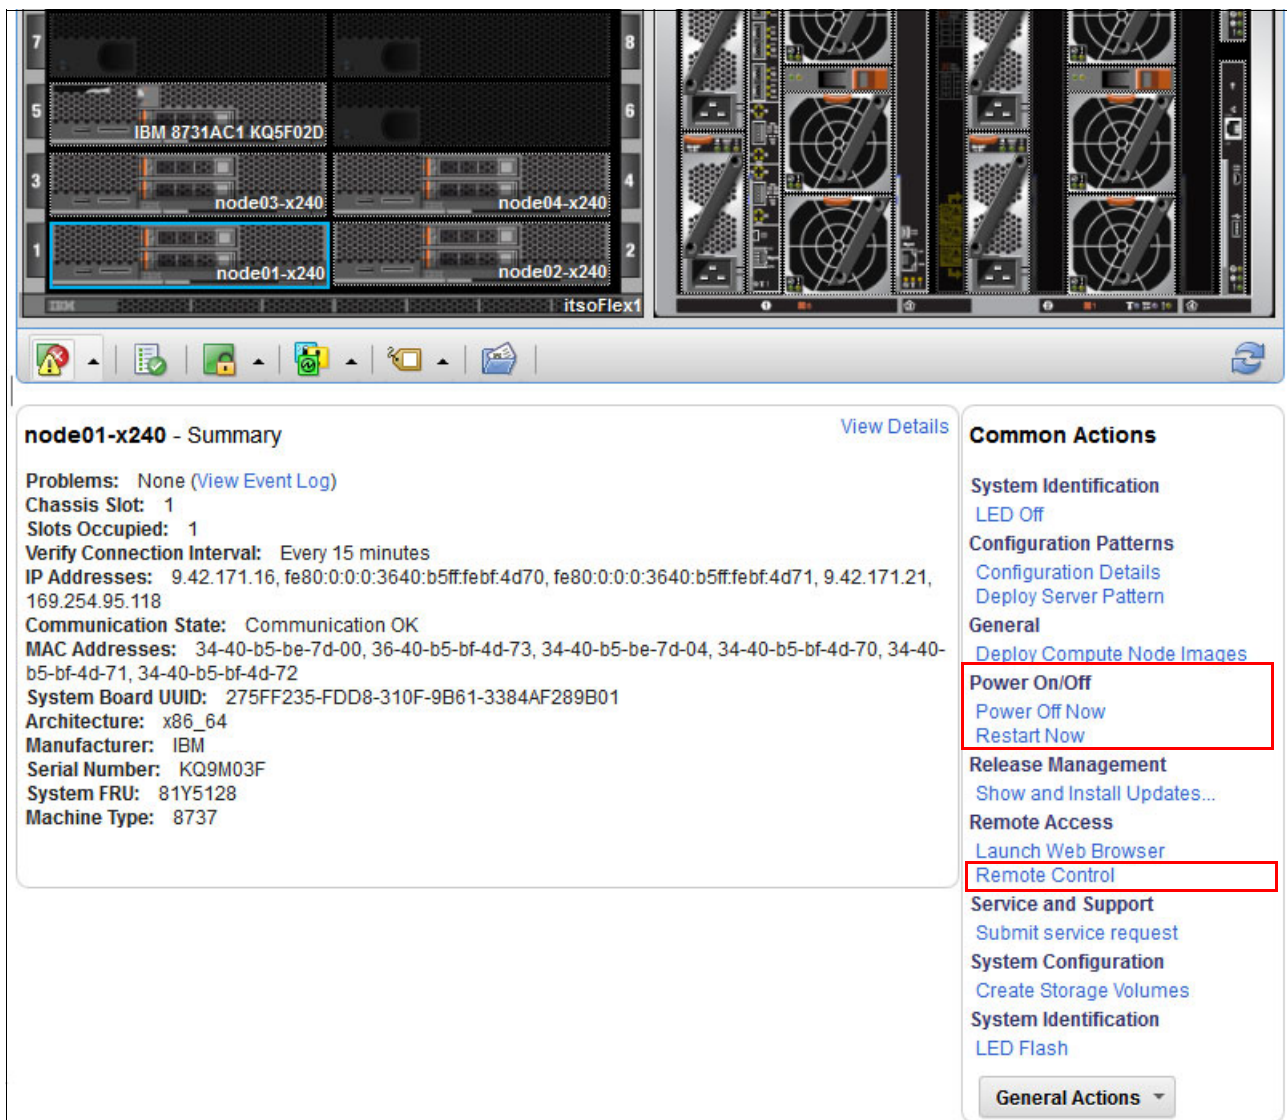


Figure 7-57 Chassis Map with selected compute node and available Common Actions menu

If you need to perform hardware maintenance, you can power off and power on or restart the compute node.

2. Under Common Actions, click the **Remote Control** task to start the Remote Control Java application. When you start the Remote Control application from FSM, you are prompted to save a shortcut to the application on your system. You can then use this shortcut to open the Remote Control session to the specified compute node without using the FSM user interface. However, your computer must have access to FSM because the application validates the user ID with the management software user registry.
3. Select one of the connection types, as shown in Figure 7-58:
 - If you need to give exclusive access to the remote control session, select **Single-user** access. All other Remote Control sessions to the selected compute node are blocked until you disconnect from the selected compute node.
 - Select **Multi-user** to allow multiple users to connect to the remote session.

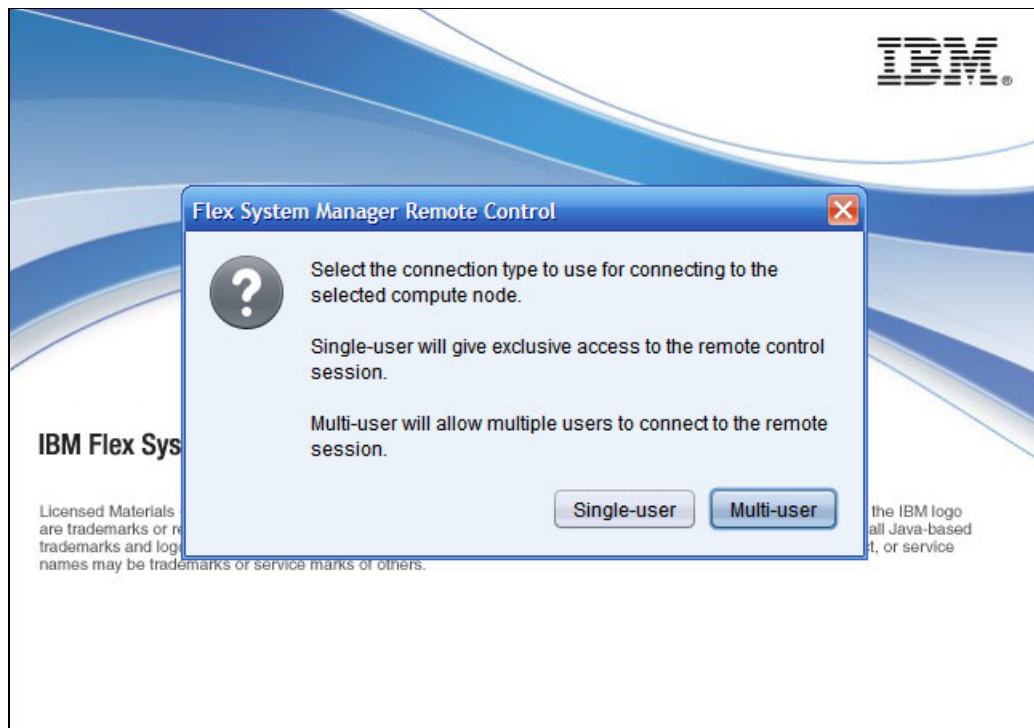


Figure 7-58 FSM Remote Control connection type window

4. The console window opens as shown in Figure 7-59. You can now use the mouse and keyboard to operate with the server as though you were at the local console. Click the arrow on top of the console window to open the Remote Control toolbar.

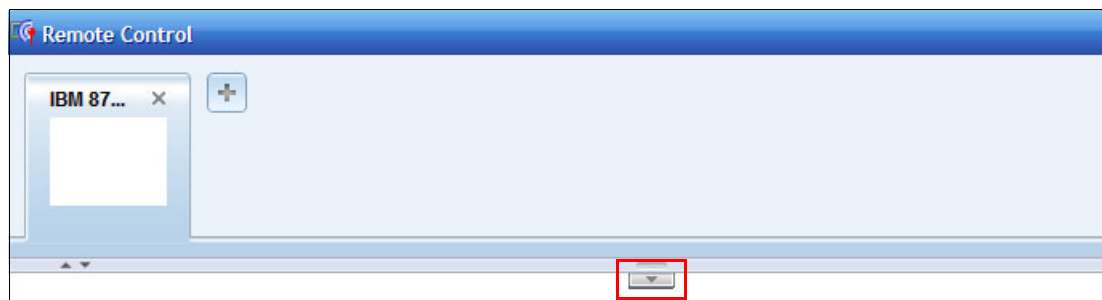


Figure 7-59 Remote Control console window

5. The Remote Control toolbar (Figure 7-60) offers functions, such as screen capture, compute node power controls, defining custom key sequences, sticky keys, mounting media, and Remote Control preferences. Click the **Mount Media** icon and select **Mount Remote Media**.

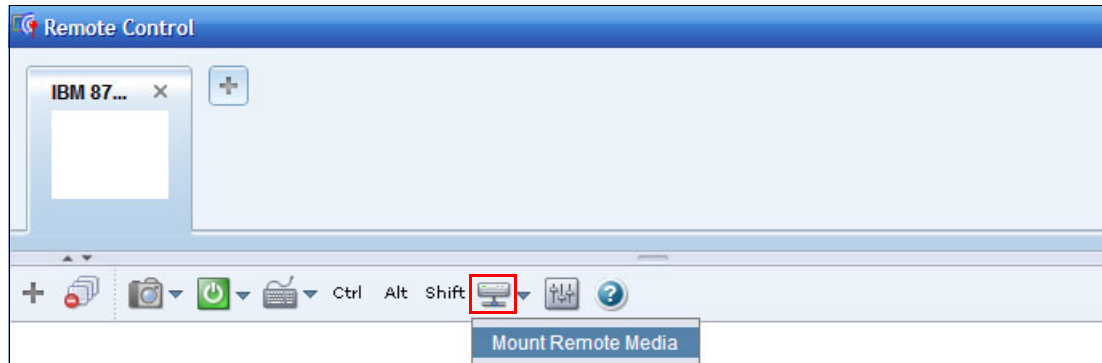


Figure 7-60 Remote Control toolbar

6. You can mount a local CD-ROM drive, upload an image to the integrated management module (IMM) (up to 50 MB), or select an image. Highlight **Select an image** and click **Add** as shown in Figure 7-61.

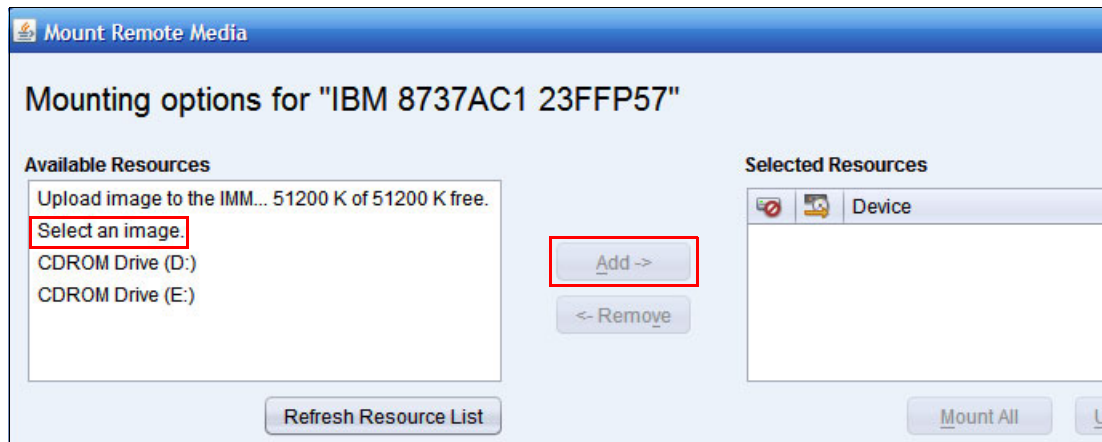


Figure 7-61 Mount Remote Media window

7. You are prompted to select an image from your local computer. Select the image that you want to mount and click **Open**, as shown in Figure 7-62.




Figure 7-62 Open dialog box

8. The image file path is displayed in the list of devices available for mounting, as shown in Figure 7-63. Click **Mount All**.




Figure 7-63 Mount Remote Media window that shows Selected Resources

9. The image is now mounted as shown in Figure 7-64. If you need to unmount it at any time, click **Unmount All**. Leave the image mounted and click **Close**.

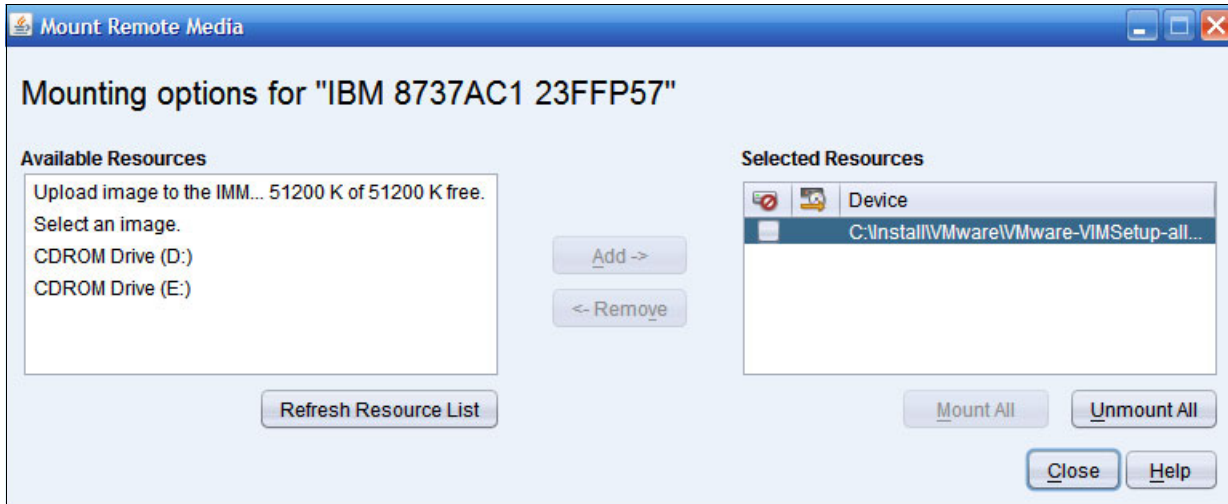


Figure 7-64 Mount Remote Media window that shows mounted resources

10. Open **My Computer** in Windows to ensure that you can access the mounted image as a CD drive.

You can also use the power control options from the Remote Control toolbar. You can choose to restart/power off the compute node immediately (**Restart Immediately/Hard Power Off**) or gracefully shut down the OS before restarting or powering off the compute node (**Shut Down OS and Restart/Shut Down OS and Power Off**), as shown in Figure 7-65.

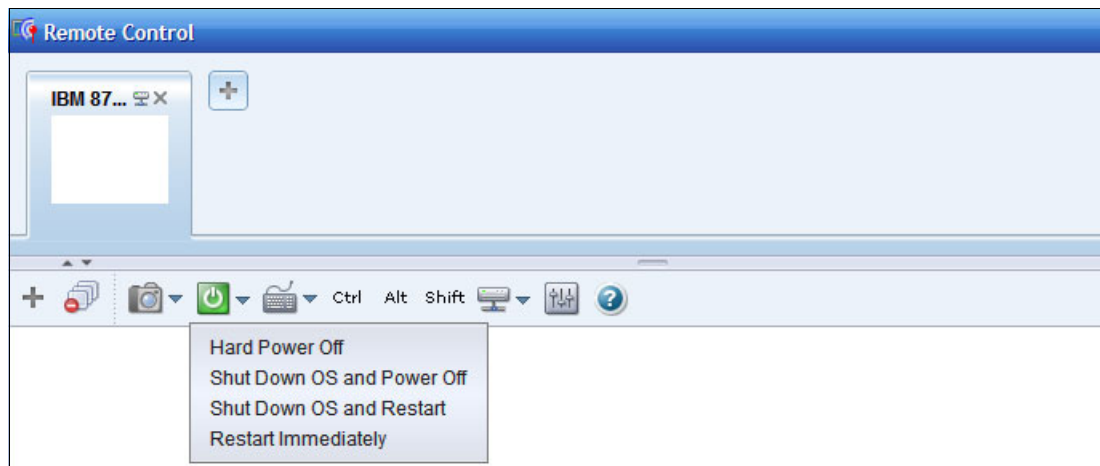


Figure 7-65 Power Controls menu on the Remote Control toolbar

11. You can quickly open a Remote Control session to another compute node by clicking the plus sign (+) icon next to your Remote Control session thumbnail. See Figure 7-66.




Figure 7-66 Window showing available compute nodes for Remote Control

12. Click **Cancel** and close the Remote Console.

The user ID that is used to start the Remote Control application must be a valid user ID that is defined in the FSM user registry. The user ID must also have sufficient user authority to access and manage a compute node. You can assign the role of SMAdministrator to the user ID. Or, you can define a custom role for compute node access and management, and assign that role to the user ID.

Restriction: Remote Control in FSM is available only for X-Architecture compute nodes. You can establish a terminal console session to any virtual server on a Power Systems compute node. For more information, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8731.doc/using_remote_access.html



IBM Fabric Manager

IBM Fabric Manager enables you to quickly replace and recover compute nodes in your IBM Flex System or to preconfigure your network and storage infrastructure in the IBM Flex System chassis.

The following topics are covered:

- ▶ 8.1, “IBM Fabric Manager overview” on page 306
- ▶ 8.2, “Starting the IBM Fabric Manager interface” on page 306
- ▶ 8.3, “Adding devices” on page 307
- ▶ 8.4, “Adding a device pool” on page 310
- ▶ 8.5, “Adding a boot target template” on page 310
- ▶ 8.6, “Adding a profile” on page 312
- ▶ 8.7, “Profile deployment” on page 313
- ▶ 8.8, “Pushing a deployment” on page 313
- ▶ 8.9, “Verifying an IBM Fabric Manager deployment” on page 314
- ▶ 8.10, “Adding and starting a monitor” on page 316

8.1 IBM Fabric Manager overview

IBM Fabric Manager enables you to use virtual Ethernet Media Access Control (MAC), Fibre Channel worldwide name (WWN), and serial-attached SCSI (SAS) WWN addresses to preconfigure your network and storage environments before any compute nodes are inserted into the chassis.

IBM Fabric Manager also monitors server health and can automatically replace a failed server from a designated pool of servers without user intervention. After a failure alert occurs, IBM Fabric Manager will attempt to power off the failing node. IBM Fabric Manager will then apply the failed node's virtualized addresses and boot parameters to the next node in the standby pool and power on that standby server.

In order to take full advantage of IBM Fabric Manager, you must set up your server environment to boot from SAN.

Note: IBM Fabric Manager is not supported with SAS adapters on Power compute nodes. Management of Fibre Channel boot targets is also not supported. You must pre-assign Fibre Channel boot targets in the Software Management Services (SMS) menus on the Power compute node.

8.2 Starting the IBM Fabric Manager interface

Perform the following steps to start the IBM Fabric Manager user interface from the IBM Flex System Manager:

1. Log in to the IBM Flex System Manager, click the **Home** tab, and select the **Applications** tab, as shown in Figure 8-1.



Figure 8-1 Flex System Manager: Applications tab

2. Select **IBM Fabric Manager**. This opens a new browser tab and presents the IBM Fabric Manager interface login window, as shown in Figure 8-2. Enter your user name and password.

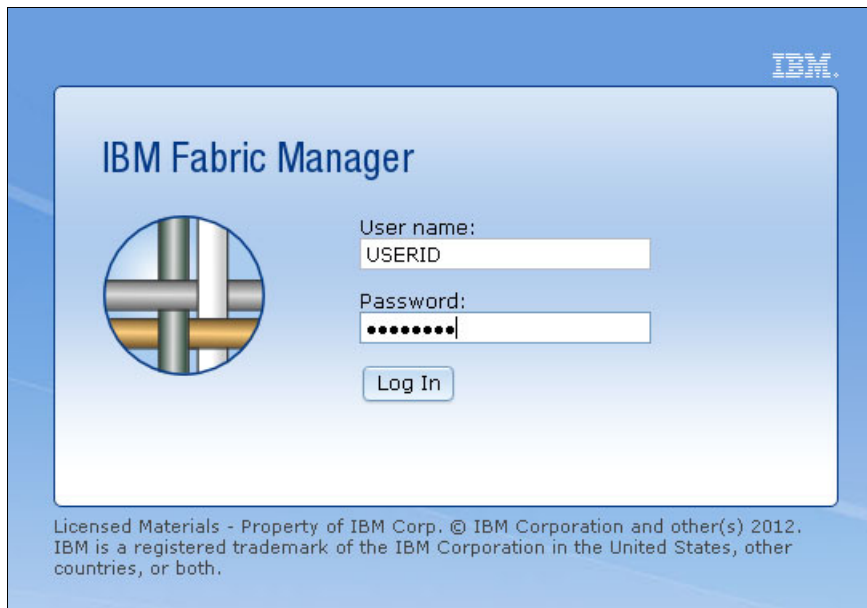


Figure 8-2 IBM Fabric Manager login window

Default credentials: IBM Fabric Manager has a default user name of USERID and password of PASSWORD (with a zero, not an O). After you log in to IBM Fabric Manager for the first time, you are required to change the password.

3. The IBM Fabric Manager main window opens.

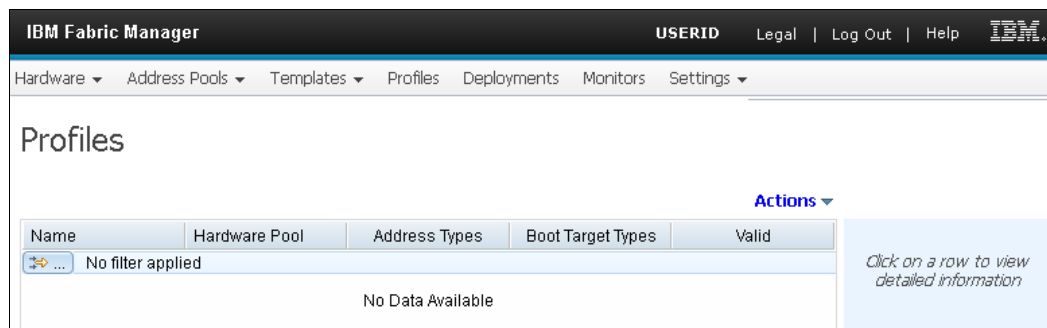


Figure 8-3 IBM Fabric Manager main window

8.3 Adding devices

This section describes how to add hardware devices from your Flex System chassis and define pools into which to group your compute nodes.

Before adding pools, you must add each chassis with hardware that you want to manage with IBM Fabric Manager.

Perform the following steps to add a chassis:

1. From the IBM Fabric Manager main window, click **Hardware** and select **Devices**.
2. Click the **Actions** menu and select **Add**, then select one of the following choices:

- Add a Single Chassis

Select this choice to add a single chassis by its IP address.

- Add Multiple Chassis by Range

Select this choice to add multiple chassis in a range of IP addresses.

- Add Chassis from file or URL

Select this choice to import IP addresses from a file or URL. There is a maximum of 100 IP addresses that can be imported. Any additional IP addresses will be ignored without warning. The file must be a text file with one valid IPv4 or IPv6 address per line.

For the example in this book, add only a single chassis and click **Next**. Figure 8-4 shows the Add Individual Hardware window that you use to add a single chassis.

The screenshot shows a dialog box titled "Add Individual Hardware" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- IP/Hostname: [Text input field]
- Username: [Text input field]
- Password: [Text input field]
- SNMPv3 receiver IP: [Text input field]
- SNMPv3 password: [Text input field]
- Backup SNMPv3 receiver IP: [Text input field]
- Backup SNMPv3 password: [Text input field]
- TCP command mode**
- Order: [Dropdown menu showing "Secure then unsecure"]
- Secure port: [Text input field containing "6091"]
- Unsecure port: [Text input field containing "6090"]
- Generic chassis
- Buttons: [Back] [Cancel] [Start]

Figure 8-4 Adding a single chassis

Enter the following information:

- IP/Hostname

Enter the IP address or Domain Name System (DNS) host name of the Chassis Management Module (CMM) for the chassis you are adding.

- Username and Password

Enter the user ID and password for the CMM for the chassis you are adding.

- SNMPv3 receiver IP and SNMPv3 password

These fields are required fields to monitor the chassis for failover events. The Flex System Manager IP address must be entered here. The password must be at least eight characters long. It does not have to be the same as any other SNMPv3 password for the environment. The Backup SNMPv3 fields are left blank.

- Generic chassis

Select this option to add the chassis but not log in or gather inventory.

All other listed fields can be left at their defaults. When you have completed the required fields, click **Start** and the Hardware Discovery Progress dialog box (see Figure 8-5) lists all chassis that it discovers and their progress. The time for this task to complete depends upon the number of chassis being discovered. After all the chassis are discovered, click **OK** to complete this process.

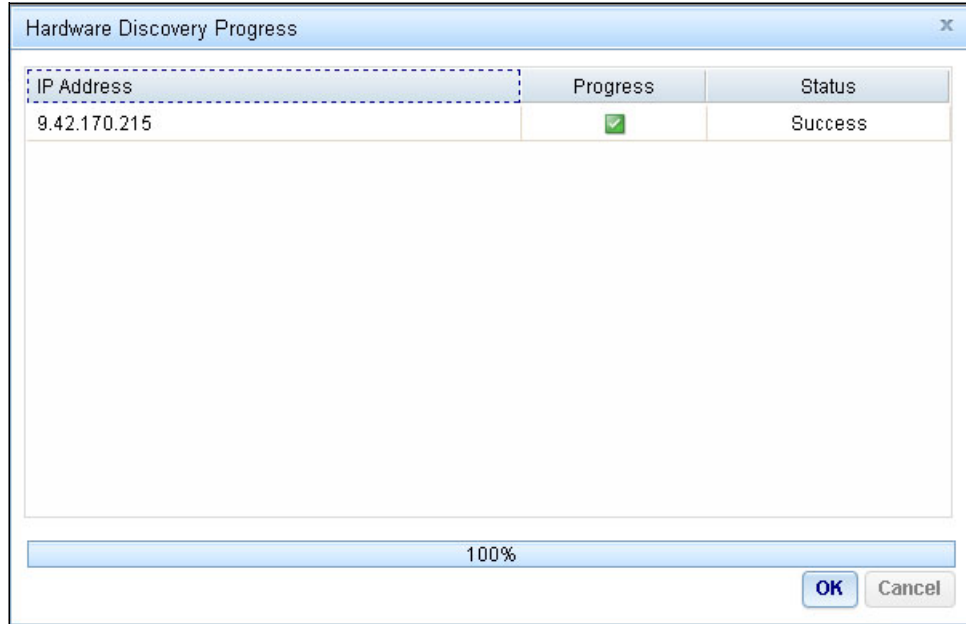


Figure 8-5 Hardware Discovery Progress window

After your chassis is added, the chassis appears on the Hardware Devices section of the IBM Fabric Manager interface as shown in Figure 8-6.

Hardware Devices

Actions ▾

Name	Type	Bay	IP Address
<input type="checkbox"/> SN#Y011BG24H0BB	Flex System	...	9.42.170.215
<input type="checkbox"/> [1] node01-x240	8737-AC1	1	...
<input type="checkbox"/> [2] node02-x240	8737-AC1	2	...
<input type="checkbox"/> [3] node03-x240	8737-AC1	3	...
<input type="checkbox"/> [4] node04-x240	8737-AC1	4	...
<input type="checkbox"/> [5] node05-FSM	8731-AC1	5	...
<input type="checkbox"/> [10] node06-p270	7954-24X	10	...
[1] EN4093 10Gb Ethernet Switc...	...	1	9.42.171.8
[3] FC3171 8Gb SAN Switch	...	3	9.42.171.9

Figure 8-6 IBM Fabric Manager Hardware Devices with chassis added

8.4 Adding a device pool

A *device pool* groups compute nodes together to define a standby pool of hardware that can be used to quickly recover from a failed compute node. All compute nodes used in a hardware pool must be of the same size (single-wide or double-wide) and configuration.

Perform the following steps to add a device pool:

1. From the IBM Fabric Manager Interface, click **Hardware** and then select **Pools**.
2. Click the **Actions** menu and select **Add**. Figure 8-7 shows the Add a New Pool dialog box.

Name	Type	Add / Remove
[-] SN#Y011BG24H0BB	Flex System	<input type="checkbox"/>
bay 1	...	<input type="checkbox"/>
bay 2	...	<input type="checkbox"/>
bay 3	...	<input checked="" type="checkbox"/>
bay 4	...	<input checked="" type="checkbox"/>
bay 5	...	<input type="checkbox"/>
bay 6	...	<input type="checkbox"/>
bay 7	...	<input type="checkbox"/>
bay 8	...	<input type="checkbox"/>

Figure 8-7 Adding a new hardware pool

Enter the following information:

- Pool name

Enter the name of the hardware pool.

- Select

You can select the **All Slots** check box to enable management of all the compute nodes or slots in the chassis, even if there is not a compute node in every slot. To choose individual slots or compute nodes, expand the list of slots by clicking to the left of the chassis name. Select the check box for each slot to add to the pool.

After you enter a name for your pool and select your slots, click **Save** to create the new pool. The pool that you just created now shows in the list of pools on the Hardware Pools page.

8.5 Adding a boot target template

This section describes how to add a boot target template to IBM Fabric Manager.

A *boot target template* enables IBM Fabric Manager to quickly recover from compute node hardware failure by reconfiguring a standby server in the hardware pool to boot from the primary server's disk.

Boot order: Although IBM Fabric Manager will configure the boot list on the SAS or FC adapter, it does not configure the Unified Extensible Firmware Interface (UEFI) boot list on System x compute nodes. You must manually configure your compute nodes to boot from the FC or SAS adapters.

Follow these steps to add a new boot template to IBM Fabric Manager:

1. From the IBM Fabric Manager interface, click **Template** and select **Boot Target**.
2. Click **Actions** and select **New**. Figure 8-8 shows the Create Boot Targets dialog box.

Order	Address	LUN
1	50:05:07:68:05:0C:03:70	0
2	50:05:07:68:05:0C:03:71	0

Order	Address	LUN
1	<Enter address>	0
2	<Enter address>	0

Figure 8-8 Creating a boot target template

Configure the following settings:

- Name
Type the name of your new boot target pool.
- Type
You can define either a Fibre Channel or SAS boot target. For this example, we will boot from SAN and will not use SAS.
- Primary boot targets
Enter the worldwide name (WWN) addresses of the target SAN storage system and logical unit number (LUN) ID of the target device to use to boot the server.
- Secondary boot targets
Enter the WWN addresses of the target SAN storage system and the LUN ID of the secondary target device, if applicable, to use to boot the server.

After completing the required fields, click **Save** to finish creating the boot target template. After it is completed, the new template shows in the list of defined templates.

8.6 Adding a profile

Profiles tie hardware pools and templates together into a single entity that can be deployed to a chassis.

Perform the following steps to add a profile:

1. From the IBM Fabric Manager interface, click **Profiles**.
2. Click the **Actions** menu and select **Add**. Figure 8-9 shows the Create Profile window.

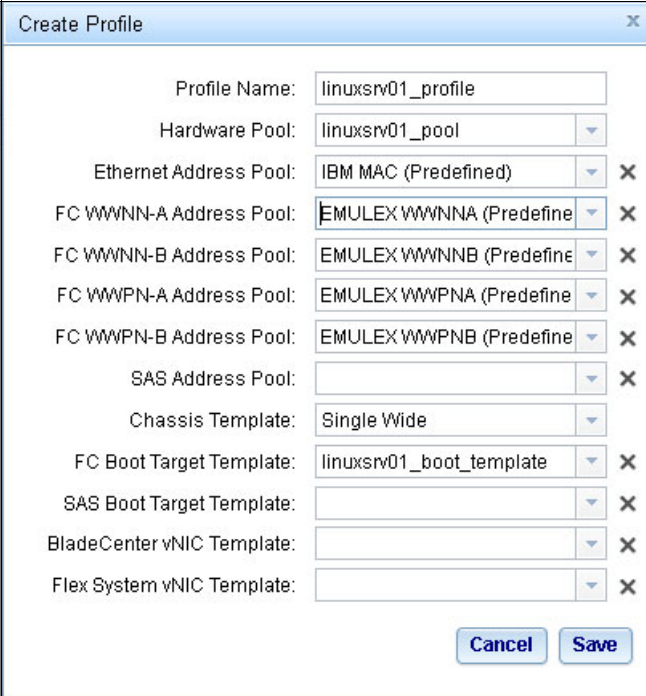


Figure 8-9 Creating a profile

Enter the following parameters:

- Profile Name

Enter the name of the profile.

- Hardware Pool

Select a previously defined hardware pool from the list.

- Address pools and templates

Profiles use Ethernet, FC, and SAS addresses and templates to configure a standby server with a failed server's virtual addresses so that a quick recovery from hardware failure can be performed automatically. There are predefined address pools and templates that can be used for simpler and faster configuration. Select the options from each list that best fit your environment and needs for this particular profile.

- Chassis Template

Select **Single Wide** or **Double Wide**. This must match the size of the compute nodes in the pool selected in the Hardware Pool field.

- FC (or SAS, if applicable) Boot Target Template

Select the previously created boot target template.

Click **Save** to finish creating the profile. The profile now appears in the list of defined profiles in the IBM Fabric Manager interface.

8.7 Profile deployment

Before IBM Fabric Manager can manage the hardware defined by a profile, it must be deployed.

Deployment: Deployment of a profile must be performed from the *Profiles* window of IBM Fabric Manager and not from the *Deployments* window.

Perform the following steps to deploy a profile:

1. From the IBM Fabric Manager interface, click **Profiles**.
2. Click the name of a profile to select it.
3. Click the **Actions** menu and select **Deploy**, as shown in Figure 8-10.



Figure 8-10 Deploying a profile

4. Type a name for the deployment in the Deploy Profile dialog box shown in Figure 8-11.




Figure 8-11 Entering a name for profile deployment

5. Click **Deploy** to deploy the profile.
6. Click **Close** on the Success dialog box to finish.

The deployed profile now is listed in the Deployments window of the IBM Fabric Manager interface. Click **Deployments** to view the list.

8.8 Pushing a deployment

The newly created deployment must now be edited and pushed out to the chassis and hardware.

Perform the following steps to push a deployment in IBM Fabric Manager:

1. Click **Deployments** in the IBM Fabric Manager interface.
2. Select the name of the deployment that you want to push.
3. Click the **Actions** menu and select **Push**, as shown in Figure 8-12.



Figure 8-12 Pushing a deployment

4. The Pre-deployment Options window opens, as shown in Figure 8-13.




Figure 8-13 Options dialog box before pushing the deployment

The Pre-deployment Options dialog box allows you to select which compute nodes in the hardware pool receive the deployment being pushed. Click **Continue**.

5. A confirmation box asking whether you want to push this deployment appears. Click **OK**.
6. A progress window showing the status of all chassis and node deployment and overall progress opens. Click **OK** when it is complete.

8.9 Verifying an IBM Fabric Manager deployment

This section describes the necessary steps to verify that a deployment was successfully pushed to the hardware in the profile defined in the hardware pool used in the deployment.

Perform the following steps to verify deployment:

1. From the IBM Fabric Manager interface, click **Hardware** and select **Devices**.

- Expand the hardware list by clicking the plus sign (+) to the left of the chassis name or IP address.
- Verify that the IBM Fabric Manager (IFM) Mode column entries for the nodes used in this deployment are ENABLED, as shown in Figure 8-14.

Hardware Devices											Actions ▾
Name	Type	Bay ▲	IP Address	License	Inventoried	Power	Network	IFM Mode	Source	IFM Status	
SN#Y011BG24H0BB	Flex System	...	9.42.170.215	✓	✓	...	🌿	
[1] node01-x240	8737-AC1	1	🟢	...	DISABLED	CMM	not available	
[2] node02-x240	8737-AC1	2	🟢	...	DISABLED	CMM	not available	
[3] node03-x240	8737-AC1	3	🟡	...	ENABLED	CMM	not available	
[4] node04-x240	8737-AC1	4	🟡	...	ENABLED	CMM	not available	

Figure 8-14 Verify IFM mode

- For nodes that are standby nodes, select each one, click the **Actions** menu, then select **Toggle IFM Mode**. The IFM Mode of each standby server must be DISABLED.
- Power on the primary node or nodes in the profile that you deployed. Select each node and click the **Actions** menu and select **Toggle Power**. Click **OK** to confirm that you want to power on the node.
- An information dialog box appears stating that it can take some time for the Devices window to reflect changes in status and recommends waiting 1 - 3 minutes before refreshing the window. Click **OK** to continue.
- When a primary node has booted and has been acquired by IBM Fabric Manager, the IFM Status column shows a green box with a white check mark inside it and a Details link, as shown in Figure 8-15. The details column to the right of the hardware table also shows an IFM Status of NORMAL.

Hardware Devices											Actions ▾
Name	Type	Bay ▲	IP Address	License	Inventoried	Power	Network	IFM Mode	Source	IFM Status	
SN#Y011BG24H0BB	Flex System	...	9.42.170.215	✓	✓	...	🌿	
[1] node01-x240	8737-AC1	1	🟢	...	DISABLED	CMM	not available	
[2] node02-x240	8737-AC1	2	🟢	...	DISABLED	CMM	not available	
[3] node03-x240	8737-AC1	3	🟢	...	ENABLED	CMM	<div style="border: 1px solid green; padding: 2px; display: inline-block;"> ✓ Details... </div>	
[4] node04-x240	8737-AC1	4	🟡	...	DISABLED	CMM	not available	
[5] node05-FSM	8731-AC1	5	🟢	...	DISABLED	CMM	not available	
[10] node06-p270	7954-24X	10	🟢	...	DISABLED	CMM	not available	
[1] EN4093 10Gb Ethernet Switc...	...	1	9.42.171.8	🟢	
[3] FC3171 8Gb SAN Switch	...	3	9.42.171.9	🟢	

Name: [3] node03-x240

Device Type: Server

IFM Mode: ENABLED

IFM Status: NORMAL

IFM Applied: true

Power State: On

Figure 8-15 Verifying IFM status for a primary node

8. Click **Details**. The Review IFM Status window opens, as shown in Figure 8-16.

Offset	Mezz	Port	WWNN	WWPN	WWPN Boot Order	WWNN Boot Order	Status
0	1	1	2F:FE:00:00:C9:00:00:00	2F:FC:00:00:C9:00:00:00	1	1	n/a
0	1	2	2F:FF:00:00:C9:00:00:00	2F:FD:00:00:C9:00:00:00	1	1	n/a
0	2	1	2F:FE:00:00:C9:00:00:01	2F:FC:00:00:C9:00:00:01	1	1	Normal
0	2	2	2F:FF:00:00:C9:00:00:01	2F:FD:00:00:C9:00:00:01	1	1	Normal

Figure 8-16 Viewing the IFM status details

9. Review the addresses applied to the node by clicking each tab and verifying that the addresses are within the range of the defined pools used in the hardware pool with which the node is associated. To close this dialog box, click the small (x) in the upper-right corner.

8.10 Adding and starting a monitor

Failover monitors can be configured to watch for particular events that will then trigger an automatic failover to a standby node in a pool.

Adding a monitor

Perform the following steps to add a monitor to IBM Fabric Manager:

1. From the IBM Fabric Manager interface, click **Monitors**.
2. Click the **Actions** menu and select **Add**. Figure 8-17 shows the Create Monitor dialog box.

Monitor name: linuxsrv01_pool

Monitored pool: linuxsrv01_pool

Standby pool: linuxsrv01_pool

Failover settings:

Ignore model Ignore type

Ignore width Ignore power state

Ignore partition Failover VLAN

Triggering events:

Power off Removal

CPU failure Hard drive failure

Memory failure Communication error

No power Voltage warning

PFA

Cancel Save

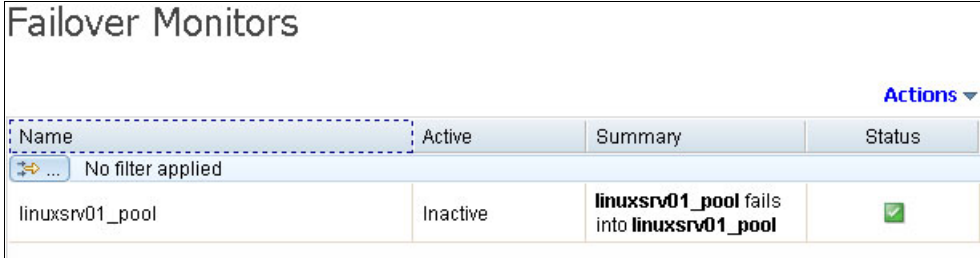
Figure 8-17 Creating a monitor with sample events selected

Enter the following parameters:

- Monitor name
Enter a name for this monitor.
- Monitored pool
Select the pool to monitor for events.
- Standby pool
Select the pool of standby hardware to which to fail over when a monitored event occurs.
- Failover settings
Select the settings for this failover event. For more detailed information about each of these settings, see the Failover Settings section of the Failover Monitors help that can be accessed by using the Help link in the upper-right corner of the IBM Fabric Manager interface.
- Triggering events
Select the events that will trigger a failover event. For more detailed information about each of these events, see the Triggering Events section of the Failover Monitors help that can be accessed by using the Help link in the upper-right corner of the IBM Fabric Manager interface.

3. Click **Save**.

The newly created monitor now appears in the list of defined monitors as shown in Figure 8-18.




Name	Active	Summary	Status
No filter applied			
linuxsrv01_pool	Inactive	linuxsrv01_pool fails into linuxsrv01_pool	

Figure 8-18 Viewing the list of defined monitors

Starting a failover monitor

Before IBM Fabric Manager can automatically fail over to standby hardware, you must start your defined monitors.

Perform the following steps to start a failover monitor:

1. To select a monitor, click the name of the monitor that you want to start.
2. Click the **Actions** menu and select **Start**.

When a failover monitor is started, it registers all the chassis associated with the pools used in the monitor. During chassis registration, you might see a status icon as shown in Figure 8-19 on page 318.

Name	Active	Summary	Status
No filter applied			
linuxsrv01_pool	Active	linuxsrv01_pool fails into linuxsrv01_pool	⚠

Figure 8-19 Chassis registration status

To refresh the monitor status, click **Monitors** in the IBM Fabric Manager interface. It can take several minutes for all of the chassis to be registered. After the chassis registration process is complete and the monitor starts, you see a green box with a white check mark in it as the status icon shown in Figure 8-20.

Name	Active	Summary	Status
No filter applied			
linuxsrv01_pool	Active	linuxsrv01_pool fails into linuxsrv01_pool	✔

Figure 8-20 Monitor startup successful

The monitor that you just created is now started and waiting for events to trigger a failover.



Managing the KVM environment with IBM Flex System Manager

This chapter addresses how to manage the Red Hat Enterprise Linux (RHEL) KVM-based virtualization environment with IBM Flex System Manager (FSM). It covers how to enable Kernel-based Virtual Machine (KVM) to be managed by FSM, and how to perform typical virtualization management tasks. These tasks include virtual machine lifecycle management, automation capabilities, and maintenance.

This chapter includes the following section:

- ▶ 9.1, “KVM management architecture” on page 320
- ▶ 9.2, “KVM platform agent installation” on page 320
- ▶ 9.3, “Image repository for KVM” on page 328
- ▶ 9.4, “Creating KVM storage system pools” on page 350
- ▶ 9.5, “Creating KVM network system pools” on page 352
- ▶ 9.6, “Creating KVM server system pools” on page 365
- ▶ 9.7, “Add host to an existing server system pool” on page 372
- ▶ 9.8, “Operating a KVM virtual infrastructure” on page 375

9.1 KVM management architecture

If you want to install an environment by using the Network File System (NFS) storage-based solution, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_kvm.html

The SAN storage configuration looks more complex than the NFS solution, but the block storage-based model offers better performance, and more functionality and flexibility.

Figure 9-1 illustrates a KVM virtualization environment with SAN storage.

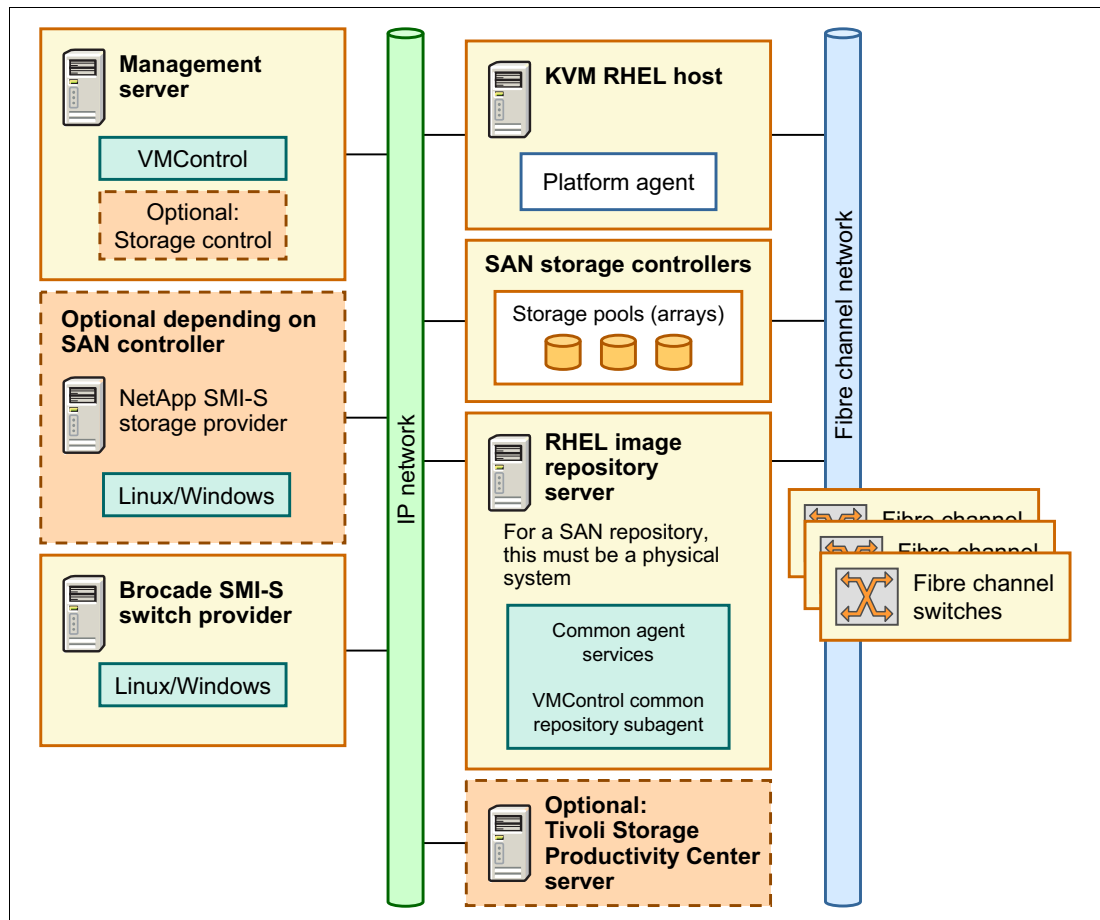


Figure 9-1 Select option to install KVM host

Before you start the implementation of the SAN storage-based solution, see “SAN storage-based model” on page 107.

9.2 KVM platform agent installation

To manage the KVM host from IBM Flex System Manager VMControl, you must manually install an agent that is called the KVM Platform Agent. You cannot use the Deploy Agent wizard on the IBM Flex System Manager to deploy this agent.

9.2.1 Preparation

Install and configure RHEL 6.2 on the compute node using the Virtualization Host role. RHEL installation is not described in this book.

To allow FSM communications to the KVM platform agent, perform the following steps:

1. Disable SUSE Linux (SELinux) as shown in Figure 9-2.

```
[root@KVM01 selinux]# pwd
/etc/selinux
[root@KVM01 selinux]# ls
config config.ori restorecond.conf restorecond_user.conf semanage.conf targeted
[root@KVM01 selinux]# cat config

# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#   enforcing - SELinux security policy is enforced.
#   permissive - SELinux prints warnings instead of enforcing.
#   disabled - No SELinux policy is loaded.
SELINUX=disabled
# SELINUXTYPE= can take one of these two values:
#   targeted - Targeted processes are protected,
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted

[root@KVM01 selinux]#
```

Figure 9-2 Disable SELinux

Tip: You can also configure SELinux in “permissive” mode if required for security reasons.

2. Configure iptables as shown in Figure 9-3.

```
[root@KVM01 selinux]# iptables -A INPUT -p tcp --dport 427 -j ACCEPT
[root@KVM01 selinux]# iptables -A INPUT -p udp --dport 427 -j ACCEPT
[root@KVM01 selinux]# iptables -A INPUT -p tcp --dport 22 -j ACCEPT
[root@KVM01 selinux]# iptables -A INPUT -p tcp --dport 15988 -j ACCEPT
[root@KVM01 selinux]# iptables -A INPUT -p tcp --dport 15989 -j ACCEPT
[root@KVM01 selinux]# service iptables save
iptables: Saving firewall rules to /etc/sysconfig/iptables:[ OK ]
[root@KVM01 selinux]#
```

Figure 9-3 Configuring iptables

You might face some issues during the inventory collection because of the iptables configuration on a KVM host. If so, remove the REJECT statement at the end of the INPUT chain in the filter table and reappend it to the end of the chain. You can also temporarily disable iptables for troubleshooting purposes.

3. Configure Yum on your system. For more information, see this website:

http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/sec-Configuring_Yum_and_Yum_Repositories.html

Tip: Configure Yum on your system because during the KVM Platform Agent (PA) installation, you might face Red Hat Package Manager (RPM) dependency requirements. You can save time if Yum is configured.

4. Check that the date on your KVM node is the same than the other hosts and the FSM as shown in Figure 9-4.

```
[root@KVM01 selinux]# date
Tue Jun 19 12:28:55 EDT 2012
[root@KVM01 selinux]# █
```

Figure 9-4 Check date on KVM host

Remember: If a Red Hat Package Manager (RPM) server is configured in your network, configure it on all your KVM hosts and on the FSM, as well.

9.2.2 KVM Platform Agent installation

To install KVM Platform Agent, perform the following steps:

1. Remove some packages by using the following command:

```
yum -y erase tog-pegasus libcmptutil libvirt-cim sblim-cmpi-nfsv3
sblim-cmpi-fsvo1 sblim-gather-provider sblim-gather sblim-cmpi-base openslp
```

2. Download KVM Platform Agent from the following URL (Figure 9-5):

```
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
```



Figure 9-5 KVM Platform Agent download

3. Put the downloaded agent in /tmp of your KVM host by using Secure Copy Protocol (SCP) and an SCP tool, then uncompress the archive as shown in Figure 9-6.

```
[root@KVM01 SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64]# pwd
/tmp/SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64
[root@KVM01 SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64]# ls
SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64.tar.gz
[root@KVM01 SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64]# tar -xvzf SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64.tar
.gz
dir6.3.1_platform_agent_linux_rhel6kvm_x86_64
platform.rsp
[root@KVM01 SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64]# ls
dir6.3.1_platform_agent_linux_rhel6kvm_x86_64 platform.rsp SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64.tar.gz
[root@KVM01 SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64]# █
```

Figure 9-6 Uncompressed agent that was previously copied in /tmp

4. Start the KVM Platform Agent installation as shown in Figure 9-7.

```
[1-Agree|0-Disagree]: 1
IBM Systems Director Platform Agent 6.3.1 installation.

Extracting RPM files to /tmp/platform.jXfCp2q6Nf
.....
114033 blocks
Preparing packages for installation...
The system you are installing on does not have the IPMI (Intelligent Platform
Management Interface) utilities installed. To install the IPMI utilities,
install the OpenIPMI package using your distribution's package management
system.

Preparing...
1:ibmcim-instrumentation ##### [100%]
2:ibmcim-baseserver ##### [ 3%]
3:openslp ##### [ 8%]
4:tog-pegasus ##### [11%]
5:ibmcim-pegasus-utils ##### [14%]
6:ibmcim-baseos ##### [16%]
7:ibmcim-baseserver-mof ##### [19%]
8:ibmcim-baseos-mof ##### [22%]
9:sblim-cmpi-base ##### [24%]
10:ibmcim-serviceprocessor##### [27%]
11:ibmcim-network ##### [30%]
12:libcmptutil ##### [32%]
Installing....
13:Lib_Utills ##### [35%]
14:sblim-gather ##### [38%]
15:libvirt-cim ##### [41%]
16:ibmcim-network-mof ##### [43%]
17:ibmcim-serviceprocessor##### [46%]
18:ibmcim-pegasus-provider##### [49%]
19:ibm-key-xchange-cmpi ##### [51%]
20:sblim-gather-provider ##### [54%]
21:lsi_mr_hhr ##### [57%]
22:sblim-cmpi-fsvol ##### [59%]
23:sblim-cmpi-nfsv3 ##### [62%]
24:ibmcim-vlan-mof ##### [65%]
25:serveraid-platformagent^[[2~##### [68%]
26:tog-pegasus-cim-tools ##### [70%]
27:qlogic_cna_providers-pe##### [73%]
```

Figure 9-7 KVM Platform Agent installation panel (1 of 2)

Figure 9-8 shows the second half of the installation panel.

```
Saving /etc/libvirt/libvirtd.conf as /etc/libvirt/libvirtd.conf.old
Changes made to /etc/libvirt/libvirtd.conf
249c249
< max_clients = 200
---
> #max_clients = 20
258c258
< max_workers = 200
---
> #max_workers = 20
274c274
< max_requests = 200
---
> #max_requests = 20
280c280
< max_client_requests = 100
---
> #max_client_requests = 5

 33:openslp-server          ##### [ 89%]
 34:ibmcim-vlan            ##### [ 92%]
 35:log4cxx                ##### [ 95%]
 36:ibm-watchdog64        ##### [ 97%]
 37:dsa                   ##### [100%]
Starting libvirtd daemon:
Starting slpd: Multicast Route Enabled          [ OK ]
Starting gatherd:                               [ OK ]
Starting reposed:                               [ OK ]
Shutting down CIM server: cimserver stop        [ OK ]

tog-pegasus: Generating cimserver SSL certificates...SSL certificates generated

Starting up CIM server: cimserver start          [ OK ]

Starting tier1slp:
please wait .....                               [ OK ]

Starting IBM platform agent watchdog service...

Installation of the IBM Systems Director Platform Agent 6.3.1 succeeded.

[root@KVM01 SysDir6_3_1_Platform_Agent_Linux_RHEL6KVM_x86_64]# █
```

Figure 9-8 KVM Platform Agent installation panel (2 of 2)

If you receive error messages similar to the error shown in Figure 9-9, use `yum` to solve the dependencies issue and restart the KVM Platform Agent installation.

```
error: Failed dependencies:
    libconfig.so.8()(64bit) is needed by libvirt-cim-0.5.14-7.el6.x86_64
The main set of RPMs will not install due to unsatisfied dependencies.
```

Figure 9-9 Dependencies error example

Your KVM Platform Agent installation is complete.

Requirement: Repeat the steps in 9.2.1, “Preparation” on page 321 and 9.2.2, “KVM Platform Agent installation” on page 322 for each host that you want to manage with FSM.

9.2.3 KVM host discovery, granting access, and inventory collection

To enable host discovery and run inventory collection, perform these steps:

1. Go to Discovery Manager as shown in Figure 9-10 and click **System Discovery**.




Figure 9-10 Discovery Manager window

2. Enter the IP address of the KVM host as shown in Figure 9-11 or the range of KVM hosts IP addresses that you want to discover. Click **Discover Now**.

The screenshot shows the 'System Discovery' configuration page. At the top, there are browser tabs for 'Chassis Man...', 'Home', 'Discovery M...', and 'System Disc...'. The main heading is 'System Discovery'. Below the heading, there is a descriptive paragraph: 'Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a resource for a single IP address, or use a discovery profile. Discovery profiles enable you to customize discovery for a range of IP addresses, and requesting access to and collecting inventory for the discovered resources.' There is a link to 'Learn more about using discovery'. The 'Select a discovery option:' dropdown is set to 'Single IPv4 address'. The 'IP address:' field contains '9', '.27', '.15', and '.75'. The 'Select the resource type to discover:' dropdown is set to 'All'. At the bottom, there are two buttons: 'Discover Now' (highlighted with a red box) and 'Schedule...'. On the right side, there is an 'Advanced Tasks' sidebar with links for 'Create new profile', 'Manage discovery profiles', and 'Discovery jobs'.

Figure 9-11 Entering the KVM host IP address for discovery

Wait until the job is completed as shown in Figure 9-12.

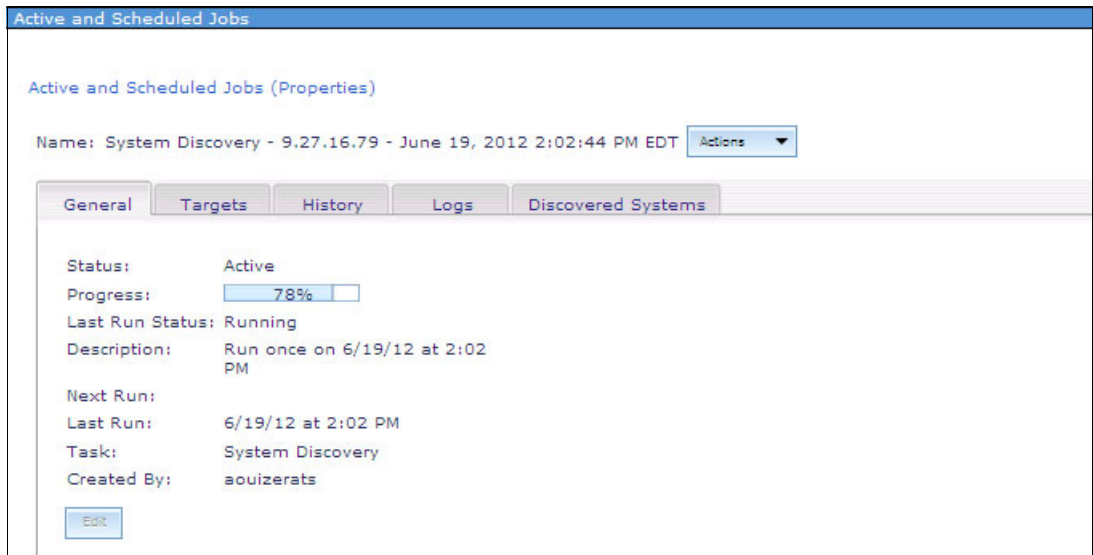


Figure 9-12 KVM discovery progress

3. When the job is complete, click the discovered system. You can see that an operating system object type has been discovered as shown in Figure 9-13. Click **No access**.

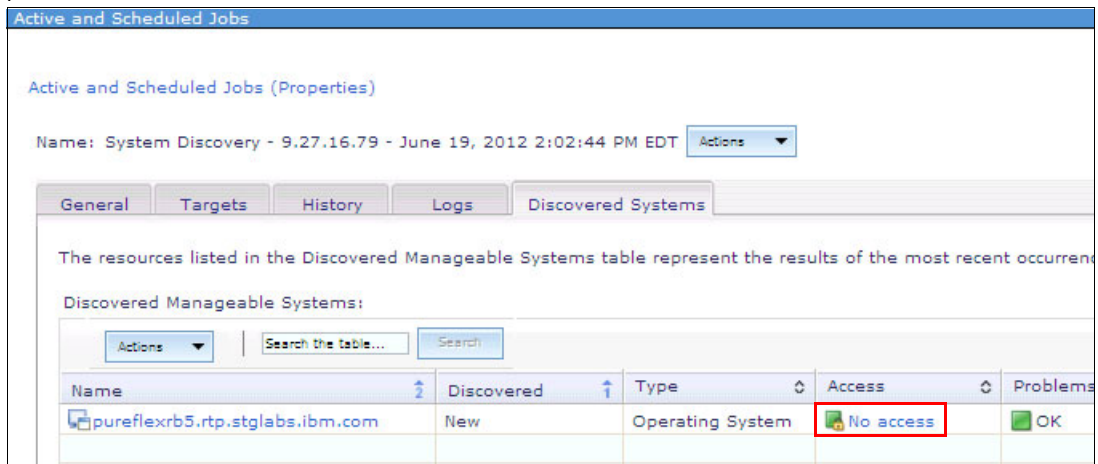


Figure 9-13 Discovered KVM host system

- Grant access to your KVM host by using root credentials as shown in Figure 9-14, then click your host.

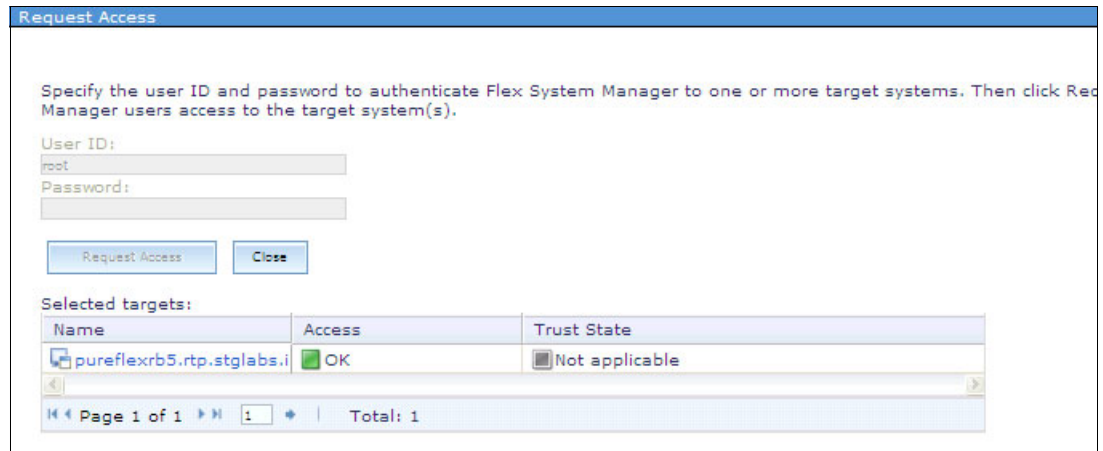


Figure 9-14 KVM host operating system access granted

- Click **Actions** → **Inventory** → **Collect Inventory** as shown in Figure 9-15.

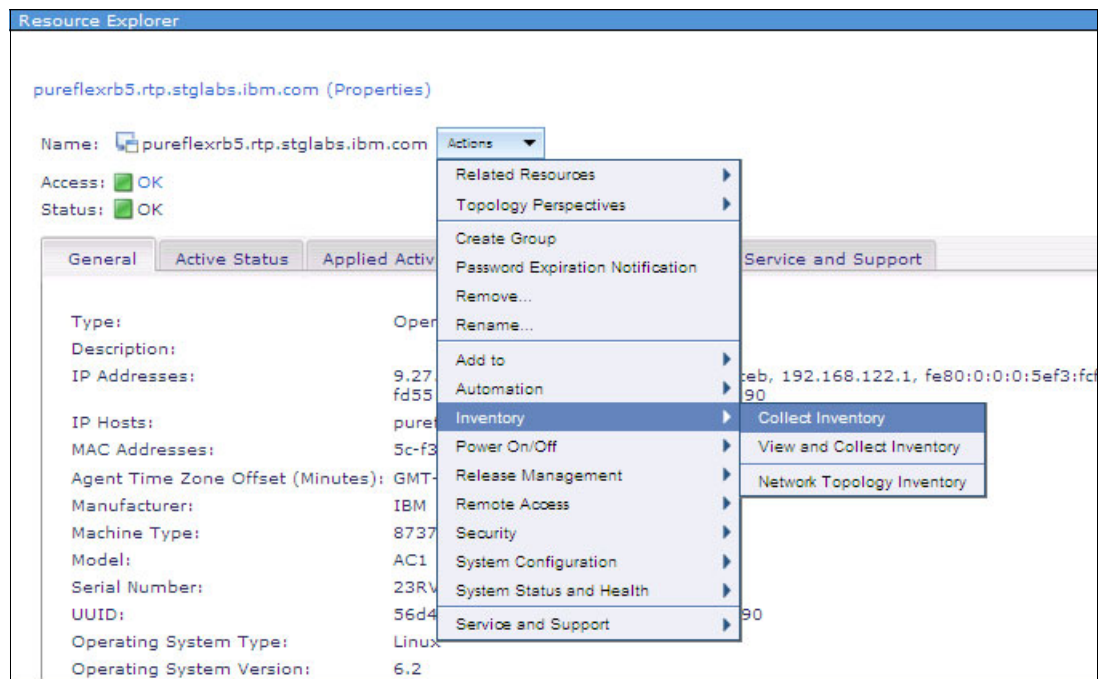


Figure 9-15 Collect Inventory on KVM host

After processing, your collection is complete as shown in Figure 9-16 and your KVM host is ready to be managed by FSM.




Figure 9-16 Collect inventory on KVM host completed

9.3 Image repository for KVM

To set up the image repository for IBM FSM/VMControl with SAN storage, several pieces must be in place. The sections describe the steps that you must perform.

This section addresses the implementation of the KVM SAN storage-based solution. If you want to install an environment using the NFS storage-based solution, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_c_learnmore_repositories_kvm.html

Figure 9-17 illustrates an image repository for a KVM virtual environment with SAN storage.




Figure 9-17 KVM image repository with SAN storage

9.3.1 Preparation

Perform these steps to establish an image repository for KVM using the SAN-based storage:

Remember: The images that are shown in the image repository are not created automatically when a new repository is created. Images must be imported or captured to deploy the virtual server and workload.

1. Install your KVM host by using the steps in 9.2.1, “Preparation” on page 321.

Tip: Stop at the end of 9.2.1, “Preparation” on page 321.

2. Install the prerequisite RPMs that are required to install Linux x86 Common Agent as shown in Figure 9-18 and Figure 9-19 on page 330.

```
[root@KVM03 tmp]# yum install libcrypt.so.1 libc.so.6 libdl.so.2 libstdc++.so.5 libgcc_s.so.1 libm.so.6 libnsl.so.1 libpam.so.0 libpthread.so.0 librt.so.1 unzip bind-utils net-tools libstdc++.so.6 libuuid.so.1 libcrypt.so.1 libxpat.so.0
Loaded plugins: product-id, security, subscription-manager
Updating certificate-based repositories.
Setting up Install Process
Package unzip-6.0-1.el6.x86_64 already installed and latest version
Package 32:bind-utils-9.7.3-8.P3.el6.x86_64 already installed and latest version
Package net-tools-1.60-109.el6.x86_64 already installed and latest version
Resolving Dependencies
--> Running transaction check
----> Package compat-expat1.i686 0:1.95.8-8.el6 will be installed
----> Package compat-libstdc++-33.i686 0:3.2.3-69.el6 will be installed
----> Package glibc.i686 0:2.12-1.47.el6 will be installed
--> Processing Dependency: libfreebl3.so(NSSRAWHASH_3.12.3) for package: glibc-2.12-1.47.el6.i686
--> Processing Dependency: libfreebl3.so for package: glibc-2.12-1.47.el6.i686
----> Package libgcc.i686 0:4.4.6-3.el6 will be installed
----> Package libstdc++.i686 0:4.4.6-3.el6 will be installed
----> Package libuuid.i686 0:2.17.2-12.4.el6 will be installed
----> Package pam.i686 0:1.1.1-10.el6 will be installed
--> Processing Dependency: libcrack.so.2 for package: pam-1.1.1-10.el6.i686
--> Processing Dependency: libaudit.so.1 for package: pam-1.1.1-10.el6.i686
--> Processing Dependency: libselinux.so.1 for package: pam-1.1.1-10.el6.i686
--> Processing Dependency: libdb-4.7.so for package: pam-1.1.1-10.el6.i686
--> Running transaction check
----> Package audit-libs.i686 0:2.1.3-3.el6 will be installed
----> Package cracklib.i686 0:2.8.16-4.el6 will be installed
----> Package db4.i686 0:4.7.25-16.el6 will be installed
----> Package libselinux.i686 0:2.0.94-5.2.el6 will be installed
----> Package nss-softokn-freebl.i686 0:3.12.9-11.el6 will be installed
--> Finished Dependency Resolution
```

Figure 9-18 Linux x86 Common Agent RPM prerequisites

```

Total                                                                 103 MB/s | 6.5 MB   00:00
warning: rpmts_HdrFromFdno: Header V3 RSA/SHA256 Signature, key ID fd431d51: NOKEY
Retrieving key from file:///mnt/rhel62iso/RPM-GPG-KEY-redhat-release
Importing GPG key 0xFD431D51:
  Userid: "Red Hat, Inc. (release key 2) <security@redhat.com>"
  From   : /mnt/rhel62iso/RPM-GPG-KEY-redhat-release
Is this ok [y/N]: y
Importing GPG key 0x2FA658E0:
  Userid: "Red Hat, Inc. (auxiliary key) <security@redhat.com>"
  From   : /mnt/rhel62iso/RPM-GPG-KEY-redhat-release
Is this ok [y/N]: y
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : libgcc-4.4.6-3.el6.i686                1/12
  Installing : glibc-2.12-1.47.el6.i686              2/12
  Installing : nss-softokn-freebl-3.12.9-11.el6.i686 3/12
  Installing : db4-4.7.25-16.el6.i686                4/12
  Installing : audit-libs-2.1.3-3.el6.i686           5/12
  Installing : cracklib-2.8.16-4.el6.i686            6/12
  Installing : libselinux-2.0.94-5.2.el6.i686        7/12
  Installing : pam-1.1.1-10.el6.i686                 8/12
  Installing : libuuid-2.17.2-12.4.el6.i686          9/12
  Installing : libstdc++-4.4.6-3.el6.i686           10/12
  Installing : compat-expat1-1.95.8-8.el6.i686       11/12
  Installing : compat-libstdc++-33-3.2.3-69.el6.i686 12/12
rhel62-cdrom/productid | 1.7 kB   00:00 ...
Installed products updated.

Installed:
  compat-expat1.i686 0:1.95.8-8.el6          compat-libstdc++-33.i686 0:3.2.3-69.el6          glibc.i686 0:2.12-1.47.el6
  libgcc.i686 0:4.4.6-3.el6                libstdc++.i686 0:4.4.6-3.el6                libuuid.i686 0:2.17.2-12.4.el6
  pam.i686 0:1.1.1-10.el6

Dependency Installed:
  audit-libs.i686 0:2.1.3-3.el6          cracklib.i686 0:2.8.16-4.el6          db4.i686 0:4.7.25-16.el6
  libselinux.i686 0:2.0.94-5.2.el6      nss-softokn-freebl.i686 0:3.12.9-11.el6

Complete!

```

Figure 9-19 Linux x86 Common Agent RPM prerequisites

9.3.2 Common Agent installation on a KVM host image repository

To install Common Agent, perform these steps:

1. Allow Transmission Control Protocol (TCP) ports that are required by the Common Agent on the firewall as shown in Figure 9-20.

```

[root@KVM03 ~]# iptables -A INPUT -p tcp --dport 5988 -j ACCEPT
[root@KVM03 ~]# iptables -A INPUT -p tcp --dport 5989 -j ACCEPT
[root@KVM03 ~]# service iptables save
iptables: Saving firewall rules to /etc/sysconfig/iptables:[ OK ]
[root@KVM03 ~]# █

```

Figure 9-20 Open ports for CAs

2. Download KVM Common Agent from the following URL (Figure 9-21):

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




Figure 9-21 Linux Common Agent x86

3. Copy the Common Agent into the /tmp folder of your KVM host by using the SCP protocol and an SCP tool. Uncompress the archive with the following command:

```
tar -xvf IBM Systems Director 6.3 Common Agent for Linux on IBM x86
```

4. Start the Common Agent installation as shown in Figure 9-22.

```
[root@KVM03 SysDir6_3_Common_Agent_Linux_x86_a]# clear
[root@KVM03 SysDir6_3_Common_Agent_Linux_x86_a]# ls
dir6.3_commonagent_linux_x86_a  diragent.rsp  SysDir6_3_Common_Agent_Linux_x86_a.tar.gz
[root@KVM03 SysDir6_3_Common_Agent_Linux_x86_a]# ./dir6.3_commonagent_linux_x86_a
[1-Agree|0-Disagree]: 1
IBM Systems Director Common Agent 6.3.0 installation.

Extracting RPM files to /tmp/agent.zNJPK7TpH7
.....
974221 blocks
Preparing packages for installation...
The system you are installing on does not have the IPMI (Intelligent Platform
Management Interface) utilities installed. To install the IPMI utilities,
install the OpenIPMI package using your distribution's package management
system.

Preparing...
ibmcim-icu
ibmcim-ssl
Creating a SSL private key...
Generating RSA private key, 1024 bit long modulus
.....+++++
.....+++++
e is 65537 (0x10001)
Generating a SSL certificate request...
Self-signing an SSL certificate based on system information...
Signature ok
subject=/C=US/ST=NORTH CAROLINA/L=RTP/O=IBM/OU=STG/CN=KVM03
Getting Private key
ibmcim-instrumentation
ibmcim-openslp
Starting slpd: Multicast Route Enabled[ OK ]
ibmcim-objectmanager
Starting IBM CIM indication listener CIM Listener 2.11.0
Starting CIMListener with the following options
    listenerPort 6988
    httpsConnection 0
    sslKeyFilePath
    sslCertificateFilePath
    consumerDir /opt/ibm/icc/lib
    consumerConfigDir /var/opt/ibm/icc/data/indication
    enableConsumerUnload 1
```

Figure 9-22 Common Agent installation on the KVM future image repository server

Wait until the Common Agent installation is complete as shown in Figure 9-23.

```
Installing Tivoli Common Agent Services.
Preparing...
ISDCommonAgent
[Wed Jun 20 01:05:41 EDT 2012]: Installing feature: com.ibm.usmi.agent.cassocketsubagent.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:42 EDT 2012]: Installing feature: com.ibm.sysmgmt.uim.provider.base.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:42 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:42 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.monitors.xlinux.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:42 EDT 2012]: Installing feature: com.ibm.director.commonagent.manager.xlinux32_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:42 EDT 2012]: Installing feature: com.ibm.director.cimlegacy.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:43 EDT 2012]: Installing feature: smallpatch.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:43 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.java.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:43 EDT 2012]: Installing feature: org.sblim.cim.client.agent.feature_2.0.2... Result: SUCCESSFUL
[Wed Jun 20 01:05:43 EDT 2012]: Installing feature: com.tivoli.twg.legacylibs.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:43 EDT 2012]: Installing feature: org.sblim.cim.client.legacy.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:44 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.monitors.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:44 EDT 2012]: Installing feature: com.ibm.usmi.agent.tpmsubagent.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:44 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.procman.xlinux.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:44 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.evtsub.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:44 EDT 2012]: Installing feature: com.ibm.sysmgmt.utils.updateinstaller.agent_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:44 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.procman.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:45 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.rcshd.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:45 EDT 2012]: Installing feature: com.ibm.usmi.slp.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:45 EDT 2012]: Installing feature: com.ibm.usmi.agent.coreagent.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:45 EDT 2012]: Installing feature: com.ibm.tivoli.remoteaccess.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:45 EDT 2012]: Installing feature: com.ibm.usmi.client.ipc.ft_agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:46 EDT 2012]: Installing feature: com.ibm.director.core.common.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:46 EDT 2012]: Installing feature: com.ibm.sysmgmt.uim.provider.software.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:46 EDT 2012]: Installing feature: com.ibm.director.mgr.discovery.common.agent.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:46 EDT 2012]: Installing feature: com.ibm.director.hw.bcx.agent_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:46 EDT 2012]: Installing feature: com.ibm.director.commonagent.manager.feature_6.3.0... Result: SUCCESSFUL
[Wed Jun 20 01:05:47 EDT 2012]: Installing feature: com.ibm.director.hw.bcx.agent.manager_6.3.0... Result: SUCCESSFUL
213:ACTIVE:InstallOrderedFeatures Plug-in:6.3.0SUCCESS
Installation of the IBM Systems Director Common Agent 6.3.0 succeeded.
[root@KVM03 SysDir6_3_Common_Agent_Linux_x86_a]#
```

Figure 9-23 Common Agent installation on the KVM future image repository server

5. Go to Discovery Manager, select the **System Discovery** task, enter the IP address of the host, and then, click **Discover Now**, as shown in Figure 9-24.




Figure 9-24 KVM hosts with Common Agent discovery

- Click **OK** to run the job now and click **Display Properties** to check the job progress status as shown in Figure 9-25.




Figure 9-25 KVM hosts with Common Agent discovery complete

- Click the **Discovered Systems** tab as shown in Figure 9-26.




Figure 9-26 Discovered systems

- Click **No access** and use root credentials to grant access to the system, as shown in Figure 9-27, then click the system name.

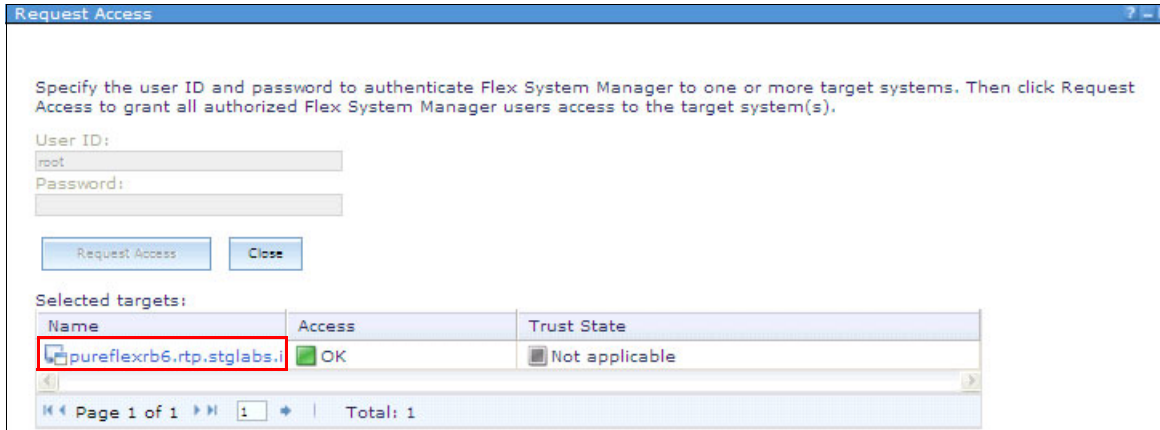


Figure 9-27 Access granted on KVM image repository

- Select **Collect Inventory** on the system as shown in Figure 9-28, then click **OK** to run the job now.

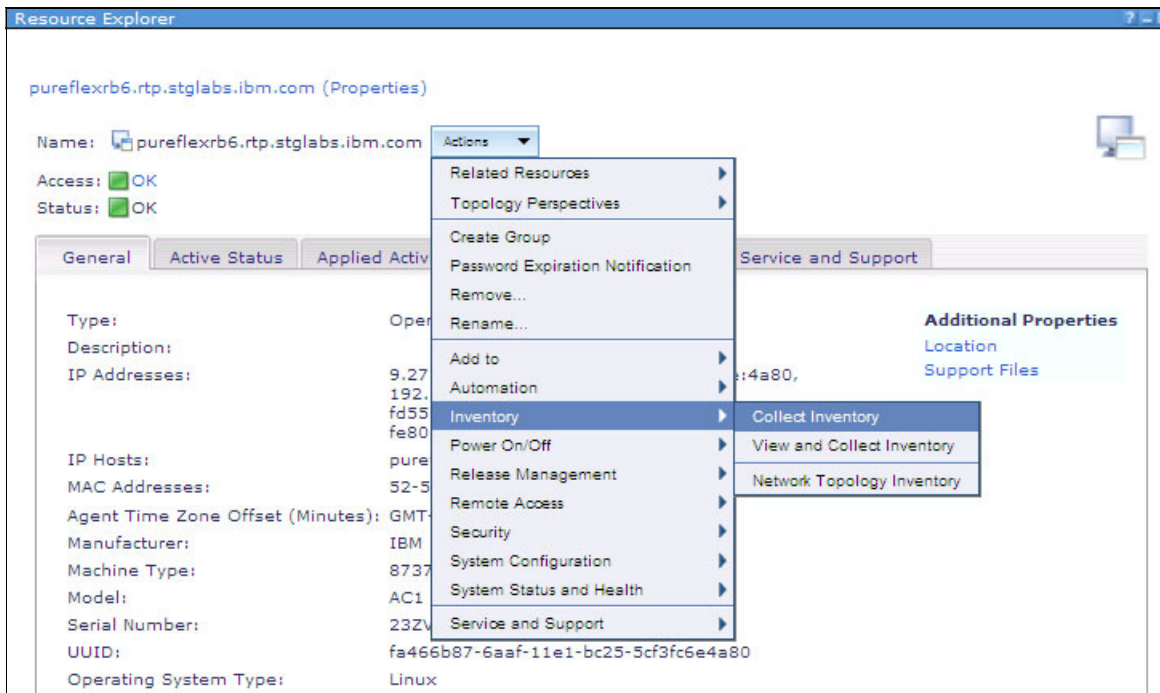


Figure 9-28 Collect Inventory window

10. Click **Display Properties** as shown in Figure 9-29 to check the progress.




Figure 9-29 Information blue box

Wait until the collection is complete as shown in Figure 9-30.

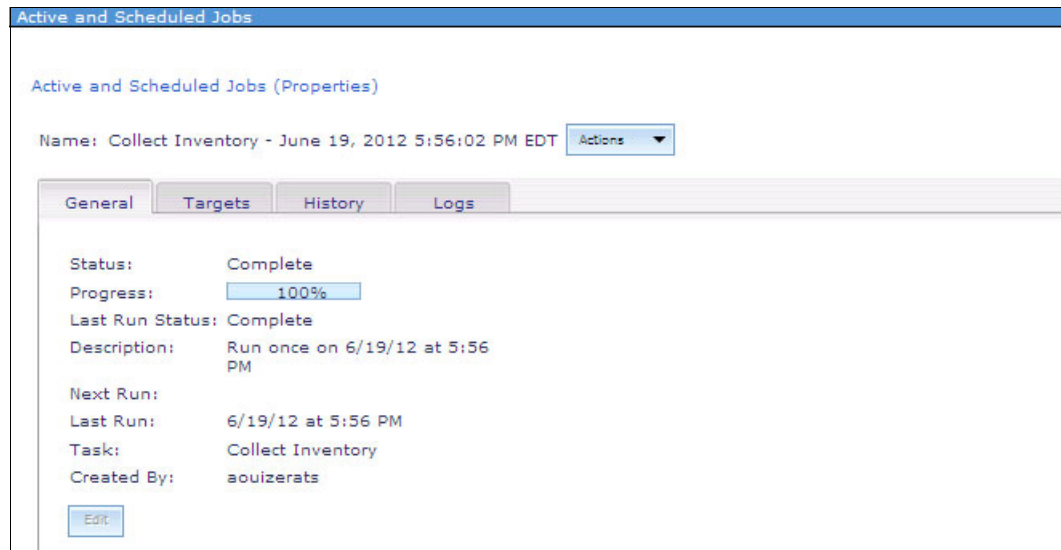


Figure 9-30 Collect inventory is complete

You successfully completed Common Agent installation, discovery, and inventory collection on the KVM repository host.

9.3.3 Subagent installation on a KVM image repository host

You can access and install IBM Flex System Manager VMControl subagents by using the IBM Flex System Manager Agent Installation Wizard. To install subagents, perform these steps:

1. Start the wizard from the VMControl Summary window or from the IBM Flex System Manager Release Management task. You can also start it from Resource Explorer by selecting your system as shown in Figure 9-31.

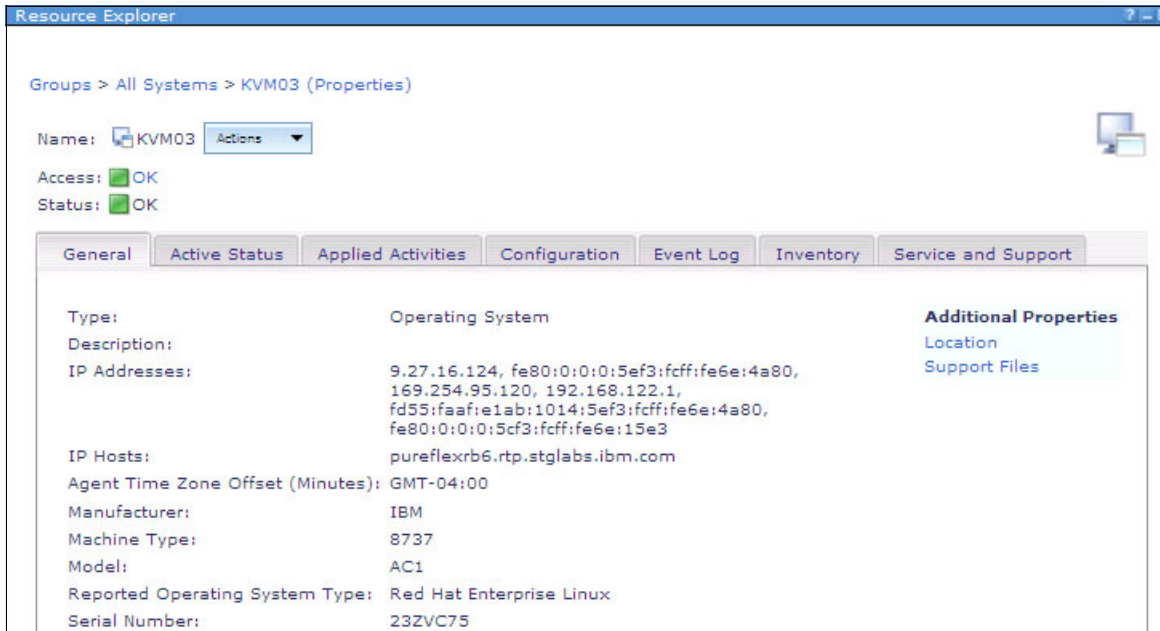


Figure 9-31 KVM image repository: Operating system

2. Click **Actions** → **Release Management** → **Install Agent** as shown in Figure 9-32.

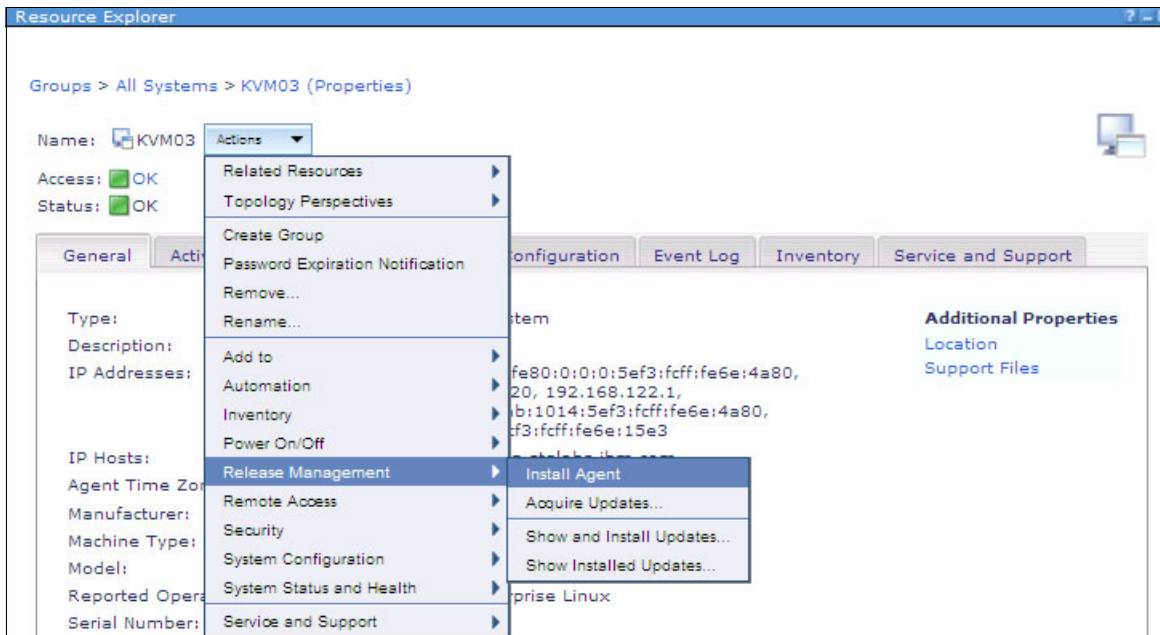


Figure 9-32 Installing the VMControl image repository subagent

3. A welcome window for Agent Installation opens as shown in Figure 9-33. Click **Next**.




Figure 9-33 Install VMControl image repository subagent

4. Click **Common Agent Subagent Packages** as shown in Figure 9-34.

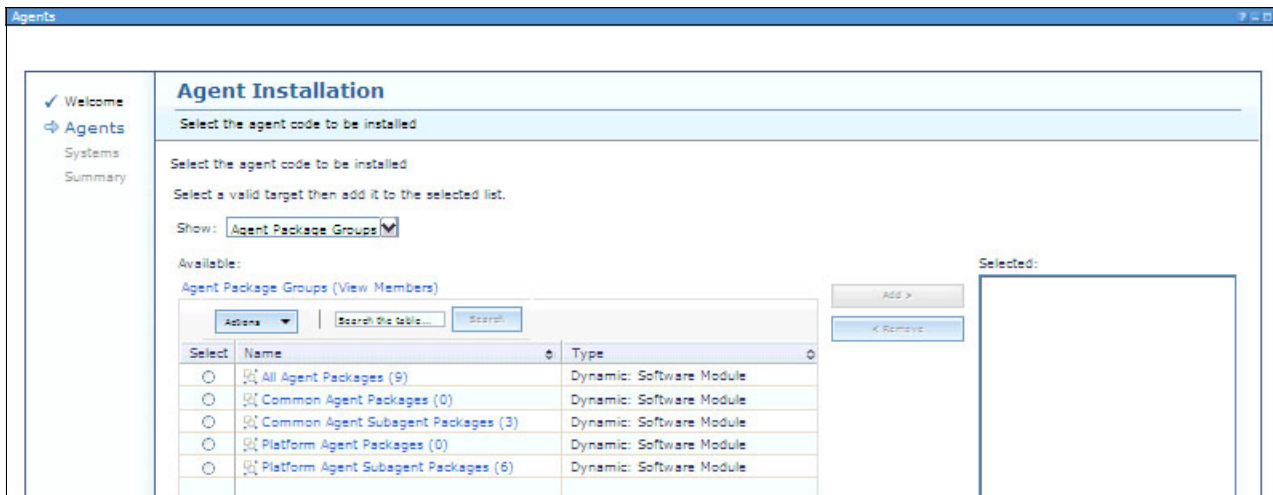


Figure 9-34 Selecting the Common Agent Subagent Packages group

5. Select **CommonAgentSubagent_VMControl_CommonRepository-2.4.1**, as shown in Figure 9-35, then click **Next**.

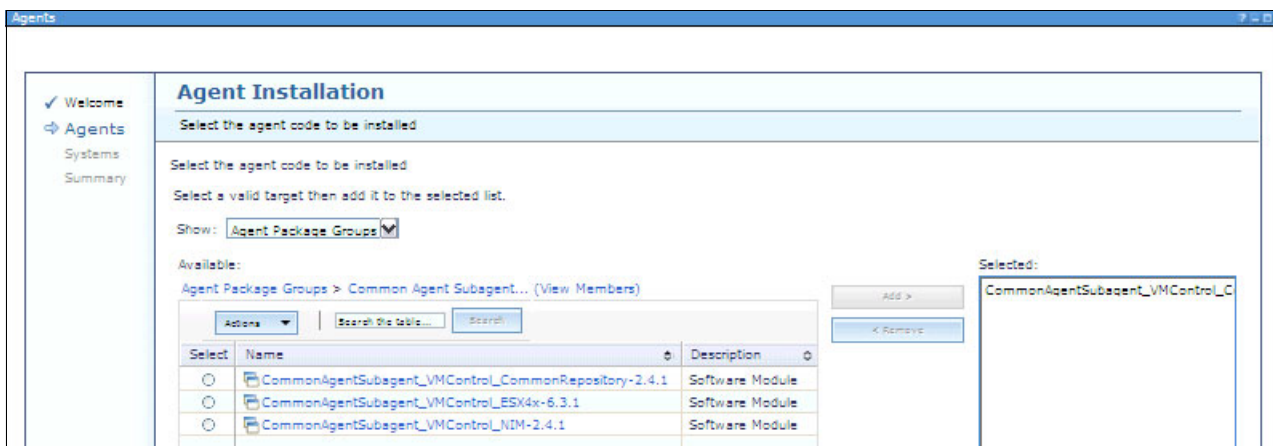


Figure 9-35 CommonAgentSubagent_VMControl_CommonRepository-2.4.1 window

- Select your future image repository system, which is an object type Operating System Linux as shown in Figure 9-36, then click **Next**.

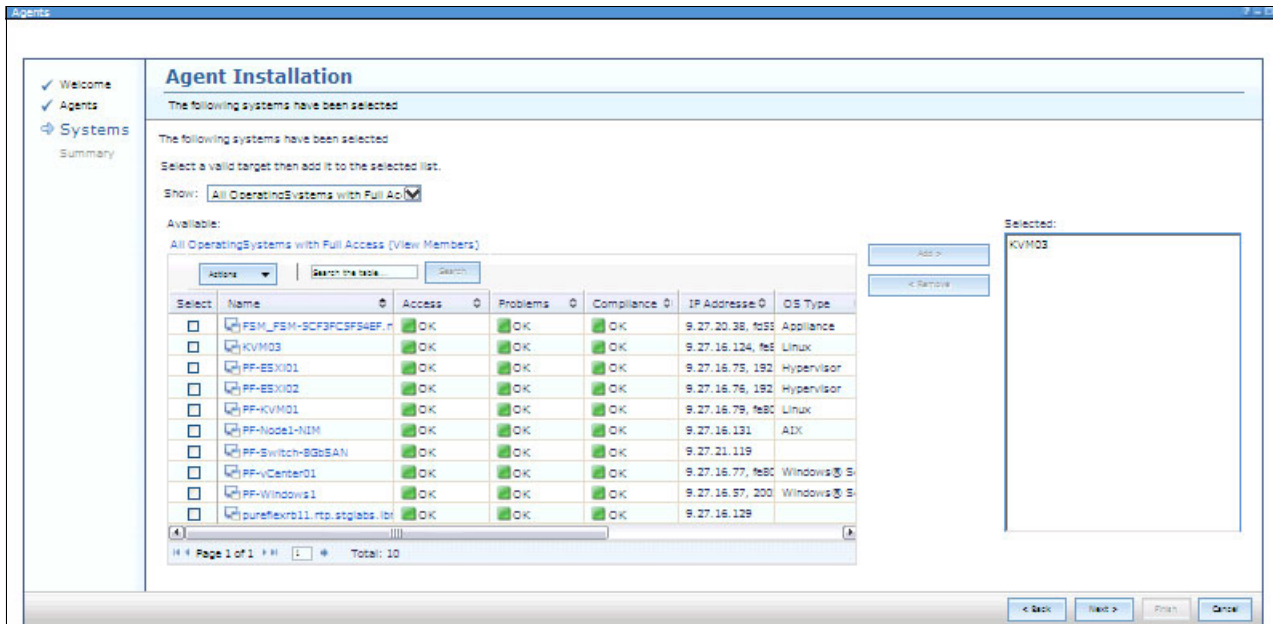


Figure 9-36 Select image repository host operating system target

- Check the summary that shows you the agent and the target for installation as shown in Figure 9-37, then click **Finish**.




Figure 9-37 Agent installation summary

8. Click **OK** to run the job now as shown in Figure 9-38.




Figure 9-38 Run agent installation job

9. Click **Display Properties** as shown in Figure 9-39 to check the job status and wait until status is complete.




Figure 9-39 Display Properties window

9.3.4 Host mappings

When you deploy to a new virtual server on Linux KVM, the disks in the virtual appliance are assigned to the disks in the virtual server based on disk order.

Important: The KVM host WWN must be visible from the V7000 storage system. If it is not, check your SAN zoning or your storage adapters.

To create host mappings, perform these steps:

1. Log on to the Storwize V7000 Storage Management GUI as shown in Figure 9-40.



Figure 9-40 V7000 GUI

2. Click the **Hosts** icon as shown in Figure 9-41.

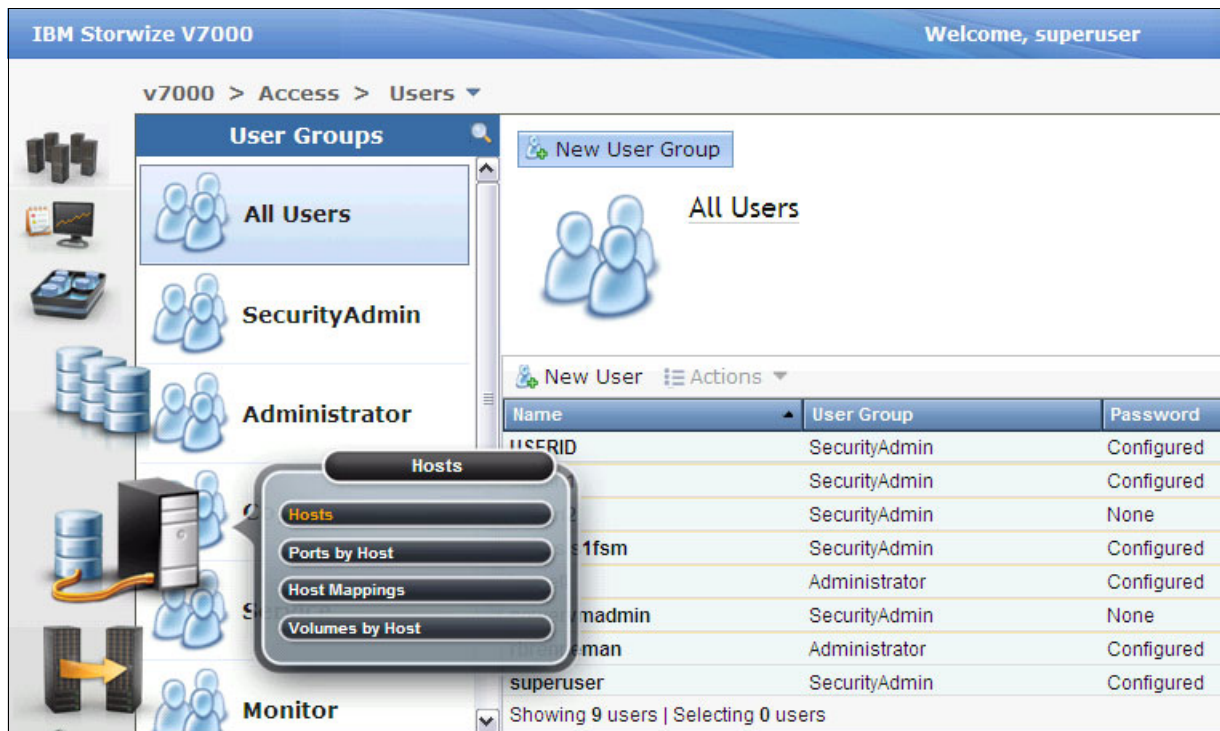
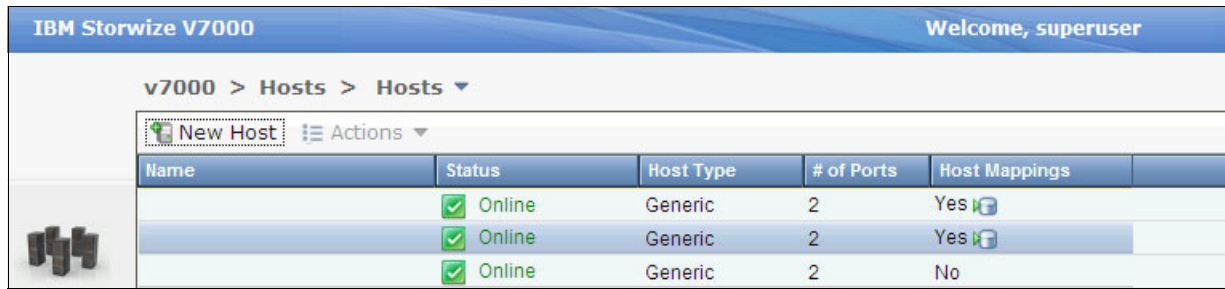


Figure 9-41 V7000 host view

You see that there is no host that is defined for your KVM hosts as shown in Figure 9-42.



The screenshot shows the IBM Storwize V7000 management interface. At the top, it says "IBM Storwize V7000" and "Welcome, superuser". Below that, the breadcrumb navigation is "v7000 > Hosts > Hosts". There is a "New Host" button and an "Actions" dropdown menu. A table lists the current hosts:

Name	Status	Host Type	# of Ports	Host Mappings
	Online	Generic	2	Yes
	Online	Generic	2	Yes
	Online	Generic	2	No

Figure 9-42 Hosts view before you define the KVM host mapping

3. Click **Create Host**, then click **Fibre Channel Host** as shown in Figure 9-43.

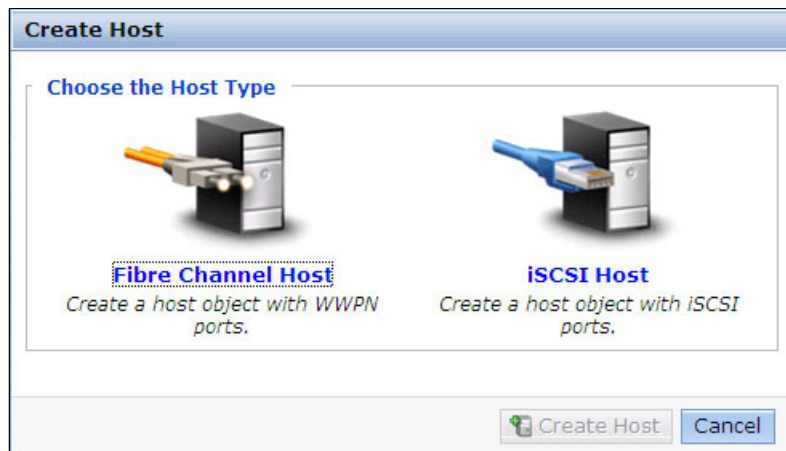


Figure 9-43 Creating the Fibre Channel host

4. Select your WWN and specify a host name as shown in Figure 9-44.

Create Host

Host Name (optional): KVM01

Fibre Channel Ports

21000024FF35EB10 Add Port to List Rescan

Port Definitions
You have not added any WWPNs yet.

Advanced Settings

I/O Group	Port Mask	Host Type
<input checked="" type="checkbox"/> io_grp0	<input checked="" type="checkbox"/> Port 1	<input checked="" type="radio"/> Generic (default)
<input checked="" type="checkbox"/> io_grp1	<input checked="" type="checkbox"/> Port 2	<input type="radio"/> HP/UX
<input checked="" type="checkbox"/> io_grp2	<input checked="" type="checkbox"/> Port 3	<input type="radio"/> OpenVMS
<input checked="" type="checkbox"/> io_grp3	<input checked="" type="checkbox"/> Port 4	<input type="radio"/> TPGS

Advanced Create Host Cancel

Figure 9-44 Selecting the KVM host WWN

5. Click **Add Port to List** as shown in Figure 9-45.

Create Host

Host Name (optional): KVM01

Fibre Channel Ports

Add Port to List Rescan

Port Definitions
21000024FF35EB10

Advanced Settings

I/O Group	Port Mask	Host Type
<input checked="" type="checkbox"/> io_grp0	<input checked="" type="checkbox"/> Port 1	<input checked="" type="radio"/> Generic (default)
<input checked="" type="checkbox"/> io_grp1	<input checked="" type="checkbox"/> Port 2	<input type="radio"/> HP/UX
<input checked="" type="checkbox"/> io_grp2	<input checked="" type="checkbox"/> Port 3	<input type="radio"/> OpenVMS
<input checked="" type="checkbox"/> io_grp3	<input checked="" type="checkbox"/> Port 4	<input type="radio"/> TPGS

Advanced Create Host Cancel

Figure 9-45 Add KVM WWN to Port Definitions list

6. Click **Create Host** to start the host mapping creation as shown in Figure 9-46.




Figure 9-46 Host map creation processing

Wait until the creation task is complete as shown in Figure 9-47.




Figure 9-47 Host map creation completed

7. Go back to the host view to check that your host was created as shown in Figure 9-48.




Figure 9-48 Host map creation checking

Remember: You must repeat this process on each KVM host.

For more information about storage, see the IBM Flex System Information Center at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Fsd0_vim_c_learnmore_storage_paths.html

9.3.5 Discover and manage V7000 storage system

For more information about how to discover and manage V7000 storage system, see 6.12, “Discover and manage external Storwize V7000” on page 234.

9.3.6 Discover and manage SAN switches

For more information about how to discover and manage SAN switches, see 6.14, “External Fibre Channel SAN switch discovery” on page 247.

9.3.7 Discover and configure an image repository server for SAN storage

To configure an image repository server, perform these steps:

1. From the VMControl main page, click the **Virtual Appliances** tab.
2. Click **Create Image Repository** under Common tasks as shown in Figure 9-49.




Figure 9-49 Create image repository

3. You are redirected to the Welcome window as shown in Figure 9-50. Click **Next**.

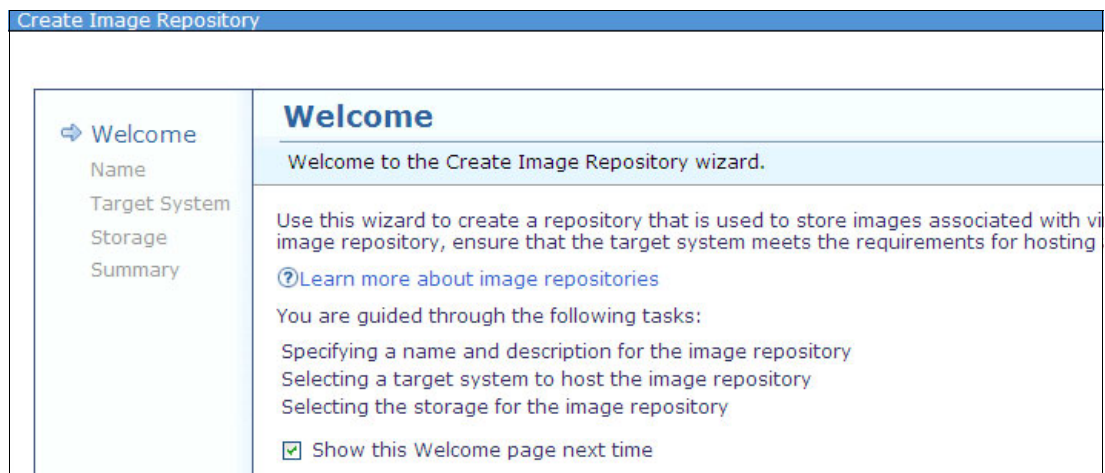


Figure 9-50 Create Image Repository Welcome window

4. Specify the name of your KVM image repository, as shown in Figure 9-51, then click **Next**.

The screenshot displays a configuration window for creating a KVM image repository. On the left, a sidebar lists the steps: 'Welcome' (checked), 'Name' (selected), 'Target System', 'Storage', and 'Summary'. The main area is titled 'Name' and contains the instruction 'Specify a name and description for the image repository you want to create'. Below this, there are two input fields: '*Name:' with the text 'KVMimagesrepo' and 'Description:' which is currently empty. A note below the description field states 'Limit of 256 characters'. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Finish'.

Figure 9-51 Giving a name to the KVM image repository

5. Select the **Target System** that was prepared before with the Common Agent and image repository subagent to create an image repository system (Figure 9-52). Click **Next**.




Figure 9-52 Selecting the image repository system

6. Select the storage that you want to use for the image repository on which virtual appliances will be stored as shown in Figure 9-53.




Figure 9-53 Selecting virtual appliance storage

- Review the summary and click **Finish** to complete the image repository creation as shown in Figure 9-54.




Figure 9-54 Image repository creation summary

- A blue information window opens as shown in Figure 9-55. Click **Display Properties**.

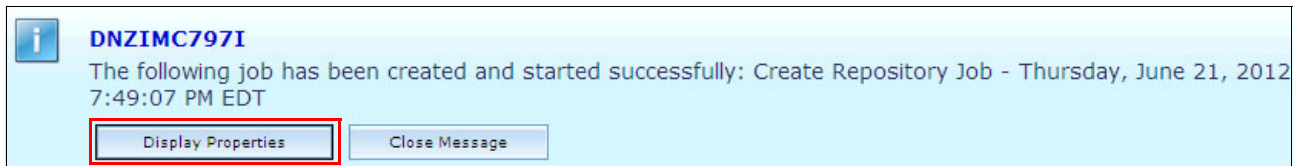


Figure 9-55 Display Properties window

Wait until the job is complete as shown in Figure 9-56.

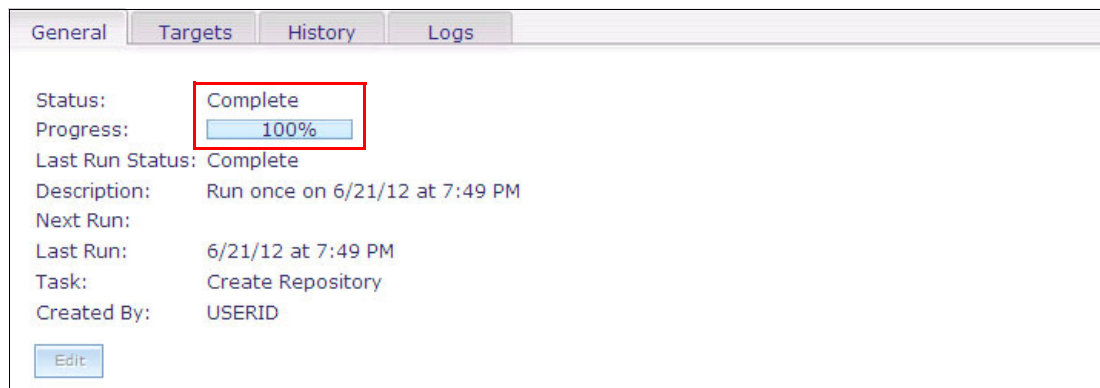


Figure 9-56 Image repository creation complete

9.4 Creating KVM storage system pools

Storage system pools provide the ability to group similar storage subsystems and automate placement within the storage system pool to simplify workload deployment operations.

To work with storage pools, perform these steps:

1. Open the VMControl main window and click **Storage system pools** as shown in Figure 9-57.

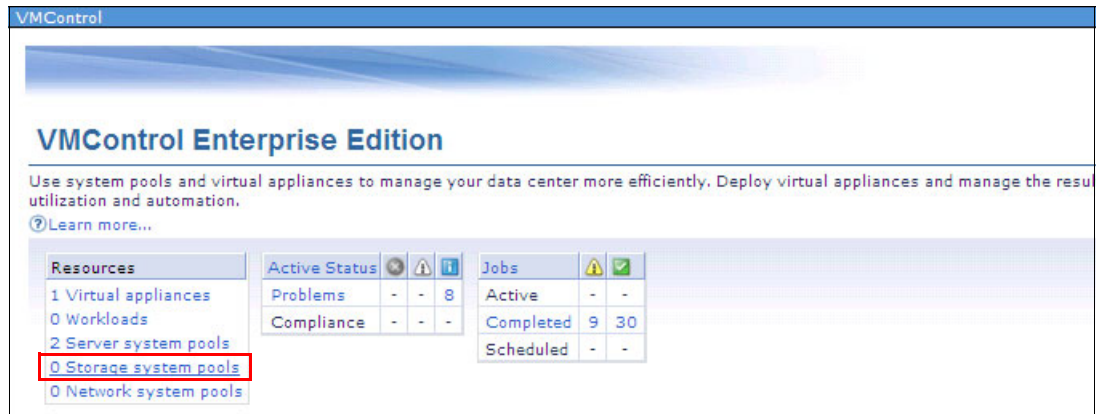


Figure 9-57 VMControl main window

2. The View storage system pools window opens as shown in Figure 9-58. There are no storage system pools available at the moment. Click **Create**.




Figure 9-58 No storage system pools available

The Welcome window opens to create the storage system pool as shown in Figure 9-59.




Figure 9-59 Welcome window to create storage system pools

3. Specify the name to assign to the storage system pool as shown in Figure 9-60, then click **Next**.

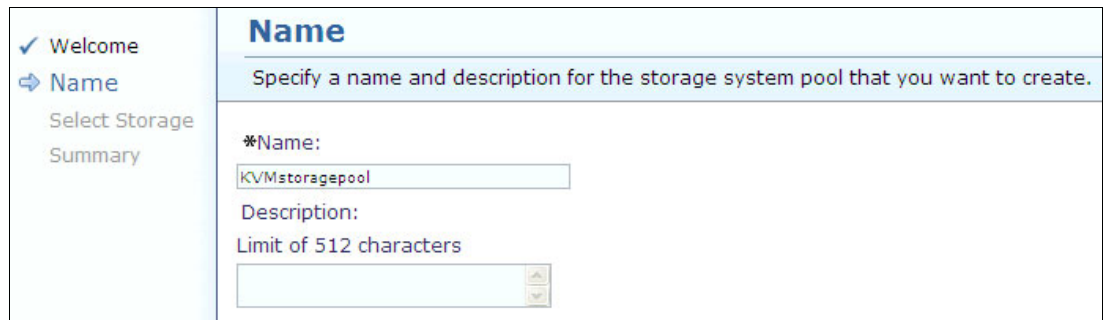


Figure 9-60 Storage system pool name

4. Select the storage subsystem that you want to assign to the storage system pool, click **Add**, as shown in Figure 9-61, and then click **Next**.




Figure 9-61 Select Storage Subsystem window

5. Review the summary as shown in Figure 9-62 and click **Finish**.

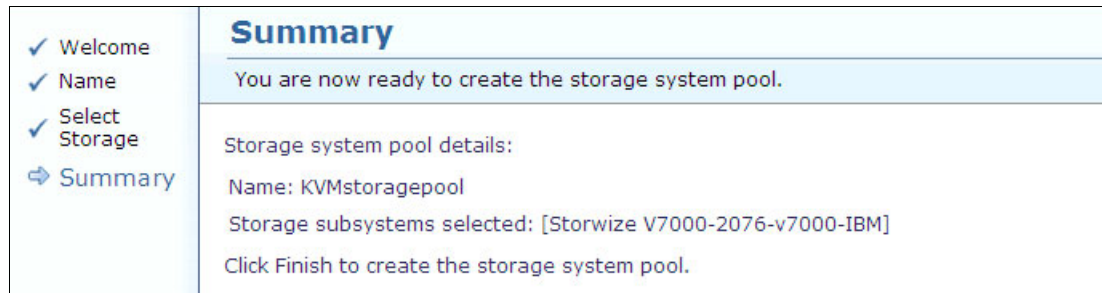


Figure 9-62 Storage system pool creation summary

6. Wait until the storage system pool is created as shown in Figure 9-63.




Figure 9-63 Storage system pool creation

7. A blue information window opens (Figure 9-64). Click **Display Properties**.

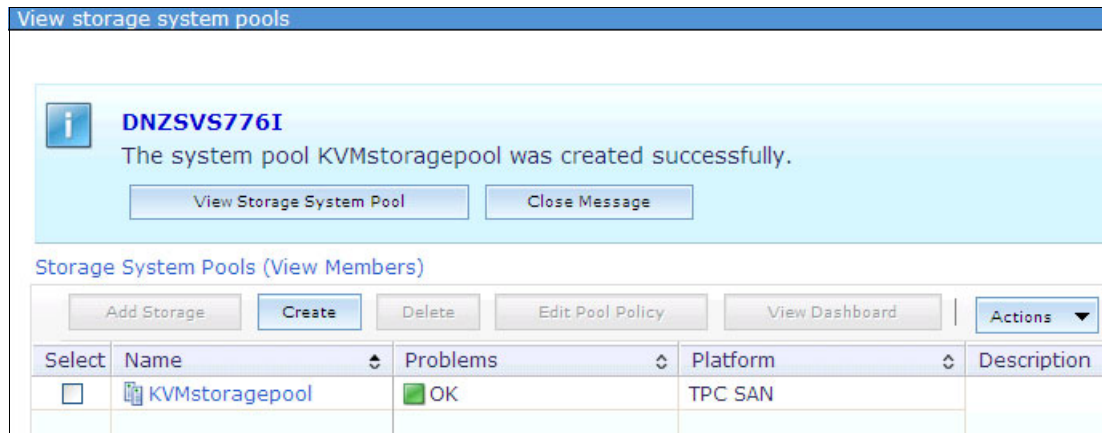


Figure 9-64 Display job properties window

At the end of this process, your storage system pool is created.

9.5 Creating KVM network system pools

By using network system pools with logical network profiles, you can partition and provision your network into separate logical networks. The network system pool functions will automatically provision both the physical and virtual switch devices to ensure connectivity between the devices, and will also provide status when connectivity is broken.

Network system pools simplify and automate network configuration tasks for virtual servers. You can manage the network connections of the pooled network systems to ensure network connectivity across a set of network switches.

Consideration: Network system pools manage Layer 2 network configuration. However, additional Layer 3 IP configuration might be required after VM deployment or migration.

Network system pools, combined with logical networks and server system pools, provide flexibility and control over how network resources are used. As an administrator, you can perform these tasks:

- ▶ Define larger network system pools to allow more efficient use of network resources
- ▶ Define logical networks within a network system pool for shaping and isolation purposes
- ▶ Define a server system pool with some or all of the servers managed by a network pool

To achieve full mobility, server system pools need to use both network system pools and shared storage. A server system pool can be associated with a network system pool when you create a server system pool. This provides physical network connectivity between all resources in the server system pool and ensures workload and virtual server relocation across any resource in the server system pool.

Logical networks within the network system pool define the logical connectivity between the systems; only systems on the same logical networks can connect to each other. Before creating associated server system pools, the network system pool must exist, and it defines the scope of the physical network mobility domain that you want.

A *logical network* defines how a set of workloads are connected together in the network. This logical network is isolated from other logical networks using virtual LANs (VLANs). You can use this architecture to define unique networks for various workloads across a single physical network. The *logical network profile* is the entity that defines how a logical network needs to be used within a network system pool. You can use a single logical network profile across multiple network system pools (physically independent networks).

Use a logical network profile to define a set of attributes defining how the workload or virtual machines will use the network. At a minimum, define the VLAN for the workload to use. Other attributes, such as Virtual Ethernet Bridging (VEB) attributes or Virtual Ethernet Port Aggregator (VEPA) attributes, can also be defined as part of this profile. After you create the logical network profile, you must associate the logical network profile with a network system pool to ensure that all virtual servers that are deployed using this profile will have connectivity to each other.

To create a configuration template that will be used in the network pool, perform the following steps:

1. On the Plug-ins tab of the Home page, click **Configuration Templates** under Configuration Manager, as shown in Figure 9-65.

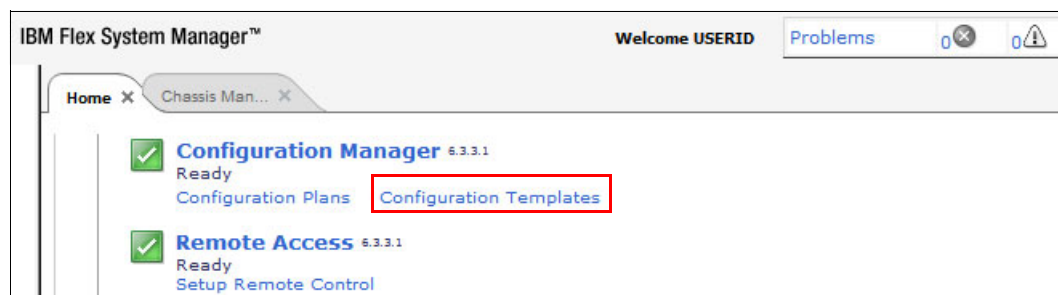


Figure 9-65 Launching Configuration Templates

2. On the Configuration Templates page, shown in Figure 9-66, click **Create**.




Figure 9-66 Configuration Templates window

3. On the Create template page, for the template type, select **System Pool**. For the configuration to create a template, select **Logical Network Configuration**. Provide a configuration template name and optionally a configuration template description, as shown in Figure 9-67. Click **Continue**.




Figure 9-67 Creating a template

- On the Logical Network Configuration Profiles page, click **Add Profile** as shown in Figure 9-68.

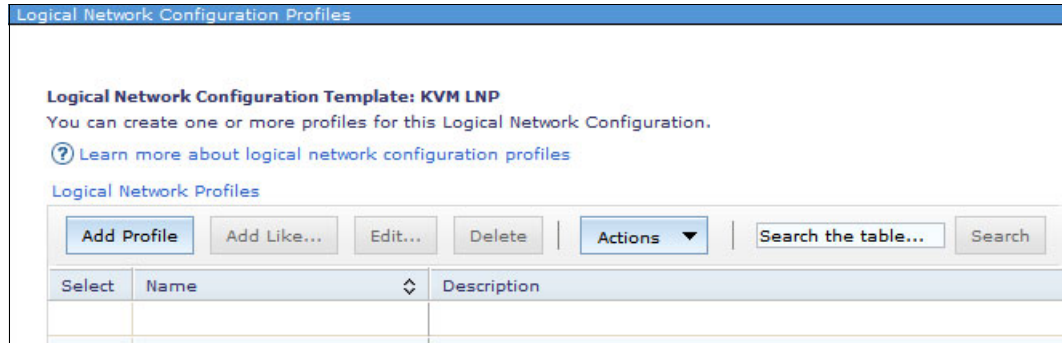


Figure 9-68 Adding a network profile

This launches the Create Logical Network Profile wizard. You can use the wizard to create multiple logical network profiles for each template.

- The Welcome page appears, as shown in Figure 9-69. Click **Next**.




Figure 9-69 Logical network profile: Welcome page

- On the Profile Name page, enter profile name and, optionally, a profile description, as shown in Figure 9-70. Click **Next**.

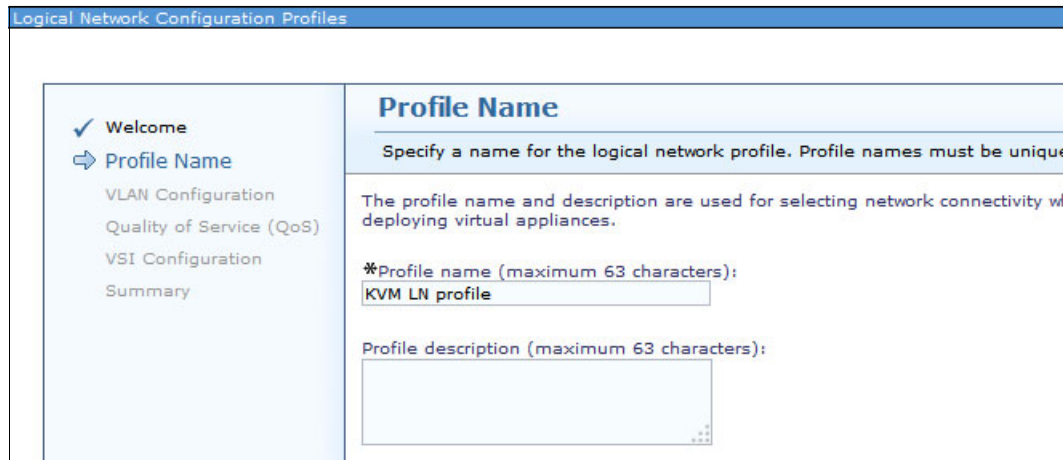


Figure 9-70 Logical network profile: Profile Name

7. On the VLAN Configuration page, enter a VLAN ID for this profile, as shown in Figure 9-71. Click **Next**.

The screenshot shows the 'VLAN Configuration' page. On the left, a navigation pane lists: Welcome (checked), Profile Name (checked), VLAN Configuration (active), Quality of Service (QoS), VSI Configuration, and Summary. The main content area has a title 'VLAN Configuration' and a subtitle 'Specify the switch port VLAN configuration needed to support the connected endpoints'. Below this, it states 'This VLAN setting applies to both VEB and VEPA configurations.' The '*VLAN ID:' field contains the value '42' in a text box, with '(1 - 4094)' indicating the valid range.

Figure 9-71 Logical network profile: VLAN Configuration

8. Enter Quality of Service (QoS) parameters if required, as shown in Figure 9-72. Click **Next**.

The screenshot shows the 'Quality of Service' page. The navigation pane on the left lists: Welcome (checked), Profile Name (checked), VLAN Configuration (checked), Quality of Service (QoS) (active), VSI Configuration, and Summary. The main content area has a title 'Quality of Service' and a subtitle 'Select Quality of Service (QoS) settings for the logical network profile.' Below this, it states 'Quality of service parameters control network performance relative to other endpoints.' There is a link '? Learn more about quality of service settings'. Under 'Port priority:', there are two radio buttons: 'Use default port priority' (selected) and 'Select a port priority'. Below the second radio button is a dropdown menu showing '1' and the text '(1 lowest priority - 7 highest priority)'. Under 'Bandwidth allocation:', there are two radio buttons: 'Use default bandwidth allocation' (selected) and 'Enter bandwidth allocation'. Below the second radio button is a text box followed by '(Kbps)'.

Figure 9-72 Logical network profile: Quality of Service

- Enter the Virtual Switch Interface (VSI) configuration settings, as shown in Figure 9-73. Click **Next**.

The screenshot shows the 'VSI Configuration' step in a wizard titled 'Logical Network Configuration Profiles'. On the left, a navigation pane lists steps: Welcome, Profile Name, VLAN Configuration, Quality of Service (QoS), VSI Configuration (highlighted with a right-pointing arrow), and Summary. The main content area is titled 'VSI Configuration' and contains the following text: 'Specify Virtual Switch Interface (VSI) settings for the network connection.' Below this, it says: 'Use a VSI interface with VEPA to communicate with multiple controllers. VSI definitions with the VSI Manager ID. If you are not using VEPA, you can skip this step.' There are three input fields: 'VSI Manager ID: Identifies the VSI Manager with the database that holds the detailed VSI manager ID can be used to obtain the IP address and other connectivity and access', 'VSI type ID (VTID): The integer identifier of the VSI type.', and 'VSI type ID version: The integer identifier designating the desired version of the VTID.'

Figure 9-73 Logical network profile: VSI Configuration

- On the Summary page, check the configuration parameters, as shown in Figure 9-74. Click **Finish**.

The screenshot shows the 'Summary' step in the 'Logical Network Configuration Profiles' wizard. The left navigation pane shows 'Summary' as the active step, indicated by a right-pointing arrow. The main content area is titled 'Summary' and contains the text: 'You have specified the following settings for this logical network profile.' Below this, a list of settings is displayed: Profile Name: KVM LN profile; Description: (empty); VLAN Configuration: VLAN ID: 42; Quality of Service: Port priority: Default, Bandwidth allocation: Default; VSI Configuration: VSI Manager ID: (empty), VSI Type ID: (empty), VSI Type version: (empty).

Figure 9-74 Logical network profile: Summary

- The newly created profile is displayed in the Template Configuration window, as shown in Figure 9-75 on page 358. If you need to configure more logical network profiles in the configuration template, repeat steps 4 - 10. Click **Save Template**.




Figure 9-75 Saving configuration template

12. The newly created template is shown in the Configuration Templates window, as shown in Figure 9-76.

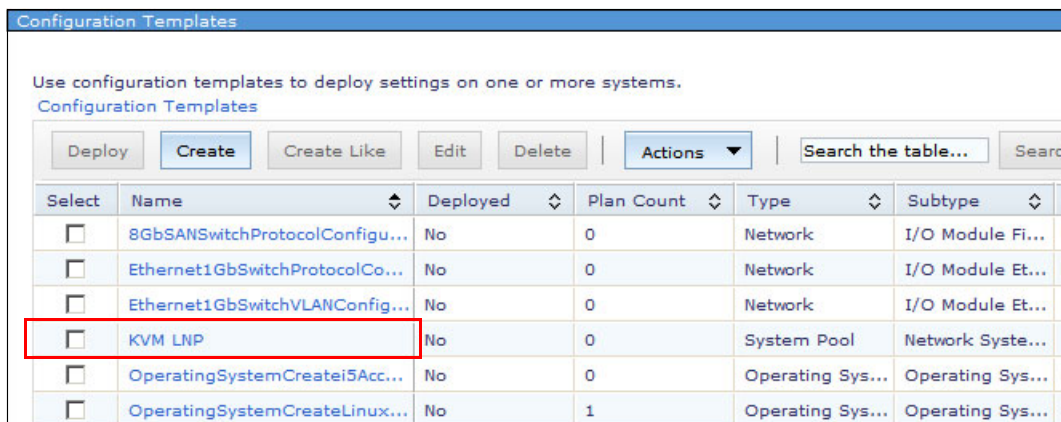


Figure 9-76 Configuration template created

The configuration template is created and is ready to be assigned to the network pool.

To create the network system pool, perform the following steps:

1. In the VMControl main window, click the **System pools** tab, then click **Network System Pools and Members**, as shown in Figure 9-77 on page 359.

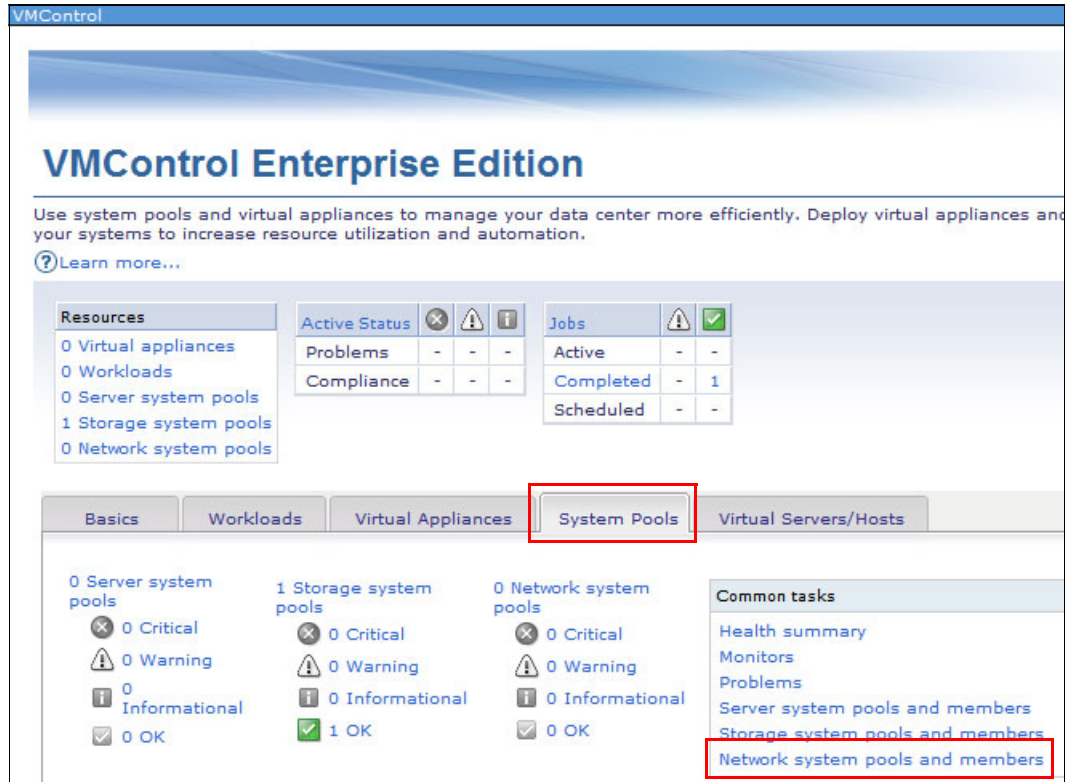


Figure 9-77 VMControl main window

2. In the Network System Pools and Members window, click **Create** to create the new network system pool, as shown in Figure 9-78.




Figure 9-78 Network System Pools and Members window

3. The Create Network System Pool wizard opens, as shown in Figure 9-79 on page 360. Click **Next**.

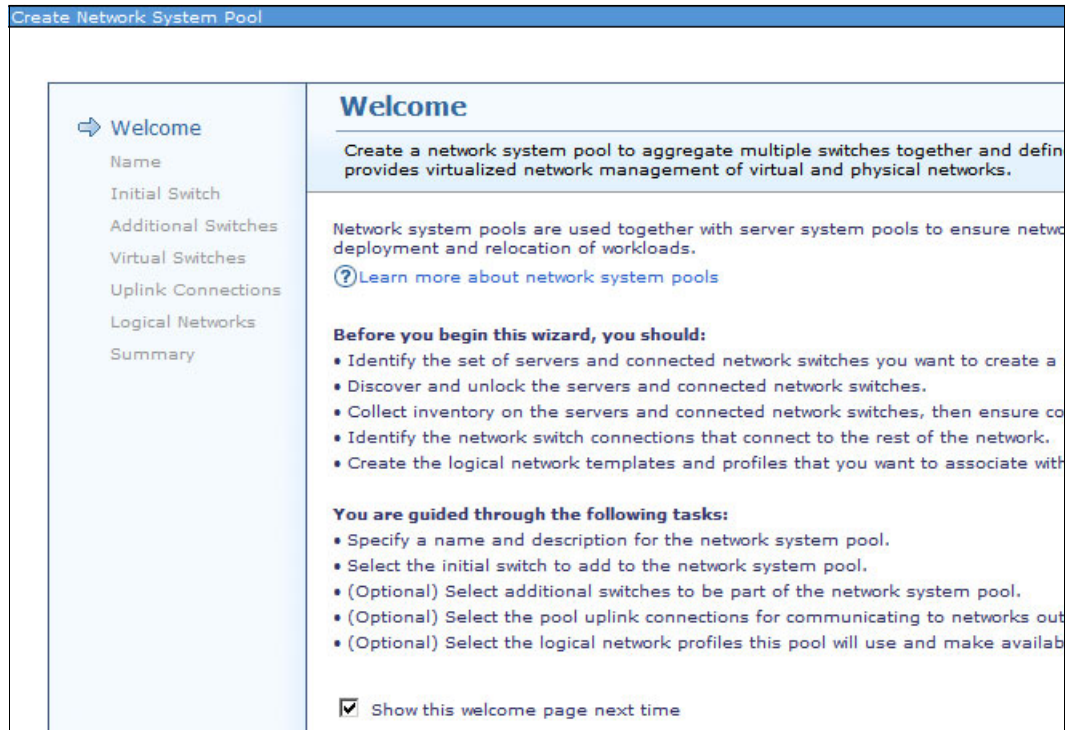


Figure 9-79 Create Network System Pool: Welcome

4. On the Name page, enter the network pool name and description, as shown in Figure 9-80. Click **Next**.




Figure 9-80 Create Network System Pool: Name

5. On the Initial Switch page, select the first switch in the mobility domain, as shown in Figure 9-81 on page 361. Click **Next**.



Figure 9-81 Create Network System Pool: Initial Switch

6. Optional: Select additional switches, as shown in Figure 9-82. Click **Next**.

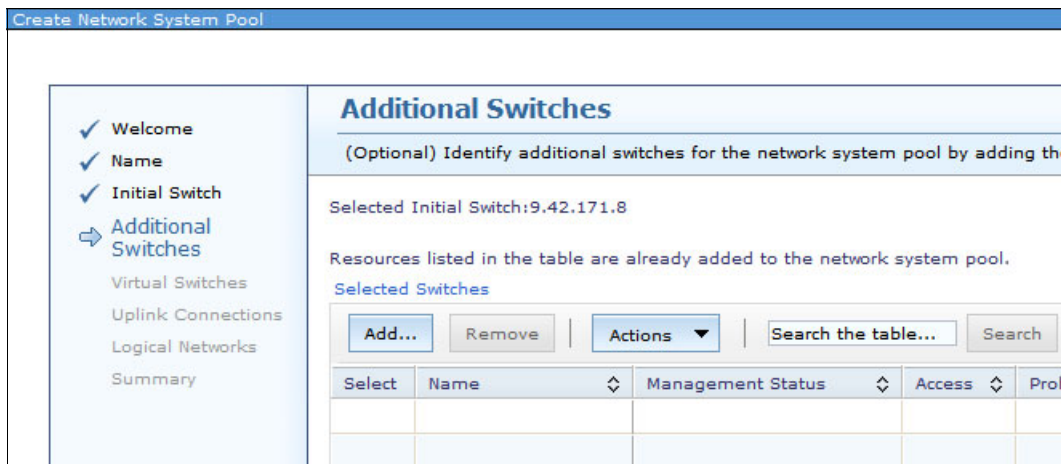


Figure 9-82 Create Network System Pool: Additional Switches

7. The wizard automatically adds virtual switches to the pool, as shown in Figure 9-83 on page 362. Click **Next**.



Figure 9-83 Create Network System Pool: Virtual Switches

- On the Pool Uplink Connections page, click **Add** to add a new uplink, as shown in Figure 9-84.



Figure 9-84 Create Network System Pool: Pool Uplink Connections

- Select the uplink port as shown in Figure 9-85 on page 363. Click **OK**.

Add Pool Uplink Connections - KVM network pool						
Based on your previous switch selections, the following list represents candidate pool uplink connections. Select the LAN network system pool to the rest of the network.						
Candidate Pool Uplink Connections						
<input type="button" value="Actions"/> <input type="text" value="Search the table..."/> <input type="button" value="Search"/>						
Select	LAN Connection	Switch	VLAN ID Set	Default VLAN ...	Description	
<input type="checkbox"/>	EXT20	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT21	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT10	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT22	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT15	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT16	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT17	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT18	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT19	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXTM	9.42.171.8	4095	PVID - 4095		
<input checked="" type="checkbox"/>	EXT1	9.42.171.8	1	Default - 1		
<input type="checkbox"/>	EXT2	9.42.171.8	1	Default - 1		

Figure 9-85 Add Uplink Connections - KVM network pool

10. The uplink port that was added appears in the list of uplinks, as shown in Figure 9-86. Click **Next**.

Create Network System Pool

- Welcome
- Name
- Initial Switch
- Additional Switches
- Virtual Switches
- Uplink Connections
- Logical Networks
- Summary

Pool Uplink Connections

(Optional) Identify the pool uplink connections for the network system pool by adding uplink connections.

For each pool uplink connection, you can define what VLAN traffic is allowed between the pool and the rest of the network. Pool uplink connections can also be modified after creating the network system pool.

[Learn more about pool uplink connections](#)

Resources listed in the table are already added to the network system pool.

Pool Uplink Connections

Select	LAN Connection	Switch	VLAN ID Set	Default VLAN
<input type="checkbox"/>	EXT1	9.42.171.8	1	Default - 1

Figure 9-86 Create Network System Pool: Pool Uplink Connections

11. Assign the logical network profile created earlier by clicking **Add**, as shown in Figure 9-87 on page 364.




Figure 9-87 Create Network System Pool: Logical Network Profiles

12. Select the logical network profile, as shown in Figure 9-88. Click **OK**.

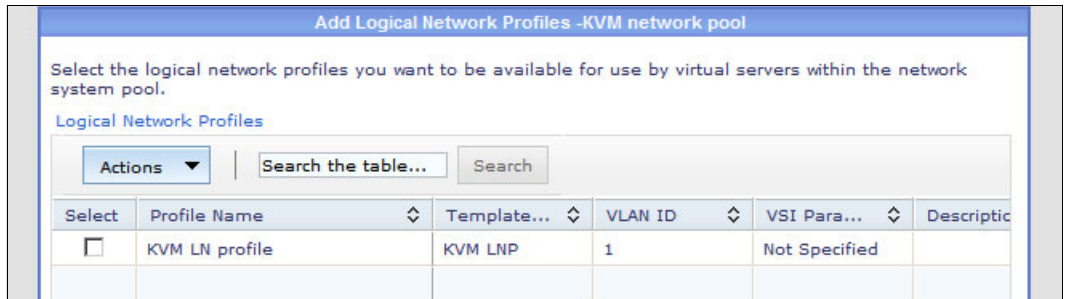


Figure 9-88 Create Network System Pool: Add Logical Network Profiles

13. The added profile is listed on the Logical Network Profiles page (see Figure 9-89). Click **Next**.

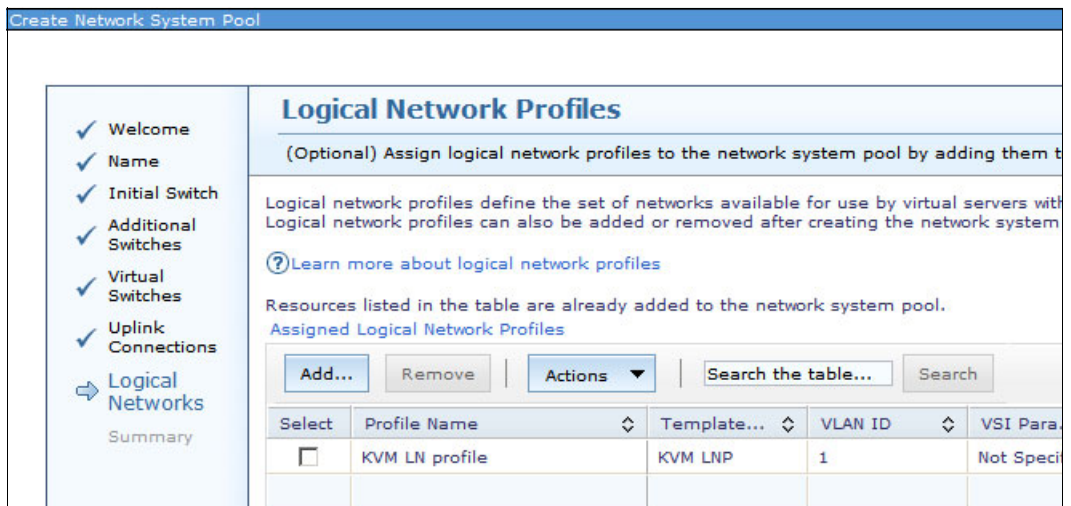


Figure 9-89 Create Network System Pool: Logical Network Profiles

14. Check the configuration settings on the Summary page, as shown in Figure 9-90 on page 365. Click **Finish**.

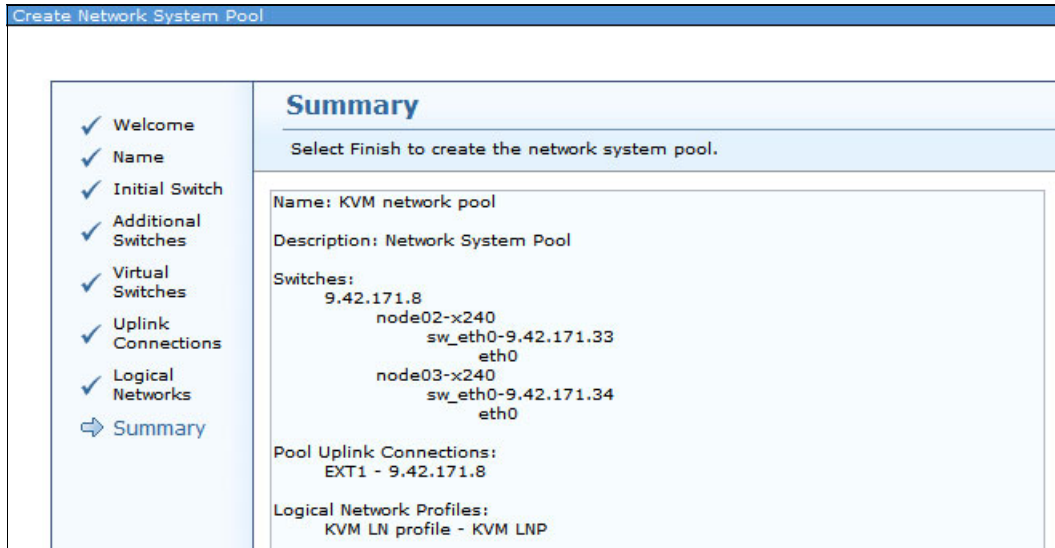


Figure 9-90 Create Network System Pool: Summary

15. The newly created network system pool is shown in the Network System Pools and Members window (see Figure 9-91).




Figure 9-91 Network System Pools and Members

For more information about implementing network system pools, see this website:

http://pic.dhe.ibm.com/infocenter/flexsys/information/topic/com.ibm.acc.8731.doc/com.ibm.sdnm.adv.helps.doc/fnc0_t_network_ctrl_managing_nsps_and_lmps.html

9.6 Creating KVM server system pools

Before you create Server system pools, review the information at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_c_learnmore_getting_started_system_pools.html

For more information about the prerequisites, see this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fcom.ibm.acc.commontasks.doc%2Fcommontasks_navigating_fsm.html

To create KVM server system pools, perform these steps:

1. Go to the VMControl main window as shown in Figure 9-92, and click **Server system pools**.



Figure 9-92 VMControl main window

2. The View Server System Pools window opens as shown in Figure 9-93. Click **Create**.




Figure 9-93 Server system pool member list

3. The Welcome window to create a server system pool opens as shown in Figure 9-94 on page 367. Click **Next**.

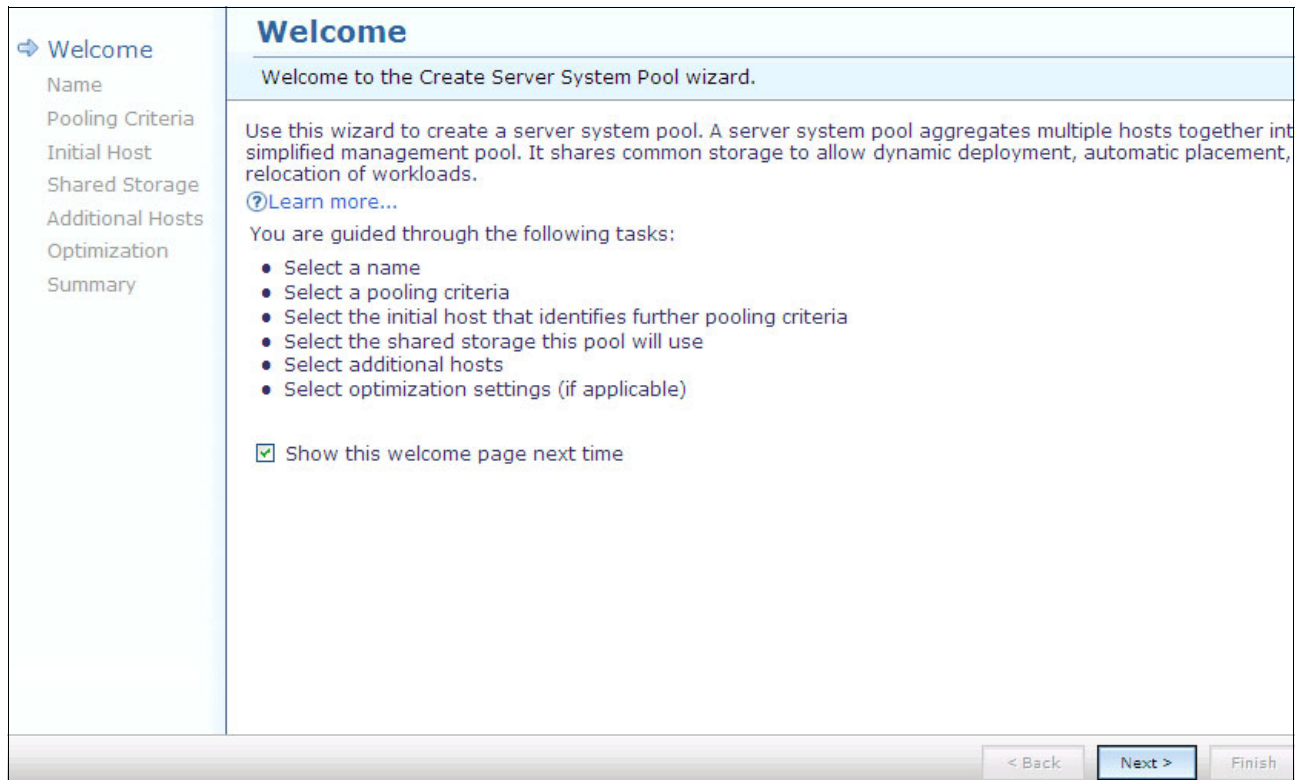


Figure 9-94 Server system pool Welcome window

- Specify a name to assign to the server system pool as shown in Figure 9-95, then click **Next**.

<ul style="list-style-type: none"> ✓ Welcome ⇒ Name Pooling Criteria Initial Host Shared Storage Additional Hosts Optimization Summary 	<h2 style="margin: 0;">Name</h2> <p style="margin: 0;">Specify a name and description for the server system pool that you want to create.</p> <p>*Name: <input type="text" value="KVMpool"/></p> <p>Description (limit of 512 characters): <input type="text"/></p>
--	--

Figure 9-95 Server system pool name

- Under Pooling Criteria, check **Only add hosts capable of live virtual server relocation** as shown in Figure 9-96, then click **Next**.

<ul style="list-style-type: none"> ✓ Welcome ✓ Name ⇒ Pooling Criteria Initial Host Shared Storage Additional Hosts Optimization Summary 	<h2 style="margin: 0;">Pooling Criteria</h2> <p style="margin: 0;">Select the pooling criteria to use for this server system pool.</p> <p>Resilience criteria: <input checked="" type="checkbox"/> Only add hosts capable of live virtual server relocation</p> <p>Network deployment criteria: <input type="checkbox"/> Only add hosts connected by a network system pool and capable of automating network system pool management.</p> <p style="margin-left: 20px;">There are no available network system pools. Learn more about network system pools</p> <p>Note: When adding hosts that contain existing virtual servers, the existing virtual servers will still run on the host, but not be managed by the server system pool.</p> <p>Learn about server system pool capabilities</p>
--	--

Figure 9-96 Pooling criteria

6. Select a host as the initial host to initiate the creation of the server system pool as shown in Figure 9-97, then click **Next**.

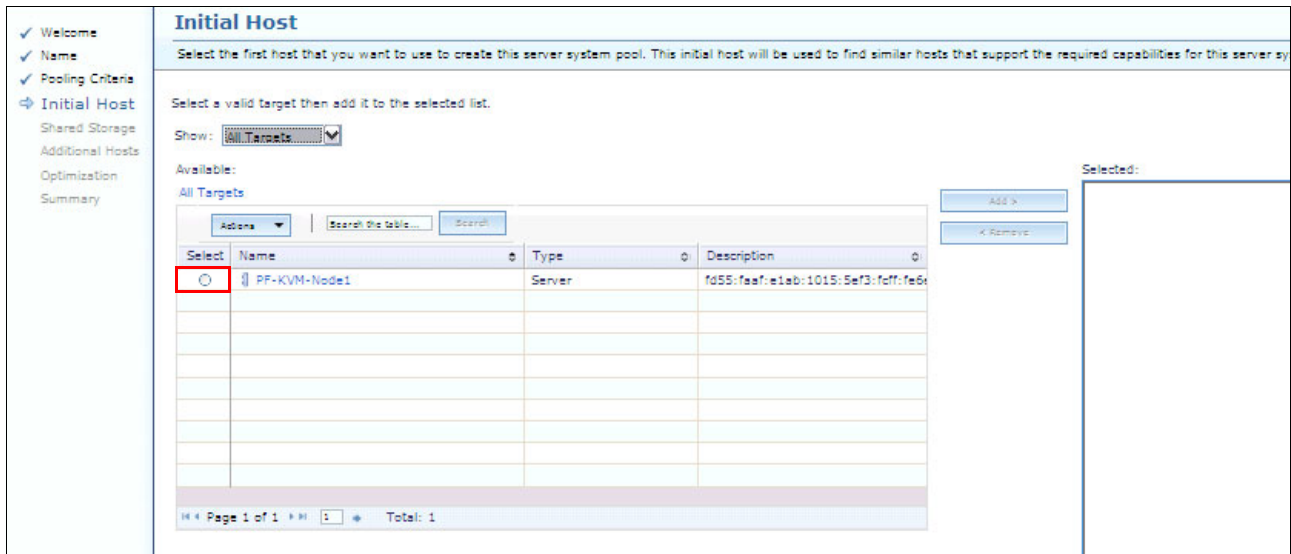


Figure 9-97 Initial server system pool host

7. Select the shared storage system that you want to assign to the server system pool as shown in Figure 9-98, then click **Next**.

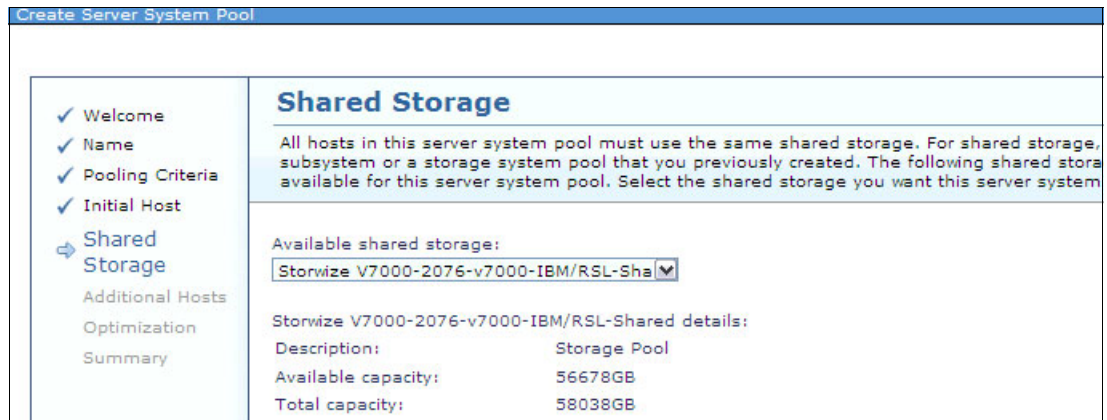


Figure 9-98 Assigning a storage system to the server system pool

8. If you have several hosts, you can add more hosts to your server system pool. This configuration increases high availability and the amount of resources that are shared in your system pool as shown in Figure 9-99. Click **Next**.

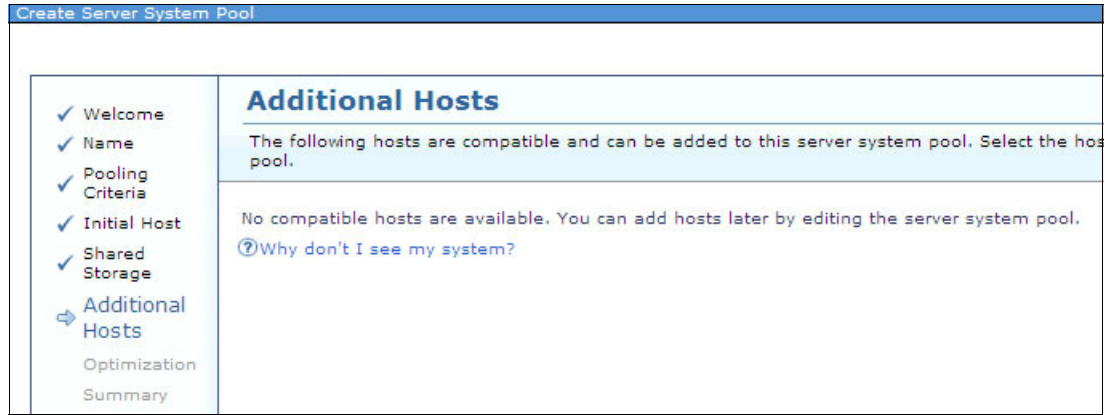


Figure 9-99 Additional Hosts window

9. Select the optimization mode that you want as shown in Figure 9-100, then click **Next**.

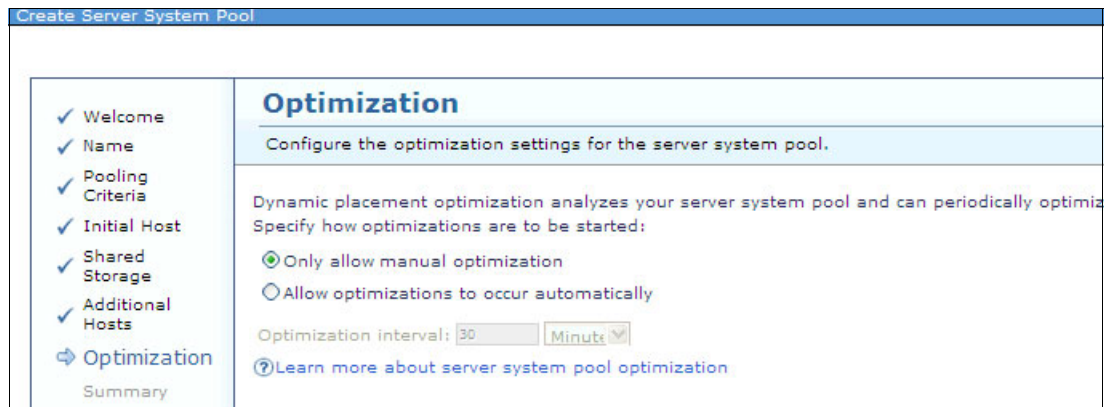


Figure 9-100 Server system pool optimization choice

10. Check the summary as shown in Figure 9-101.

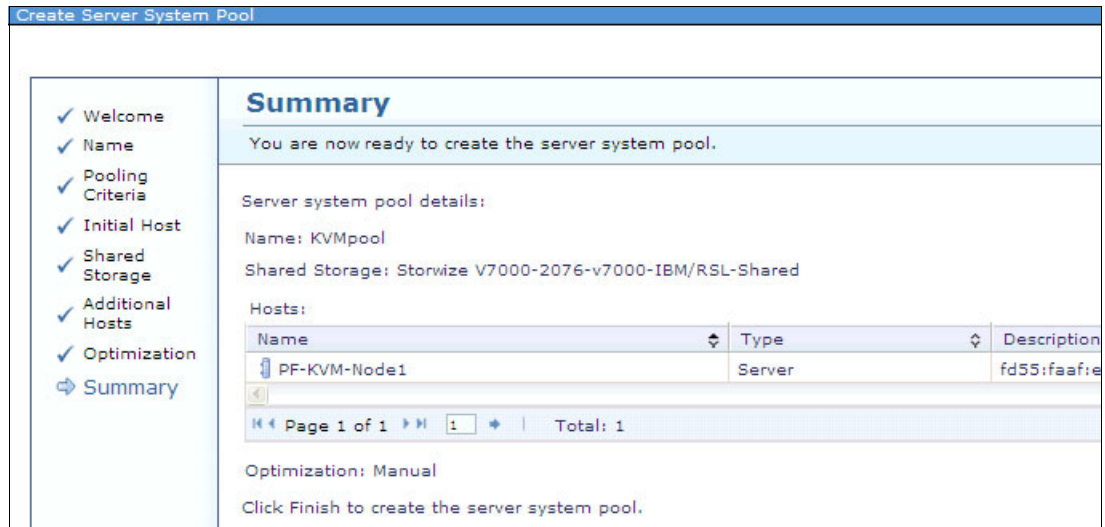


Figure 9-101 Checking the summary

Click **Display Properties** to check the creation progress as shown in Figure 9-102.

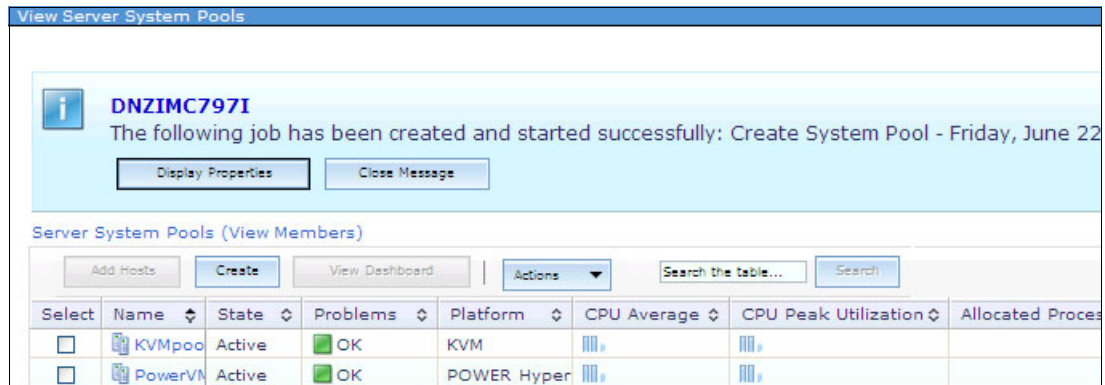


Figure 9-102 Display Properties window

When the creation is complete, go back to **View Server System Pools** and check that the new server system pool is available as shown in Figure 9-103.

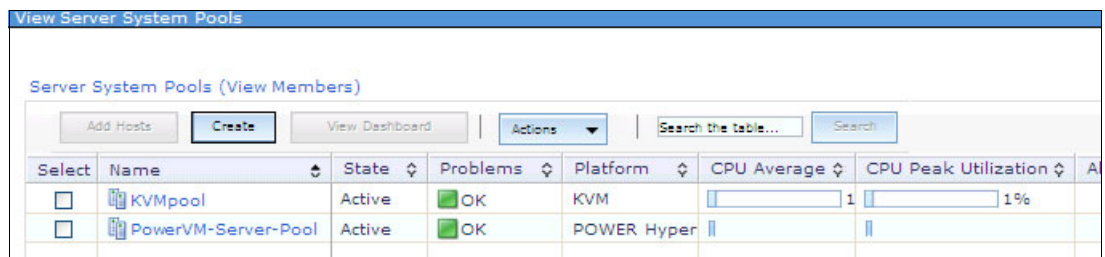
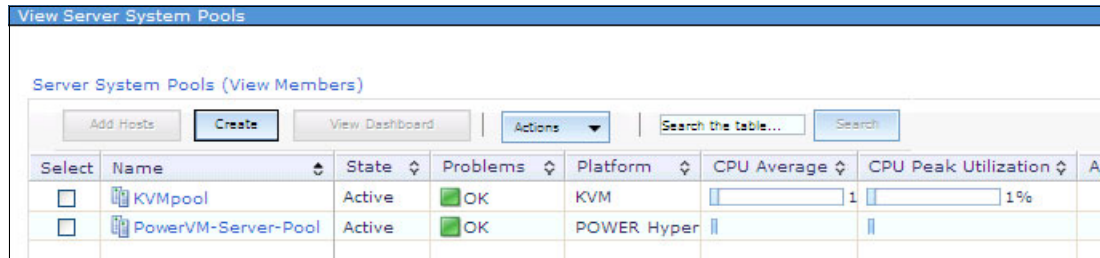


Figure 9-103 Server system pools view

9.7 Add host to an existing server system pool

To increase your server system pool availability or to increase the amount of resources in your pool, add processors and memory. To add memory and processors, add more KVM hosts in your server system pool. To add more hosts, perform these steps:

1. Go to the server system pools view as shown in Figure 9-104.

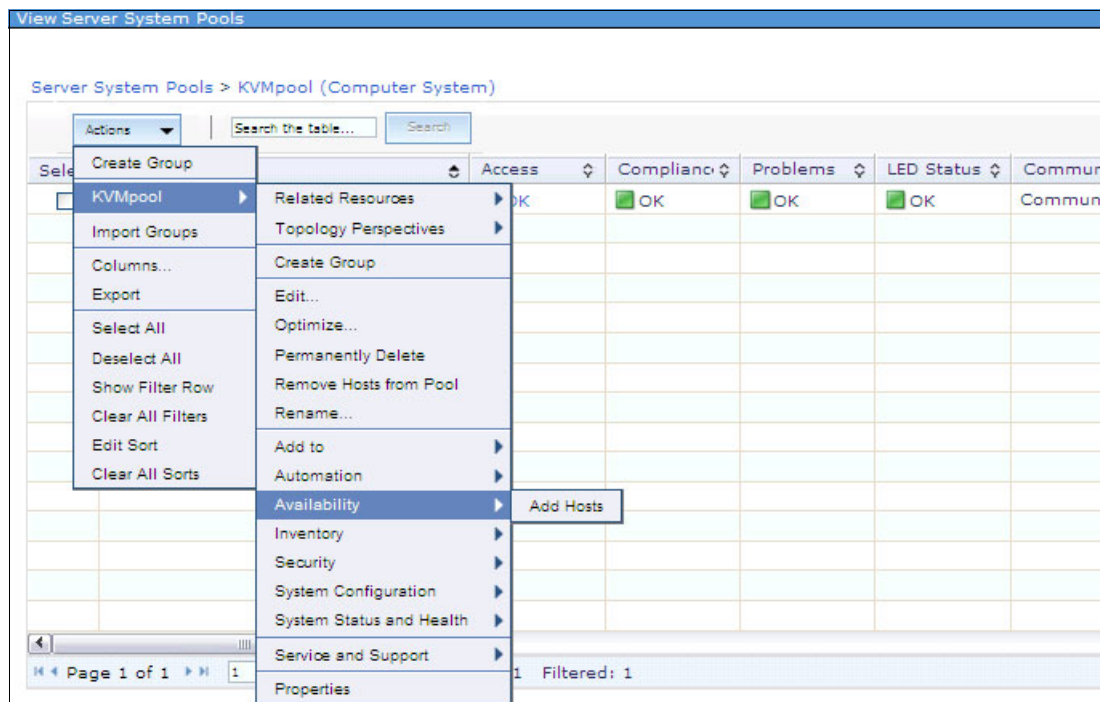


The screenshot shows the 'View Server System Pools' interface. At the top, there are buttons for 'Add Hosts', 'Create', and 'View Dashboard'. Below these is a search bar and an 'Actions' dropdown menu. The main part of the interface is a table with the following columns: Select, Name, State, Problems, Platform, CPU Average, and CPU Peak Utilization. There are two rows of data:

Select	Name	State	Problems	Platform	CPU Average	CPU Peak Utilization
<input type="checkbox"/>	KVMpool	Active	OK	KVM	1	1%
<input type="checkbox"/>	PowerVM-Server-Pool	Active	OK	POWER Hyper		

Figure 9-104 Add hosts to a server system pool

2. Select your server system pool, and click **KVMpool** → **Availability** → **Add Hosts** as shown in Figure 9-105.



The screenshot shows the 'View Server System Pools' interface with the 'KVMpool' row selected. A context menu is open over the 'KVMpool' row, showing various actions. The 'Availability' option is selected, and a sub-menu is open showing the 'Add Hosts' option. The table has the following columns: Select, Name, Access, Compliance, Problems, LED Status, and Community. The 'KVMpool' row has the following values: in the Select column, 'KVMpool' in the Name column, 'OK' in the Access column, 'OK' in the Compliance column, 'OK' in the Problems column, 'OK' in the LED Status column, and 'Commun' in the Community column.

Select	Name	Access	Compliance	Problems	LED Status	Community
<input type="checkbox"/>	KVMpool	OK	OK	OK	OK	Commun

Figure 9-105 Selecting Add Hosts

3. The Welcome window for adding hosts opens as shown in Figure 9-106. Click **Next**.




Figure 9-106 Welcome to the Add Host wizard window

4. Select the host that you want to add to your server system pool as shown in Figure 9-107, then click **Next**.




Figure 9-107 Select hosts to add to an existing server system pool

5. Review the summary and click **Finish** as shown in Figure 9-108.




Figure 9-108 Add hosts summary

6. Click **Display Properties** as shown in Figure 9-109.

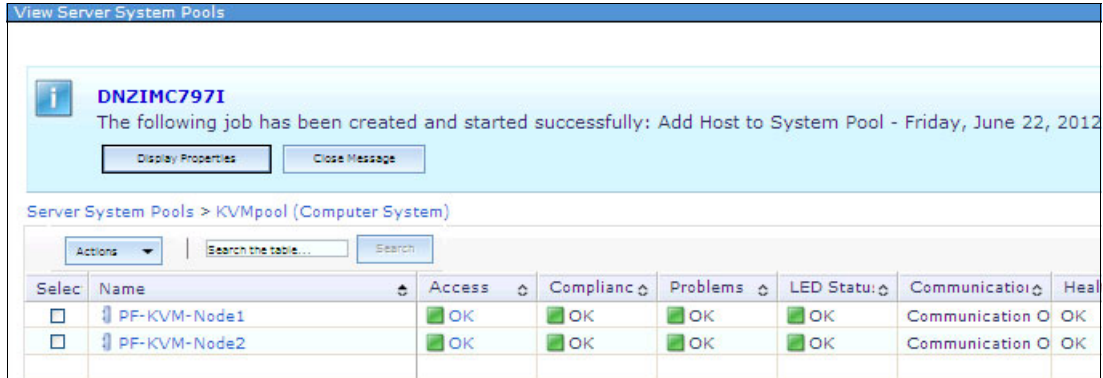


Figure 9-109 Display job properties

Wait until the job is complete as shown in Figure 9-110.

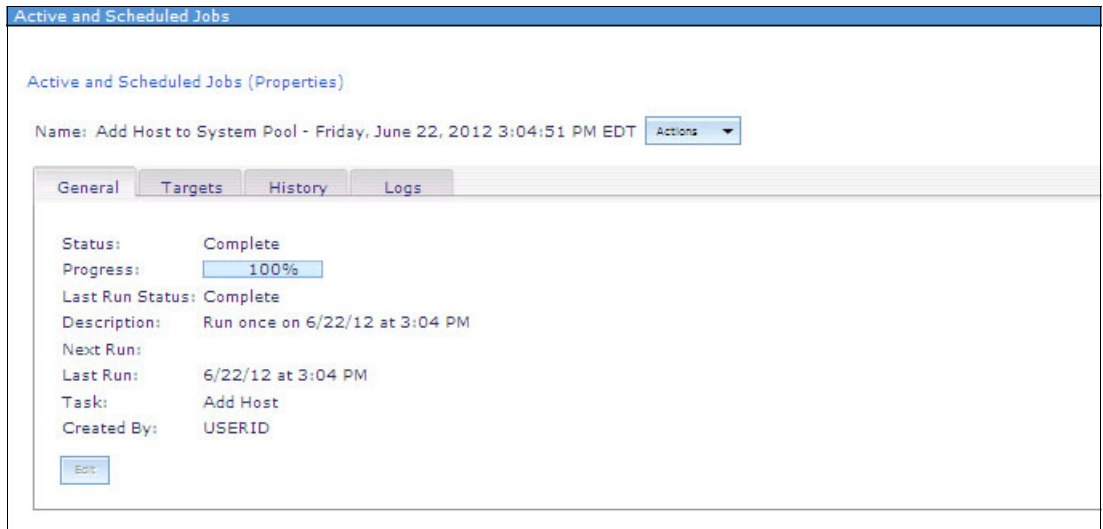


Figure 9-110 Job complete

7. Go back to **Server system pools** view and click your KVM pool. You can see now that you have an additional KVM host in your server system pool as shown in Figure 9-111.

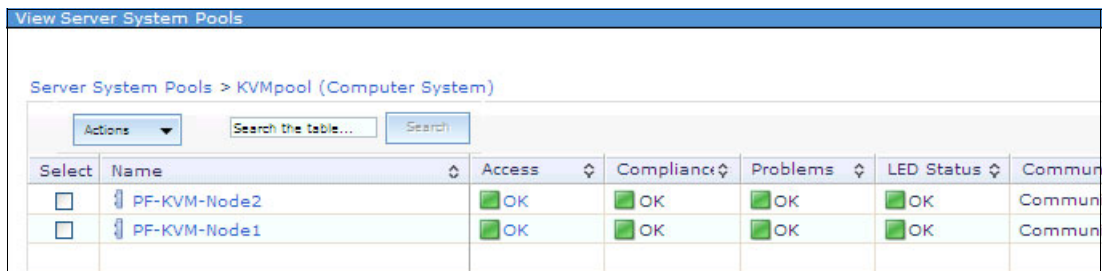


Figure 9-111 KVM server system pool

9.8 Operating a KVM virtual infrastructure

The following tasks can be performed when you are operating the KVM infrastructure:

- ▶ Importing a virtual appliance
- ▶ Deploy a virtual appliance to create a virtual server
- ▶ Capturing a virtual appliance
- ▶ Relocate virtual servers

9.8.1 Importing a virtual appliance

To import virtual appliances, perform these steps:

1. Go to the VMControl plug-in tab as shown in Figure 9-112.



Figure 9-112 VMControl plug-in main window

2. Click the **Virtual Appliances** tab as shown in Figure 9-113, then click **Import** in the Common tasks list.



Figure 9-113 Virtual Appliances tab

3. You are redirected to the import appliances Welcome window as shown in Figure 9-114. Click **Next** to start the import process.

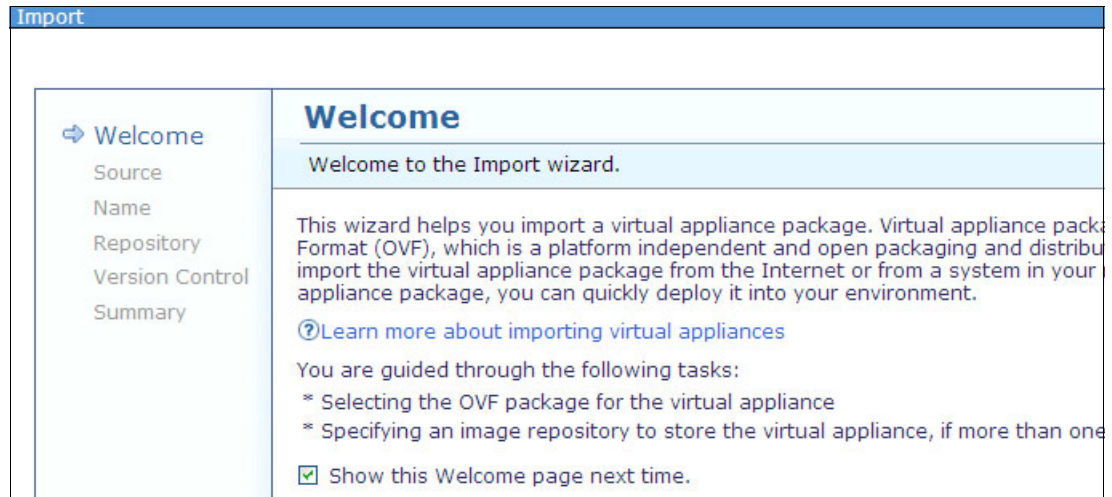


Figure 9-114 Import appliances Welcome window

4. Enter the path to import your appliance. In this example, import the appliance from an http server as shown in Figure 9-115. Click **Next**.

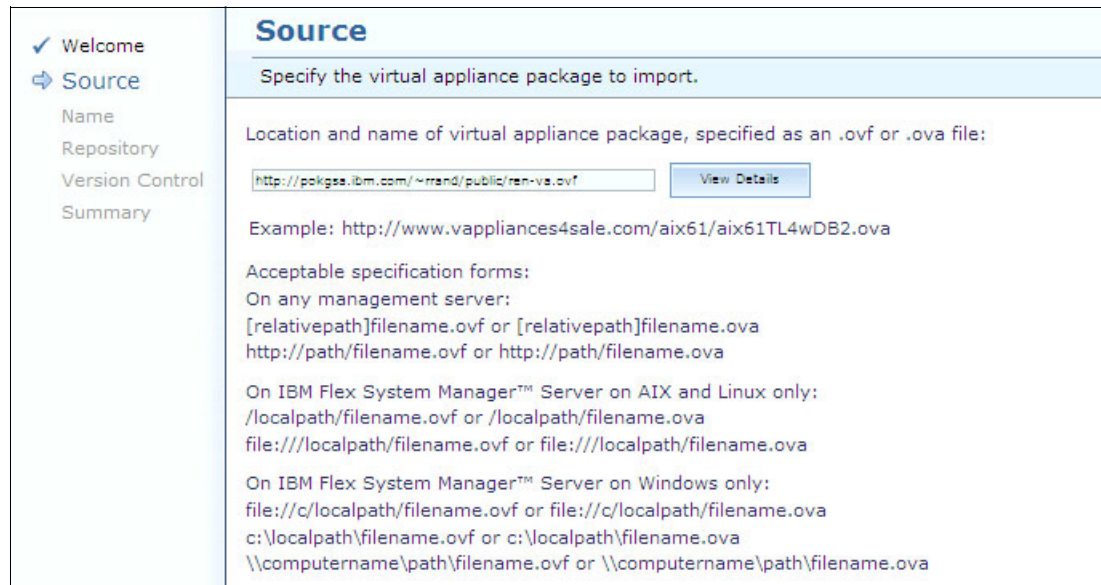


Figure 9-115 Import appliance source path file

- If you do not want to import the digital signature, select the **Import without digital signature** check box as shown in Figure 9-116, then click **Next**.



Figure 9-116 Digital Signature window

- By default, the original appliance name is assigned, as shown in Figure 9-117, but you can specify another one. Click **Next**.



Figure 9-117 Assign a name to a virtual appliance

- Select the image repository where you want to import the virtual appliance as shown in Figure 9-118, then click **Next**.



Figure 9-118 Appliance repository

- You can create a version tree for the imported appliance as shown in Figure 9-119. Or, you can add it under an existing tree as a child appliance of an existing one. Click **Next** to continue.



Figure 9-119 Version Control window

- Review the summary as shown in Figure 9-120, then click **Next**.




Figure 9-120 Import appliance summary

- Select **Run Now** and click **OK** to start the virtual appliance import process as shown in Figure 9-121.




Figure 9-121 Run appliance import now

11. Click **Display Properties** to open the Active and Scheduled Jobs window as shown in Figure 9-122.

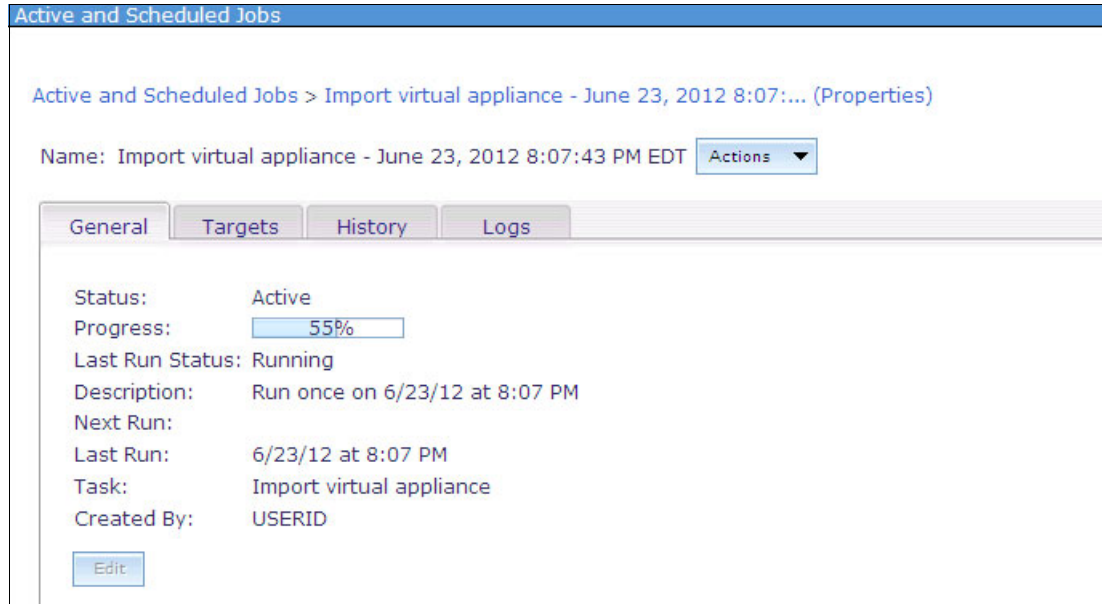


Figure 9-122 Import appliance job progress

Remember: Appliance import times can vary depending on where the appliance is in the network.

12. When the import process is complete, go back to the **Virtual Appliances** tab as shown in Figure 9-123. Check that the new appliance is available for deployment.

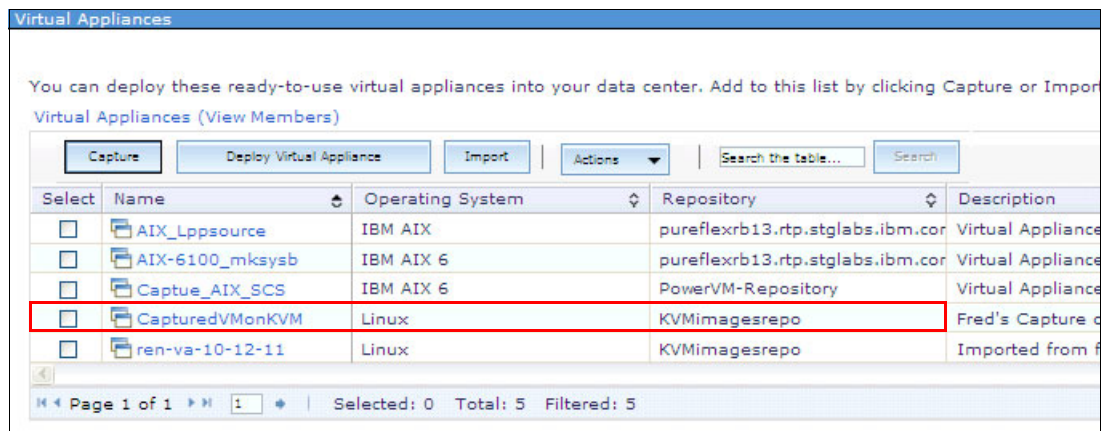


Figure 9-123 Virtual Appliances view

9.8.2 Deploy a virtual appliance to create a virtual server

You can deploy virtual appliances in the Linux KVM virtualization environment on IBM Flex System Manager VMControl to new or existing virtual servers, or to server system pools.

For more information about deployment requirements and limitations, see the IBM Flex System Information Center at this website:

http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp?topic=%2Fco.ibm.director.vim.helps.doc%2Ffsd0_vim_r_kvm_deploy_reqs.html

To deploy a virtual appliance, perform these steps:

1. Go to the main page of the VMControl Plug-in as shown in Figure 9-124. Click **Deploy virtual appliance** in the Common tasks list.



Figure 9-124 VMControl plug-in main window

2. The Welcome window for the Deploy Virtual Appliance opens as shown in Figure 9-125. Click **Next**.




Figure 9-125 Deploy Virtual Appliance Welcome window

3. Select the appliance that you want to deploy as shown in Figure 9-126, then click **Next**.

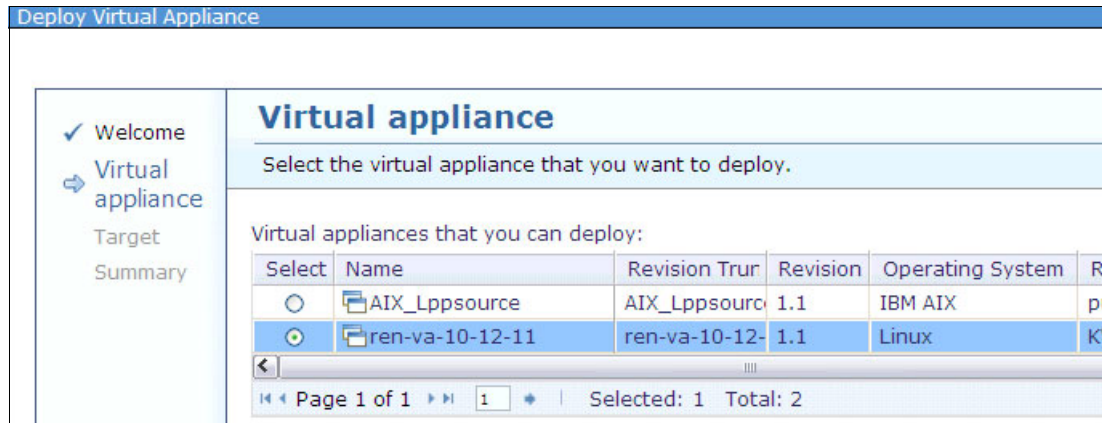


Figure 9-126 Virtual appliances catalog

Remember: You can see the appliances that are available from different image repositories, as shown in Figure 9-126.

4. Select the target location where you want to deploy the new virtual server as shown in Figure 9-127 and click **Next**.




Figure 9-127 Target to deploy a virtual server

Tip: You can choose to deploy an appliance on an existing virtual server. A virtual server is a virtual machine with processor, RAM, and hard disk drive (HDD) resources on which you can install an OS or deploy a virtual appliance.

5. Specify a workload name as shown in Figure 9-128 and click **Next**.

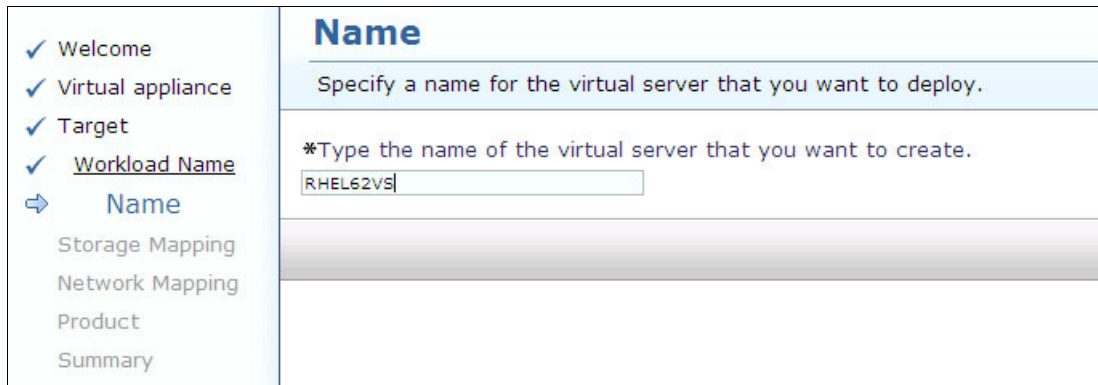


The screenshot shows a configuration window titled "Workload Name". On the left is a navigation pane with the following items: "Welcome" (checked), "Virtual appliance" (checked), "Target" (checked), "Workload Name" (selected with a blue arrow), "Name", "Storage Mapping", "Network Mapping", "Product", and "Summary". The main content area has a header "Workload Name" and a sub-header "A workload is created as a result of deploying the virtual appliance." Below this is a text input field with the value "RHEL62deployed" and a label "*Specify a unique name for the workload." The input field is highlighted with a light blue border.

Figure 9-128 Workload Name window

Clarification: A Workload in FSM is a group that contains one or several virtual servers.

6. Specify a name for your virtual server as shown in Figure 9-129, then click **Next**.



The screenshot shows a configuration window titled "Name". On the left is a navigation pane with the following items: "Welcome" (checked), "Virtual appliance" (checked), "Target" (checked), "Workload Name" (checked), "Name" (selected with a blue arrow), "Storage Mapping", "Network Mapping", "Product", and "Summary". The main content area has a header "Name" and a sub-header "Specify a name for the virtual server that you want to deploy." Below this is a text input field with the value "RHEL62VS" and a label "*Type the name of the virtual server that you want to create." The input field is highlighted with a light blue border.

Figure 9-129 Virtual server naming

- Assign a disk from a storage pool or from a storage volume, as shown in Figure 9-130. Select **Assign to Storage Pool**, then click **Next**.

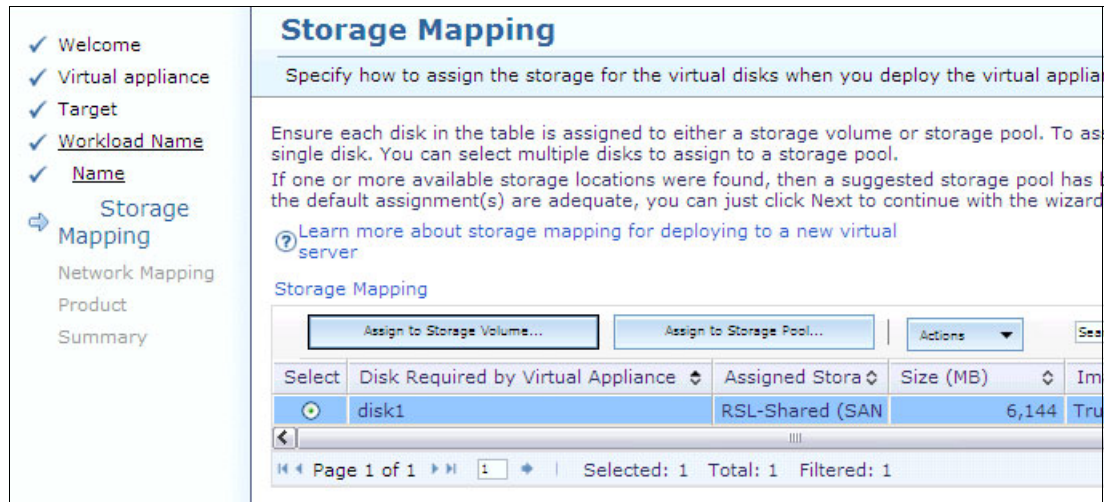


Figure 9-130 Assigning a disk

Tip: By default, the storage volume (disk1) is selected, as shown in Figure 9-130.

- Select the storage pool that you configured previously, as shown in Figure 9-131, then click **OK**.




Figure 9-131 Selecting the storage pool

- The Storage Mapping GUI is displayed as shown in Figure 9-132. Note that disk1 is not selected now because a disk is created automatically in your storage pool. Click **Next**.




Figure 9-132 Storage Mapping window

- Select the virtual network adapter to create a virtual Ethernet adapter on your virtual server, as shown in Figure 9-133, then click **Next**.

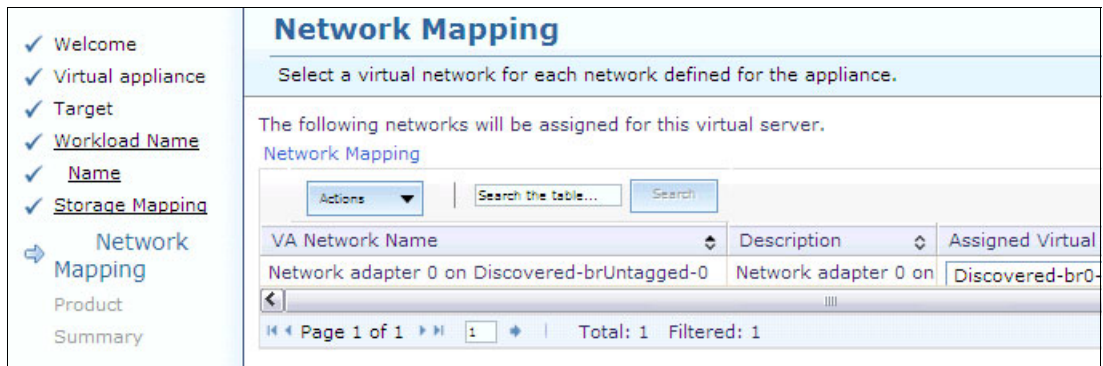


Figure 9-133 Network Mapping window

11. You can preconfigure several parameters that include the name and the network configuration of your virtual server as shown in Figure 9-134. Click **Next**.

Figure 9-134 Virtual server preconfiguration

12. Review the summary and click **Finish** as shown in Figure 9-135.

Figure 9-135 Virtual server creation summary

13. Click **OK** to start the creation of your virtual server workload member as shown in Figure 9-136.




Figure 9-136 Run job

14. Click **Display Properties** as shown in Figure 9-137.

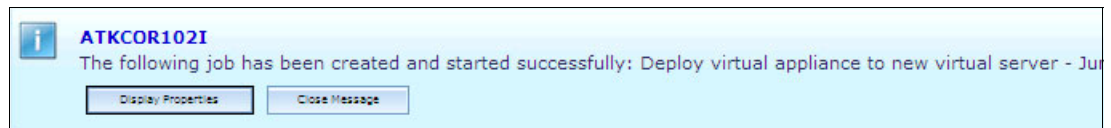


Figure 9-137 Display Properties window

Virtual server creation is complete as shown in Figure 9-138.

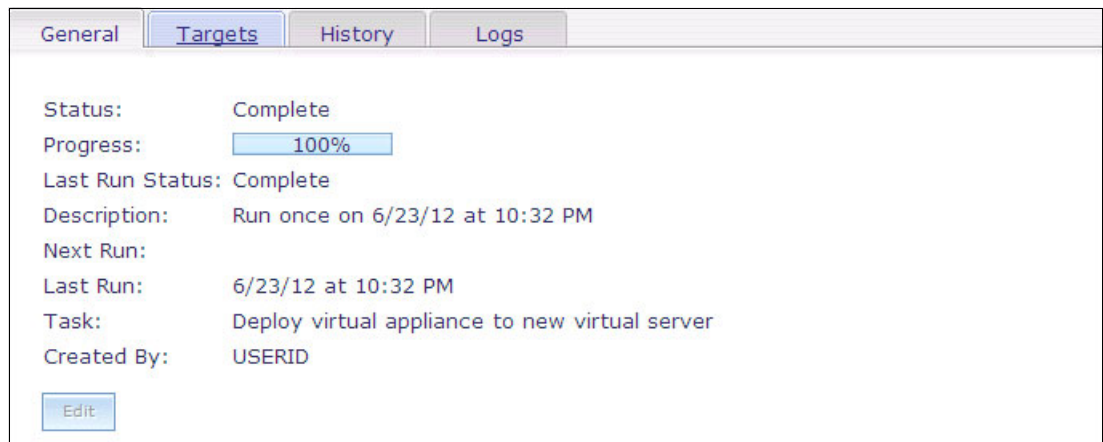


Figure 9-138 Virtual server creation complete

15. Go back to the **Virtual Servers and Hosts** view to check that your new server is deployed as shown in Figure 9-139.

Select	Name	State	OS Name	OS Type and Version	Access
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008	OK
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008	OK
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK
<input type="checkbox"/>	vm003	Stopped			OK
<input type="checkbox"/>	PF-HyperV-Node1	Started	PF-HyperV01	Windows® Server 2008	OK
<input type="checkbox"/>	VMWindowsHyperV	Started			OK
<input type="checkbox"/>	PF-KVM-Node1	Started	PF-KVM01	Linux 6.2	OK
<input type="checkbox"/>	RHEL62VS	Started			OK
<input type="checkbox"/>	VMRHEL62x86temp	Started			OK
<input type="checkbox"/>	PF-KVM-Node2	Started	PF-KVM02	Linux 6.2	OK
<input type="checkbox"/>	RHEL62vmForVNC	Started			OK
<input type="checkbox"/>	vmRHEL62	Stopped			OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK
<input type="checkbox"/>	PF-Node1-NIM	Started	PF-Node1-NIM	AIX 6.1	OK

Figure 9-139 Virtual Servers and Hosts view

9.8.3 Capturing a virtual appliance

To capture a virtual appliance, perform these steps:

1. Click the **Virtual Appliances** tab in the VMControl plug-in main window, then click **Capture** in the Common tasks list as shown in Figure 9-140.

Basics	Workloads	Virtual Appliances	System Pools	Virtual Servers/Hosts
What to deploy: 2 Virtual appliances		Where to deploy: 8 Existing virtual servers 2 Hosts and 2 server system pools		Common tasks Deploy virtual appliance Capture Import View active and scheduled jobs View virtual appliance versions Create image repository
What to capture: 1 Workloads 6 Virtual servers and operating systems		Where to store: 3 Image repositories		

Figure 9-140 VMControl main page

2. The Welcome capture window opens as shown in Figure 9-141. Click **Next**.




Figure 9-141 Capture Welcome window

3. Specify a name for the virtual appliance as shown in Figure 9-142, then click **Next**.



Figure 9-142 Appliance capture name

4. Select the source type to capture as shown in Figure 9-143, then click **Next**.



Figure 9-143 Source type to capture

5. Select the virtual server that you want to capture as shown in Figure 9-144.

Select the virtual server to capture.

Select	Name	State	Access	Problems	C
<input type="radio"/>	PF-Node1-NIM	Started	OK	OK	
<input type="radio"/>	PF-Node1-Test02	Stopped	Offline	Information	
<input type="radio"/>	RHEL62vmForVNC	Started	OK	OK	
<input checked="" type="radio"/>	RHEL62VS	Started	OK	OK	
<input type="radio"/>	vmRHEL62	Stopped	OK	OK	
<input type="radio"/>	VMRHEL62x86template	Started	OK	OK	

Page 1 of 1 | Selected: 1 Total: 6

Figure 9-144 Select server to capture

6. Select the image repository where you want to put the appliance that is generated by the capture process as shown in Figure 9-145.

Select the repository where you want to store the image that is associated with the

Repositories that are capable of storing the image associated with the new virtual a

Image Repositories

Actions | Search the table... Search

Select	Name	Image Count	Managed By	Des
<input type="radio"/>	PowerVM-Repository	0	SN101D88B_VIOS1	Ima
<input checked="" type="radio"/>	KVMimagesrepo	1	PF-KVM03	Ima

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

Figure 9-145 Repository to capture

- Select the disk that you want to capture from your existing virtual server as shown in Figure 9-146.

Disks

Specify the disks and disk images to be captured. Selecting a disk captures info disk image additionally captures the disk contents.

By default all compatible disks and their associated disk image contents are select choose to exclude a disk or disk image from the capture. The resulting virtual app needs to create an operational virtual server when it is deployed. For example, th image is required.

[Learn more about capturing disks and disk images](#)

Disks and Images to Capture

Actions | Search the table... Search

Select	Disk Name	Storage Server	Size (MB)
<input checked="" type="checkbox"/>	RHEL62VS2895D0	Storwize V7000-2076-v7000-IBM	6144

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

Figure 9-146 Select the disk for capture

- Select the network mapping for your future appliance as shown in Figure 9-147.

Network Mapping

Specify a description to use for each virtual network

Network Mapping

Actions | Search the table... Search

Network	Description
Discovered-br0-0	Network adapter 0 on Discov

Page 1 of 1 | 1 | Total: 1 Filtered: 1

Figure 9-147 Network Mapping window

9. If no operating system was discovered from the original virtual server, you must specify the type of operating system as shown in Figure 9-148.




Figure 9-148 Virtual server captured operating system

10. Select the version control type for your future virtual appliance as shown in Figure 9-149, then click **Next**.




Figure 9-149 Version Control window

11. Review the summary as shown in Figure 9-150. You cannot capture a running server.

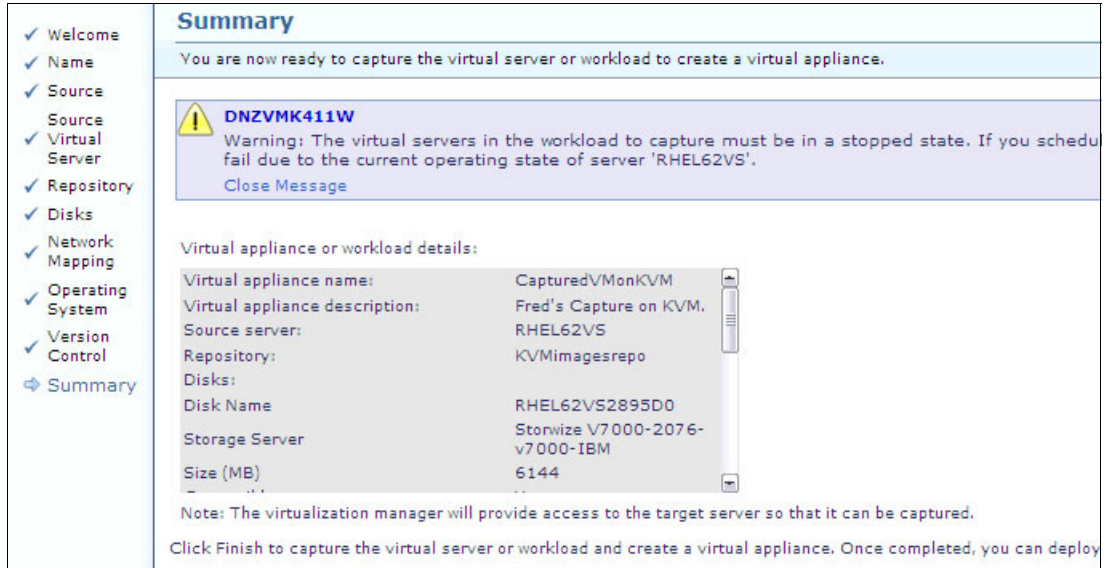


Figure 9-150 Virtual server capture summary

12. Go back to the Virtual Servers and Hosts window. Power off the virtual server that you want to capture by right-clicking the name and clicking **Power On/Off** → **Power Off Now**, as shown in Figure 9-151.




Figure 9-151 Power off the server for capture

13. Check that the virtual server is stopped as shown in Figure 9-152.

Select	Name	State	OS Name	OS Type and Version	Access	Problems	Compliance
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	Information	OK
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008	OK	Information	OK
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008	OK	Information	OK
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	Information	OK
<input type="checkbox"/>	vm003	Stopped			OK	Information	OK
<input type="checkbox"/>	PF-HyperV-Node1	Started	PF-HyperV01	Windows® Server 2008	OK	Information	OK
<input type="checkbox"/>	VMWindowsHyperV	Started			OK	Information	OK
<input type="checkbox"/>	PF-KVM-Node1	Started	PF-KVM01	Linux 6.2	OK	OK	OK
<input checked="" type="checkbox"/>	RHEL62VS	Stopped			OK	OK	OK
<input type="checkbox"/>	VMRHEL62x86tem	Started			OK	OK	OK
<input type="checkbox"/>	PF-KVM-Node2	Started	PF-KVM02	Linux 6.2	OK	OK	OK
<input type="checkbox"/>	RHEL62vmForVNC	Started			OK	OK	OK
<input type="checkbox"/>	vmRHEL62	Stopped			OK	OK	OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK	OK	OK
<input type="checkbox"/>	PF-Node1-NIM	Started	PF-Node1-NIM	AIX 6.1	OK	OK	OK

Figure 9-152 Virtual server stopped for capture

14. Go back to the Summary window and close the Warning information window as shown in Figure 9-153, then click **Finish**.

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:	CapturedVMonKVM
Virtual appliance description:	Fred's Capture on KVM.
Source server:	RHEL62VS
Repository:	KVMimagesrepo
Disks:	
Disk Name	RHEL62VS2895D0
Storage Server	Storwize V7000-2076-v7000-IBM
Size (MB)	6144
Compatible	Yes
Include Image	Yes

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 complete your environment.

Figure 9-153 Capture summary

15. Click **OK** to run the job as shown in Figure 9-154 and begin to capture your virtual server.

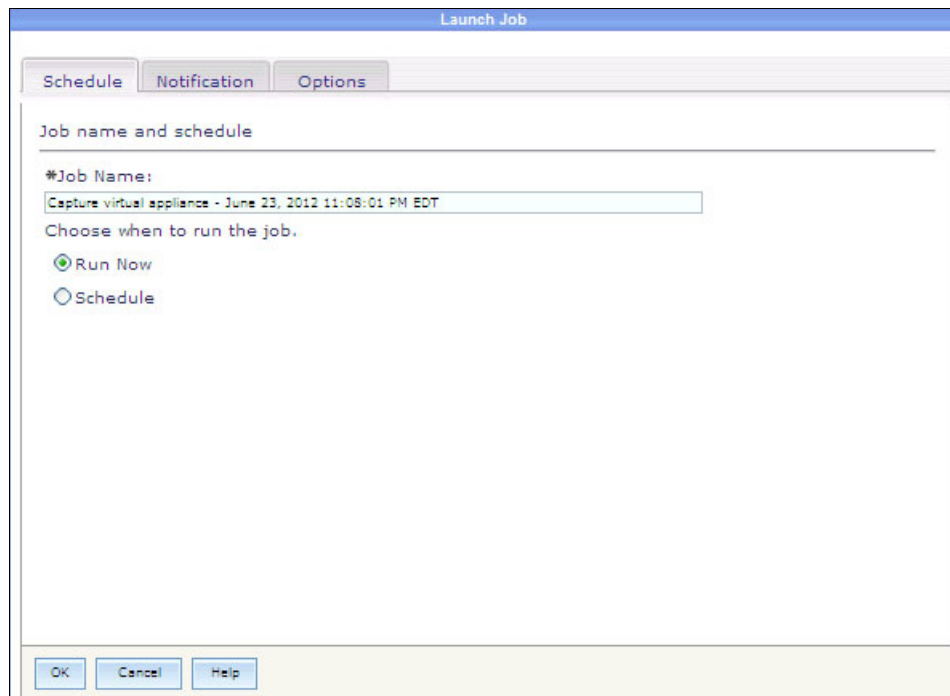


Figure 9-154 Run job now

16. Click **Display Properties** to check the job status as shown in Figure 9-155.

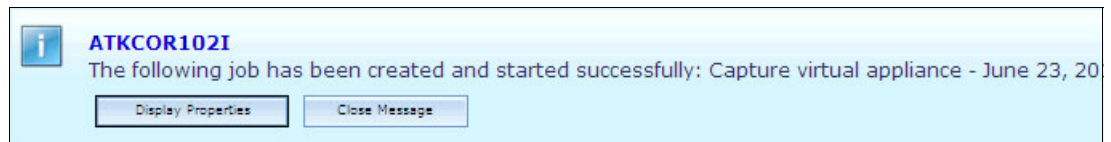


Figure 9-155 Capture job status

Wait until the job is complete as shown in Figure 9-156.




Figure 9-156 Capture complete

17. Go back to the VMControl main tab and check that your appliance count is incremented as shown in Figure 9-157. Click **Virtual Appliances**.




Figure 9-157 VMControl main window

In the Virtual Appliances window, you can see the new virtual appliance as shown in Figure 9-158.




Figure 9-158 Virtual Appliances window

9.8.4 Relocate virtual servers

You can migrate your virtual server from a physical KVM host to another physical KVM host. To do so, perform these steps:

1. From the VMControl plug-in main page, click the **Virtual Servers/Hosts** tab.
2. Click the **Virtual servers and hosts** link under the Common tasks, as shown in Figure 9-159.

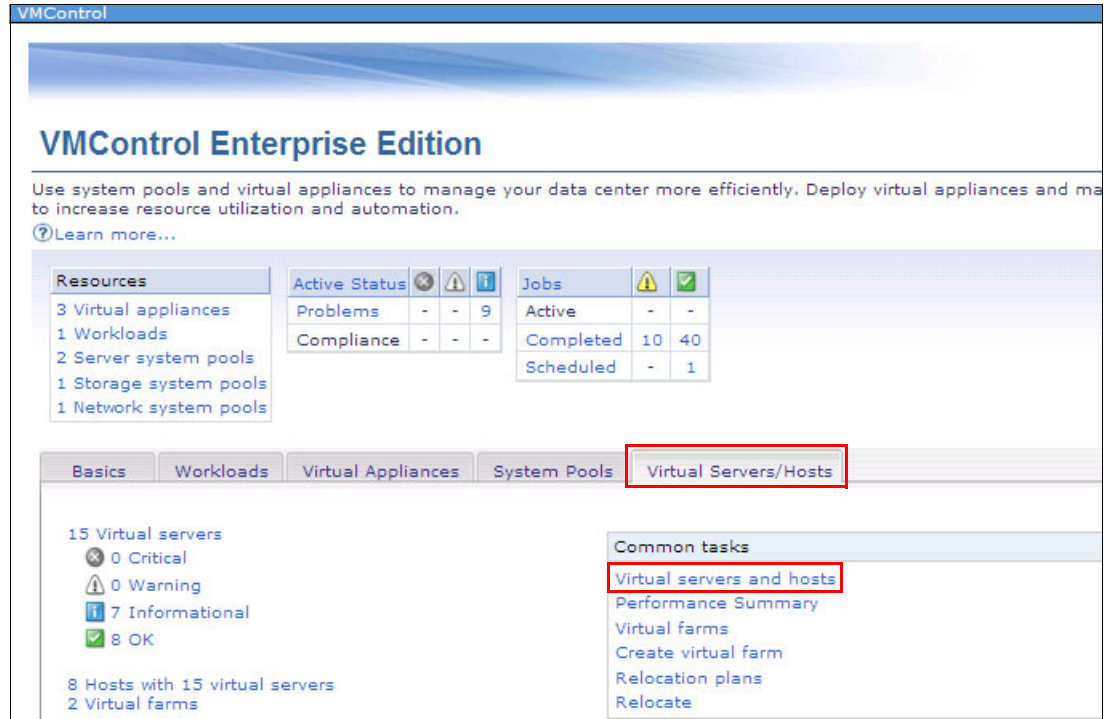


Figure 9-159 VMControl plug-in main page: Virtual Servers/Hosts tab

The Virtual Servers and Hosts window opens as shown in Figure 9-160.

Virtual Servers and Hosts (View Members)

Select	Name	State	OS Name	OS Type and Version	Access	Problems
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008 6	OK	Information
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008 6	OK	Information
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm003	Stopped			OK	Information
<input type="checkbox"/>	PF-HyperV-Node1	Started	PF-HyperV01	Windows® Server 2008 6	OK	Information
<input type="checkbox"/>	VMWindowsHyperV	Started			OK	Information
<input type="checkbox"/>	PF-KVM-Node1	Started	PF-KVM01	Linux 6.2	OK	OK
<input type="checkbox"/>	RHEL62VS	Stopped			OK	OK
<input type="checkbox"/>	VMRHEL62x86temp	Started			OK	OK
<input type="checkbox"/>	PF-KVM-Node2	Started	PF-KVM02	Linux 6.2	OK	OK
<input type="checkbox"/>	RHEL62vmForVNC	Started			OK	OK
<input type="checkbox"/>	vmRHEL62	Stopped			OK	OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK	OK
<input type="checkbox"/>	PF-Node1-NIM	Started	PF-Node1-NIM	AIX 6.1	OK	OK

Page 1 of 2 Selected: 0 Total: 23 Filtered: 23

Figure 9-160 Virtual Servers and Hosts window

3. Right-click your virtual server and select **Availability** → **Relocate** as shown in Figure 9-161.

Virtual Servers and Hosts (View Members)

Select	Name	State	OS Name	OS Type and Version	Access	Problems
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008 6	OK	Information
<input type="checkbox"/>	vm002			Windows® Server 2008 6	OK	Information
<input type="checkbox"/>	PF-ESXi-Node2			Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm003				OK	Information
<input type="checkbox"/>	PF-HyperV-Node1			Windows® Server 2008 6	OK	Information
<input type="checkbox"/>	VMWindow				OK	Information
<input type="checkbox"/>	PF-KVM-Node1			Linux 6.2	OK	OK
<input checked="" type="checkbox"/>	RHEL62VS				OK	OK
<input type="checkbox"/>	VMRHEL62				OK	OK
<input type="checkbox"/>	PF-KVM-Node2			Linux 6.2	OK	OK
<input type="checkbox"/>	RHEL62vm				OK	OK
<input type="checkbox"/>	vmRHEL62				OK	OK
<input type="checkbox"/>	PF-PowerVM-N				OK	OK
<input type="checkbox"/>	PF-Node1			AIX 6.1	OK	OK

Page 1 of 2 Selected: 23

Figure 9-161 Select server for relocation

4. A confirmation window opens to confirm that your virtual server can be relocated as shown in Figure 9-162. Click **OK** to start the server relocation.




Figure 9-162 Relocate virtual server information window

5. Click **Display Properties** to check the relocation status as shown in Figure 9-163.

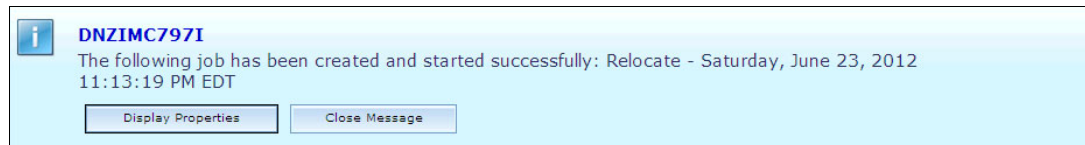


Figure 9-163 Display Properties window

- Go back to the Virtual Servers and Hosts window. Make sure that your virtual server is “hosted” by another KVM member of the system pool as shown in Figure 9-164.

Virtual Servers and Hosts

Virtual Servers and Hosts (View Members)

Performance Summary | Actions | Search the table... Search

Select	Name	State	OS Name	OS Type and Version	Access	Problems
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008 6.	OK	Information
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008 6.	OK	Information
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm003	Stopped			OK	Information
<input type="checkbox"/>	PF-HyperV-Node1	Started	PF-HyperV01	Windows® Server 2008 6.	OK	Information
<input type="checkbox"/>	VMWindowsHyperV	Started			OK	Information
<input type="checkbox"/>	PF-KVM-Node1	Started	PF-KVM01	Linux 6.2	OK	OK
<input type="checkbox"/>	VMRHEL62x86templa	Started			OK	OK
<input type="checkbox"/>	PF-KVM-Node2	Started	PF-KVM02	Linux 6.2	OK	OK
<input type="checkbox"/>	RHEL62vmForVNC	Started			OK	OK
<input type="checkbox"/>	RHEL62VS	Stopped			OK	OK
<input type="checkbox"/>	vmRHEL62	Stopped			OK	OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK	OK
<input type="checkbox"/>	PF-Node1-NIM	Started	PF-Node1-NIM	AIX 6.1	OK	OK

Page 1 of 2 | Selected: 0 Total: 23 Filtered: 23

Figure 9-164 Check relocation



Managing the PowerVM environment with IBM Flex System Manager

This chapter addresses how to manage the PowerVM infrastructure through IBM Flex System Manager (FSM). IBM Flex System p260 and p460 compute nodes have the same capabilities as the rack POWER7 system family, which includes Power 770 and Power 750. You can create logical partition profiles, and modify, delete, and activate them. You can use p260 and p460 compute nodes for full partition purposes, and you can set up the Virtual I/O Server (VIOS) environment on the p260 and p460 compute nodes. You can manage the Virtual I/O Controller (VIOC) client image through VMControl, which is a feature in the Flex System Manager. The VMControl feature can capture, deploy, and relocate virtual servers. The included example scenario describes the steps to set up the PowerVM virtualization environment.

This chapter includes the following sections:

- ▶ 10.1, “Initial deployment of virtual machine” on page 402
- ▶ 10.2, “Capturing virtual machines” on page 416
- ▶ 10.3, “Deploying virtual machines” on page 471
- ▶ 10.4, “Relocating virtual machines” on page 486

10.1 Initial deployment of virtual machine

This section describes how to deploy the PowerVM virtual machine through Flex System Manager. For more information about planning for PowerVM infrastructure management, see 5.2.3, “Planning for PowerVM virtualization” on page 111.

10.1.1 Solution architecture

The overall architecture of the PowerVM infrastructure with Flex System Manager is shown in Figure 10-1.

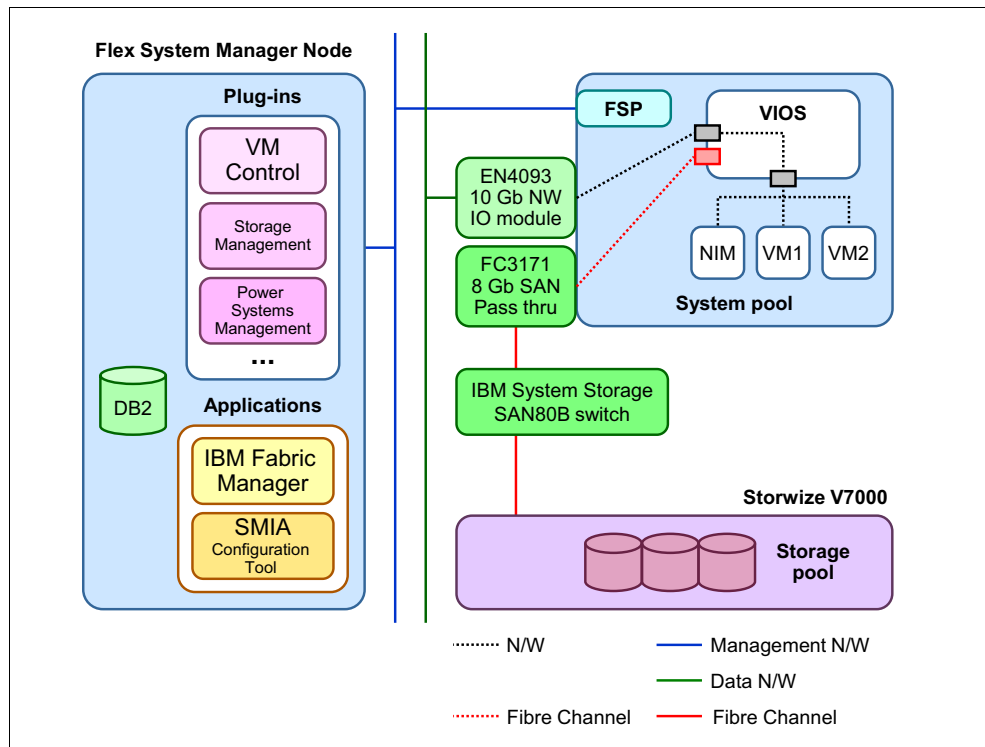


Figure 10-1 Architectural overview with Flex System Manager

The diagram in Figure 10-1 shows the relationships between PowerVM infrastructure components and Flex System Manager. The blue boxes represent physical compute nodes. The purple box stands for Storwize V7000 and the storage pools that are defined on it. The green boxes are I/O modules in the chassis (10 GbE switch, 8 Gb FC Pass-thru module, and external 8 Gb FC SAN switch). A Flex System Manager node box describes components that are needed for the PowerVM implementation. A p260 compute node box represents logical partitions (LPARs), the flexible service processor (FSP), and the Shared Ethernet Adapter (SEA).

10.1.2 Setting up VIOS and Network Installation Manager server

This section shows the initial deployment of PowerVM on the p260 compute node.

You can see physical compute nodes and I/O modules in the Chassis Manager view of Flex System Manager, as shown in Figure 10-2. For more information about managing the chassis, see Chapter 4, “Chassis Management Module operations” on page 41.

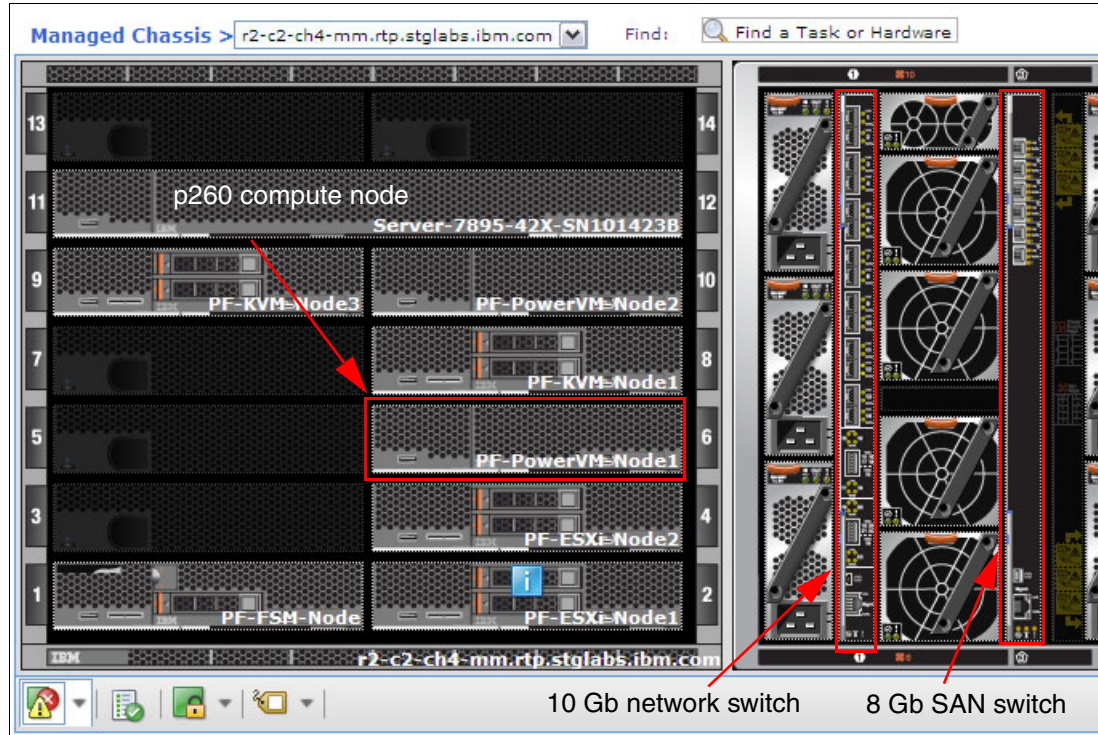


Figure 10-2 Chassis components view

First, install VIOS on a p260 compute node.

If you want more information about how to implement the p260 and p460 compute nodes, see *IBM Flex System p260 and p460 Planning and Implementation Guide*, SG24-7989, at this website:

<http://www.redbooks.ibm.com/abstracts/sg247989.html?Open>

There are two options to install VIOS on the p260 compute node for the first time:

- ▶ Install VIOS by using a DVD through the supported USB optical DVD drive, which is connected to the p260 front panel as shown in Figure 10-3.




Figure 10-3 Front panel of the IBM Flex System p260 compute node

- ▶ Use the Network Installation Manager (NIM) method.

Restriction: This option is only for a NIM server that exists at the site.

Creating and activating the VIOS profile

One of the strengths of FSM is the capability to provide a single point of management. To activate the VIOS profile, perform these steps:

1. Click **Chassis Manager** on the Initial Setup tab.
2. Click the chassis name as shown in Figure 10-4.




Figure 10-4 Installed chassis name

3. Click **Manage Power Systems Resources** in the menu at the far left of the chassis graphical view as shown in Figure 10-5.




Figure 10-5 *Manage Power Systems Resources*

- If you discovered and collected inventory on p260 compute nodes earlier, you see the physical p260 compute nodes in the opened window (Figure 10-6).

Tip: If you do not see any compute node in the chassis, go to 6.1, “IBM Flex System Manager Setup Wizard” on page 123.

Right-click the discovered compute node, and select **System Configuration** → **Create Virtual Server** (Figure 10-6).

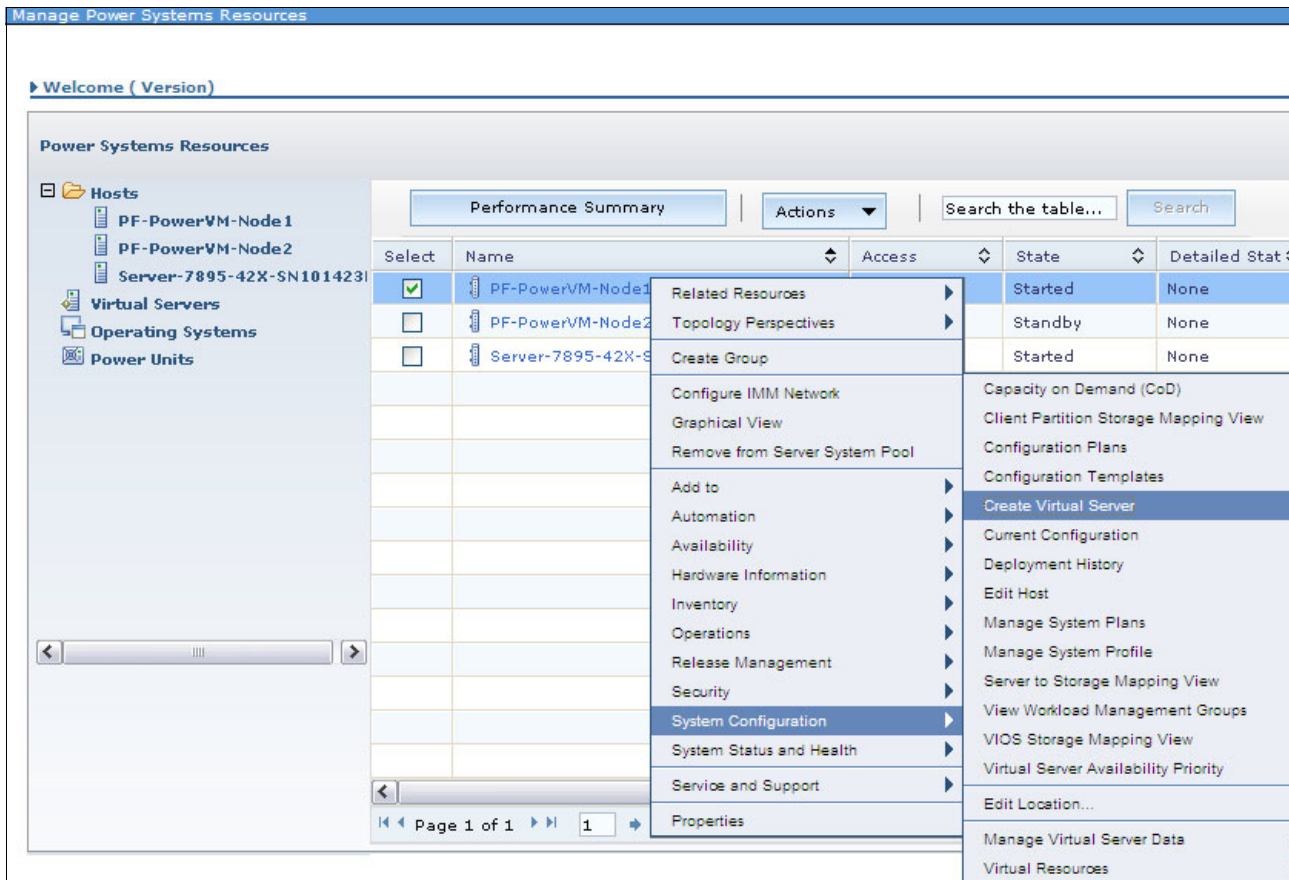


Figure 10-6 Manage Power Systems Resources main window

5. Create the VIOS profile based on your system requirements. Enter the VIOS name, as shown in Figure 10-7.

The screenshot shows the 'Name' configuration step of a wizard. The left sidebar lists steps: Name (selected), Memory, Processor, Ethernet, Physical I/O, and Summary. The main area is titled 'Name' and contains the following fields and options:

- Host name: PF-PowerVM-Node1
- *Virtual server name: SN101D88B_VIOS1
- Virtual server ID: 1
- Environment: AIX/Linux (dropdown)
- Suspend capable
- Remote Restart capable
- Assign all resources to this virtual server.
- Enable virtual trusted platform module (VTPM)
- Warning: The VTPM key is set to default key.

Figure 10-7 Creating the VIOS profile: Name

6. Define the VIOS memory size, as shown in Figure 10-8.

The screenshot shows the 'Memory' configuration step of a wizard. The left sidebar lists steps: Name (checked), Memory (selected), Processor, Ethernet, Virtual Storage Adapters, Physical I/O, and Summary. The main area is titled 'Memory' and contains the following information:

- Select the memory mode and assigned memory for the virtual server.
- Dedicated Memory**
- Total system memory: 32.0 GB
- Memory available: 31.0 GB
- *Assigned memory (GB): 10

Figure 10-8 Creating the VIOS profile: Memory

- Define the processing mode and how many processors are used for VIOS, as shown in Figure 10-9.




Figure 10-9 Creating the VIOS profile: Processor

- Define the virtual Ethernet adapter for the Shared Ethernet Adapter (SEA), as shown in Figure 10-10. Click **Add** to create new adapter or **Edit** to modify an existing adapter.

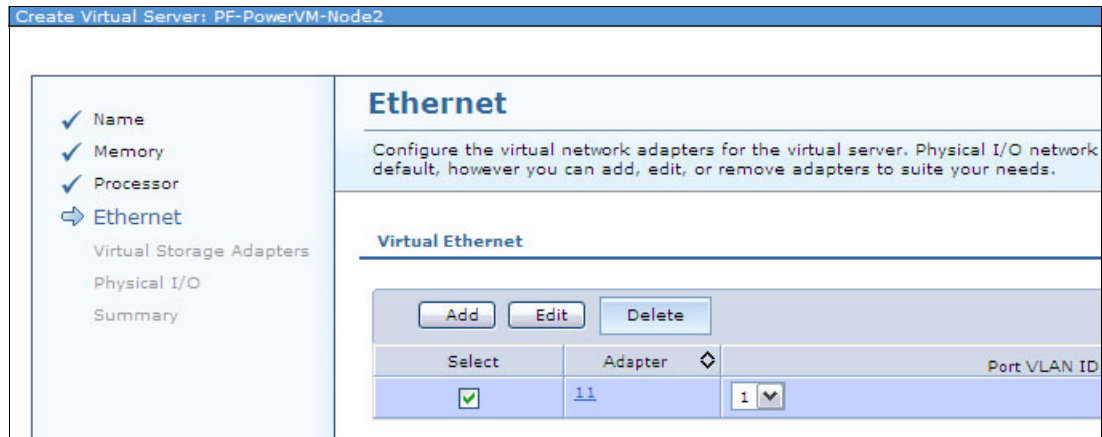


Figure 10-10 Creating the VIOS profile: Ethernet

9. Create one virtual Ethernet adapter for the Shared Ethernet Adapter and assign priority 1 to the VIOS, as shown in Figure 10-11.

Virtual Ethernet - Modify Adapter

Specify an adapter ID and virtual Ethernet for this adapter.

*Adapter Id
11

*Port Virtual Ethernet
1

VSI Type Id
[]

VSI Type Version
[]

VSI Manager Id
[]

IEEE Settings
Select this option to allow additional virtual LAN IDs for the adapter.

IEEE 802.1q compatible adapter

Maximum number of VLANs: 20

Additional VLAN IDs:
[] 1,20,48,...

Shared Ethernet Settings
Select Ethernet bridging to link (bridge) the virtual Ethernet to a physical network.

Use this adapter for Ethernet bridging

Priority:
1 (1 or 2)

Figure 10-11 Creating the VIOS profile: Virtual Ethernet: Modify Adapter

10. Click **Create Adapter** as shown in Figure 10-12.

Virtual Storage Adapters

Specify the virtual storage adapters required for this virtual server.

*Maximum number of virtual adapters : 100

No adapters configured. Select "Create Adapter.." button to create a new virtual adapter.

Create Adapter..

*Note: 1) You can use the Virtual Storage Management task to define the physical block storage Fibre Channel server adapters that the client virtual servers will use for storage access

Navigation menu:
✓ Name
✓ Memory
✓ Processor
✓ Ethernet
Virtual Storage Adapters (selected)
Physical I/O
Summary

Figure 10-12 Creating the VIOS profile: Virtual Storage Adapters

11. Create an adapter for virtual SCSI, as shown in Figure 10-13.




Figure 10-13 Creating the VIOS profile: Create Virtual Adapter

12. Define the connecting VIOC ID and an adapter ID, as shown in Figure 10-14. After you enter parameters for virtual storage adapters, click **OK**.




Figure 10-14 Creating the VIOS profile: Create Virtual Adapter continued

13. You can see the newly created virtual SCSI adapter in the Virtual Storage Adapters pane, as shown in Figure 10-15. If you need to create more virtual storage adapters, click **Add**.

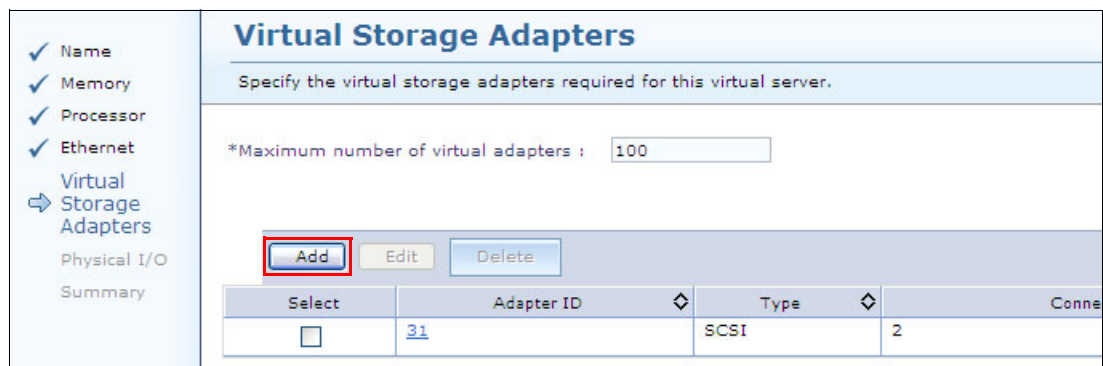


Figure 10-15 Creating the VIOS profile: Virtual Storage Adapters

14. Define the physical I/O adapters, as shown in Figure 10-16. Assign all physical adapters to the VIOS.

Select	Location Code	
<input checked="" type="checkbox"/>	U78AE.001.WZS00T2-P1-C18-L1	Ethernet controller
<input checked="" type="checkbox"/>	U78AE.001.WZS00T2-P1-C19-L1	Fibre Channel Serial Bus
<input checked="" type="checkbox"/>	U78AE.001.WZS00T2-P1-T2	PCI-E SAS Controller
<input checked="" type="checkbox"/>	U78AE.001.WZS00T2-P1-T1	PCI-to-PCI bridge
<input checked="" type="checkbox"/>	U78AE.001.WZS00T2-P1-C18-L2	Ethernet controller

Figure 10-16 Creating the VIOS profile: Physical I/O Adapters

15. Review the information in the Summary window, as shown in Figure 10-17.

Server Name:	PF-PowerVM-Node1
Virtual server name:	SN101D88B_VIOS1
Virtual server ID:	1
Environment:	VIOS
Memory:	10 GB [Dedicated]
Processors:	1 [Shared, DefaultPool(0)]
Virtual Ethernets:	11 [1, Bridge]
Virtual Adapters:	31 [SCSI, 2:1]
Physical adapters:	U78AE.001.WZS00T2-P1-C18-L1 U78AE.001.WZS00T2-P1-C18-L2

Figure 10-17 Creating the VIOS profile: Summary

16. Right-click the created VIOS, and select **Operations** → **Activate** → **Profile** as shown in Figure 10-18.

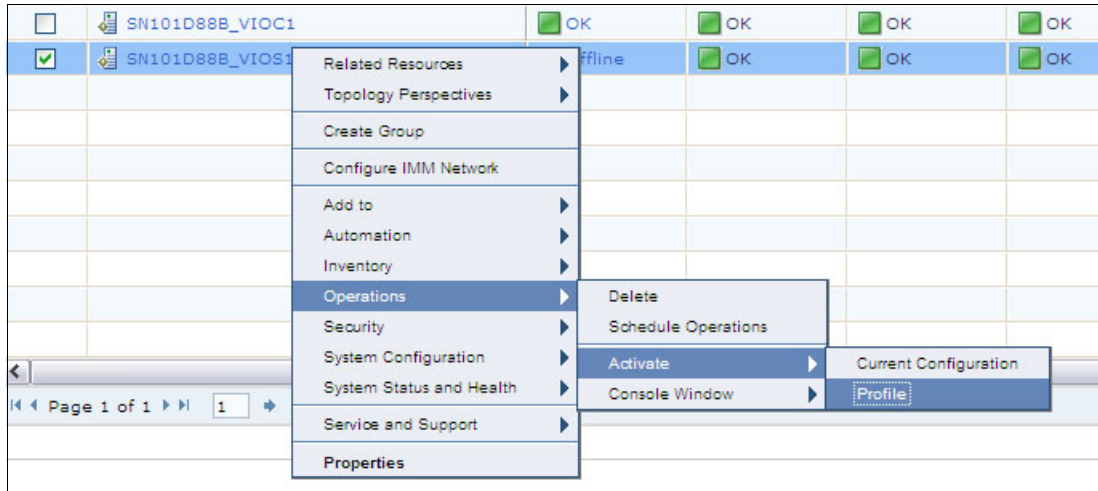


Figure 10-18 Activating the VIOS profile

17. Click **Advanced**, as shown in the Figure 10-19.

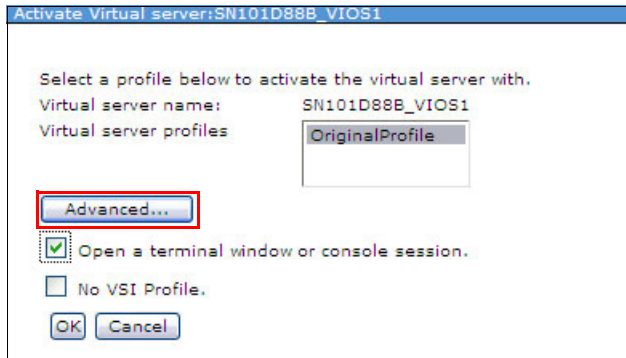


Figure 10-19 Activating the virtual server

18. Set the Boot mode value to **SMS** mode, as shown in the Figure 10-20. Click **OK**.




Figure 10-20 Activating the virtual server: Advanced properties

19. The terminal console is displayed as shown in the Figure 10-21. Enter Passw0rd, then you can see the Software Management Services (SMS) mode menu.

```
In order to access the terminal, you must first authenticate with the following
management console: 9.27.20.38
----
User ID: USERID
Password: █
```

Figure 10-21 Entering the SMS mode menu

20. Select **5. Select Boot Options**, as shown in Figure 10-22.

```
Version AF743_103
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Main Menu
1. Select Language
2. Setup Remote IPL (Initial Program Load)
3. Change SCSI Settings
4. Select Console
5. Select Boot Options
```

Figure 10-22 SMS menu: Main

21. Select **1. Select Install/Boot Device**, as shown in Figure 10-23.

```
Version AF743_103
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Multiboot
1. Select Install/Boot Device
2. Configure Boot Device Order
3. Multiboot Startup <OFF>
4. SAN Zoning Support
5. Management Module Boot List Synchronization
```

Figure 10-23 SMS menu: Multiboot

22. Select **3. CD/DVD**, as shown in Figure 10-24.

```
Version AF743_103
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Select Device Type
1. Diskette
2. Tape
3. CD/DVD
4. IDE
5. Hard Drive
6. Network
7. List all Devices
```

Figure 10-24 SMS menu: Select Device Type

23. Select **6. USB**, as shown in the Figure 10-25.




Figure 10-25 SMS menu: Select Media Type

Follow the VIOS installation menu to complete the installation of VIOS. For more information, see the *IBM Flex System p260 and p460 Planning and Implementation Guide*, SG24-7989, at this website:

<http://www.redbooks.ibm.com/abstracts/sg247989.html?Open>

AIX installation

The next step is to install an AIX image. For more information about implementing a VIOC, see *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940, at this website:

<http://www.redbooks.ibm.com/abstracts/sg247940.html?Open>

Tip: When you are installing a VIOC, use a virtual CD-ROM host by VIOS.

You can check the media device configuration, as shown in Figure 10-26.

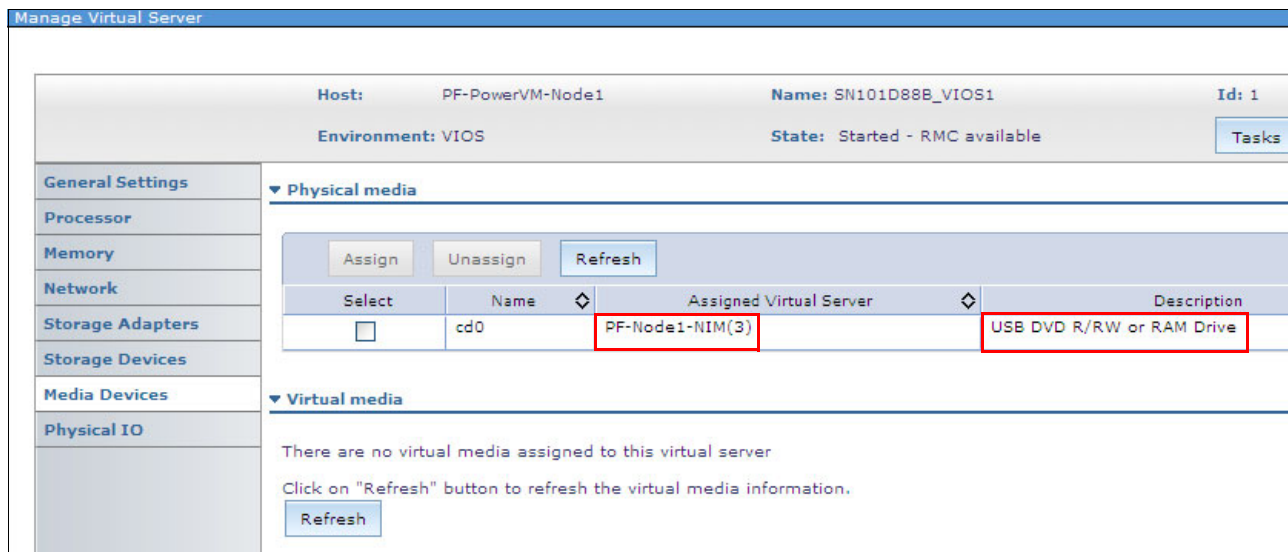


Figure 10-26 Media Devices window in the Manage Virtual Server

To do so, perform these steps:

1. Right-click **VIOS name** and select **System Configuration** → **Manage Virtual Server**.
2. Click **Media Devices** to see the current VIOS configuration. You can see which server owns the media device. In the example, PF-Node1-NIM owns the virtual CD-ROM. The description indicates which media is assigned to the virtual server, for example, CD-ROM.

There are two partitions, as shown in Figure 10-27. One is VIOS and the other one is for the NIM server to deploy AIX OS. The minimum requirement is set to deploy the PowerVM virtual machine by using VMControl.




Figure 10-27 Created two partitions

10.2 Capturing virtual machines

This section addresses how to capture virtual machines in the PowerVM environment.

There are two methods to capture virtual machines in PowerVM virtual infrastructures:

- ▶ Capturing AIX by using Network Installation Manager (NIM)
- ▶ Capturing AIX by using storage copy services (SCS)

For more information about NIM and SCS-based capture methods, see 5.2.3, “Planning for PowerVM virtualization” on page 111.

10.2.1 Capturing AIX by using Network Installation Manager (NIM)

This section describes how to capture AIX by using NIM through IBM Flex System Manager. Figure 10-28 shows the minimum configuration for NIM-based virtual machine deployment.




Figure 10-28 System diagram for NIM-based deployment

To capture AIX by using NIM, perform these steps:

1. Click the **Plug-ins** tab on the Flex System Manager Home page, as shown in Figure 10-29.




Figure 10-29 Flex System Manager main window

2. Click **VMControl Enterprise Edition** in the Plug-ins window, as shown in Figure 10-30.



Figure 10-30 VMControl Enterprise Edition menu

3. Click the **Virtual Appliances** tab to see the status of managed virtual appliances from the perspective of the virtual machine deployment, as shown in Figure 10-31 on page 418.




Figure 10-31 Virtual Appliances tab

The following information is highlighted in Figure 10-31:

- **What to deploy:** This menu shows ready-to-use virtual appliances in your data center, such as Lpp_source capture or mksysb capture. When you finish this task, **0 Virtual appliances** changes to **1 Virtual appliances**.
- **Where to deploy:** When you are deploying a virtual machine, select one of these two options: virtual server itself or server system pools that are normally physical servers.
- **What to capture:** This menu shows the partitions that are going to be gold images.
- **Where to store:** This menu shows all image repositories that VMControl has.

- From the **Virtual Servers/Hosts** tab, select the PowerVM node, then collect inventory for the physical server and for the VIO Server as shown in Figure 10-32.

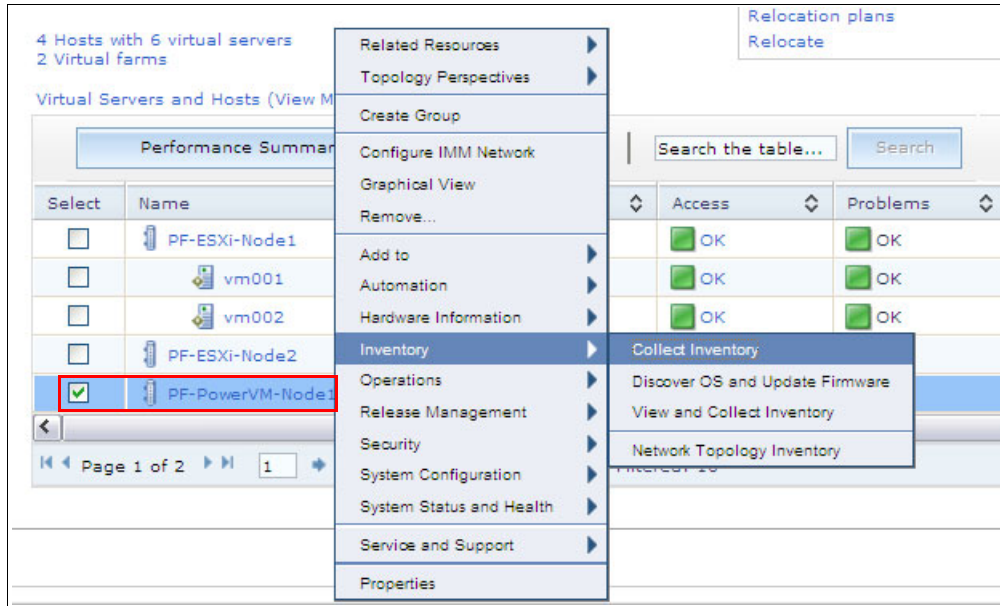


Figure 10-32 Running a Collect Inventory job

- Check the log for a job, as shown in Figure 10-33.

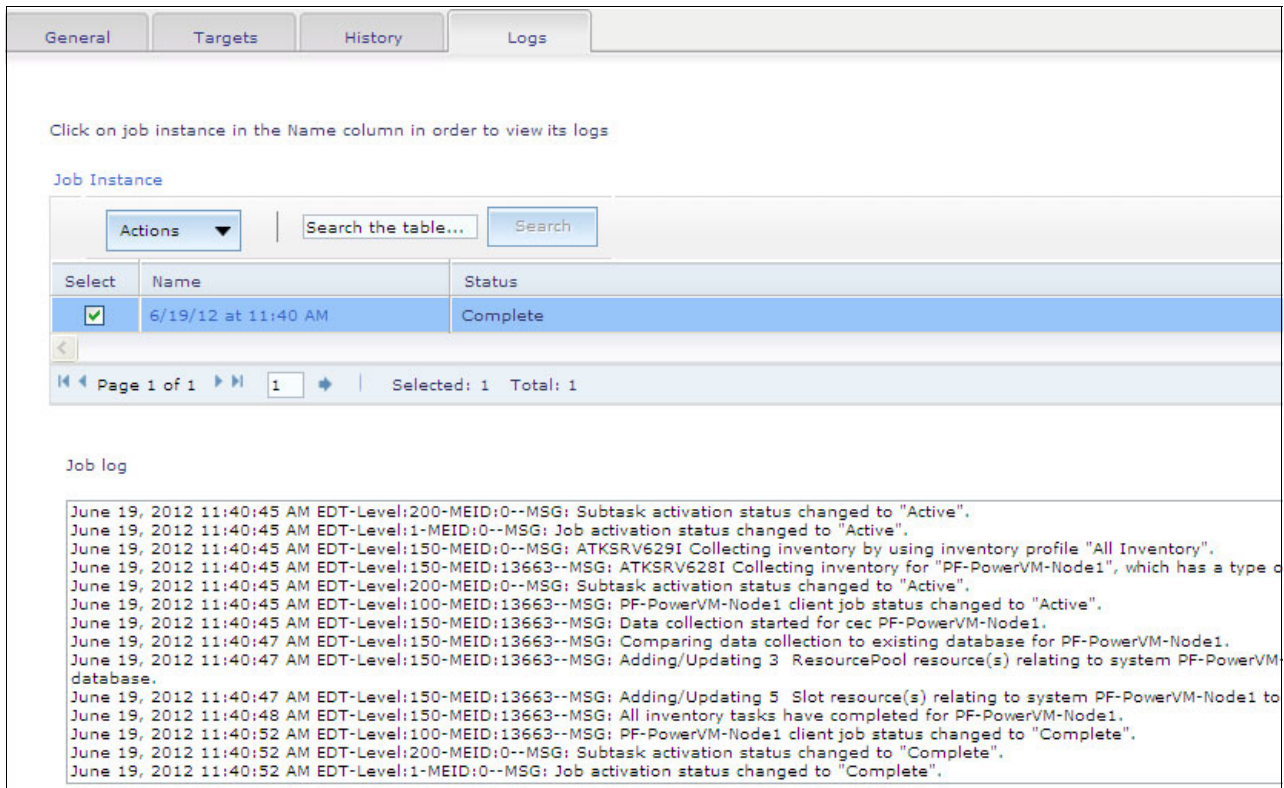


Figure 10-33 Checking the log in the Logs tab

- You can check the Resource Explorer window where three more virtual resources are added to the physical server, as shown in Figure 10-34 compared to Figure 10-27 on page 416.




Figure 10-34 Four PowerVM virtual servers in the Resource Explorer menu

- On the Virtual Servers/Hosts tab, the total number of virtual servers has increased to six, as shown in Figure 10-35.



Figure 10-35 Discovered virtual servers on the Virtual Server/Hosts tab

The second step is creating a VMControl repository by using NIM. The following characteristics are required:

- You can manage the AIX mksysb or lpp_source images.
- The VMControl repository must be configured as a NIM master system:
 - nim_master_setup
 - AIX 7.1 or later

- File sets:
 - dsm.core (not installed as part of the default file set)
 - openssh.base.client
 - openssl.base
- ▶ IBM Flex System Manager Common Agent must be installed.
- ▶ VMControl NIM subagent must be installed.
- ▶ NIM master must be discovered, accessed, and inventoried by IBM Flex System Manager.
- ▶ After /export/nim is created, make sure that it is large enough to hold appliances.

To create the repository, perform these steps:

1. Make sure that you have the SystemMgmtClient bundle installed on the NIM server, as shown in Figure 10-36.

```
# lslpp -l Director*
Fileset                Level  State   Description
-----
Path: /usr/lib/objrepos
DirectorCommonAgent    6.2.1.3  COMMITTED  All required files of Director
Common Agent, including JRE,
LWI
DirectorPlatformAgent  6.2.1.2  COMMITTED  Director Platform Agent for
IBM Systems Director on AIX

Path: /etc/objrepos
DirectorCommonAgent    6.2.1.3  COMMITTED  All required files of Director
Common Agent, including JRE,
LWI
DirectorPlatformAgent  6.2.1.2  COMMITTED  Director Platform Agent for
IBM Systems Director on AIX

# lslpp -l cas*
Fileset                Level  State   Description
-----
Path: /usr/lib/objrepos
cas.agent              1.4.2.32  COMMITTED  Common Agent Services Agent

Path: /etc/objrepos
cas.agent              1.4.2.32  COMMITTED  Common Agent Services Agent
# █
```

Figure 10-36 SystemMgmtClient file set

2. Ensure that openssh and openssl are installed and ssh is started on the NIM server, as shown in Figure 10-37.

```
# lslpp -l | grep dsm
dsm.core                6.1.7.15  COMMITTED  Distributed Systems Management
dsm.dsh                 6.1.6.15  COMMITTED  Distributed Systems Management
dsm.core                6.1.7.15  COMMITTED  Distributed Systems Management
# lslpp -l | grep openssh
openssh.base.client    5.0.0.5301  COMMITTED  Open Secure Shell Commands
openssh.base.server    5.0.0.5301  COMMITTED  Open Secure Shell Server
openssh.license        5.0.0.5301  COMMITTED  Open Secure Shell License
openssh.man.en_US      5.0.0.5301  COMMITTED  Open Secure Shell
openssh.msg.en_US      5.0.0.5301  COMMITTED  Open Secure Shell Messages -
openssh.base.client    5.0.0.5301  COMMITTED  Open Secure Shell Commands
openssh.base.server    5.0.0.5301  COMMITTED  Open Secure Shell Server
# lslpp -l | grep openssl
openssl.base           0.9.8.1800  COMMITTED  Open Secure Socket Layer
openssl.license        0.9.8.802   COMMITTED  Open Secure Socket License
openssl.man.en_US      0.9.8.1800  COMMITTED  Open Secure Socket Layer
openssl.base           0.9.8.1800  COMMITTED  Open Secure Socket Layer
# lssrc -s sshd
Subsystem      Group      PID      Status
sshd           ssh       5898428  active
#
```

Figure 10-37 Checking the ssh status

3. Discover the NIM server by using its IP address, as shown in Figure 10-38.

Requirement: This step is a prerequisite to deploying the VMControl agent on the NIM server from FSM.




Figure 10-38 Discovering the NIM server

4. Check the progress, as shown in Figure 10-39.




Figure 10-39 Checking the discovery job

5. Check the logs as shown in Figure 10-40.




Figure 10-40 Checking the log

6. Click **No access** to gain access to the server, as shown in Figure 10-41.




Figure 10-41 Discovered Systems tab

7. Enter the root password for your NIM server, as shown in Figure 10-42.




Figure 10-42 Request Access window

You can see access is granted, as shown in Figure 10-43.

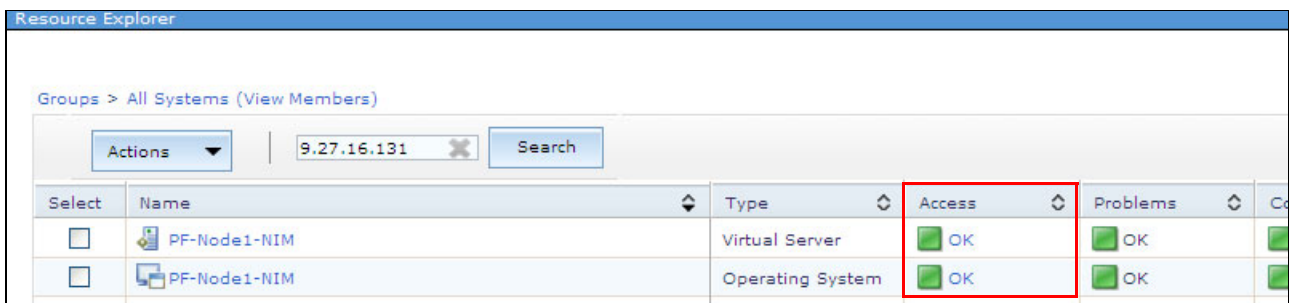


Figure 10-43 Checking accessibility in the Resource Explorer menu

8. Discover the NIM server in VMControl, as shown in Figure 10-44.




Figure 10-44 Discover virtual appliances

9. Select the NIM server that you want to discover, then click **Add** as shown in Figure 10-45 on page 426.

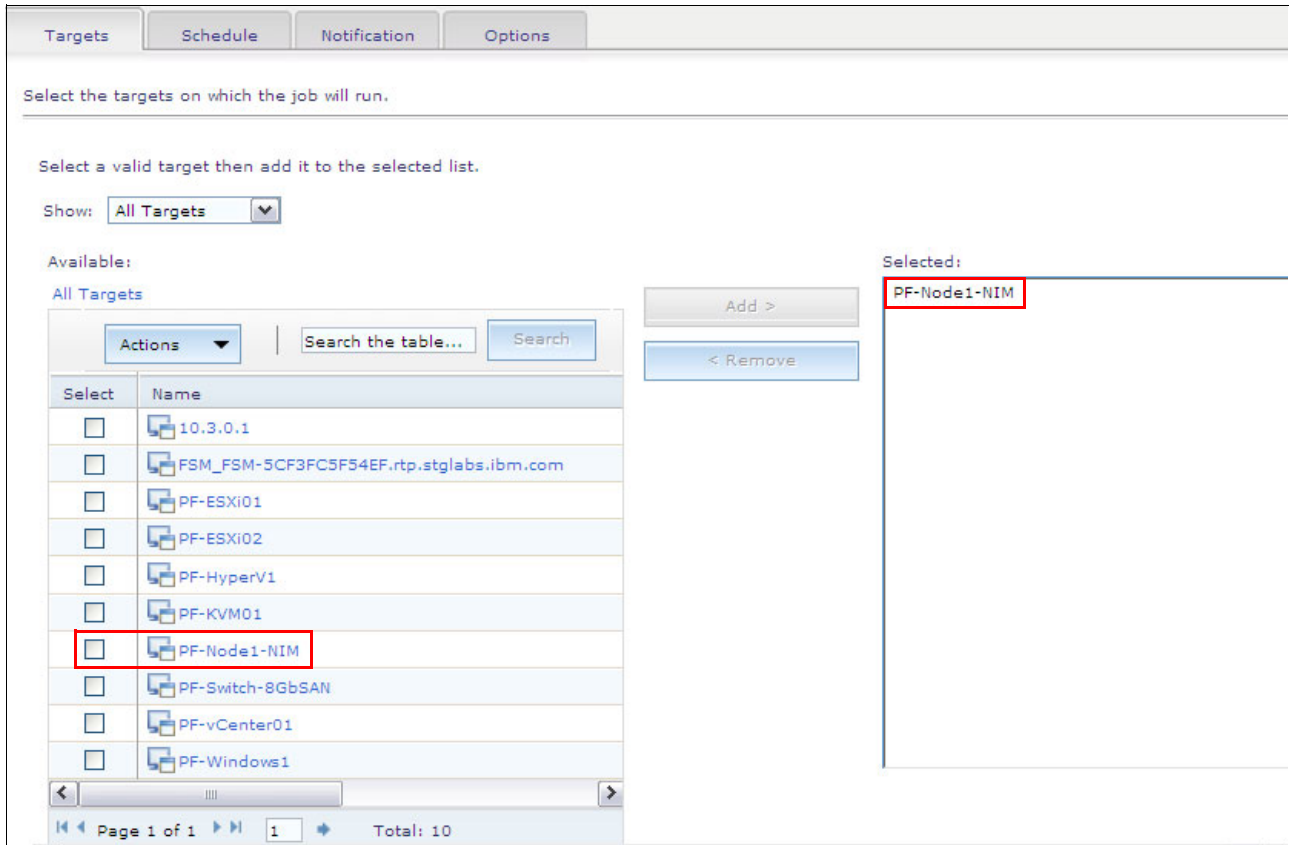


Figure 10-45 Discover virtual appliances menu

You can now deploy the VMControl agent on the NIM server.

To install the VMControl agent on the NIM server, perform these steps:

1. Click **Install Agents**, as shown in Figure 10-46 on page 427.

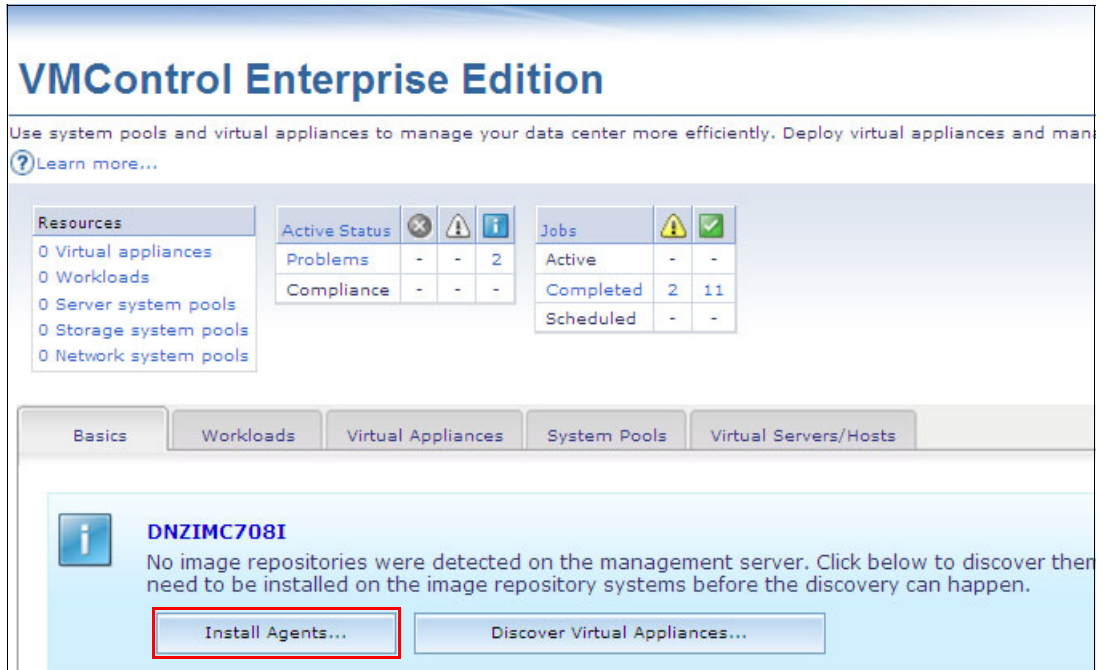


Figure 10-46 Install Agents on the Basics tab

The Welcome window opens as shown in Figure 10-47.

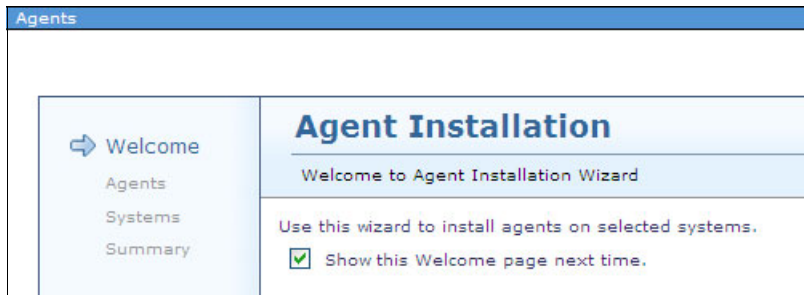


Figure 10-47 Agent Installation Welcome window

2. Select **Common Agent Subagent Packages**, as shown in Figure 10-48.

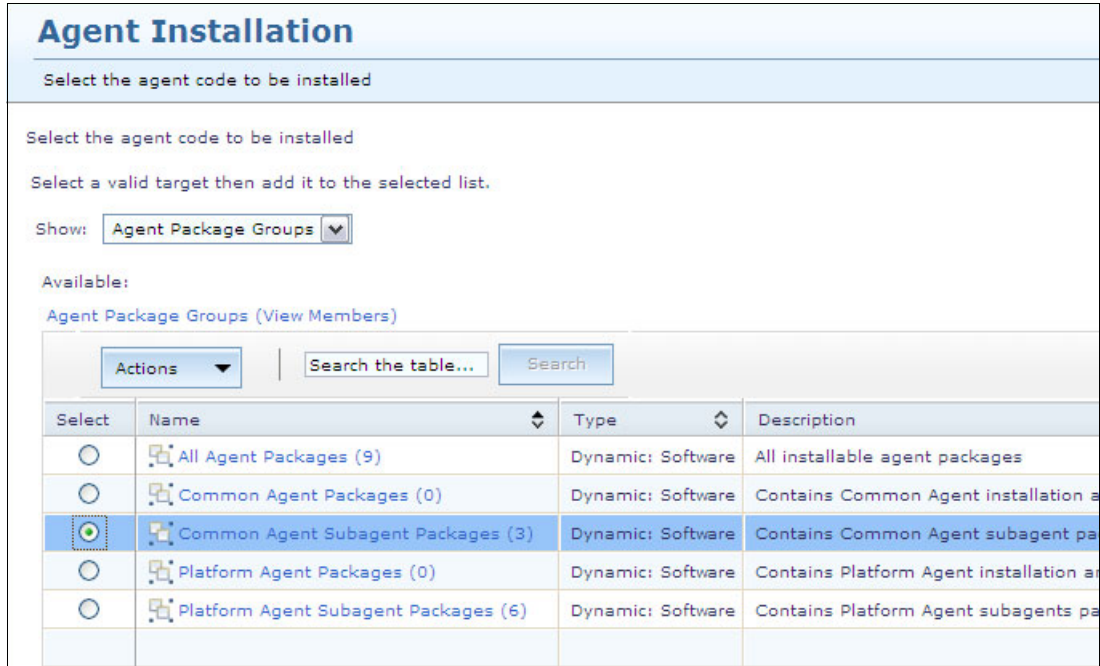


Figure 10-48 Agent Installation window

3. Select the agent for the NIM server as shown in Figure 10-49.

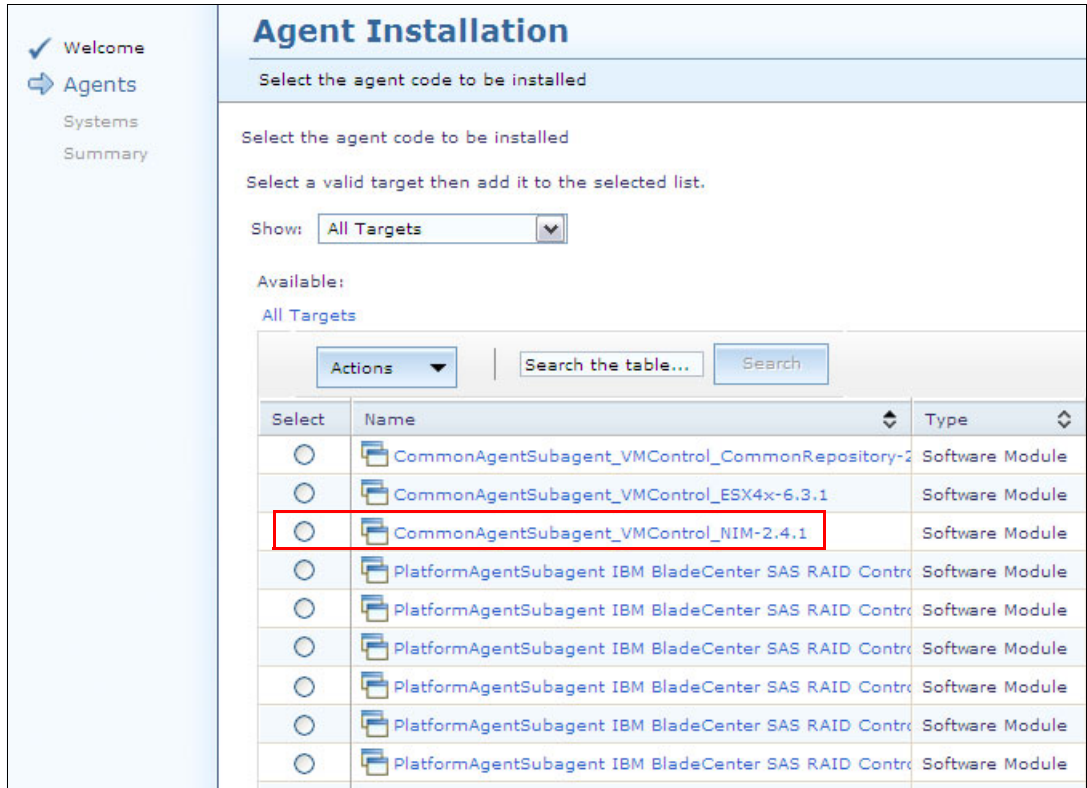


Figure 10-49 Agent Installation window

4. Click **Add** as shown in Figure 10-50.

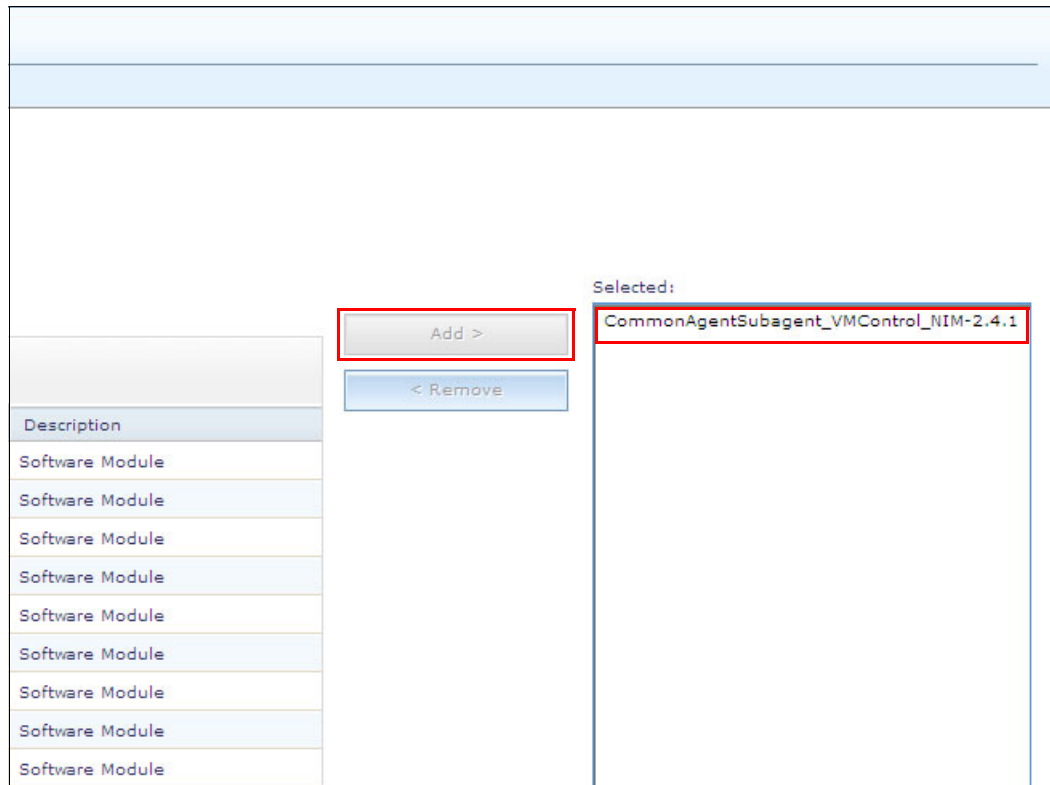


Figure 10-50 Agent Installation window

5. Select the NIM server for agent installation, then click **Add**, as shown in Figure 10-51.




Figure 10-51 Agent Installation window

6. Verify the information, then click **Finish** to install the agent for NIM, as shown in Figure 10-52.




Figure 10-52 Agent Installation

7. A menu for the installation task opens as shown in Figure 10-53. Click **OK**.




Figure 10-53 Agent Installation

8. Check the progress, as shown in Figure 10-54.

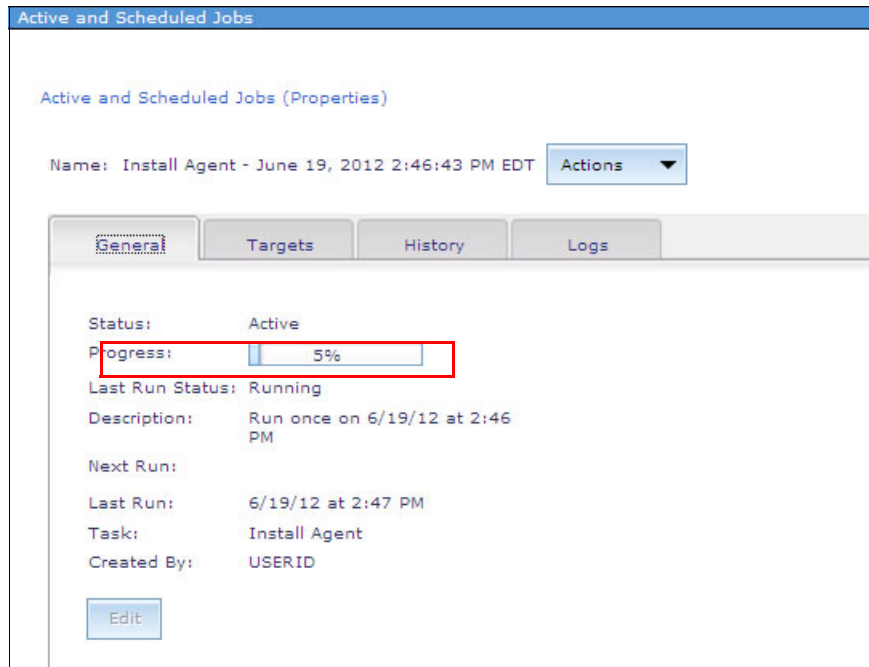


Figure 10-54 Checking your progress

9. Check the logs, as shown in Figure 10-55.




Figure 10-55 Checking the logs

10. The server object is displayed for the NIM server. Run the Inventory task against both objects, as shown in Figure 10-56.




Figure 10-56 Checking the NIM server status in the Resource Explorer menu

11. Click **Create image repository**, as shown in Figure 10-57.

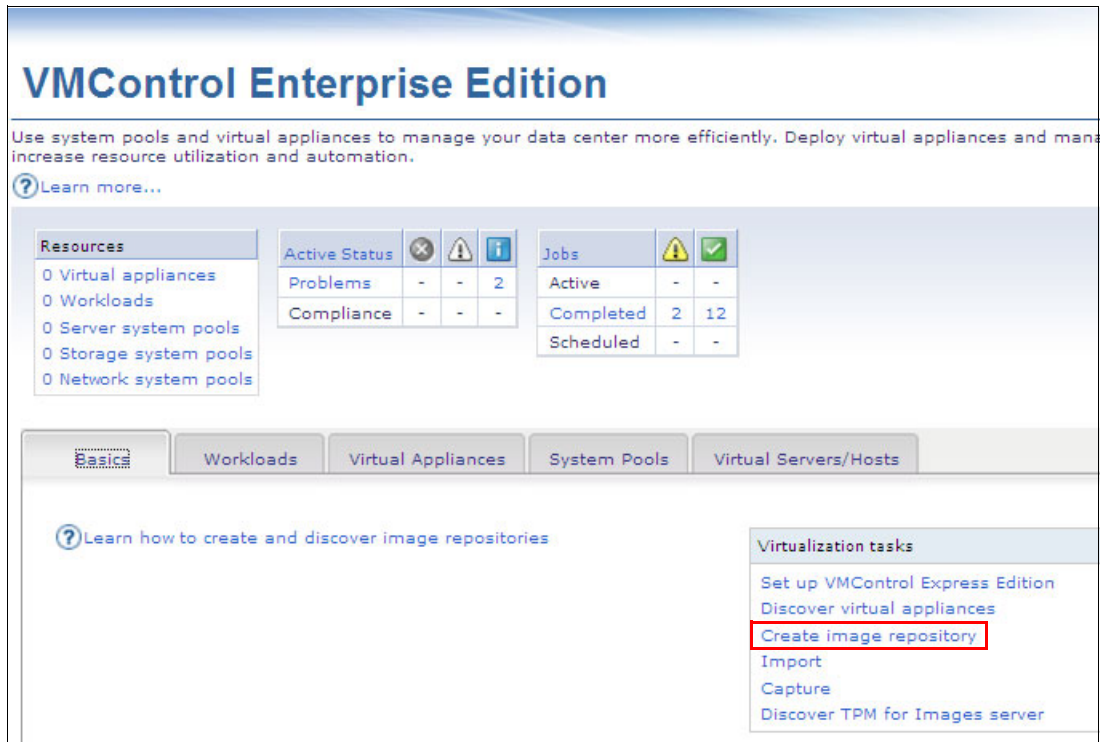


Figure 10-57 Create the image repository on the Basics tab

12. Select the NIM server to use as the repository, as shown in Figure 10-58.




Figure 10-58 Creating the image repository

13. The new NIM repository is created. Click the **Image repositories** link for more information about this repository, as shown in Figure 10-59.

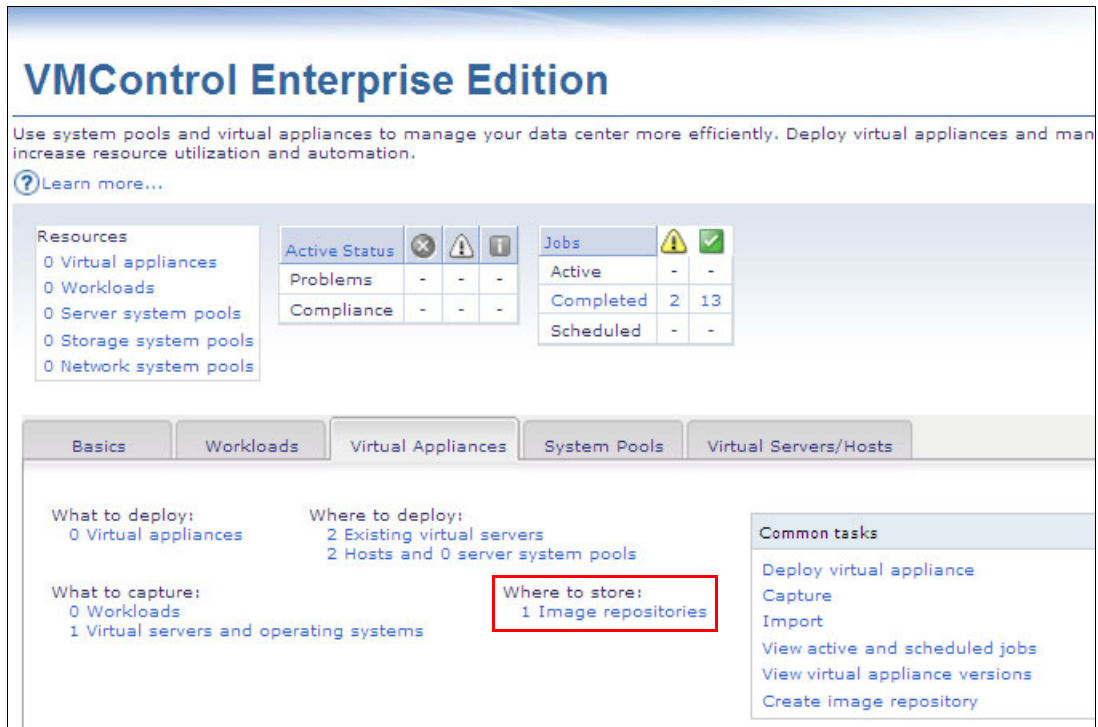


Figure 10-59 Checking the newly created image repository

14. Review the newly created image repository, as shown in Figure 10-60.




Figure 10-60 Checking the image repository

The prerequisite steps to capture the virtual server are complete.

10.2.2 Capturing the Network Installation Manager server

There are two possible capture methods for the NIM:

- ▶ Creating the LPP_source base
- ▶ Capturing the mkysyb base

Creating the LPP_source base

This section describes to how to create the LPP_source.

Figure 10-61 shows the LPP_source base capture system diagram.




Figure 10-61 System diagram for LPP_source based deployment

The LPP_source base capture involves the following basic steps:

1. Create the LPP_source file from the AIX CD/DVD on the NIM server.
2. Check that all related servers can be seen in FSM. If they cannot, discover the specific object, then run the Collect Inventory task.
3. Put the LPP_source file together with the Open Virtualization Format (OVF) format by using the **captureva** command.

The following commands take an existing lpp_source from the NIM server and make it into a virtual appliance, as shown in Figure 10-62. This function cannot be completed from the GUI.

```

USERID@FSM-5CF3FC5F54EF:~> smcli lsrepos -o
pureflexrb13.rtp.stglabs.ibm.com, 20609 (0x5081)
USERID@FSM-5CF3FC5F54EF:~> smcli captureva -r 20609 -F repos:lpp_source_6100 -n"
AIX_Lppsource" -A "cpushare=1,memsize=8192"
USERID@FSM-5CF3FC5F54EF:~> █
    
```

Figure 10-62 Create LPP_source base capture

Use the `lsrepos` command to list the repositories:

- ▶ Display all repositories: `smcli lsrepos -v`
- ▶ Display the object identifier (OID): `smcli lsrepos -o`

Use the `captureva` command to capture a virtual appliance:

- ▶ Capture a virtual appliance from a virtual server: `smcli captureva -v -s 123 -r 345 -n "XYZLpar" -D "Production server"`
- ▶ Capture a virtual appliance from existing lpp_source and set the processor and memory size: `smcli captureva -r 20609 -F repos:lpp_source_6100 -n "AIX_Lppsource" -A "cpushare=1,memsize=8192"`

Click the **Virtual Appliances** tab to see the newly created virtual appliance, as shown in Figure 10-63.




Figure 10-63 Checking the newly created virtual appliance

Capturing the mkysyb base

This section describes how to create a virtual appliance by using the mkysyb method.

The following steps are a general overview of the mkysyb base capture:

1. Create two partitions without VIOS: one for NIM server (if it exists, no need to create it) and one AIX partition where mkysyb will run.
2. Make sure that all related servers are seen in FSM. If you cannot see all of them, discover the specific object, then run the Collect Inventory task.
3. Check the file system size where the mkysyb file will be stored in the NIM.

To capture the mksysb base, perform these steps:

1. Click **Capture** in the **Virtual Appliances** tab, as shown in Figure 10-64.

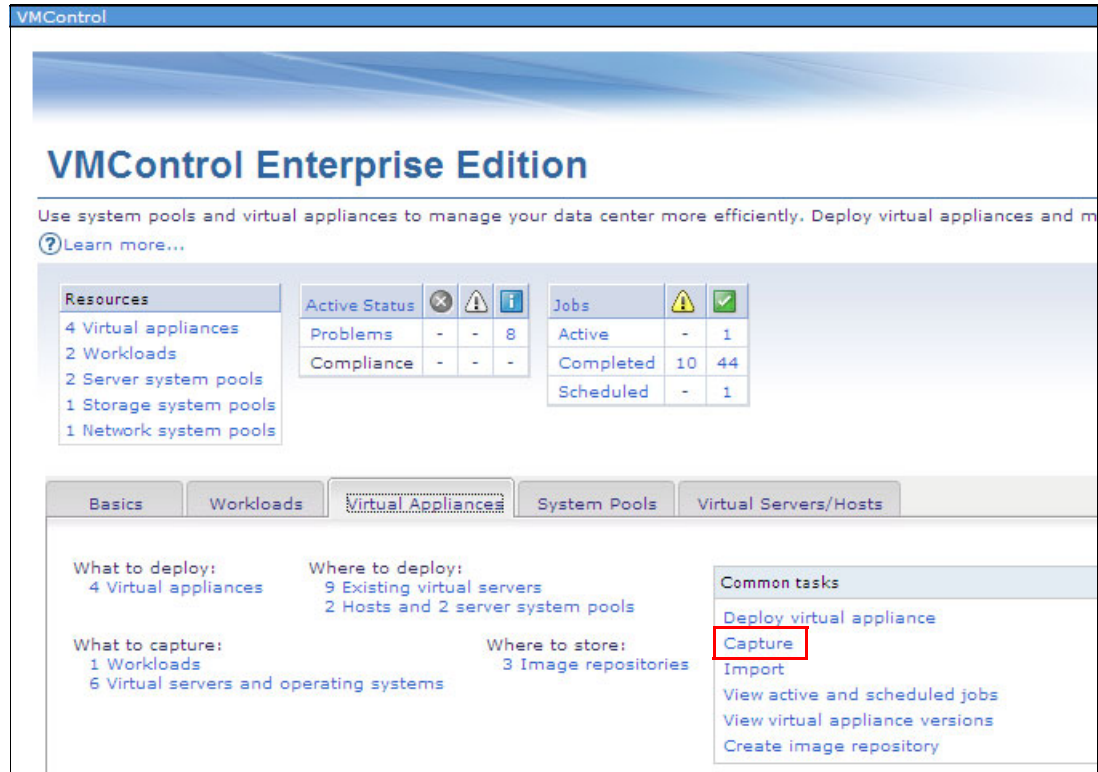


Figure 10-64 Virtual Appliances window

Figure 10-65 shows the Capture Welcome window.




Figure 10-65 Capture Welcome window

2. Enter the virtual appliance name, as shown in Figure 10-66. In this example, it is AIX-6100-mksysb.



Figure 10-66 Capture: Name

3. Select the source type to capture, as shown in Figure 10-67. In this example, it is **Virtual Server**.

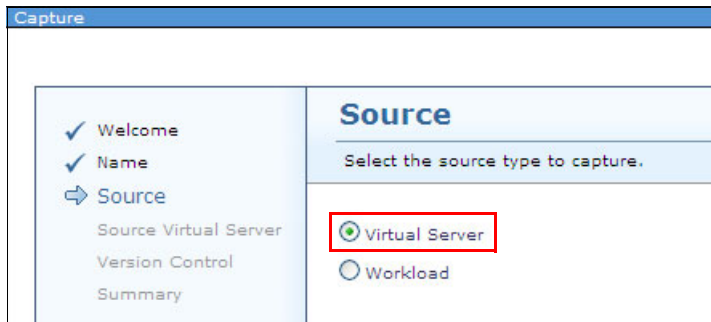


Figure 10-67 Capture: Source

4. Select the source virtual server, as shown in Figure 10-68.

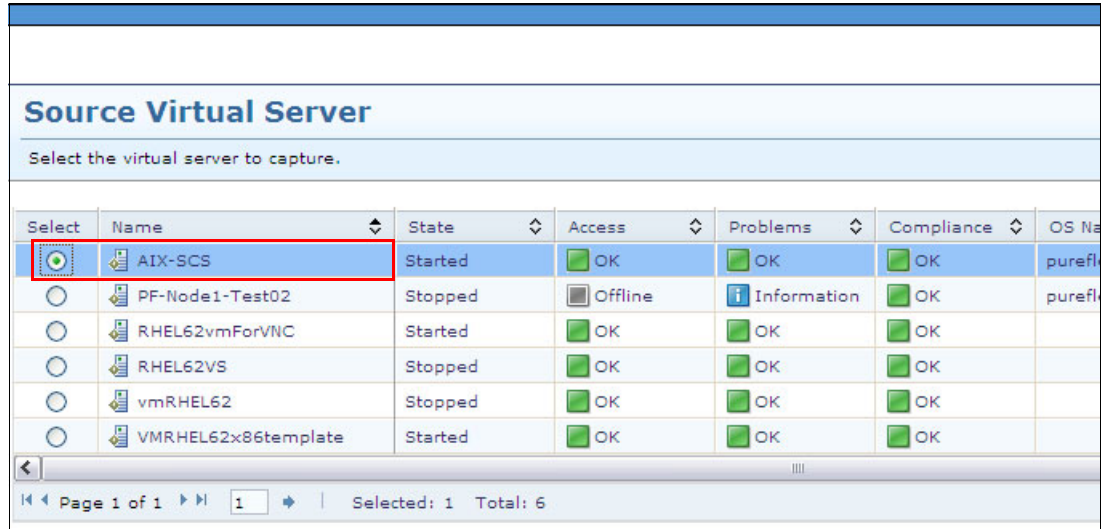


Figure 10-68 Capture: Source Virtual Server

5. Select the repository as shown in Figure 10-69. The repository is where you want to store the image that is associated with the new virtual appliance.




Figure 10-69 Capture: Repository

6. Select a network as shown in Figure 10-70. In the example, only one SEA was created, so only one network is displayed. If more than one SEA exist, select the network that you want to use.




Figure 10-70 Capture: Network Mapping

7. Figure 10-71 shows the Version Control window. Normally, mkysyb-based capture is based on a source that is already running an AIX image. Therefore, the wizard selects **Set the version based on the virtual appliance from which the virtual server was originally deployed: Capture_AIX_SCS**, by default.

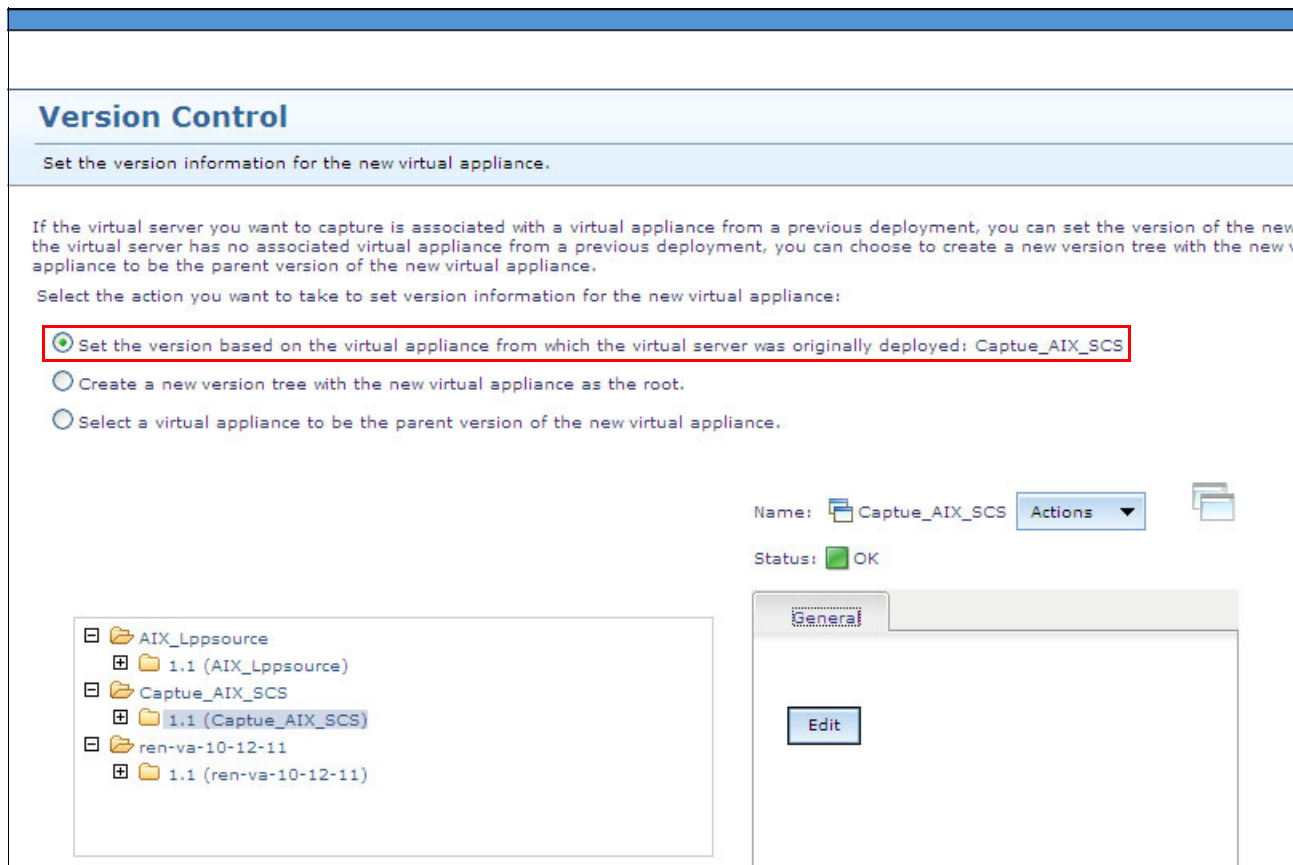


Figure 10-71 Capture: Version Control

Figure 10-72 shows a summary.

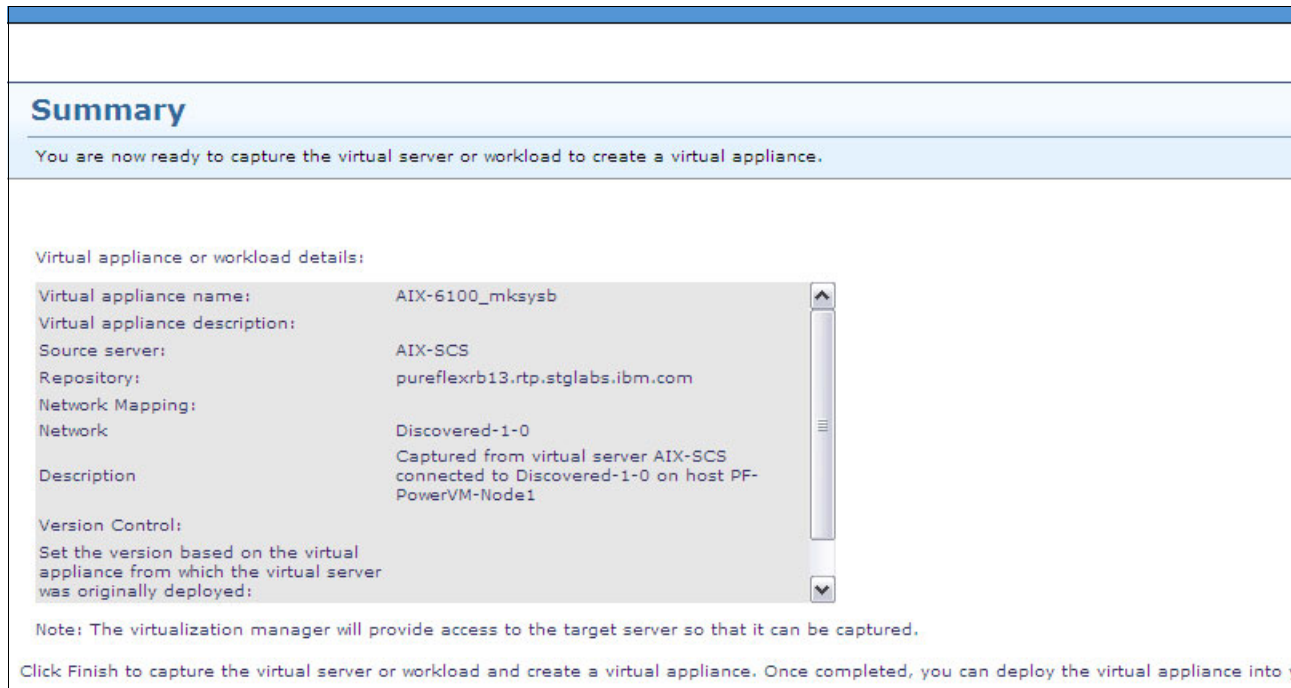


Figure 10-72 Capture: Summary

8. Figure 10-73 shows the Launch Job menu.

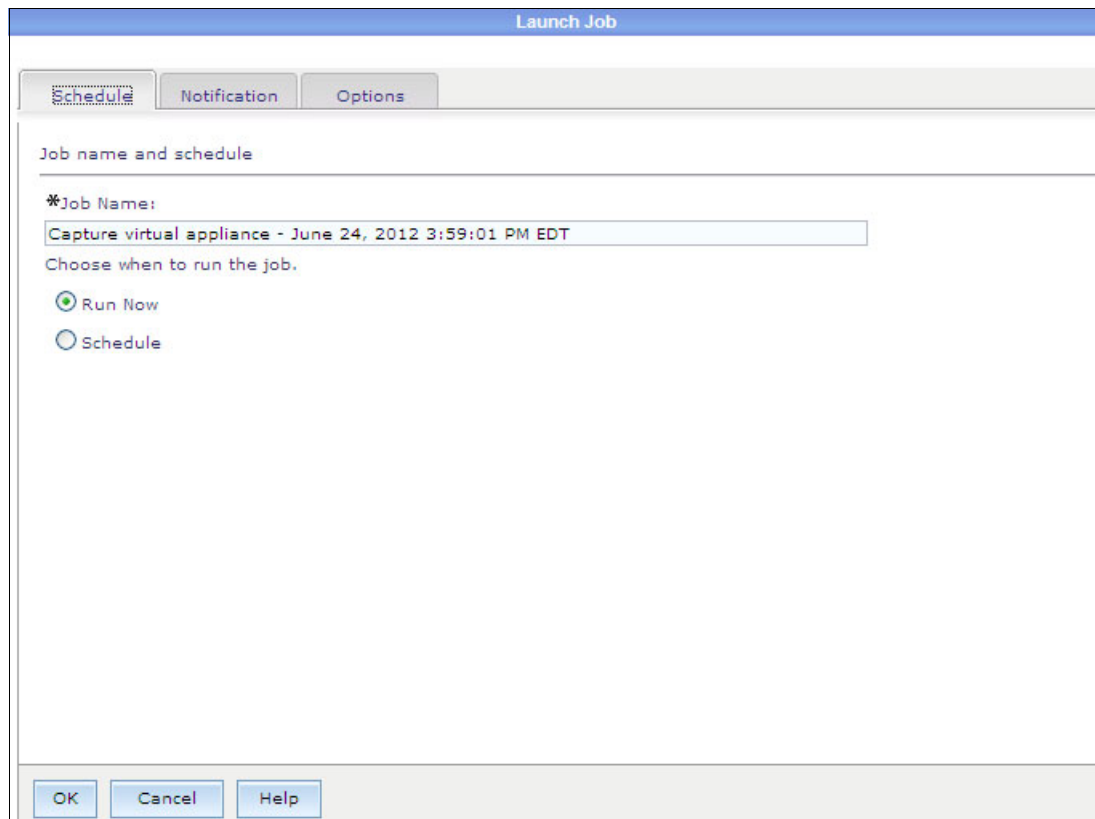


Figure 10-73 Capture: Launch Job

9. Click **OK** and check the log as shown in Figure 10-74.

Click on job instance in the Name column in order to view its logs

Job Instance

Actions | Search the table... Search

Select	Name	Status
<input checked="" type="checkbox"/>	6/24/12 at 6:11 PM	Complete

Page 1 of 1 | 1 | Selected: 1 Total: 1

Job log

```

June 24, 2012 6:11:42 PM EDT-Level:200-MEID:0--MSG: No clients to start.
June 24, 2012 6:11:42 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
June 24, 2012 6:11:42 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
June 24, 2012 6:11:42 PM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Active".
June 24, 2012 6:11:42 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
June 24, 2012 6:11:43 PM EDT-Level:150-MEID:0--MSG: DNZLOP411I Capturing virtual server AIX-SCS to virtual appliance AIX-6100_mk...
pureflexrb13.rtp.stglabs.ibm.com.
June 24, 2012 6:11:44 PM EDT-Level:150-MEID:0--MSG: DNZLOP410I Configuring the NIM capture on the virtual server and the NIM mast...
June 24, 2012 6:12:05 PM EDT-Level:150-MEID:0--MSG: DNZLOP407I Initiating capture processing on the NIM master.
June 24, 2012 6:14:20 PM EDT-Level:150-MEID:0--MSG: DNZLOP408I NIM master capture processing complete.
June 24, 2012 6:14:41 PM EDT-Level:150-MEID:0--MSG: DNZLOP409I Creating the OVF for the virtual appliance.
June 24, 2012 6:14:44 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
June 24, 2012 6:14:44 PM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
June 24, 2012 6:14:44 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
June 24, 2012 6:14:44 PM EDT-Level:100-MEID:0--MSG: Capture virtual appliance complete.
    
```

Figure 10-74 Capture: Logs

10. Review the newly captured virtual appliance, as shown in Figure 10-75.

Basics | Workloads | **Virtual Appliances** | System Pools | Virtual Servers/Hosts

What to deploy: 5 Virtual appliances | Where to deploy: 8 Existing virtual servers, 2 Hosts and 2 server system pools

What to capture: 1 Workloads, 7 Virtual servers and operating systems | Where to store: 3 Image repositories

Common tasks: Deploy virtual appliance, Capture, Import, View active and scheduled jobs, View virtual appliance versions, Create image repository

Virtual Appliances (View Members)

Capture | Deploy Virtual Appliance | Import | Actions | Search the table... Search

Select	Name	Operating System	Repository
<input type="checkbox"/>	AIX_Lppsource	IBM AIX	pureflexrb13.rtp.stglabs.ibm.com
<input checked="" type="checkbox"/>	AIX-6100_mksysb	IBM AIX 6	pureflexrb13.rtp.stglabs.ibm.com
<input type="checkbox"/>	Capture_AIX_SCS	IBM AIX 6	PowerVM-Repository
<input type="checkbox"/>	CapturedVMonKVM	Linux	KVMimagesrepo
<input type="checkbox"/>	ren-va-10-12-11	Linux	KVMimagesrepo

Figure 10-75 Capture finished

10.2.3 Capturing AIX by using storage copy services (SCS)

This section addresses how to work with the storage copy services (SCS) method. Figure 10-76 shows an SCS-based capture system diagram.

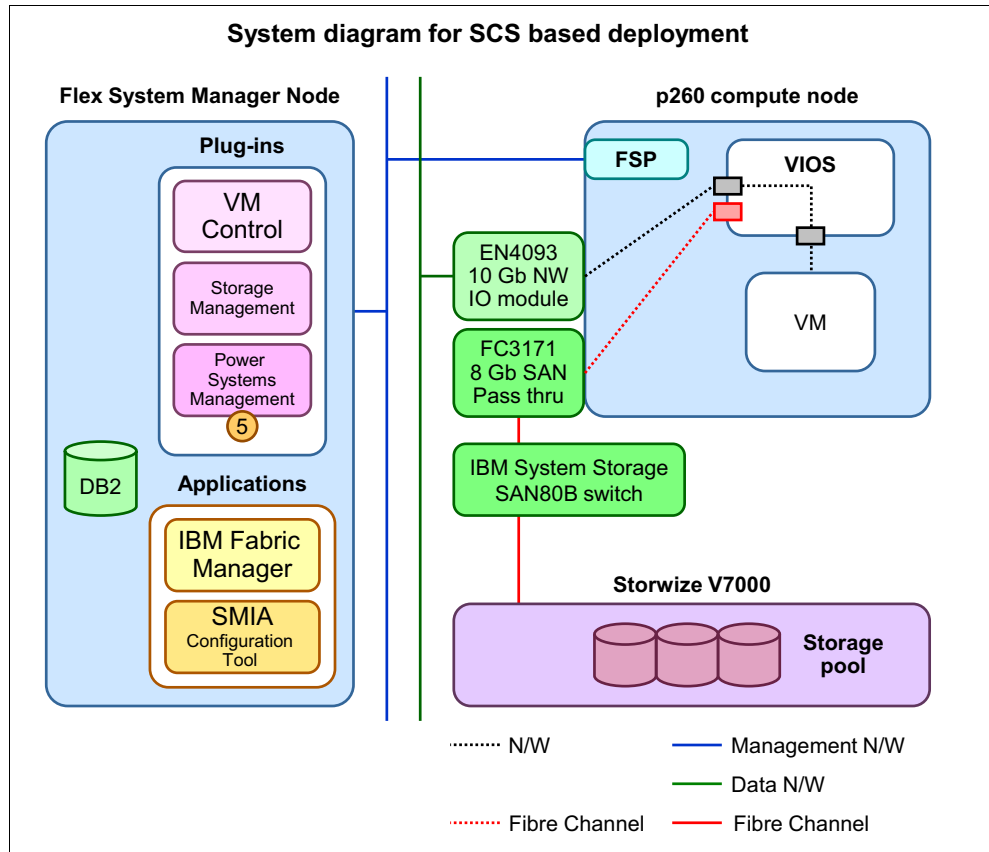


Figure 10-76 System diagram of SCS-based deployment

The following steps are a general overview of the SCS base capture:

1. Hardware preparation:
 - a. Storage configuration
 - b. Fabric zoning
 - c. Installation/configuration of VIOS on Power servers
2. Storage Management Initiative Specification (SMI-S) provider check step:
 - a. Configuration of SMI-S provider for the storage area network (SAN) switch (if needed)
3. Management check step:
 - a. VMControl
 - b. IBM Flex System Manager Storage Control
4. Infrastructure discovery:
 - a. Discovery of server infrastructure: Managed Power Servers
 - b. Discovery of storage infrastructure: V7000 and SAN fabric
 - c. Discovery of VIOS

5. Configure image repository and system pool:
 - a. Deploy Common Agent VMControl Subagent
 - b. Image repository and system pool creation
6. Preparation for capture:
 - a. Installation of activation engine
 - b. Enable activation engine
7. Functional test from FSM console

Check Storage Copy Services configuration

This section describes the checklist to configure the SCS. To do so, perform these steps:

1. Check whether IBM Flex System Manager Storage Control is running, as shown in Figure 10-77.



Figure 10-77 Storage Control main window

2. Check that related resources exist in the farm by right-clicking the farm name and selecting **Related Resources** → **Storage System**, as shown in Figure 10-78.

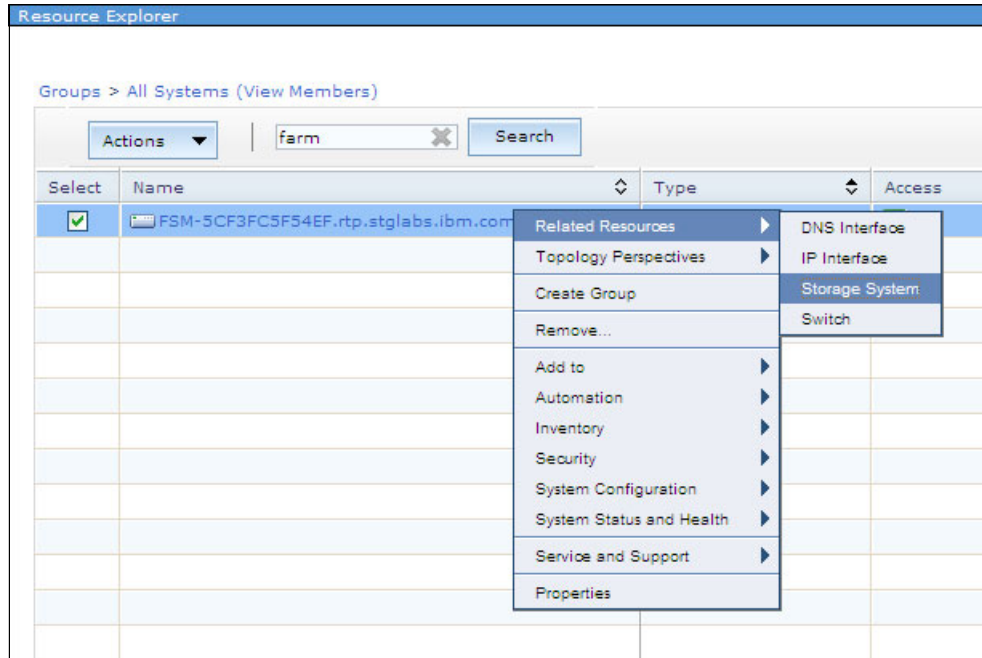


Figure 10-78 Checking the farm resources

3. Check that related resources exist in the All Systems window, as shown in Figure 10-79. In the example environment, there are three storage-related components: Storwize V7000, a Feature Code (FC) 3171 SAN pass-through, and an IBM System Storage® SANB80 switch.

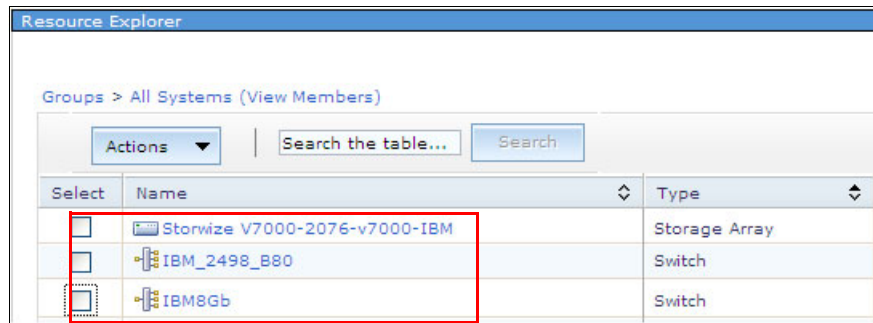


Figure 10-79 All Systems window in the Resource Explorer

4. Check ssh in the Remote Service Access Point (RSAP) configuration as shown in Figure 10-80.




Figure 10-80 Checking the RSAP configuration

5. Check that SSH access is displayed in the Access column, as shown in Figure 10-81.




Figure 10-81 Checking the RSAP configuration

For more information about how to add Storwize V7000 and third-party SAN switches, see 6.12, “Discover and manage external Storwize V7000” on page 234.

SCS configuration

To configure the SCS, perform these steps:

1. Check the auto start setting of the Common Agent file set by using the `lssvc DIRECTOR_agent` command, as shown in Figure 10-82.




Figure 10-82 Agent startup setting check step

Tip: The VIOS OS image has the Common Agent file, by default.

2. Discover VIOS and update the VIOS information in the FSM by clicking **Inventory** → **System Discovery**.
3. Right-click the discovered system in the Resource Explorer, then click **Configure Access**. Check that the Common Agent Services (CAS), Common Information Model (CIM), and Secure Shell (SSH) protocols are displayed as shown in Figure 10-83. Click **Request Access**, then enter the correct credentials.




Figure 10-83 Configure Access window

4. Check the access status, as shown in Figure 10-84.

Configure Access

Access to the identified system for all authorized Flex System Manager users is specified below in the Access field. To evaluate the overall status.

Select Request Access to enter a user ID and password to gain access to the identified system. Select Revoke Access to revoke access.

System: SN101D88B_VIOS1

Access: OK Request Access Revoke Access

For basic system management function, a remote system credential is needed to grant all authorized Flex System Manager users access to the identified system. A remote system credential is needed to grant all authorized Flex System Manager users access to the identified system. Certain types of access point credentials might not be visible.

SN101D88B_VIOS1 (Remote Service Access Point)

Select	Access Type	Access	Trust State	Access Information
<input type="checkbox"/>	SSH	No access	Not applicable	https://9.27.16.128:22/
<input type="checkbox"/>	CIM	No access	Not applicable	https://9.27.16.128:5989/
<input type="checkbox"/>	CAS	OK	Not applicable	https://9.27.16.128:9510/

Figure 10-84 Checking the access status

Agent installation

To install agents, perform these steps:

1. Click **Install agents** as shown in Figure 10-85.




Figure 10-85 Clicking Install agents

Figure 10-86 shows the Agent Installation Welcome window.

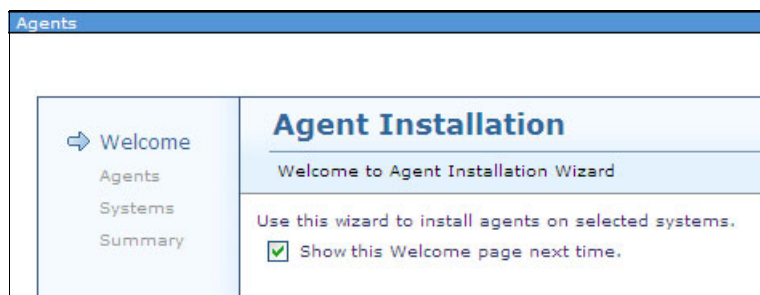


Figure 10-86 Agent Installation Welcome window

2. Select **CommonAgentSubagent_VMControl_CommonRepository-2.4.1** as shown in Figure 10-87.




Figure 10-87 Selecting the agent

3. Click **Add** as shown in Figure 10-88.

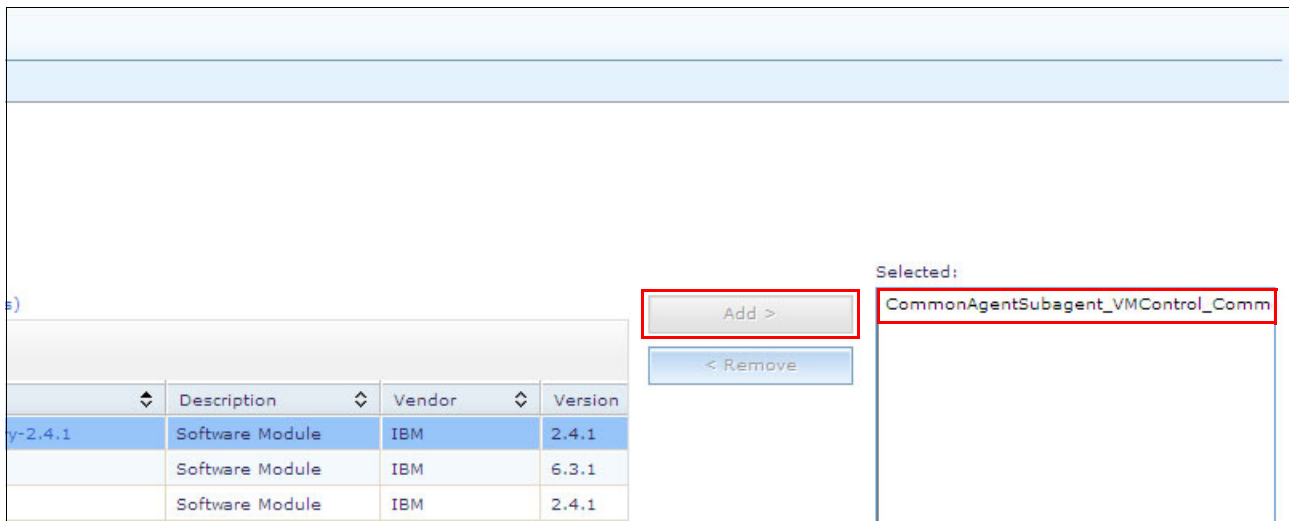


Figure 10-88 Adding the agent

4. Select a VIOS to deploy an agent, as shown in Figure 10-89.




Figure 10-89 Selecting a target VIOS

Figure 10-90 shows the summary.

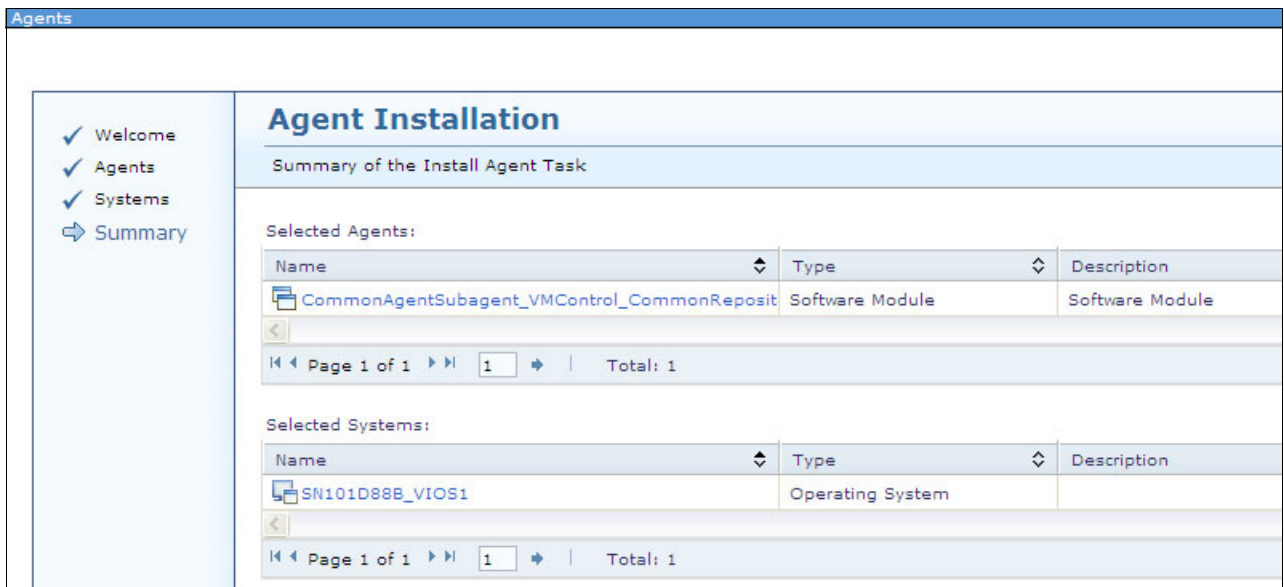


Figure 10-90 Summary window

5. Figure 10-91 shows the Launch Job menu. Click **OK**.

Launch Job

Schedule Notification Options

Job name and schedule

*Job Name:
Install Agent - June 22, 2012 9:15:46 AM EDT

Choose when to run the job.

Run Now

Schedule

OK Cancel Help

Figure 10-91 Launch Job menu

6. Check the log as shown in Figure 10-92.

Active and Scheduled Jobs (Properties)

Name: Install Agent - June 22, 2012 9:15:46 AM EDT Actions

General Targets History **Logs**

Click on job instance in the Name column in order to view its logs

Job Instance

Actions | Search the table... Search

Select	Name	Status
<input checked="" type="checkbox"/>	6/22/12 at 9:16 AM	Complete

Page 1 of 1 | 1 | Selected: 1 Total: 1

Job log

```

June 22, 2012 9:16:07 AM EDT-Level:1-MEID:0--MSG: Job "Install Agent - June 22, 2012 9:15:46 AM EDT" ac
June 22, 2012 9:16:08 AM EDT-Level:200-MEID:0--MSG: Subtask "Install Agent" activated.
June 22, 2012 9:16:08 AM EDT-Level:200-MEID:0--MSG: Starting clients
June 22, 2012 9:16:08 AM EDT-Level:100-MEID:0--MSG: Clients started for task "Install Agent"
June 22, 2012 9:16:08 AM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
June 22, 2012 9:16:08 AM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Active".
June 22, 2012 9:16:08 AM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Starting".
June 22, 2012 9:16:08 AM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
June 22, 2012 9:16:08 AM EDT-Level:100-MEID:28572--MSG: SN101D88B_VIOS1 client job status changed to
June 22, 2012 9:16:08 AM EDT-Level:100-MEID:28572--MSG: SN101D88B_VIOS1 client job status changed to
June 22, 2012 9:17:46 AM EDT-Level:100-MEID:28572--MSG: SN101D88B_VIOS1 client job status changed to
June 22, 2012 9:17:46 AM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
June 22, 2012 9:17:46 AM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
    
```

Figure 10-92 Checking the log

Image repository creation

To create an image repository, perform these steps:

1. Click **Create image repository** as shown in Figure 10-93.



Figure 10-93 Image repository creation step

Figure 10-94 shows the Welcome window.

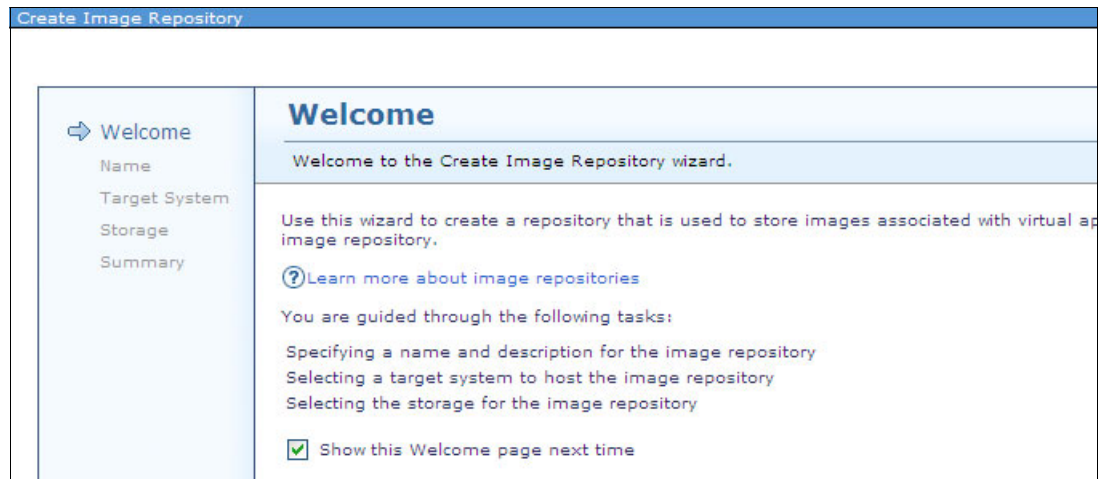


Figure 10-94 Image repository creation Welcome window

2. Enter the repository name as shown in Figure 10-95. In this example, the repository name is PowerVM-Repository.



Figure 10-95 Entering the repository name

3. Select the VIOS as shown in Figure 10-96. If the AIX deployment uses Storage Copy Services (SCS), the newly created logical unit number (LUN) is allocated to the VIOS. The VIOS assigns this LUN to the new virtual server by using vscsi mapping.

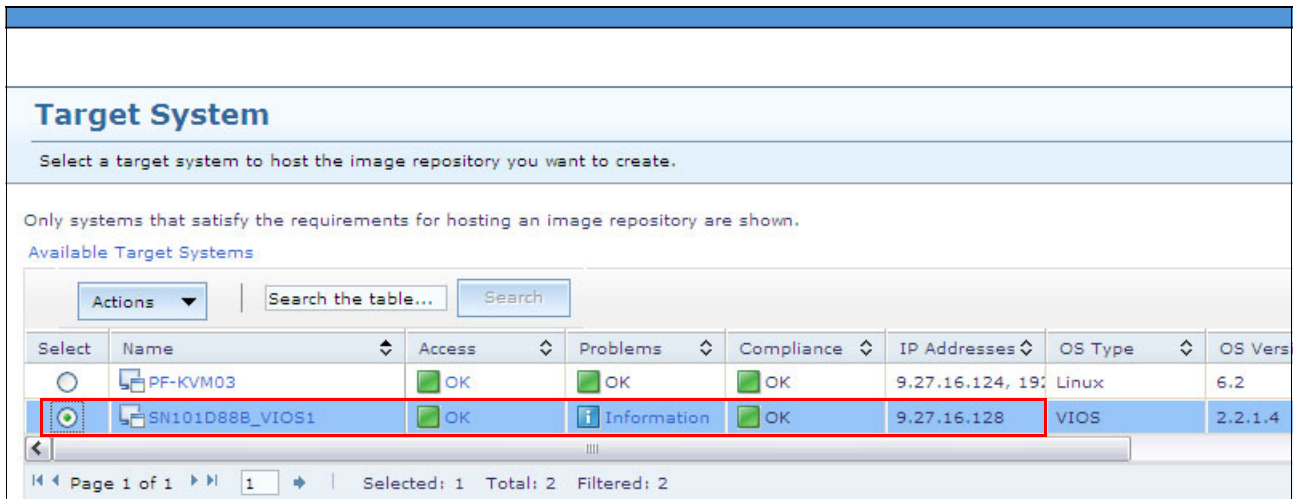


Figure 10-96 Selecting the target system

4. Select a target storage pool as shown in Figure 10-97. Whenever you create a virtual server, FSM creates a LUN in the storage pool and allocates it to VIOS.




Figure 10-97 Selecting a target storage

Figure 10-98 shows the Summary window.




Figure 10-98 Summary window

5. When the Launch Job window opens, click **OK**.
6. Check the log as shown in Figure 10-99.




Figure 10-99 Check the log

Creating the server system pool

To create the server system pool, perform these steps:

1. Click **Server system pools and members** as shown in Figure 10-100.



Figure 10-100 Server system pools creation

2. Click **Create** as shown in Figure 10-101.




Figure 10-101 Clicking Create

Figure 10-102 shows the Create Server System Pool Welcome window.



Figure 10-102 Create Server System Pool Welcome window

Figure 10-103 shows the Pooling Criteria window.

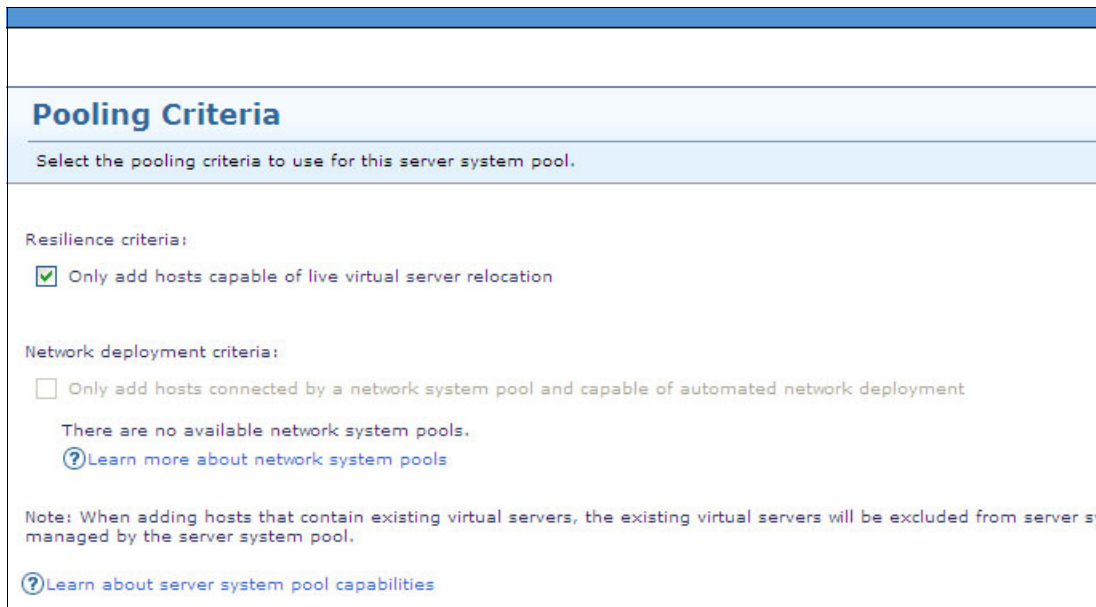


Figure 10-103 Pooling Criteria window

3. Select a physical system to use as a pool as shown in Figure 10-104.

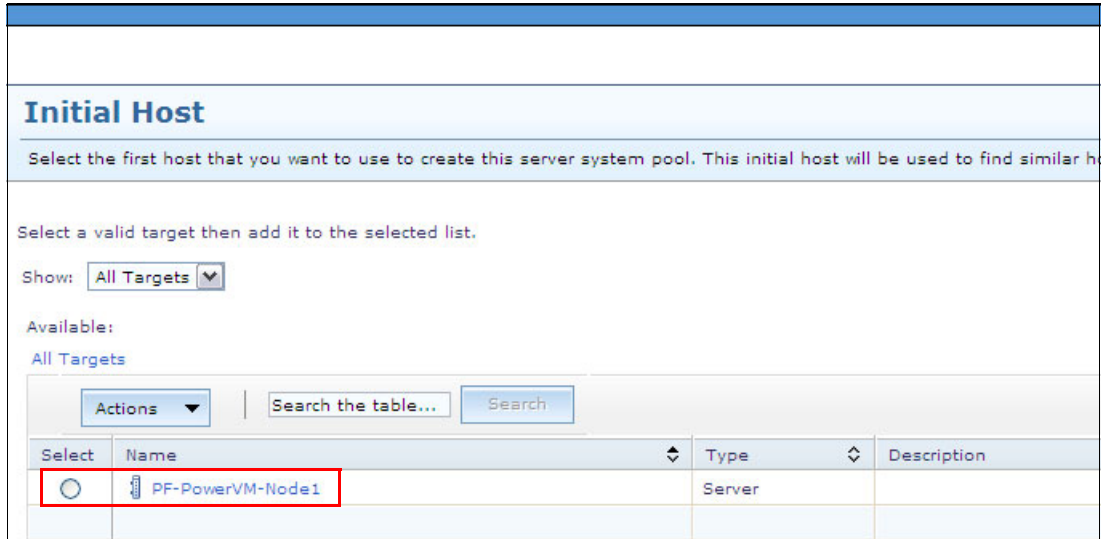


Figure 10-104 Selecting a physical server (part 1 of 2)

4. Click **Add** as shown in Figure 10-105.

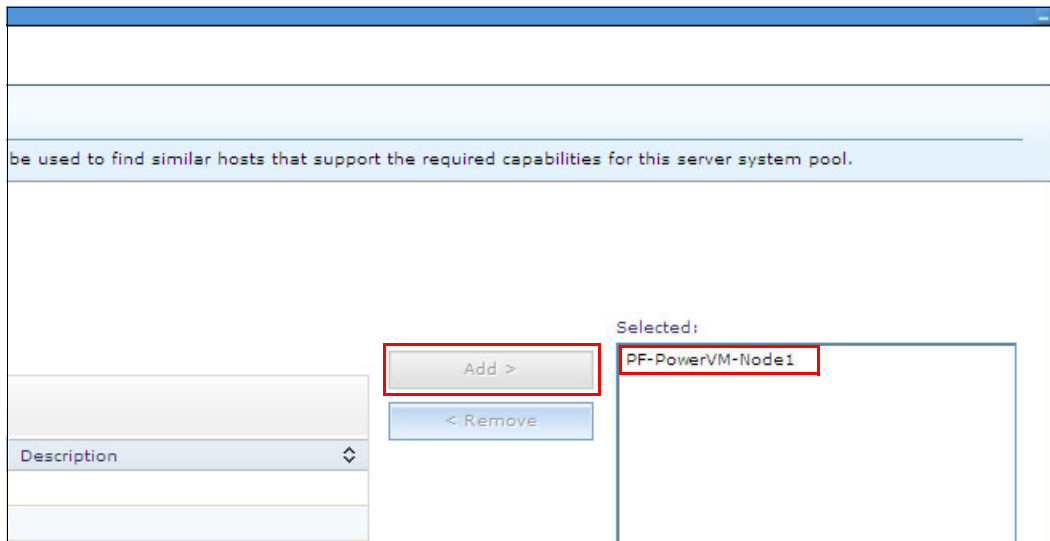


Figure 10-105 Selecting a physical server (part 2 of 2)

5. Select a storage pool that you want to use as shown in Figure 10-106. If you define more storage pools, you will see more storage pools in this window.



Figure 10-106 Selecting a storage pool

Figure 10-107 shows the Additional Hosts window.

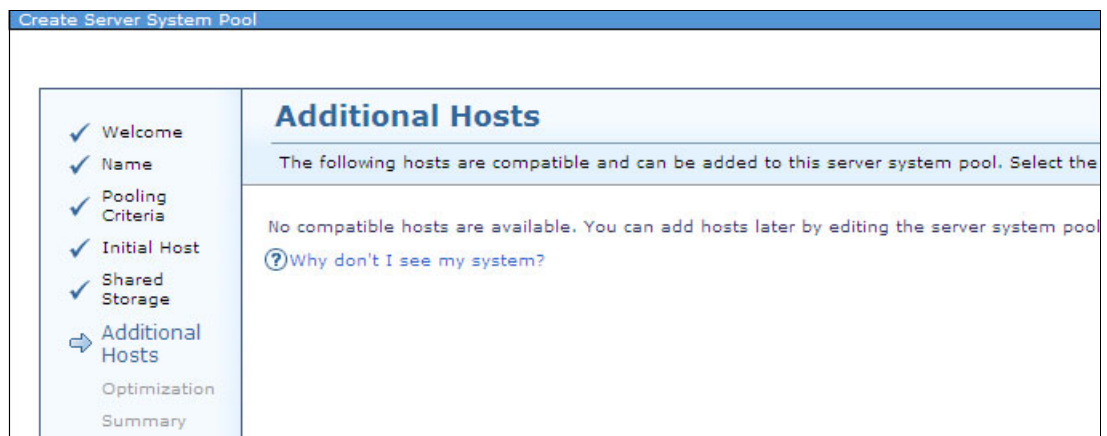


Figure 10-107 Additional Hosts window

- When you deploy a new virtual server in the server system pool, FSM deploys that server on the correct physical server automatically if you select **Allow optimizations to occur automatically**, as shown in Figure 10-108.

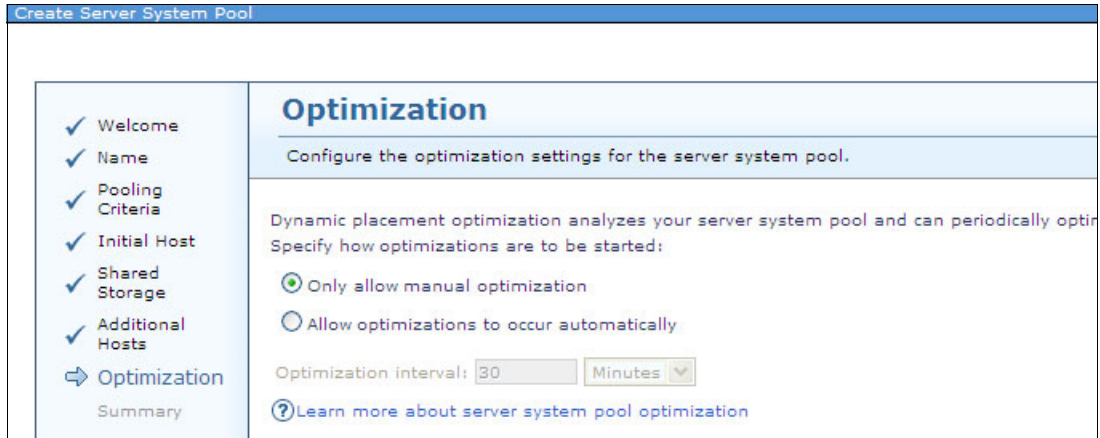


Figure 10-108 Optimization window

Figure 10-109 shows the Summary window.

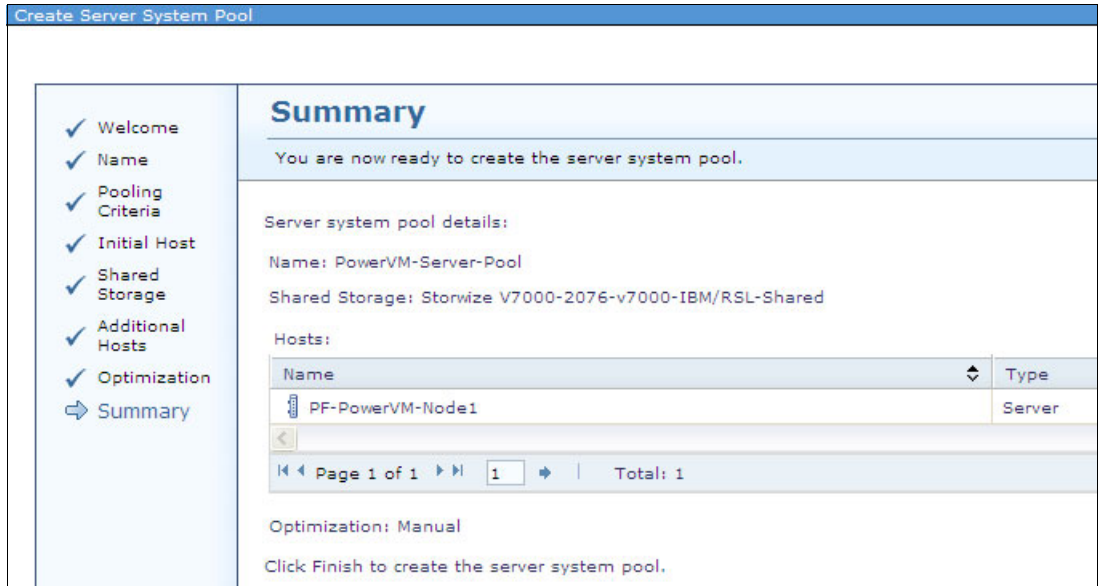


Figure 10-109 Summary window

- Click **Finish** to run the job.

8. Check the log as shown in Figure 10-110.

Click on job instance in the Name column in order to view its logs

Job Instance

Actions | Search the table... Search

Select	Name	Status
<input checked="" type="checkbox"/>	6/22/12 at 10:22 AM	Complete

Page 1 of 1 | 1 | Selected: 1 Total: 1

Job log

```

June 22, 2012 10:22:58 AM EDT-Level:1-MEID:0--MSG: Job "Create System Pool - Friday, June 22, 2012 10:22:57
June 22, 2012 10:22:58 AM EDT-Level:200-MEID:0--MSG: Subtask "Create a server system pool" activated.
June 22, 2012 10:22:58 AM EDT-Level:200-MEID:0--MSG: Starting clients
June 22, 2012 10:22:58 AM EDT-Level:100-MEID:0--MSG: Clients started for task "Create a server system pool"
June 22, 2012 10:22:58 AM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
June 22, 2012 10:22:58 AM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Active".
June 22, 2012 10:22:58 AM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
June 22, 2012 10:23:03 AM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
June 22, 2012 10:23:03 AM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
June 22, 2012 10:23:03 AM EDT-Level:100-MEID:0--MSG: Create system pool complete.

```

Figure 10-110 Checking the log

Preparation for capture

To capture the virtual server in an SCS environment, perform these steps:

1. Copy the `vmc.vsae.tar` file from Flex System Manager by using the `scp` command as shown in Figure 10-111.

```

USERID@FSM-5CF3FC5F54EF:~> scp /opt/ibm/director/proddata/activation-engine/vmc.vsae.tar root@9.27.16.129:/o
pen/
The authenticity of host '9.27.16.129 (9.27.16.129)' can't be established.
RSA key fingerprint is 44:c4:61:8d:56:b2:ba:e5:cd:7b:29:bc:43:bc:16:67.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '9.27.16.129' (RSA) to the list of known hosts.
root@9.27.16.129's password:
vmc.vsae.tar                                     100% 7900KB   7.7MB/s   00:00
USERID@FSM-5CF3FC5F54EF:~> █

```

Figure 10-111 Copying the `vmc.vsae.tar` file from FSM

2. Extract the contents of the .tar file by using the `tar -xvf vmc.vsa.e.tar` command, as shown in Figure 10-112.

```
# ls -al
total 15808
drwxr-xr-x  2 root    system      256 Jun 22 14:51 .
drwxr-xr-x  5 root    system      4096 Jun 22 14:48 ..
-rw-r--r--  1 root    system     8089600 Jun 22 14:35 vmc.vsa.e.tar
# tar -xvf vmc.vsa.e.tar
x activation-engine-2.1-1.13.aix5.3.noarch.rpm, 86482 bytes, 169 tape blocks
x activation-engine-jython-2.1-1.13.aix5.3.noarch.rpm, 7871473 bytes, 15374 tape blocks
x activation-engine-libxml2-python-2.1-1.13.noarch.rpm, 1569 bytes, 4 tape blocks
x activation-engine-libxml2-python-2.1-1.13.aix5.3.noarch.rpm, 1171 bytes, 3 tape blocks
x activation-engine-2.1-1.13.noarch.rpm, 74360 bytes, 146 tape blocks
x activation-engine-python-xml-2.1-1.13.noarch.rpm, 1553 bytes, 4 tape blocks
x activation-engine-python-xml-2.1-1.13.noarch.rpm, 1553 bytes, 4 tape blocks
x vmc-vsa.e-ext-1.1.0-1.noarch.rpm, 28958 bytes, 57 tape blocks
x linux-install.sh, 3408 bytes, 7 tape blocks
x aix-install.sh, 2233 bytes, 5 tape blocks
```

Figure 10-112 Unpacking the tar file

3. For AIX, ensure that the `JAVA_HOME` environment variable is set and points at a Java runtime environment (JRE), as shown in Figure 10-113.

```
# set JAVA_HOME=/usr/java5/jre
# echo $JAVA_HOME
```

Figure 10-113 Setting up the environment

4. Run `aix-install.sh` as shown in Figure 10-114.

```
# ./aix-install.sh
Install VSAE and VMC extensions
JAVA_HOME=/usr/java5/jre
*sys-package-mgr*: processing new jar, '/opt/ibm/ae/lib/jython/jython.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/vm.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/core.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/charsets.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/graphics.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/security.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmpkcs.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmorb.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmcfw.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmorbapi.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmjcefw.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmjgssprovider.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmjsseprovider2.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmjaaslm.jar'
*sys-package-mgr*: processing new jar, '/usr/java5/jre/lib/ibmcertpathprovider.jar'
```

Figure 10-114 Installing the VSAE file

5. Prepare the virtual server to be captured by running the `AE.sh --reset` command as shown in Figure 10-115.

```
# /opt/ibm/ae/AE.sh --reset
JAVA_HOME=/usr/java5/jre
[2012-06-22 16:16:23,391] INFO: Looking for platform initialization commands
[2012-06-22 16:16:23,398] INFO:  Version: AIX pureflexrb11 1 6 0001D88BD400

[2012-06-22 16:16:23,573] INFO: No initialization commands found....continuing
[2012-06-22 16:16:23,576] INFO: CLI parameters are ' --reset'
[2012-06-22 16:16:23,578] INFO: AE base directory is /opt/ibm/ae/
[2012-06-22 16:16:23,588] INFO: Resetting system. AP file: None. Interactive: False
[2012-06-22 16:16:23,752] INFO: In reset
[2012-06-22 16:16:23,753] INFO: Resetting products
[2012-06-22 16:16:23,755] INFO: Start to reset com.ibm.ovf.vmcontrol.system
ifconfig: error loading
/usr/lib/drivers/if_eth: No such file or directory
[2012-06-22 16:16:23,984] INFO: Reset: about to execute path /opt/ibm/ae/AS/vmc-system/resetAIX.sh
[2012-06-22 16:16:24,107] INFO: [com.ibm.ovf.vmcontrol.system] reset Not Activated
```

Figure 10-115 Preparing the virtual server to be captured

Tip: If you previously captured the virtual server and want to capture it again, run the following commands:

```
rm /opt/ibm/ae/AP/*
cp /opt/ibm/ae/AS/vmc-network-restore/resetenv /opt/ibm/ae/AP/ovf-env.xml
```

6. The virtual server shuts down automatically as shown in Figure 10-116.

```
Broadcast message from root@pureflexrb11.rtp.stglabs.ibm.com

Broadcast message from root@ (tty) at 16:16:32 ...

!!! SYSTEM BEING BROUGHT DOWN NOW !!!

JAVA_HOME=/usr/java5/jre
[2012-06-22 16:16:39,532] INFO: Looking for platform initialization commands
[2012-06-22 16:16:39,542] INFO:  Version: AIX pureflexrb11 1 6 0001D88BD400

[2012-06-22 16:16:39,799] INFO: No initialization commands found....continuing
[2012-06-22 16:16:39,803] INFO: CLI parameters are ' -d stop'
[2012-06-22 16:16:39,805] INFO: AE base directory is /opt/ibm/ae/
[2012-06-22 16:16:39,822] INFO: Stopping AE daemon.
[2012-06-22 16:16:39,841] INFO: AE daemon was not running.
Stopping The LWI Nonstop Profile...
Waiting for The LWI Nonstop Profile to exit...
```

Figure 10-116 Checking the `AE.sh` progress

Capture AIX by using Storage Copy Services

To capture AIX by using SCS, perform these steps:

1. Click **Capture** as shown in Figure 10-117.



Figure 10-117 Capturing AIX using SCS

2. Enter the name as shown in Figure 10-118.



Figure 10-118 Entering a name

3. Select a source virtual server as shown in Figure 10-119. Storwize V7000 runs a `flashcopy` command for the LUN on a virtual server that you select.

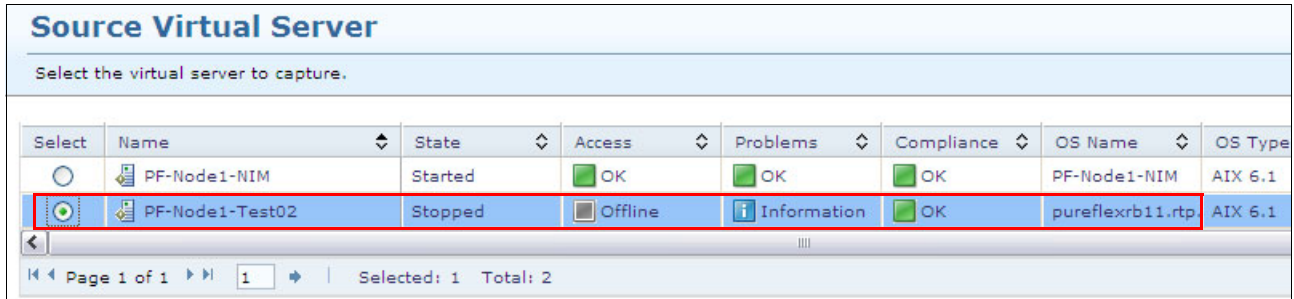


Figure 10-119 Selecting a source virtual server

4. Select **Virtual Server**, then click **Next**.
5. Select a repository as shown in Figure 10-120.




Figure 10-120 Selecting a repository

Figure 10-121 shows information about the LUN that is being captured.




Figure 10-121 Disk information to be captured

Figure 10-122 shows network mapping.

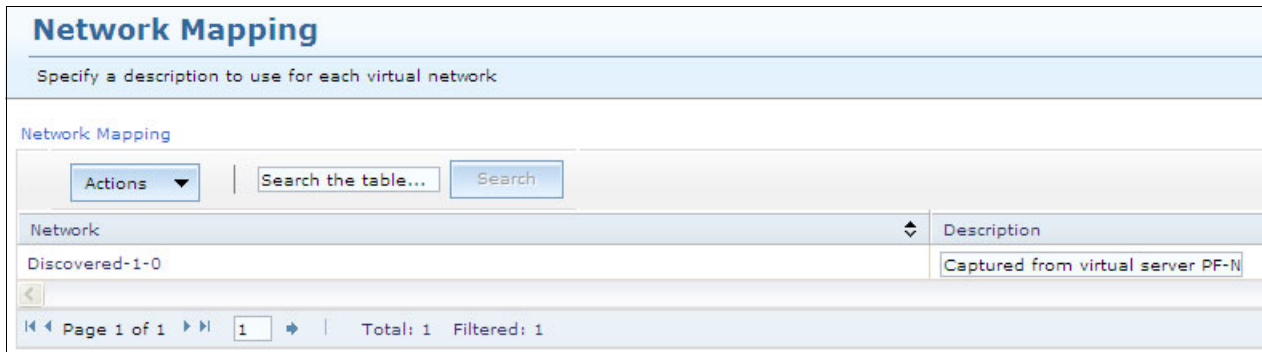


Figure 10-122 Network mapping window

6. Select the version as shown in Figure 10-123.

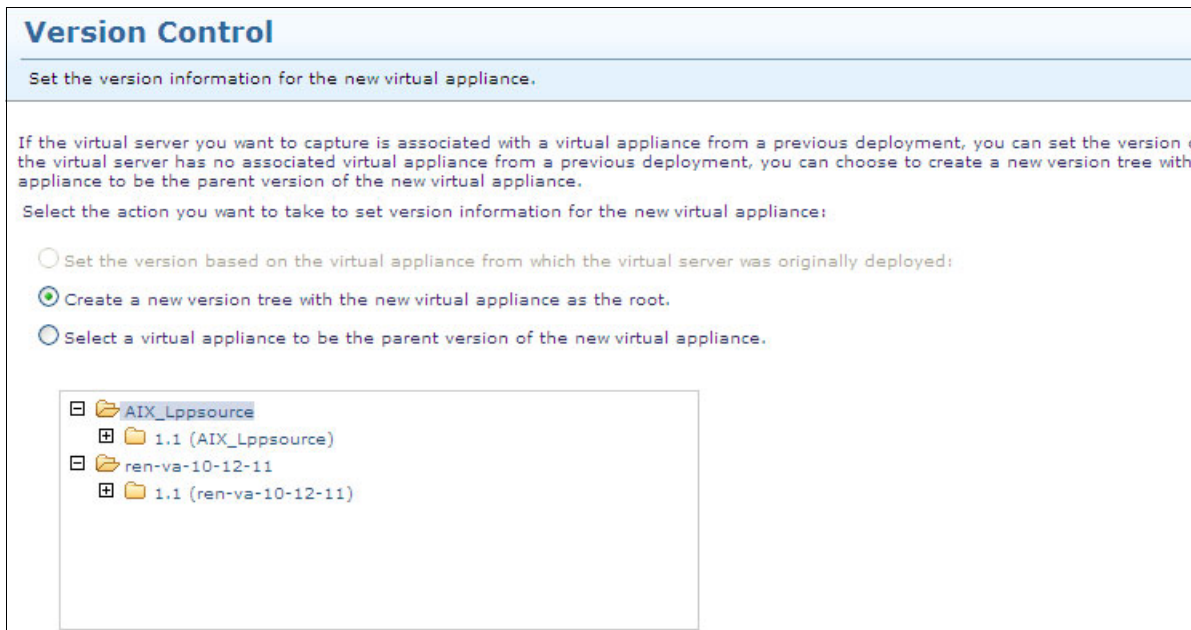


Figure 10-123 Version Control window

7. Figure 10-124 shows summary window. Click **Finish**.

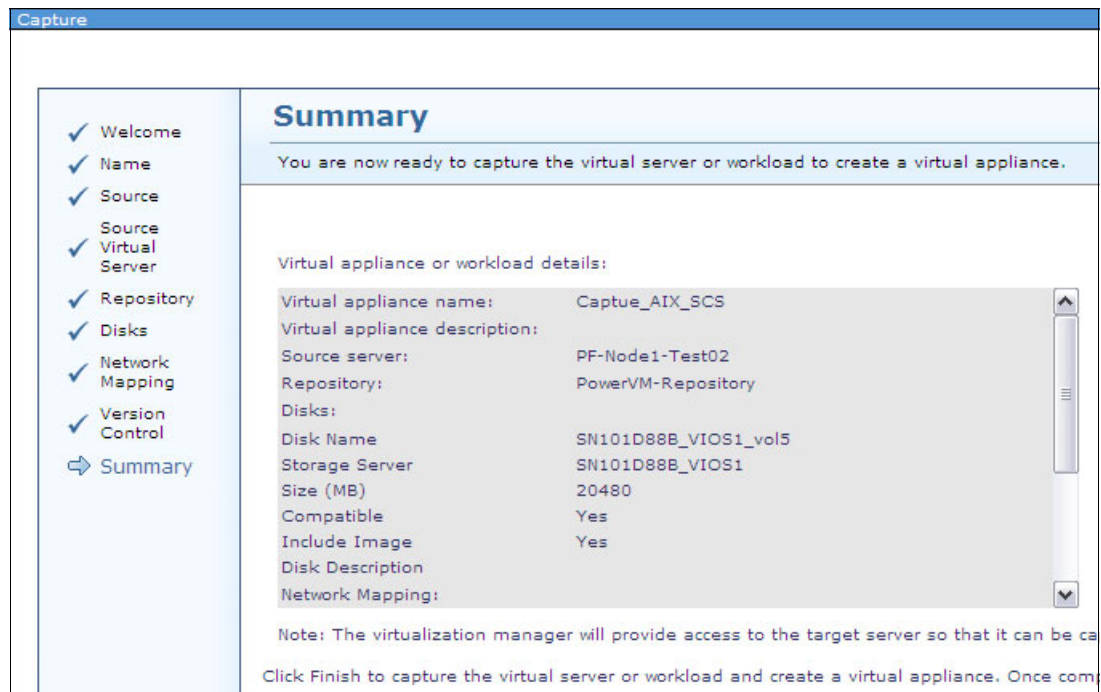


Figure 10-124 Summary window

8. Figure 10-125 shows the menu for starting jobs. Click **OK**.

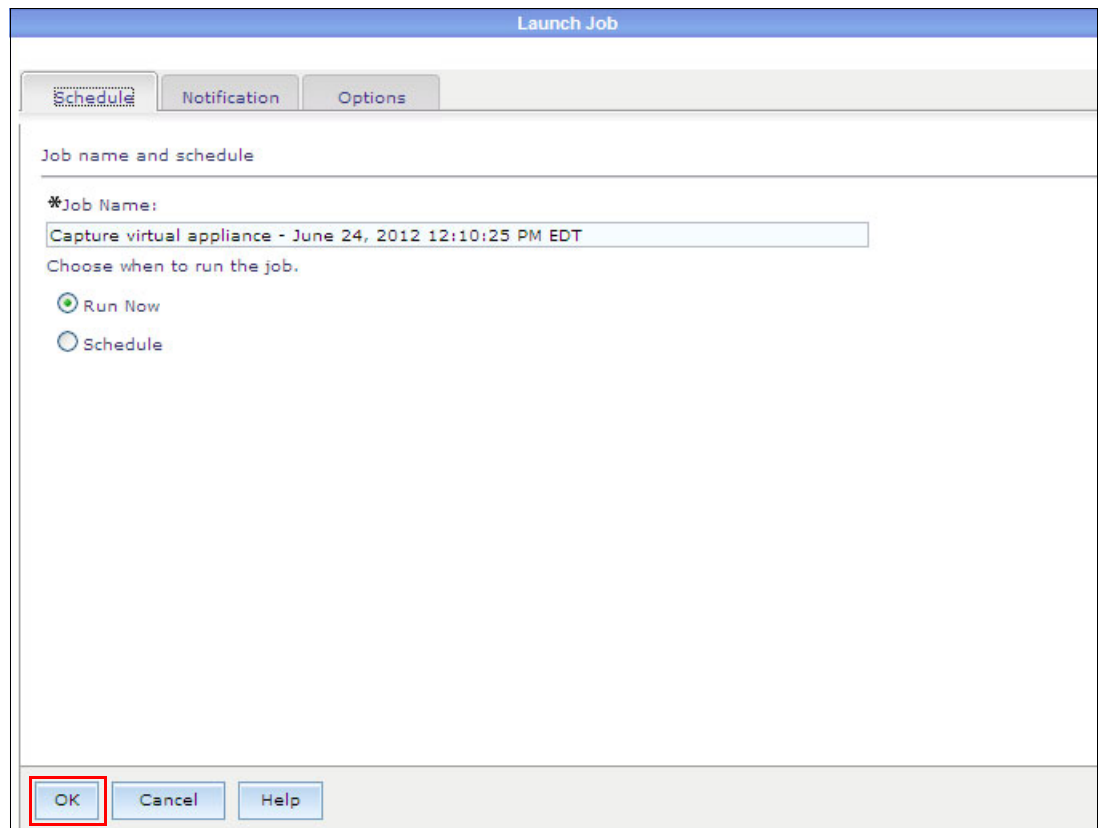


Figure 10-125 Launch Job menu

9. Check the log for the task completion, as shown in Figure 10-126.

The screenshot shows a web interface with tabs for General, Targets, History, and Logs. The 'Logs' tab is active. Below the tabs, there is a message: "Click on job instance in the Name column in order to view its logs".

Under "Job Instance", there is a table with columns "Select", "Name", and "Status". One row is visible with a checked checkbox, the name "6/24/12 at 12:10 PM", and the status "Complete".

Below the table is a pagination bar: "Page 1 of 1", "1", "Selected: 1 Total: 1".

Under "Job log", there is a list of log entries. The last few entries indicate completion:

- June 24, 2012 12:10:55 PM EDT-Level:100-MEID:0--MSG: Capture virtual appliance complete.
- June 24, 2012 12:10:55 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
- June 24, 2012 12:10:55 PM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
- June 24, 2012 12:10:55 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
- June 24, 2012 12:10:55 PM EDT-Level:150-MEID:0--MSG: Creating new container for the software image. The derby container ID is '1'
- June 24, 2012 12:10:55 PM EDT-Level:150-MEID:0--MSG: Discovering software image where the Derby ID is '1'
- June 24, 2012 12:10:54 PM EDT-Level:150-MEID:0--MSG: DNZLOP409I Creating the OVF for the virtual appliance.
- June 24, 2012 12:10:53 PM EDT-Level:150-MEID:0--MSG: DNZLOP909I Copying disk images
- June 24, 2012 12:10:53 PM EDT-Level:150-MEID:0--MSG: DNZLOP414I The virtual server is using disk group DG_06.24.2012-11:27:52:28 SAN volumes: [SN101D888_VIOS1_vol5].
- June 24, 2012 12:10:53 PM EDT-Level:150-MEID:0--MSG: DNZLOP948I New disk group: DG_06.24.2012-12:10:53:072
- June 24, 2012 12:10:53 PM EDT-Level:150-MEID:0--MSG: DNZLOP900I Requesting SAN volume(s)
- June 24, 2012 12:10:48 PM EDT-Level:150-MEID:0--MSG: DNZLOP912I Disk group to be captured: DG_06.24.2012-11:27:52:281

Figure 10-126 Check log

10. Check the captured image as shown in Figure 10-127.

The screenshot shows the "Virtual Appliances" section of a management console. It includes summary statistics for deployment and capture, a list of common tasks, and a table of virtual appliances.

Summary statistics:

- What to deploy: 4 Virtual appliances
- Where to deploy: 8 Existing virtual servers, 2 Hosts and 2 server system pools
- What to capture: 1 Workloads, 6 Virtual servers and operating systems
- Where to store: 3 Image repositories

Common tasks:

- Deploy virtual appliance
- Capture
- Import
- View active and scheduled jobs
- View virtual appliance versions
- Create image repository

Virtual Appliances (View Members) table:

Select	Name	Operating System	Repository
<input type="checkbox"/>	AIX_Lppsouce	IBM AIX	pureflexrb13.rtp.stglabs.ibm
<input type="checkbox"/>	Captue_AIX_SCS	IBM AIX 6	PowerVM-Repository
<input type="checkbox"/>	CapturedVMonKVM	Linux	KVMimagesrepo
<input type="checkbox"/>	ren-va-10-12-11	Linux	KVMimagesrepo

At the bottom, there is a pagination bar: "Page 1 of 1", "1", "Selected: 0 Total: 4 Filtered: 4".

Figure 10-127 Checking a captured image

10.3 Deploying virtual machines

This section addresses different types of deployment methods of previously captured virtual machines (VMs). It describes the following methods:

- ▶ Deploying virtual machines by using the LPP_source
- ▶ Deploying a virtual machine by using mksysb
- ▶ Deploying a virtual machine by using Storage Copy Services (SCS)

10.3.1 Deploying virtual machines by using the LPP_source

To deploy VMs by using the LPP_source, perform these steps:

1. Click the **Virtual Appliances** tab.
2. Select an LPP_source, then click the **Deploy Virtual Appliance** task, as shown in Figure 10-128.




Figure 10-128 VM deployment by using the LPP_source

3. Figure 10-129 shows the Deploy Virtual Appliance Welcome window. Click **Next**.




Figure 10-129 Welcome window

- Select a target onto which to deploy the LPP_source as shown in Figure 10-130. It can be a physical server or a partition. 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 Addresses	Installed OS	Description
<input checked="" type="radio"/>	PF-PowerVM-Node1	Started	9.27.21.44, fd55:		
<input type="radio"/>	PF-PowerVM-Node2	Standby	9.27.21.46, fe80:		

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

Deploy to an existing virtual server:

Select	Name	State	IP Addresses	Description
<input type="radio"/>	PF-Node1-NIM	Started	9.27.16.131	
<input type="radio"/>	SN101D88B_VIOC1	Started		

Page 1 of 1 | 1 | Selected: 0 Total: 2 Filtered: 2

Figure 10-130 Choosing an available target

- Enter the workload name as shown in Figure 10-131, then click **Next**.

Workload Name

A workload is created as a result of deploying the virtual appliance.

*Specify a unique name for the workload.

PF-Node1-Test1

Figure 10-131 Entering the workload name

- Figure 10-132 shows storage mapping. You can choose to assign each disk in the table to either a storage volume or storage pool. For this example, deploy a disk size of 9,537 MB. Click **Assign to Storage Pool**.

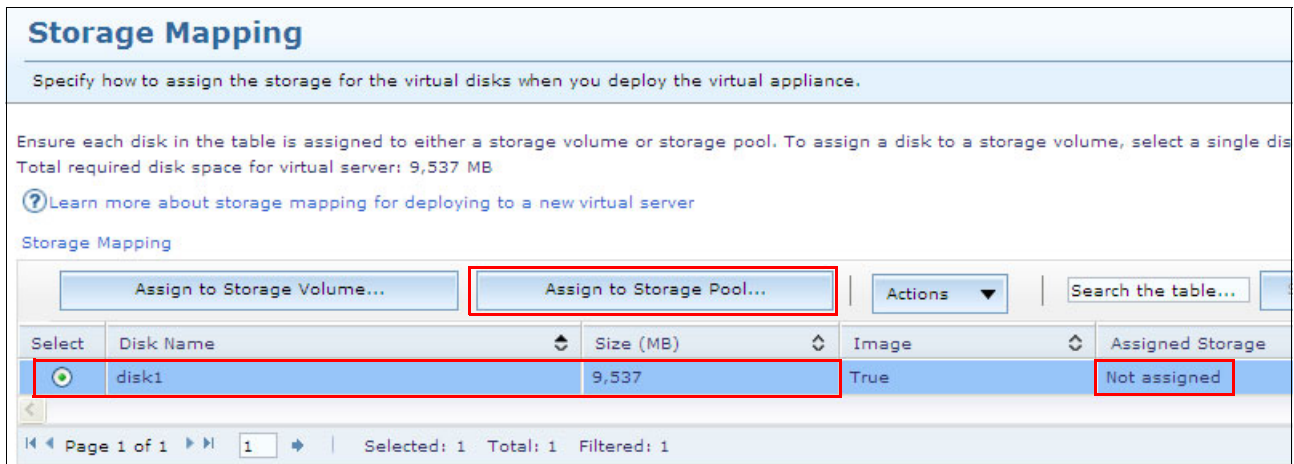


Figure 10-132 Storage Mapping window

- Assign the disk to a new virtual server from the VIOS rootvg storage pool as shown in Figure 10-133. Click **OK**.

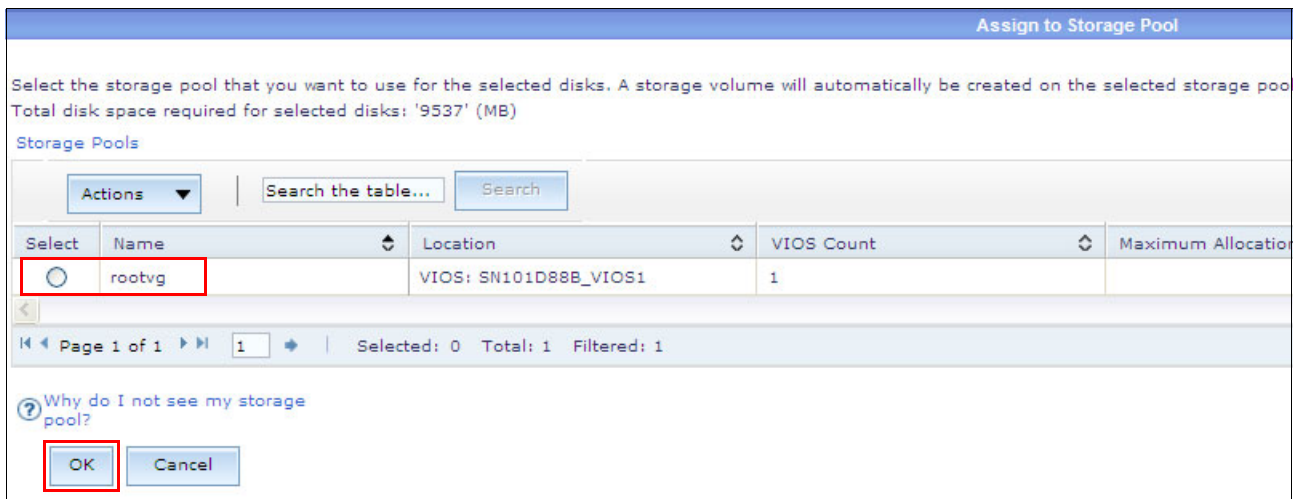


Figure 10-133 Assigning a disk to a new virtual server

Tip: There are three ways to allocate disk to a virtual server through VMControl:

- ▶ lv from vg (one of the storage pool types).
- ▶ A LUN is already assigned to VIOS by the storage subsystem.
- ▶ A LUN is allocated to VIOS upon request by using SMI-S (another storage pool type).

8. Figure 10-134 shows the Storage Mapping view. Click **Next**.

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
 Total required disk space for virtual server: 9,537 MB
 ? 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...

Select	Disk Name	Size (MB)	Image	Assigned Storage
<input type="radio"/>	disk1	9,537	True	Storage pool: rootvg

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

Figure 10-134 Storage Mapping view

9. Enter the IP address information as shown in Figure 10-135, then click **Next**.

Product

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

System Level Networking

Short host name for the system. PF-Node1-

DNS domain name for the system. rtp.stglabs

IP addresses of DNS servers for system. 9.42.242.2

Default IPv4 gateway. 9.27.16.1

Network adapter configuration for Network adapter 1 on Network 1

Internet Protocol Version 4

Static IP address for the network adapter "Network adapter 1 on Network 1". 9.27.16.12

Static network mask for network adapter "Network adapter 1 on Network 1". 255.255.255.0

Deployment use

The adapter order for network adapter "Network adapter 1 on Network 1".

Default network

NIM-specific settings

NIM-specific settings

NIM Resource or Resource Group

Figure 10-135 IP address configuration

10. Review the information in the Summary window that is shown in Figure 10-136, then click **Finish** to run the job.



Figure 10-136 Summary window

11. Figure 10-137 shows the Launch Job menu. Click **OK**.

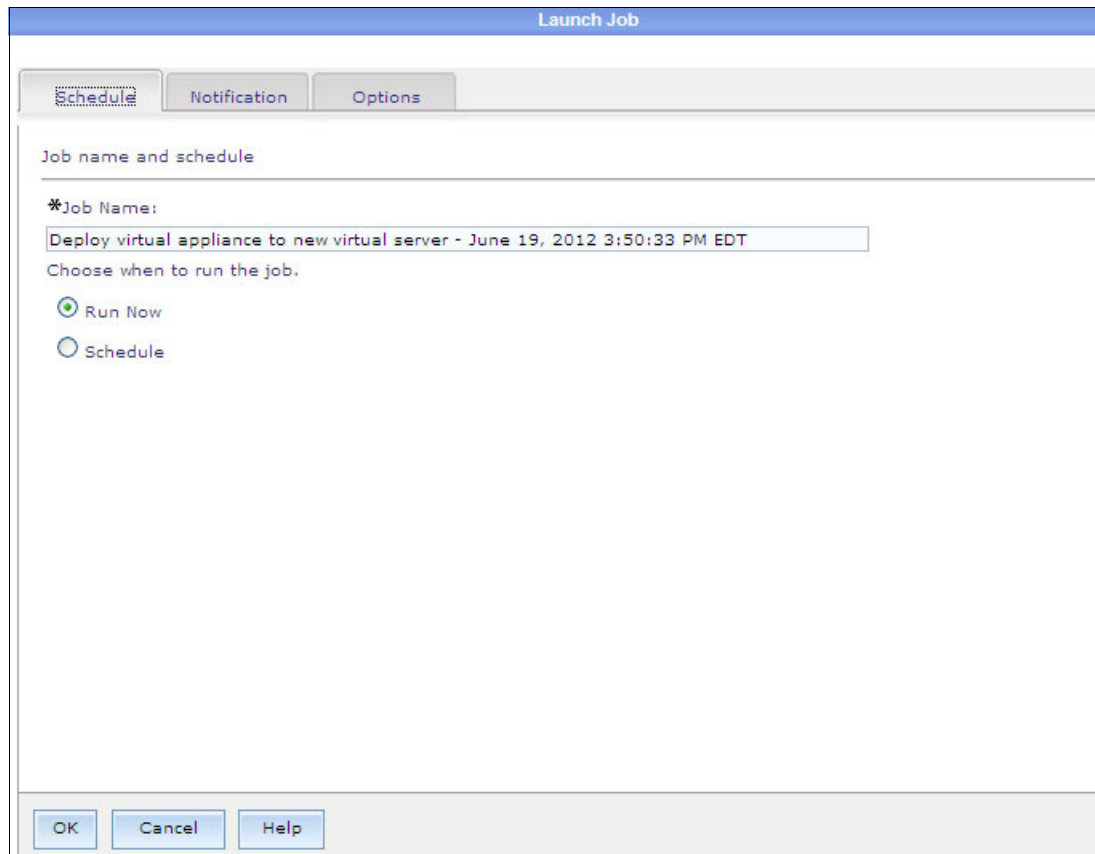


Figure 10-137 Launch Job menu

Figure 10-138 shows the newly deployed virtual server.

Select	Name	State	OS Name	OS Type and Version	Access
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008 6.	OK
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008 6.	OK
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK
<input type="checkbox"/>	vm003	Started			OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK
<input type="checkbox"/>	PF-Node1-NIM	Started	PF-Node1-NIM	AIX 6.1	OK
<input type="checkbox"/>	PF-Node1-Test01	Started			OK
<input type="checkbox"/>	SN101D88B_VIOC1	Started			OK
<input type="checkbox"/>	SN101D88B_VIOS1	Started			OK
<input type="checkbox"/>	PF-PowerVM-Node2	Standby			OK

Figure 10-138 Checking the newly deployed virtual server

12. Check the virtual server in the Resource Explorer window, as shown in Figure 10-139.

Select	Name	Access	Compliance	Problems
<input type="checkbox"/>	ETHERNET0-IBM*7895-22X*101D88B	OK	OK	OK
<input type="checkbox"/>	PF-Node1-NIM	OK	OK	OK
<input type="checkbox"/>	PF-Node1-Test01	OK	OK	OK
<input type="checkbox"/>	SN101D88B_VIOC1	OK	OK	OK
<input type="checkbox"/>	SN101D88B_VIOS1	OK	OK	OK

Figure 10-139 Virtual server in the Resource Explorer

- Discover the newly deployed virtual server as shown in Figure 10-140. Discover the virtual server and collect inventory in the Resource Explorer window. Perform this step even if you can see the virtual server in the Virtual Servers and Hosts tab in VMControl. Click the **No Access** link to get access to the newly discovered server.

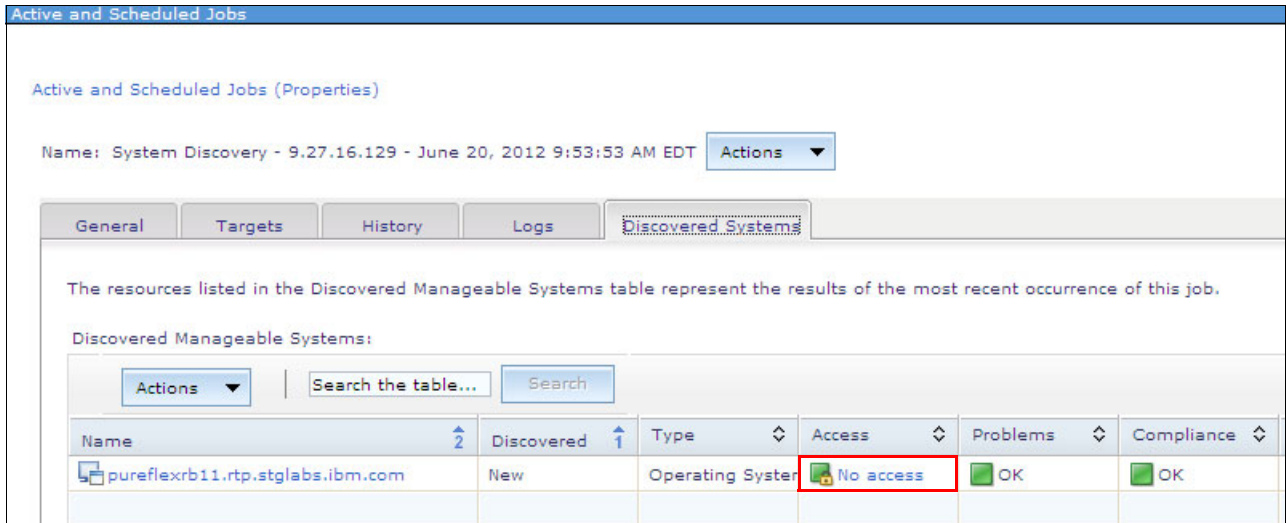


Figure 10-140 Checking the virtual server status

- Enter the credentials for the virtual server and click **Request Access** as shown in Figure 10-141.

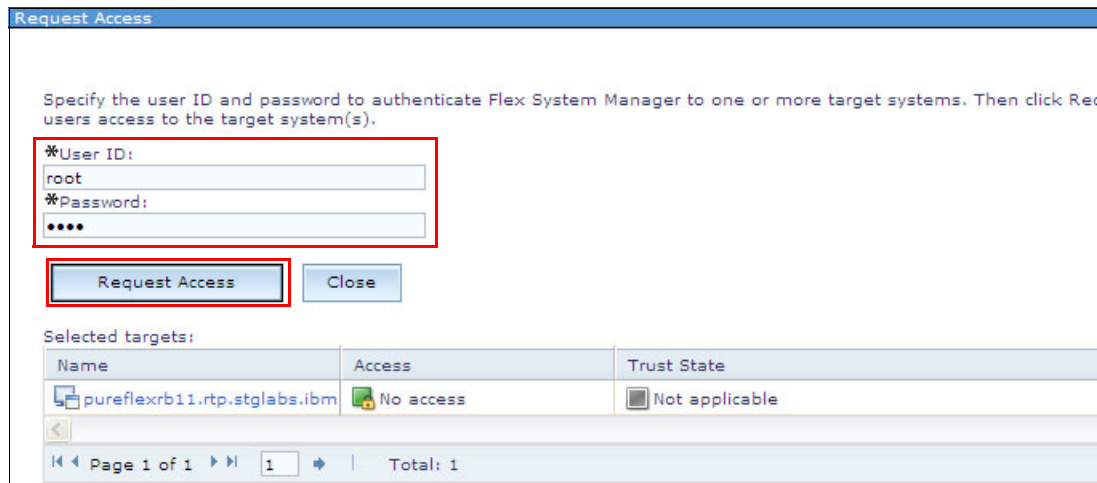


Figure 10-141 Entering the credentials for the virtual server

Figure 10-142 shows that the request access task is completed successfully.

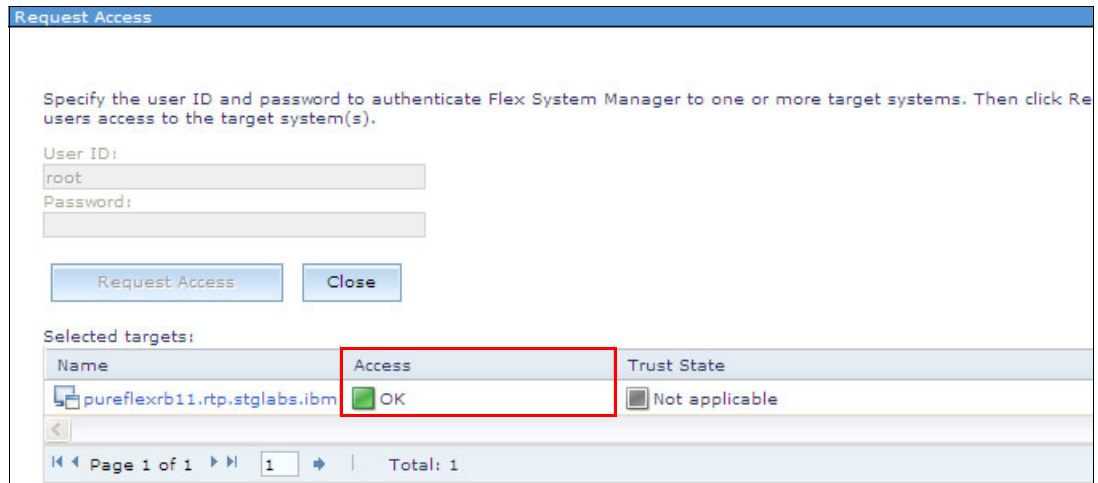


Figure 10-142 Request Access window

10.3.2 Deploying a virtual machine by using mksysb

To deploy VMs by using mksysb, perform these steps:

1. From the Virtual Appliances tab, click the **Deploy virtual appliance** task.
2. Click **Next** in the Welcome window.
3. Select the virtual appliance of the mksysb type as shown in Figure 10-143.

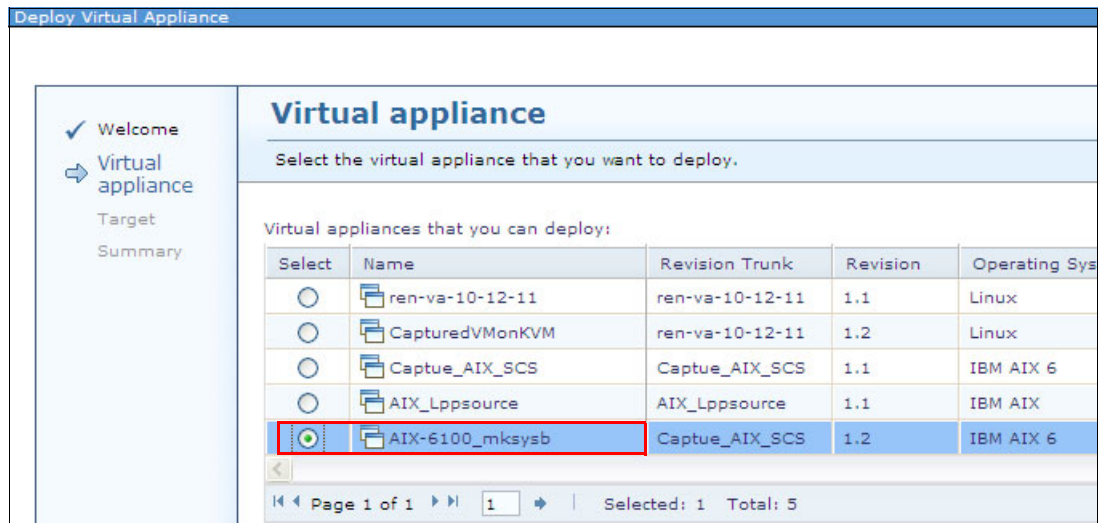


Figure 10-143 Selecting a virtual appliance of the mksysb type

4. Select a target where you want to deploy this mksysb image as shown in Figure 10-144, then click **Next**.

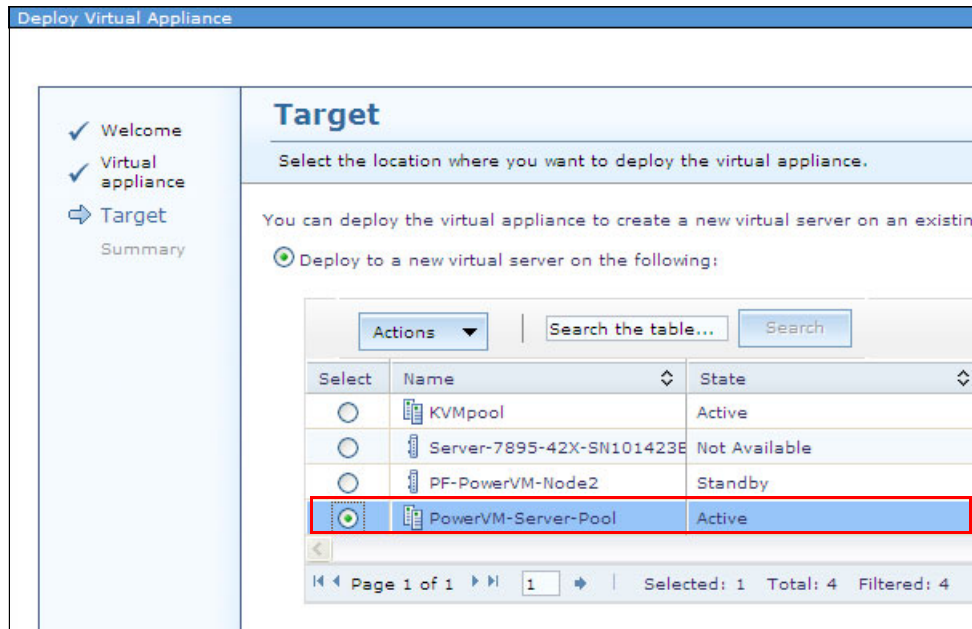


Figure 10-144 Selecting a target

5. Enter the workload name as shown in Figure 10-145 and click **Next**.




Figure 10-145 Entering the workload name

6. Figure 10-146 shows the networking mapping. Click **Next**.

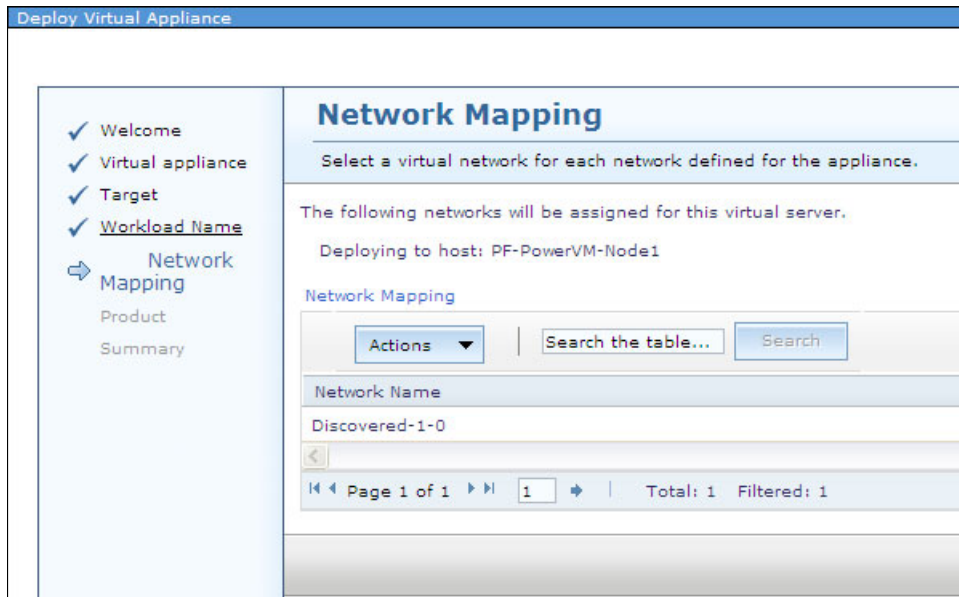


Figure 10-146 Network Mapping window

7. Enter the IP address and the Domain Name System (DNS) information as shown in Figure 10-147.

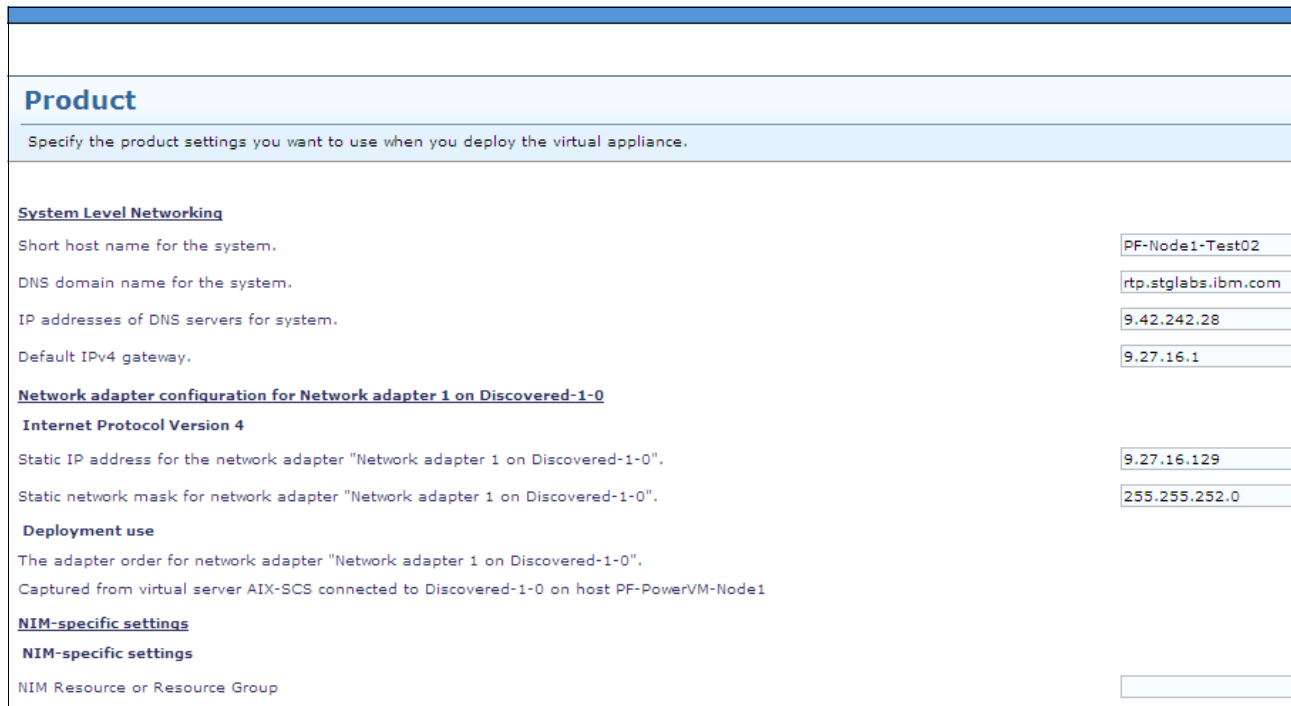


Figure 10-147 IP address configuration window

8. Click **Finish**, then click **OK** in the Launch Job window.

Figure 10-148 shows the newly deployed virtual machine that was deployed through the mksysb virtual appliance.




Figure 10-148 Newly deployed virtual server

10.3.3 Deploying a virtual machine by using Storage Copy Services (SCS)

To deploy a virtual machine by using the Storage Copy Services (SCS) method, perform these steps:

1. Click **Deploy virtual appliance** as shown in Figure 10-149.



Figure 10-149 Deploying the virtual appliance task

2. Figure 10-150 shows the Welcome window. Click **Next**.

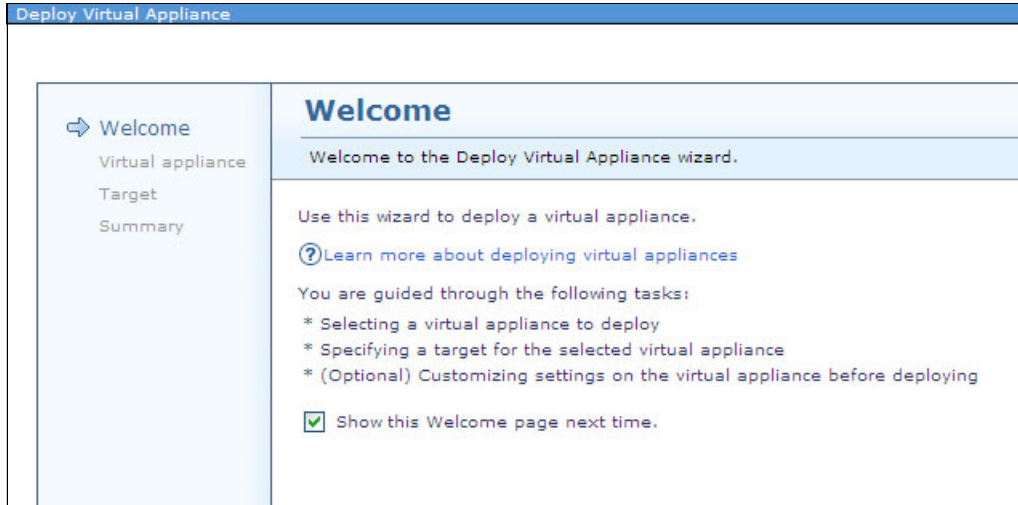


Figure 10-150 Welcome window

3. Select a virtual appliance as shown in Figure 10-151. Captue_AIX_SCS is an image that is created by the V7000 FlashCopy feature. Click **Next**.

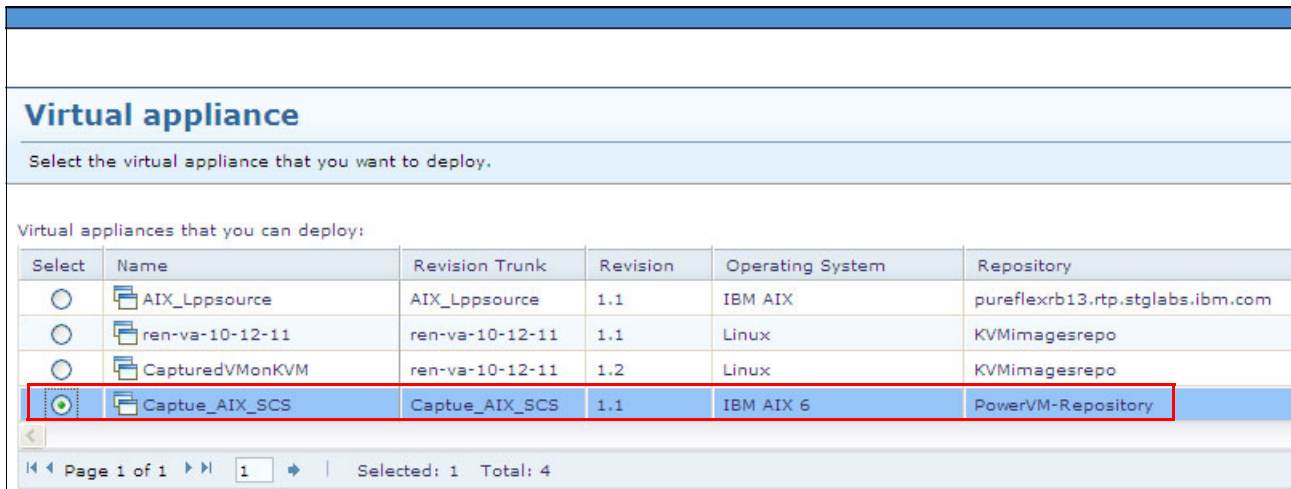


Figure 10-151 Selecting a virtual appliance

4. Select a target as shown in Figure 10-152, then 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 a

Deploy to a new virtual server on the following:

Select	Name	State	IP Addresses	Installed OS Name
<input type="radio"/>	KVMpool	Active		
<input type="radio"/>	Server-7895-42X-SN101423E	Started	169.254.3.196, fd55:faaf:e1ab:	
<input type="radio"/>	PF-PowerVM-Node2	Standby	9.27.21.46, fe80:0:0:0:5ef3:fcf	
<input checked="" type="radio"/>	PowerVM-Server-Pool	Active		

Page 1 of 1 | 1 | Selected: 1 Total: 4 Filtered: 4

Deploy to an existing virtual server:

Select	Name	State	IP Addresses	Description
<input type="radio"/>	PF-Node1-NIM	Started	0.0.0.0, 9.27.16.131	
<input type="radio"/>	test	Stopped		
<input type="radio"/>	SN101D88B_VIO1	Stopped		
<input type="radio"/>	PF-Node1-Test02	Stopped	9.27.16.129	

Figure 10-152 Select a target

5. Enter the workload name as shown in Figure 10-153, then click **Next**.

Deploy Virtual Appliance

- ✓ Welcome
- ✓ Virtual appliance
- ✓ Target
- ➔ Workload Name
- Network Mapping
- Product
- Summary

Workload Name

A workload is created as a result of deploying the virtual appliance.

*Specify a unique name for the workload.

PF-Node1-AIX-SCS

Figure 10-153 Enter the workload name

6. Figure 10-154 shows the networking mapping. Click **Next**.




Figure 10-154 Network Mapping window

7. Enter the IP address information as shown in Figure 10-155 and click **Next**.

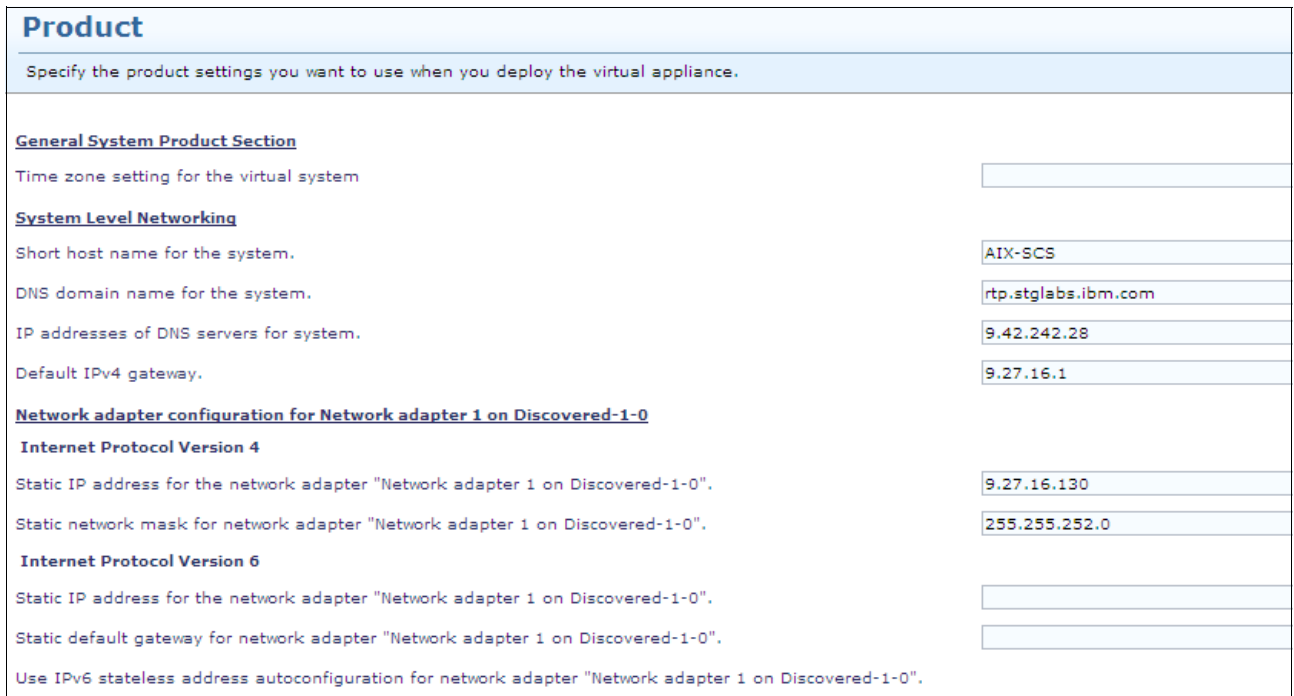


Figure 10-155 IP address configuration window

8. Figure 10-156 shows the summary. Click **Finish**.

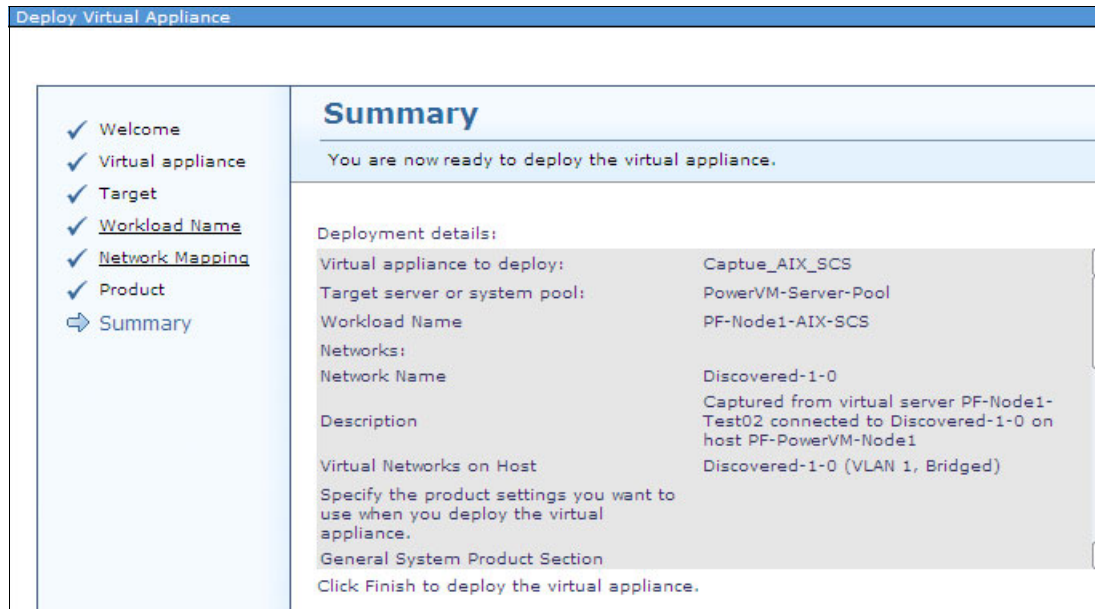


Figure 10-156 Summary window

9. Figure 10-157 shows the Launch Job window. Click **OK**.

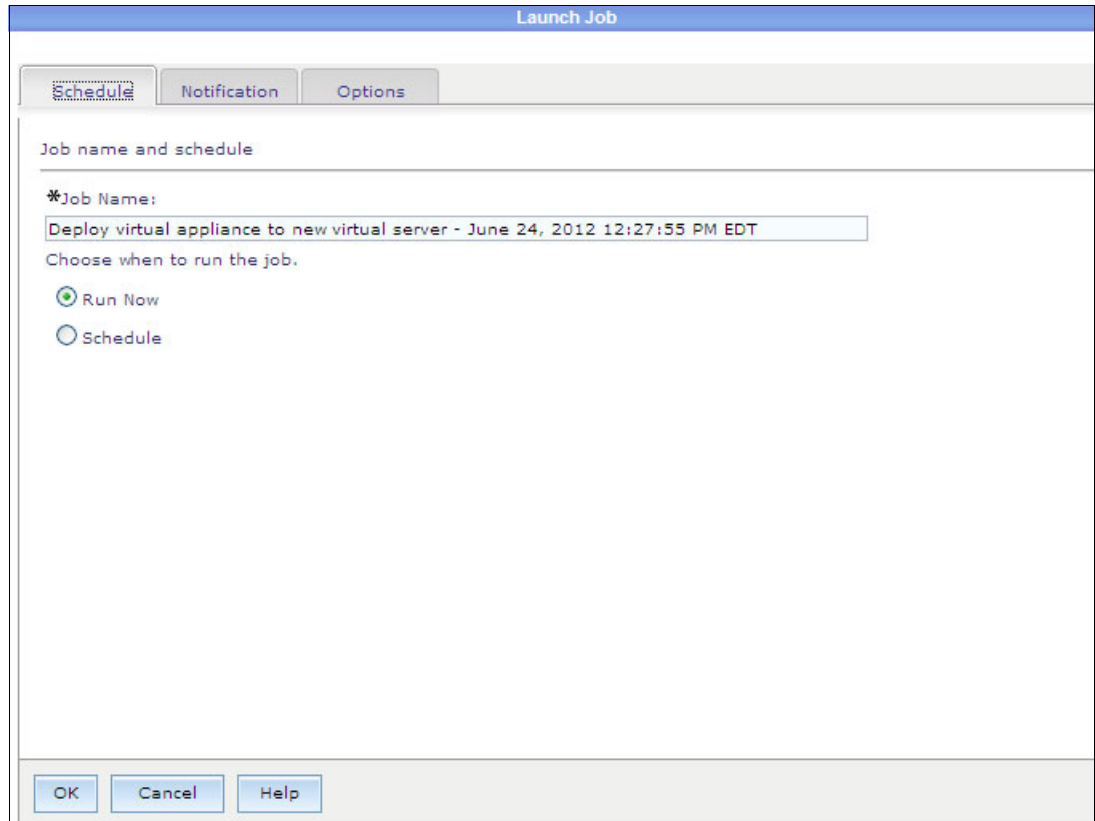
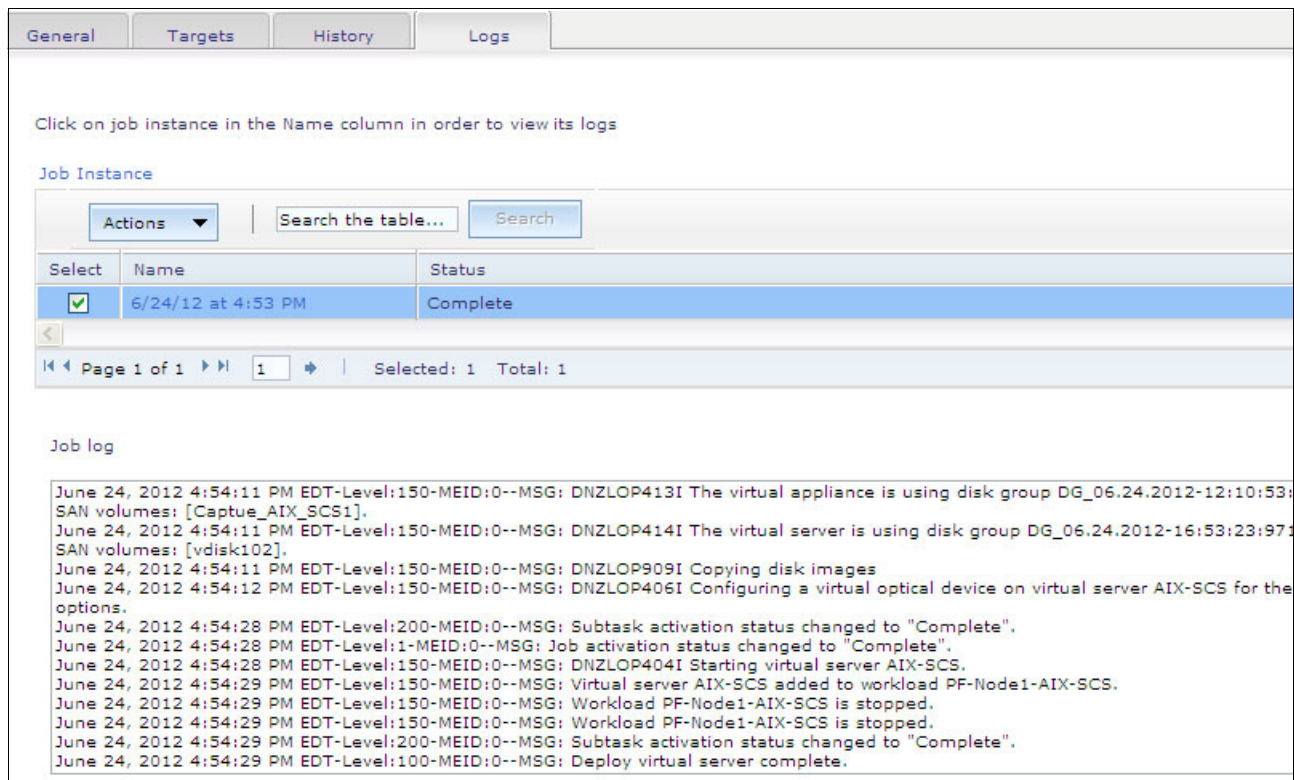


Figure 10-157 Launch Job window

Figure 10-158 shows the job log with the completed task.



Click on job instance in the Name column in order to view its logs

Job Instance

Select	Name	Status
<input checked="" type="checkbox"/>	6/24/12 at 4:53 PM	Complete

Page 1 of 1 | 1 | Selected: 1 Total: 1

Job log

```
June 24, 2012 4:54:11 PM EDT-Level:150-MEID:0--MSG: DNZLOP413I The virtual appliance is using disk group DG_06.24.2012-12:10:53:
SAN volumes: [Captue_AIX_SCS1].
June 24, 2012 4:54:11 PM EDT-Level:150-MEID:0--MSG: DNZLOP414I The virtual server is using disk group DG_06.24.2012-16:53:23:97:
SAN volumes: [vdisk102].
June 24, 2012 4:54:11 PM EDT-Level:150-MEID:0--MSG: DNZLOP909I Copying disk images
June 24, 2012 4:54:12 PM EDT-Level:150-MEID:0--MSG: DNZLOP406I Configuring a virtual optical device on virtual server AIX-SCS for the
options.
June 24, 2012 4:54:28 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
June 24, 2012 4:54:28 PM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
June 24, 2012 4:54:28 PM EDT-Level:150-MEID:0--MSG: DNZLOP404I Starting virtual server AIX-SCS.
June 24, 2012 4:54:29 PM EDT-Level:150-MEID:0--MSG: Virtual server AIX-SCS added to workload PF-Node1-AIX-SCS.
June 24, 2012 4:54:29 PM EDT-Level:150-MEID:0--MSG: Workload PF-Node1-AIX-SCS is stopped.
June 24, 2012 4:54:29 PM EDT-Level:150-MEID:0--MSG: Workload PF-Node1-AIX-SCS is stopped.
June 24, 2012 4:54:29 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
June 24, 2012 4:54:29 PM EDT-Level:100-MEID:0--MSG: Deploy virtual server complete.
```

Figure 10-158 Job log: Task completed

10.4 Relocating virtual machines

IBM Flex System Manager VMControl can relocate virtual servers in response to predicted hardware failures related to processors, memory subsystems, power source, and storage. You can also choose to relocate virtual servers for planned maintenance or downtime, or to adjust resources to improve performance.

The following relocation methods are supported in a PowerVM environment:

- ▶ Static relocation in virtual farms
- ▶ Live relocation in virtual farms
- ▶ Live relocation in server system pools

10.4.1 Manual relocation

You can choose to relocate one or more virtual servers from an existing host at any time. When you relocate virtual servers within server system pools, the relocation target is automatically identified.

10.4.2 Automatic relocation

VMControl server system pools can predict hardware failure problems and relocate virtual servers to maintain resilience. However, you might also want to monitor and adjust resources within your server system pool.

For example, you might want to monitor the hosts in your server system pool for high processor utilization. To do so, activate a threshold to monitor high and low values for processor utilization in your workloads. Then, if the threshold is reached, a message is displayed in the Server system pools dashboard, and in the Problems view.

10.4.3 Relocating virtual servers manually

Figure 10-159 shows the overall architecture for virtual server relocation.

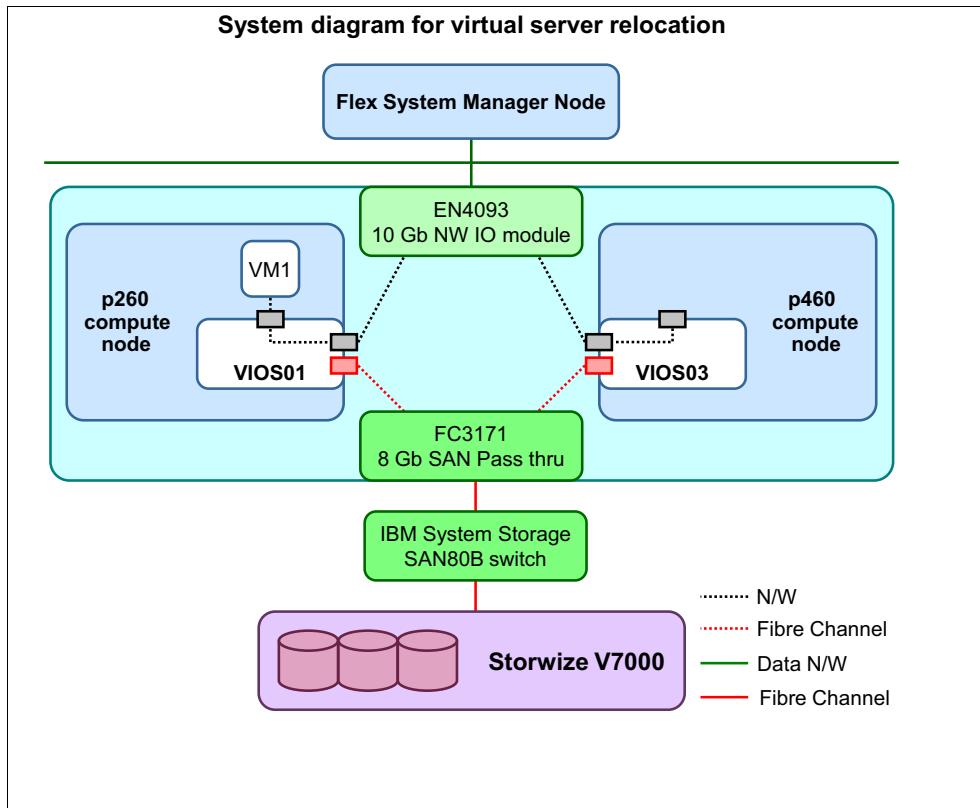


Figure 10-159 System diagram for virtual server relocation

Figure 10-160 shows the physical compute nodes. The example environment uses p260 (PF-PowerVM-Node1) and p460 (Server-7895-42X) for virtual server relocation.

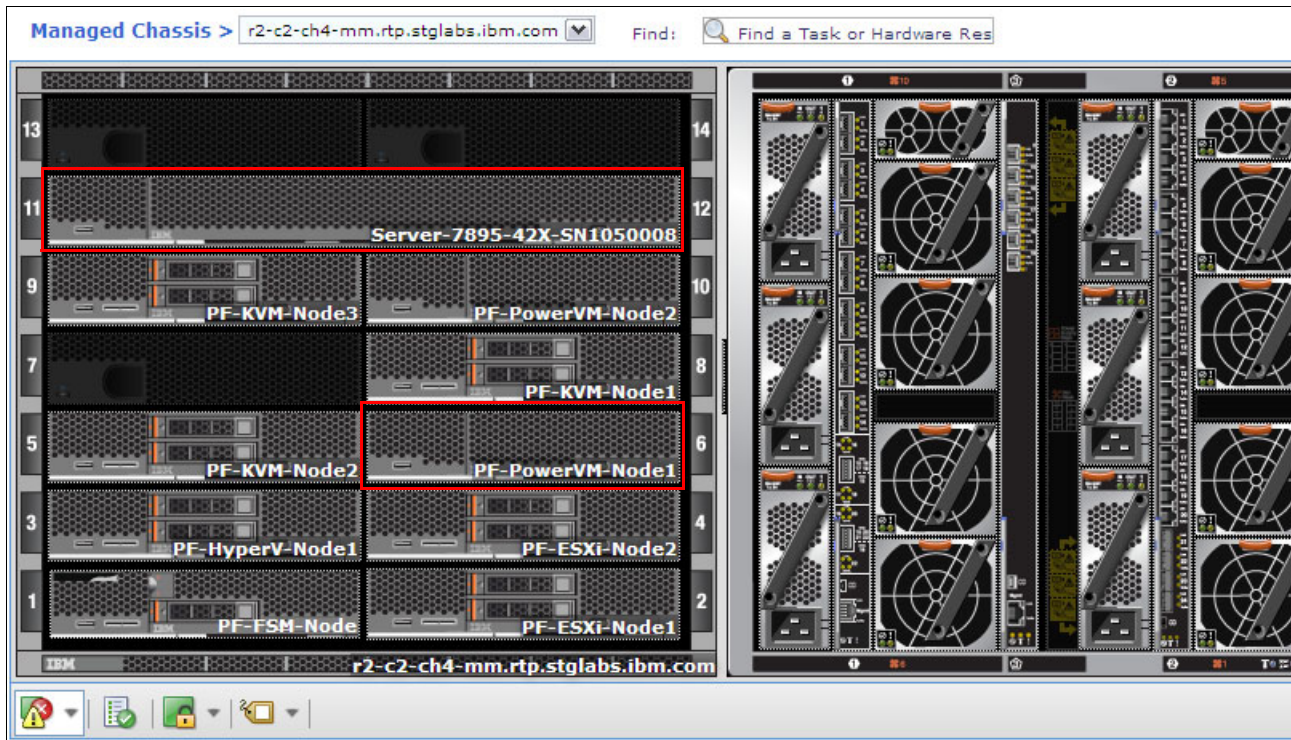


Figure 10-160 Chassis front view

This section describes several important steps that you must follow to set up virtual server relocation (Live Partition Mobility). For more information about Live Partition Mobility, see *IBM PowerVM Live Partition Mobility*, SG24-7460, at this website:

<http://www.redbooks.ibm.com/abstracts/sg247460.html?Open>

To relocate virtual servers manually, perform these steps:

1. Click **Manage Virtual Server** to check the VIOS profile settings as shown in Figure 10-161.

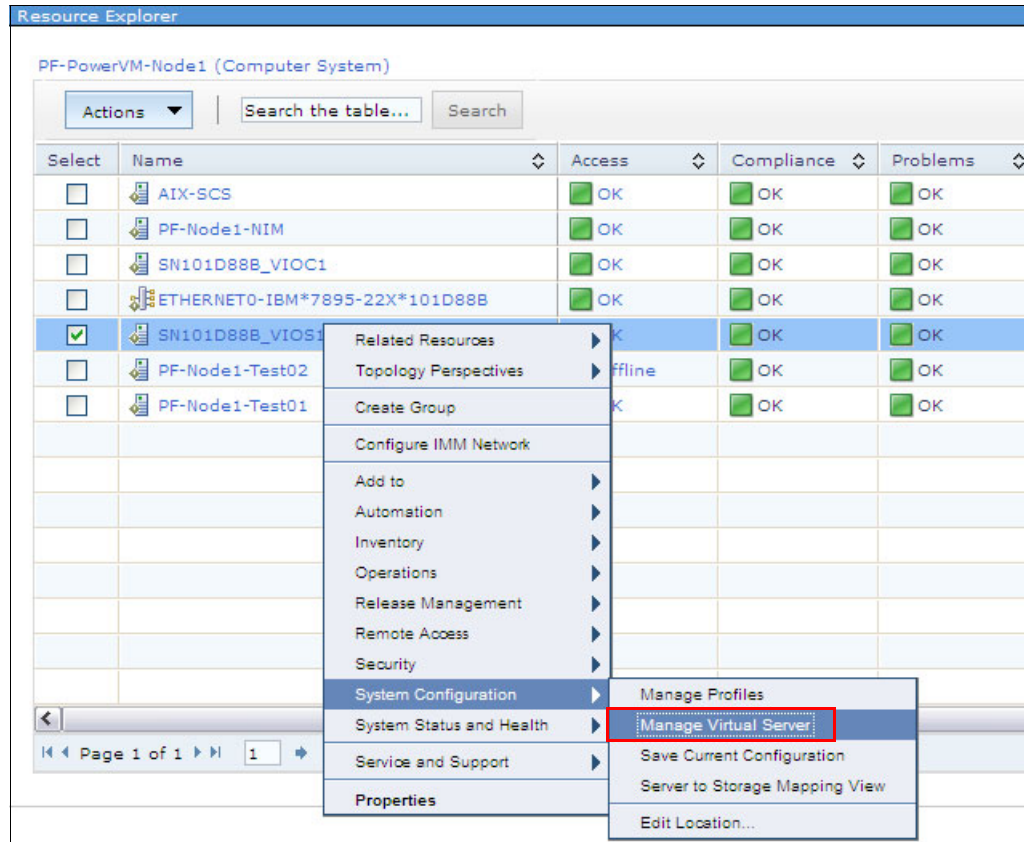


Figure 10-161 Launch Virtual Server settings

2. Select **Mover service** to perform Live Partition Mobility as shown in Figure 10-162.




Figure 10-162 Check VIOS server setting

- When you create a VIOS pair to perform Live Partition Mobility, priorities for Shared Ethernet Adapters need to be changed as shown in Figure 10-163.

Adapter Id*
11

Port Virtual Ethernet
1

VSI Type Id
[]

VSI Manager Id
[]

VSI Type Version
[]

IEEE Settings
Select this option to allow additional virtual LAN IDs for the adapter.

IEEE 802.1q compatible adapter
Maximum number of VLANs: 20
Additional VLAN IDs:
[] 1,20,48,..

Shared Ethernet Settings
Select Ethernet bridging to link (bridge) the virtual Ethernet to a physical network.

Use this adapter for Ethernet bridging*
Priority:*
1 (1 or 2)

Figure 10-163 Virtual Ethernet Adapter setting

- Figure 10-164 shows a virtual SCSI configuration. The Live Partition Mobility virtual server is AIX SCS. The AIX SCS partition has vhost4 as the vscsi in the VIOS environment.

Manage Virtual Server

Host: PF-PowerVM-Node1 Name: SN101D88B_VIOS1 Id: []

Environment: VIOS State: Started - RMC available

General Settings Virtual Storage Adapters

Processor

Memory

Network

Storage Adapters

Storage Devices

Media Devices

Physical IO

Available Virtual Slots: 92

[Add] [Remove] [Properties]

Select	Adapter(Id)	Type	Connecting virtual server
<input type="checkbox"/>	vhost2(2)	SCSI	PF-Node1-Test01(4)
<input type="checkbox"/>	vhost4(3)	SCSI	AIX-SCS(6)
<input type="checkbox"/>	vhost0(21)	SCSI	SN101D88B_VIOC1(2)
<input type="checkbox"/>	vhost1(22)	SCSI	PF-Node1-NIM(3)
<input type="checkbox"/>	vhost3(24)	SCSI	Any

Figure 10-164 Virtual SCSI configuration window

Figure 10-165 shows disk allocation information in the VIOS environment.

Host: PF-PowerVM-Node1 Name: SN101D88B_VIOS1 Id: ...
 Environment: VIOS State: Started - RMC available

Virtual Disks

Assign Unassign Refresh

Select	Name	Assigned Virtual Server
<input type="checkbox"/>	lp4vd1	PF-Node1-Test01(4)

Physical Volumes

Assign Unassign Refresh

Select	Name	Size (GB)	Assigned Virtual Server	Storage Pool	
<input type="checkbox"/>	hdisk5	20.0	Any (Virtual Slot 24)		U78AE.001.W
<input type="checkbox"/>	hdisk6	20.0	AIX-SCS(6)		U78AE.001.W
<input type="checkbox"/>	hdisk0	279.4		rootvg	U78AE.001.W
<input type="checkbox"/>	hdisk1	20.0	SN101D88B_VIOC1(2)		U78AE.001.W
<input type="checkbox"/>	hdisk2	200.0	PF-Node1-NIM(3)		U78AE.001.W
<input type="checkbox"/>	hdisk3	20.0			U78AE.001.W
<input type="checkbox"/>	hdisk4	20.0			U78AE.001.W

Figure 10-165 Disk allocation information

5. Shared disk drives on VIO Servers must have the Reserve policy set to no_reserve by using the `chdev` command, as shown in Figure 10-166.

```

$ lspv
NAME                PVID                VG                STATUS
hdisk0              0001d88be70e39c1   rootvg           active
hdisk1              0001d88bf1c70f7c   None
hdisk2              0001d88b00668e98   None
hdisk3              0001d88b01220754   None
hdisk4              none                None
hdisk5              0001d88b0f6e913a   None
hdisk6              0001d88b0f6e913a   None
$ chdev -dev hdisk6 -attr reserve_policy=no_reserve
hdisk6 changed
$
    
```

Figure 10-166 Set to no_reserve

6. From the Resource Explorer, click **Migrate** as shown in Figure 10-167.




Figure 10-167 Launch the relocation of a virtual server

7. Figure 10-168 shows the migration wizard. Go through the wizard, then click **Finish** after checking the summary. You can then observe the relocation of the virtual server.

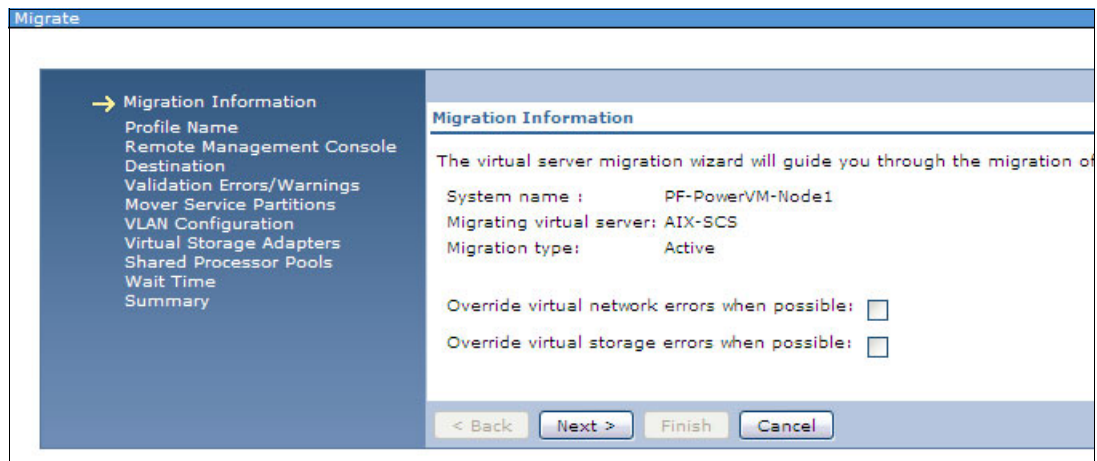


Figure 10-168 Virtual server relocation wizard

Figure 10-169 shows AIX-SCS running on the p460 (Server-7895-42X).

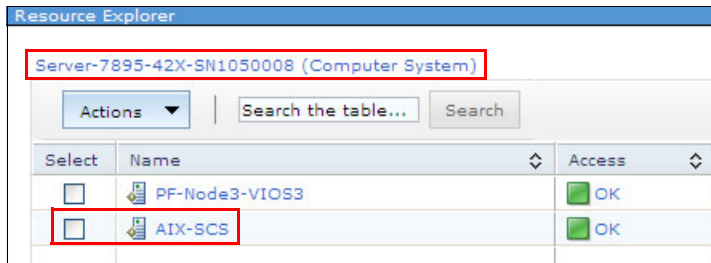



Figure 10-169 Checking the status of AIX running on p460



Managing the VMware environment with IBM Flex System Manager

This chapter addresses Flex System Manager (FSM) integration with the VMware environment. It describes common tasks that can be run on your VMware infrastructure by using FSM. These tasks include creating, editing, and relocating virtual servers, as well as reconfiguring clusters and working with maintenance mode.

This chapter also addresses how to use FSM to configure simple but powerful automation plans that can be used to proactively protect your virtual servers from hardware problems. A common use case scenario of a hardware problem is provided to illustrate the results from the automaton plan.

This chapter includes the following sections:

- ▶ 11.1, “Environment overview” on page 496
- ▶ 11.2, “Deploying a VM” on page 498
- ▶ 11.3, “Relocating a VM” on page 514
- ▶ 11.4, “Relocating all VMs from a host and saving a relocation plan” on page 519
- ▶ 11.5, “Modifying the Virtual Server resource allocation” on page 523
- ▶ 11.6, “Enabling VMware Distributed Resource Scheduler (DRS)” on page 529
- ▶ 11.7, “Putting a host in maintenance mode” on page 534
- ▶ 11.8, “Topology view” on page 539
- ▶ 11.9, “Automating preventive actions in response to hardware alerts” on page 544

11.1 Environment overview

VMware vCenter Server is the central management component for VMware ESX/ESXi hosts. vCenter is used in almost all VMware environments, and it is required for you to use VMware cluster features.

FSM uses its VMControl plug-in to interact with vCenter. FSM does not replace vCenter. In fact, VMControl uses the robust and virtualization specialized vCenter to run tasks that are targeted at the VMware vSphere infrastructure components. FSM provides an essential collection of the most commonly used tasks by a privileged administrator. Using these tasks, an enterprise administrator with full privileges can manage all platforms in your chassis from the single FSM interface. In addition, junior administrators with lower privileges can perform activities directly on vCenter, if needed.

Additionally, integrating FSM with VMware allows you to correlate events and automate tasks over the physical hardware through your hypervisor, clusters, and virtual servers. It gives you a full picture of your infrastructure end to end. By using VMware, you can operate your system from a single pane of glass from both a hardware and software perspective.

The example initial vSphere 5.0 environment used in this chapter is shown in Figure 11-1.

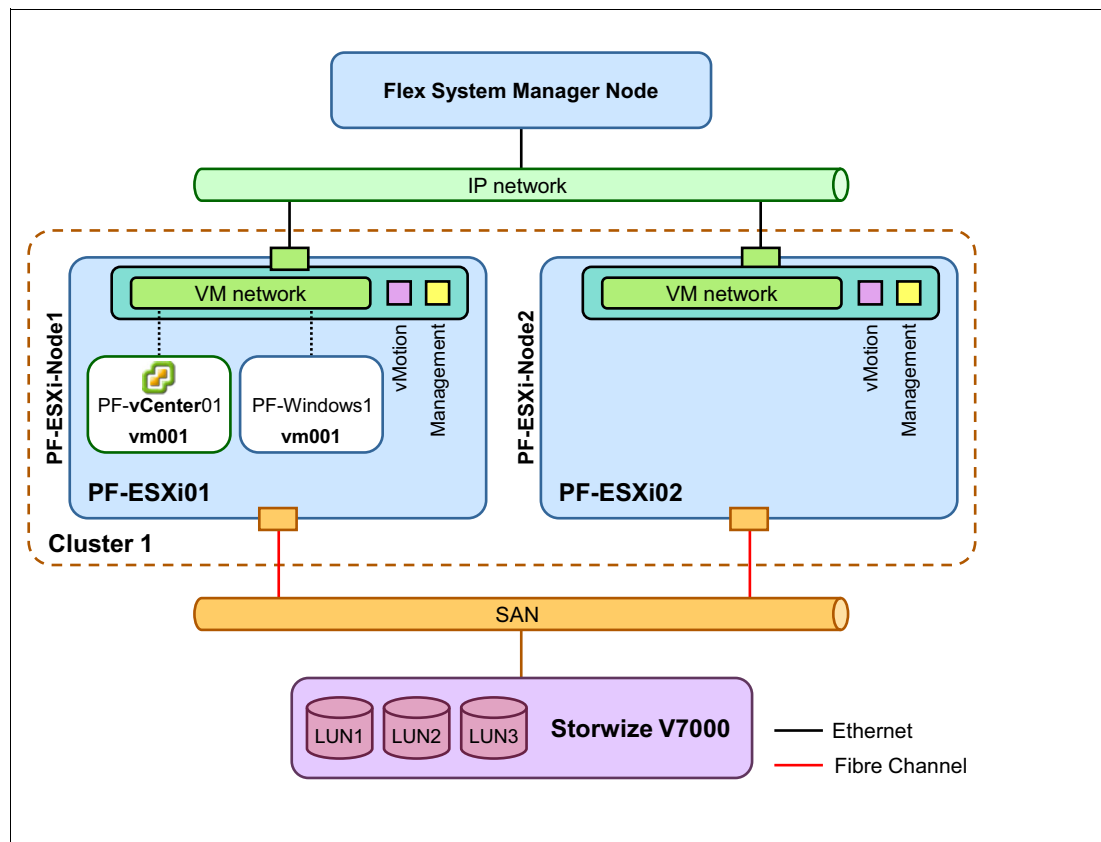


Figure 11-1 VMware environment diagram

Two physical X-Architecture compute nodes are used in a PureFlex Chassis to set up a small vSphere 5 cluster. The cluster has vCenter running in a virtual machine on the first node. Shared SAN storage is provided by Storwize V7000. A simple VM network and a vMotion network are configured for the hosts. FSM eth1 has network connectivity to vCenter.

See Table 11-1 for more information about each component.

Table 11-1 VMware environment components

Component	Description
Hosts	Two ESXi 5.0 hosts. Compute node PF-ESXi-Node1 runs ESXi with host name PF-ESXi01. Compute node PF-ESXi-Node2 runs ESXi with host name PF-ESXi02.
Virtual machines	Two VMware virtual machines Version 8: vm001 and vm002. Both run Microsoft Windows Server 2008 R2 as guest OS. Both are hosted by PF-ESXi-Node1. vm001 has host name PF-vCenter01 and runs the vCenter Server application. vm002 has host name PF-Windows1.
vCenter Server	One VMware vCenter Server 5.0 is running in a virtual machine vm001, which is hosted by PF-ESXi01. It manages both PF-ESXi01 and PF-ESXi02.
Data centers and clusters	One data center, Datacenter1, includes one cluster: Cluster1. Cluster1 has two member hosts: PF-ESXi01 and PF-ESXi02. Cluster1 does not have <i>VMware High Availability (HA)</i> or <i>Distributed Resource Scheduler (DRS)</i> enabled.
Network	Each host has one vSwitch, which has these components: <ul style="list-style-type: none"> ▶ One virtual machine port group VM Network ▶ One vMotion VMkernel port ▶ One management port The vMotion network is configured by using a non-routable network. The two virtual machines, two management ports, and FSM eth1 are in the same network. Each vSwitch has one 10 Gbit uplink.
Storage	Each host is connected through an 8-Gbit interface through the SAN fabric to Storwize V7000. Three 100 GB logical unit numbers (LUNs) are zoned and mapped to both hosts. All three LUNs are formatted with Virtual Machine File System 5 (VMFS5) and are used to store vm001, vm002, and future virtual machine files.

Tip: Generally, run vCenter Server in a virtual machine. This configuration has the following benefits:

- ▶ Easy live migration between physical hosts
- ▶ Easy backup and protection by VMware HA
- ▶ Easy to resize its allocated resources
- ▶ Reduced costs by eliminating the need for a dedicated physical host

For more information about planning for VMware, see 5.2.4, “Planning for VMware virtualization” on page 115.

After you set up the environment, discover the vCenter operating system endpoint by using FSM and request access using the Administrator local user. The Administrator user has full vCenter privileges. All ESXi hosts are discovered automatically after the vCenter compute node is accessed by FSM.

For more information about the discovery of OS, see 6.9.3, “Updating compute node firmware” on page 201.

11.2 Deploying a VM

To deploy a virtual machine (VM), perform these steps:

1. From the VMControl plug-in main window, select the **Virtual Servers/Hosts** tab and click **Virtual Servers and hosts** under Common tasks, as shown in Figure 11-2.



Figure 11-2 VMControl main window

2. The Virtual Servers and Hosts window opens as shown in Figure 11-3.

Select	Name	State	OS Name	OS Type and Version	Access
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008...	OK
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008...	OK
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK
<input type="checkbox"/>	PF-Node1-NIM	Started			OK
<input type="checkbox"/>	PF-Node1-VIOC2	Started			OK
<input type="checkbox"/>	SN101D88B_VIOC1	Stopped			OK
<input type="checkbox"/>	SN101D88B_VIOS1	Started			OK
<input type="checkbox"/>	PF-PowerVM-Node2	Standby			OK

Page 1 of 1 | 1 | Selected: 0 Total: 10 Filtered: 10

Figure 11-3 Virtual Servers and Hosts window

3. Select the first ESXi node and click **Actions** → **System Configuration** → **Create Virtual Server**, as shown in Figure 11-4.




Figure 11-4 Create Virtual Server Actions menu for selected ESXi host

4. Click **Next** on the Welcome window as shown in Figure 11-5.

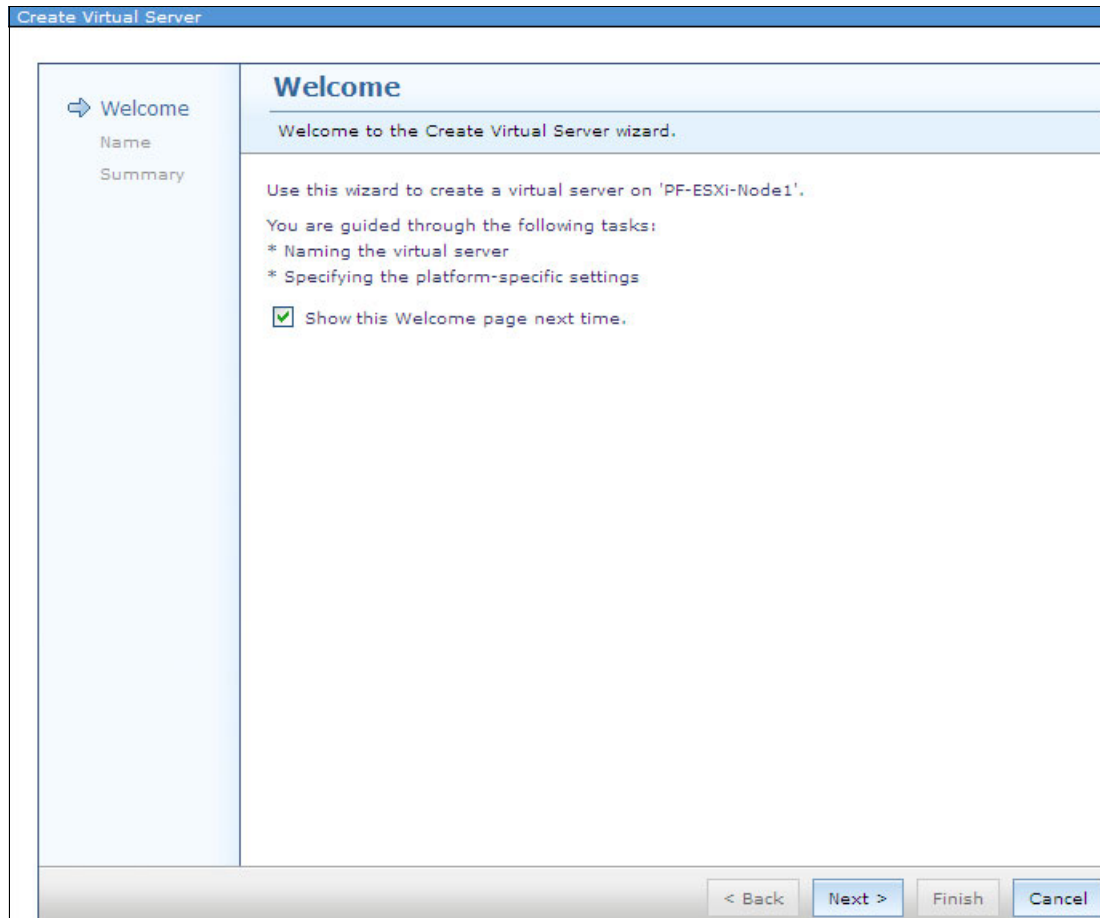


Figure 11-5 Create Virtual Server Welcome window

5. Enter a name for the virtual server that you want to create, as shown in Figure 11-6. In this example, it is vm003. Click **Next**.

✓ Welcome
⇒ Name
Summary

Name

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

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

< Back Next > Finish Cancel

Figure 11-6 Create Virtual Server Name window

6. Select the operating system that you are planning to install on this virtual server. In this example, **Windows Server 2008 R2 (64 bit)** is selected, as shown in Figure 11-7. Click **Next**.




Figure 11-7 Create Virtual Server Source window

7. Specify a number of virtual processors to assign to the virtual server. In this case, enter 2 and click **Next**, as shown in Figure 11-8.

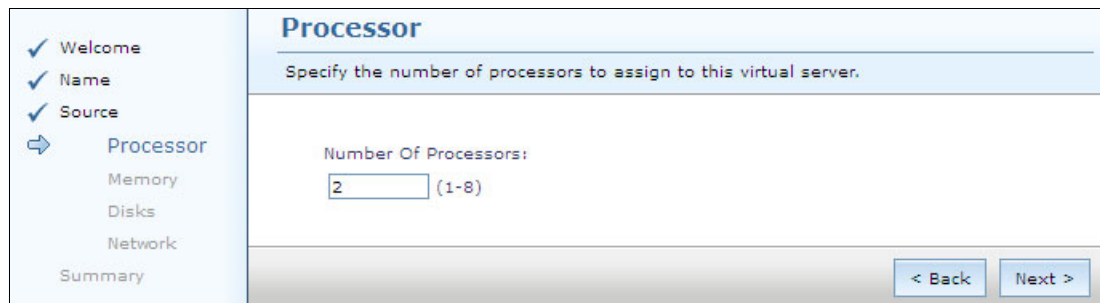


Figure 11-8 Create Virtual Server Processor window

- Enter the amount of memory to assign to this virtual server in MB (see Figure 11-9). In this case, enter 2048 and click **Next**.

Figure 11-9 Create Virtual Server Memory window

- Select a datastore from the Volume label list box where you want to store the virtual machine files. All datastores that are visible by the ESXi host that you selected initially are listed. Make sure that you select a shared datastore to take advantage of cluster features.
- Specify a virtual disk size in GB. The wizard creates one thick lazy zeroed dependent virtual disk with the size that you specify. In this example, enter 40 and click **Next** (Figure 11-10).

Figure 11-10 Create Virtual Server Disks window

- Select a virtual machine port group from the Network Label list box. For this example, select **VM Network** and click **Next** (Figure 11-11). The wizard configures the virtual machine with one virtual network card connected to the port group that you selected.

Figure 11-11 Create Virtual Server Network window

12. Review the Summary window and click **Finish** (Figure 11-12).




Figure 11-12 Create Virtual Server Summary window

13. Click **OK** in the Launch Job window to start the virtual server creation immediately as shown in Figure 11-13.




Figure 11-13 Create Virtual Server Launch Job window

14. Click **Display Properties** in the Create Virtual Server window to see the job status (see Figure 11-14).

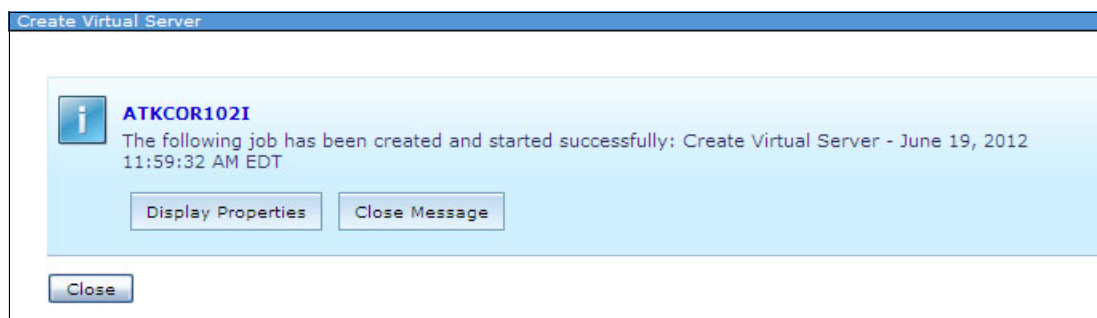


Figure 11-14 Create Virtual Server job message box

15. Ensure that the Create Virtual Server job completed successfully (Figure 11-15), and close the **Active and Scheduled Jobs** tab.

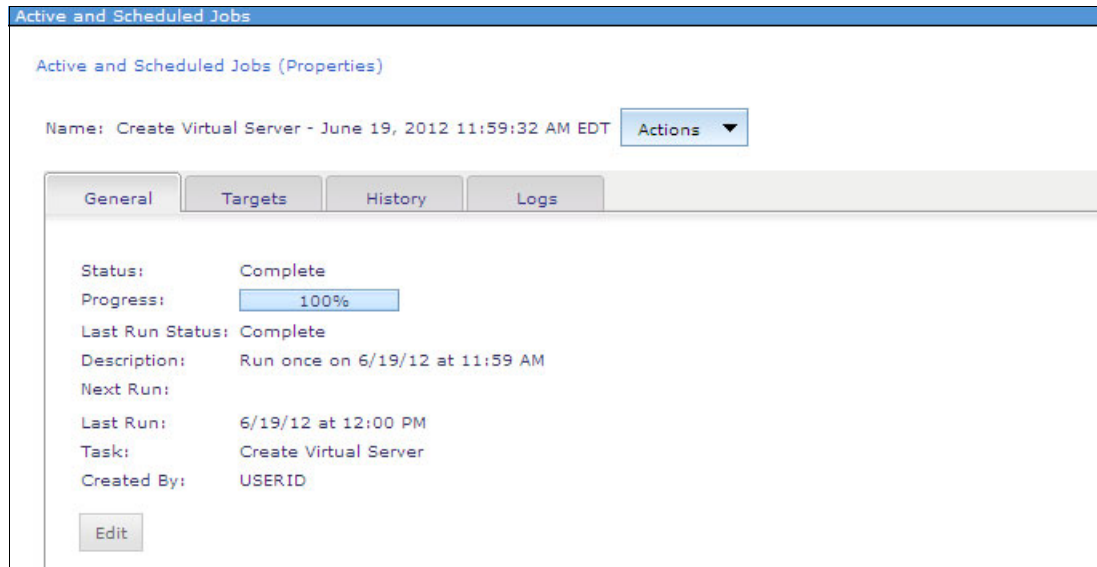


Figure 11-15 Create Virtual Server job details window

16. Return to the Virtual Servers and Hosts window to see the newly created virtual server vm003. It is in the Stopped state, as shown in Figure 11-16. The virtual machine was created on the ESXi host PF-ESXi-Node1, which is managed by the vCenter server PF-vCenter01. Click the **Information** link on the PF-ESXi-Node1 row to open the list of events for that ESXi server.

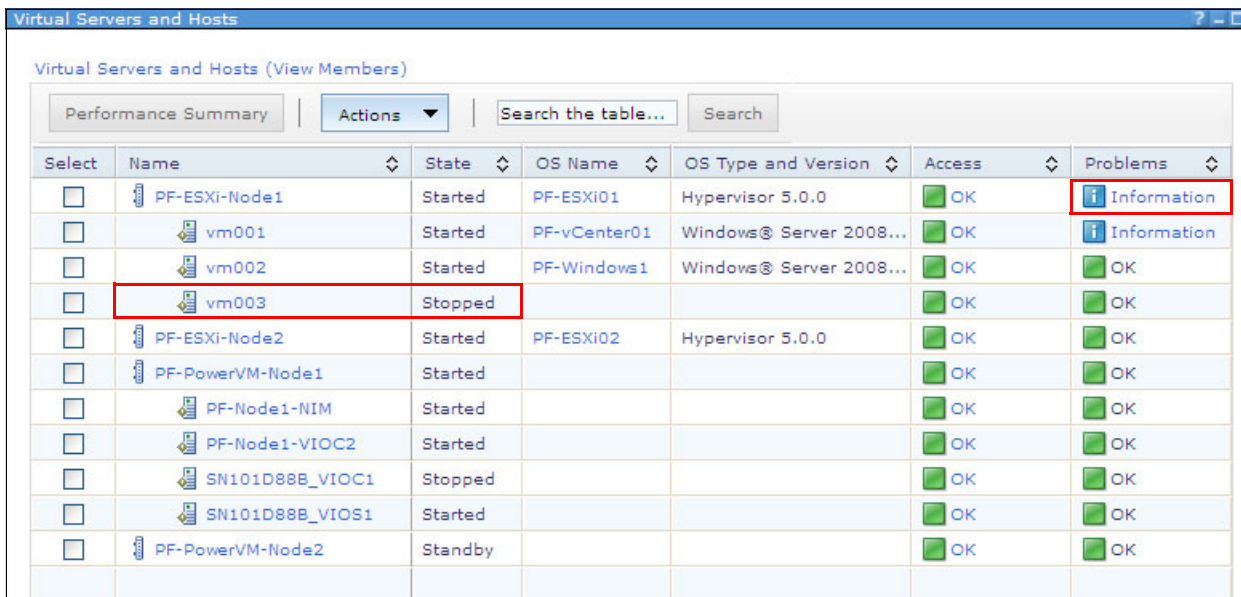


Figure 11-16 Virtual Servers and Hosts window

17. An informational event about the virtual server creation is displayed as shown in Figure 11-17. Similar informational events are also displayed under PF-vCenter01 because the virtual server was created on a host that is managed by PF-vCenter01.

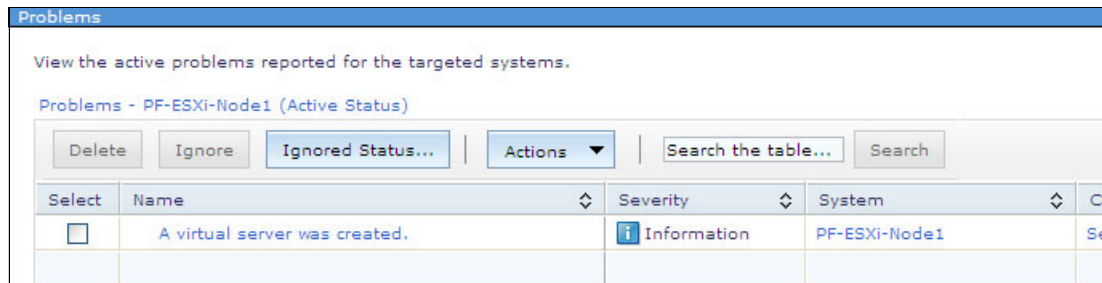


Figure 11-17 Virtual server creation informational event in the Problems window

18. If you need to delete the informational event, select it and click **Actions** → **Delete**, as shown in Figure 11-18.

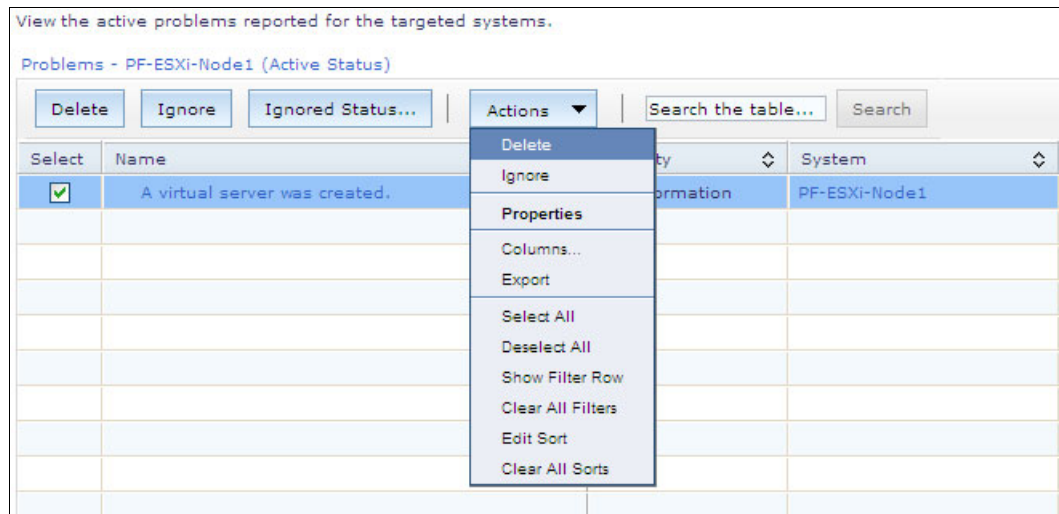


Figure 11-18 Delete menu item for selected event in the Problems window

19. For this example, connect to your vCenter server and observe its state (see Figure 11-19). The vm003 virtual machine was created. The virtual machine creation was started by Administrator, which is the user that FSM used to discover and authenticate to vCenter.




Figure 11-19 vSphere Client window that shows the newly created virtual machine

20. Return to the Virtual Servers and Hosts view to power on the newly created virtual server. Select **vm003**, and click **Actions** → **Power On/Off** → **Power On**, as shown in Figure 11-20.

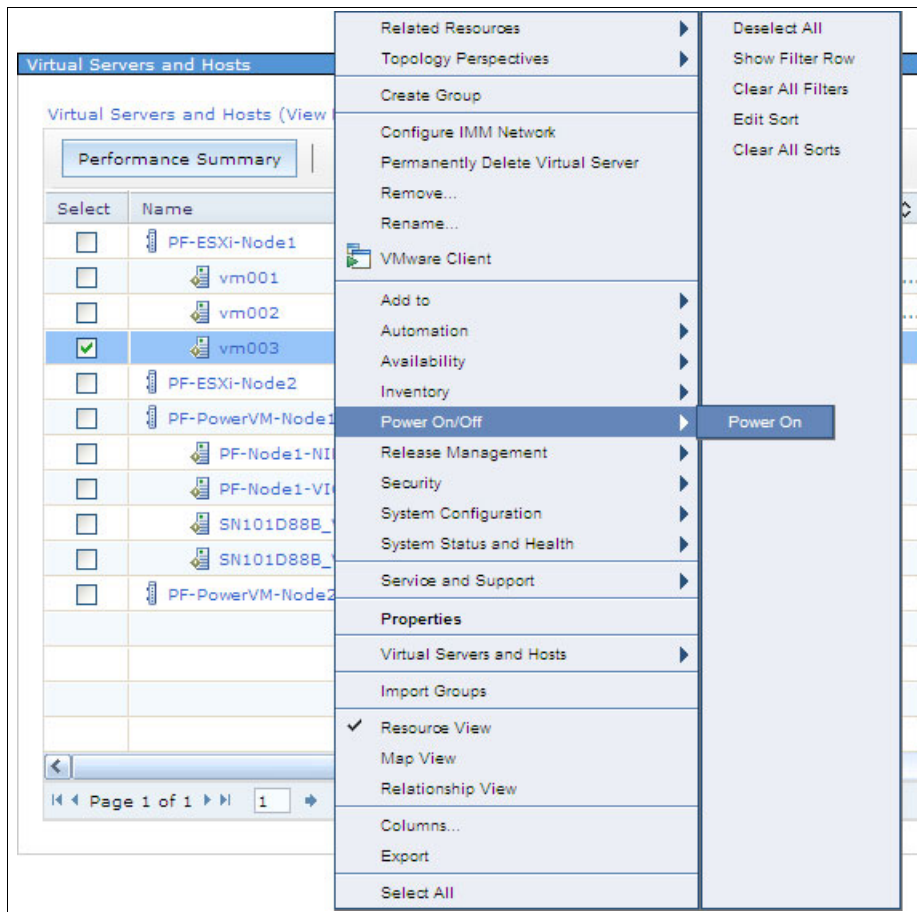


Figure 11-20 Power On menu item for selected virtual server

21. Click **OK** to start the Power On job immediately as shown in Figure 11-21.




Figure 11-21 Power On Launch Job window

22.The virtual server State changes to Started in the Virtual Servers and Hosts view, as shown in Figure 11-22.

The screenshot shows a notification at the top: "ATKCOR102I The following job has been created and started successfully: Power On - June 19, 2012 12:27:56 PM EDT". Below this is a table titled "Virtual Servers and Hosts (View Members)". The table has columns for Select, Name, State, OS Name, OS Type and Version, Access, and Prob. The row for "vm003" is selected, and its "State" is "Started", which is highlighted with a red box. Other rows include "PF-ESXi-Node1", "vm001", "vm002", "PF-ESXi-Node2", "PF-PowerVM-Node1", "PF-Node1-NIM", "PF-Node1-VIOC2", "SN101D88B_VIOC1", "SN101D88B_VIOS1", and "PF-PowerVM-Node2".

Select	Name	State	OS Name	OS Type and Version	Access	Prob
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	C
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008...	OK	I
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008...	OK	C
<input checked="" type="checkbox"/>	vm003	Started			OK	I
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	C
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK	C
<input type="checkbox"/>	PF-Node1-NIM	Started			OK	C
<input type="checkbox"/>	PF-Node1-VIOC2	Started			OK	C
<input type="checkbox"/>	SN101D88B_VIOC1	Stopped			OK	C
<input type="checkbox"/>	SN101D88B_VIOS1	Started			OK	C
<input type="checkbox"/>	PF-PowerVM-Node2	Standby			OK	C

Page 1 of 1 | 1 | Selected: 1 Total: 11 Filtered: 11

Figure 11-22 Virtual Servers and Hosts window that shows the started virtual server

In Figure 11-23, you can see the powered on virtual machine in vCenter.




Figure 11-23 vSphere Client window that shows a powered on virtual machine

FSM creates a virtual server without any guest OS installed. Normally, this task is performed by an enterprise administrator who manages the entire chassis by using FSM and has full privileges to create a virtual server. At this point, a junior administrator with virtual machine user privileges in vCenter can connect to the Virtual Machine console through a vSphere client. The junior administrator can then proceed with the guest OS installation.

11.3 Relocating a VM

To relocate the newly created virtual server to the second host while the virtual server is running, perform these steps:

1. Select the virtual server that you want to relocate to another host (in this example, **vm003**), and click **Actions** → **Availability** → **Relocate**, as shown in Figure 11-24.

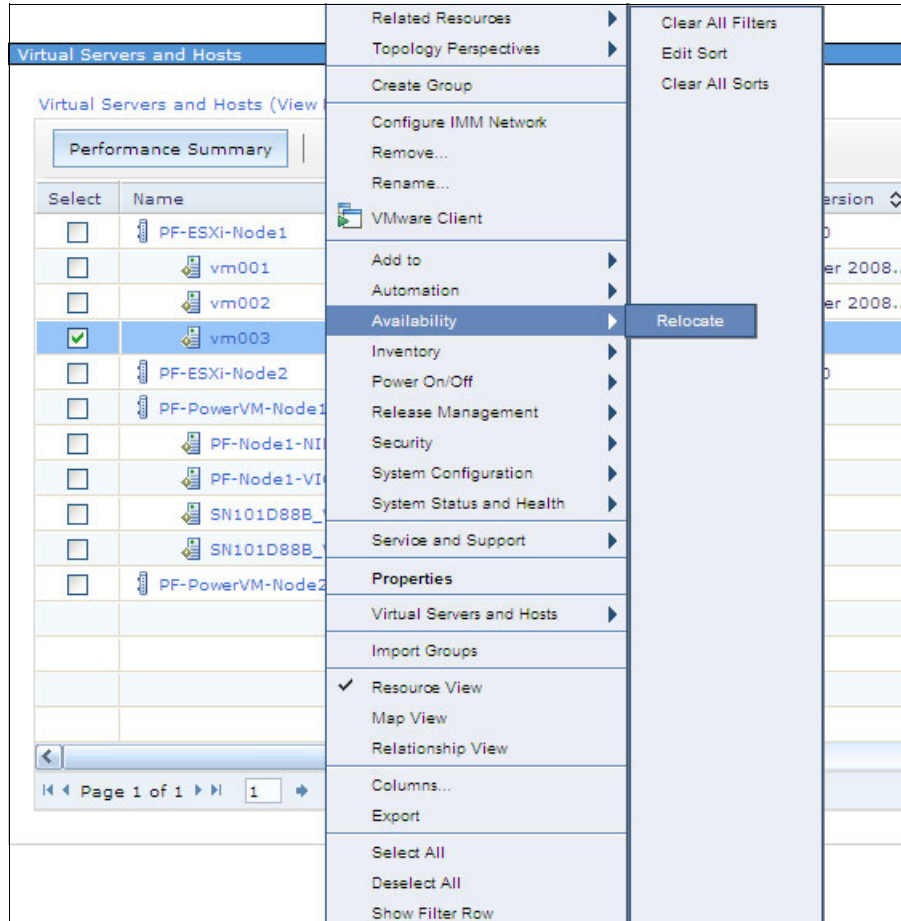


Figure 11-24 Relocate menu item for selected virtual server

2. Verify the virtual machine name and click **Next**, as shown in Figure 11-25.

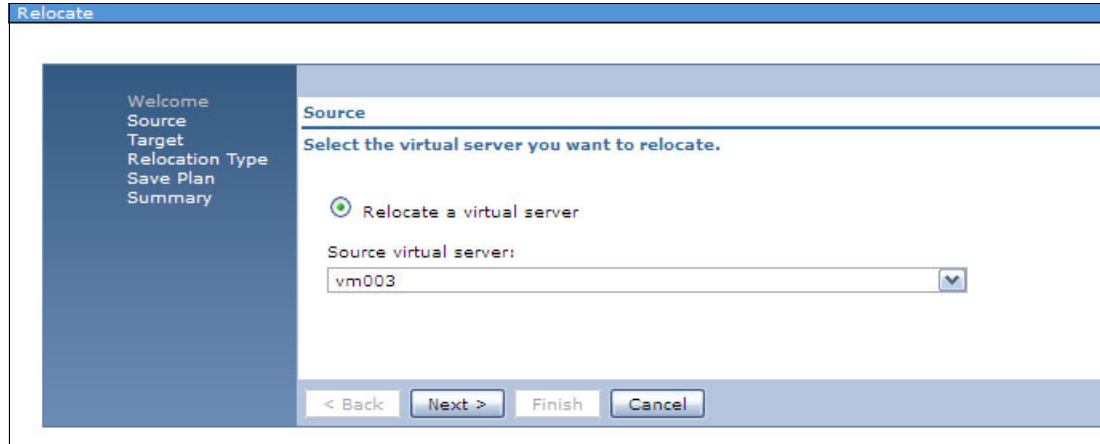


Figure 11-25 Relocate Welcome window

3. Select the target host for the virtual machine. You can also select “Relocate by CPU utilization” if you want the virtual server to be moved to the host with the lowest processor utilization. Select **PF-ESXi-Node2**, as shown in Figure 11-26.

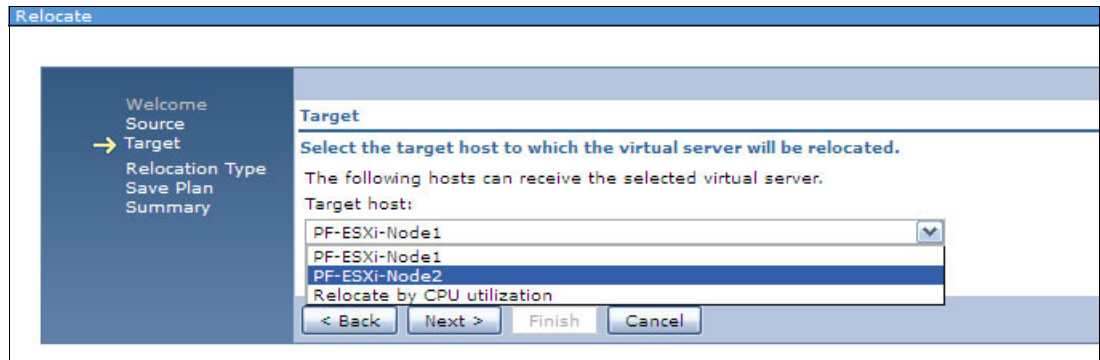


Figure 11-26 Relocate Target window

4. You can save the plan for relocation to run it again or use it later, if needed. Select **Relocate only** and click **Next** as shown in Figure 11-27.

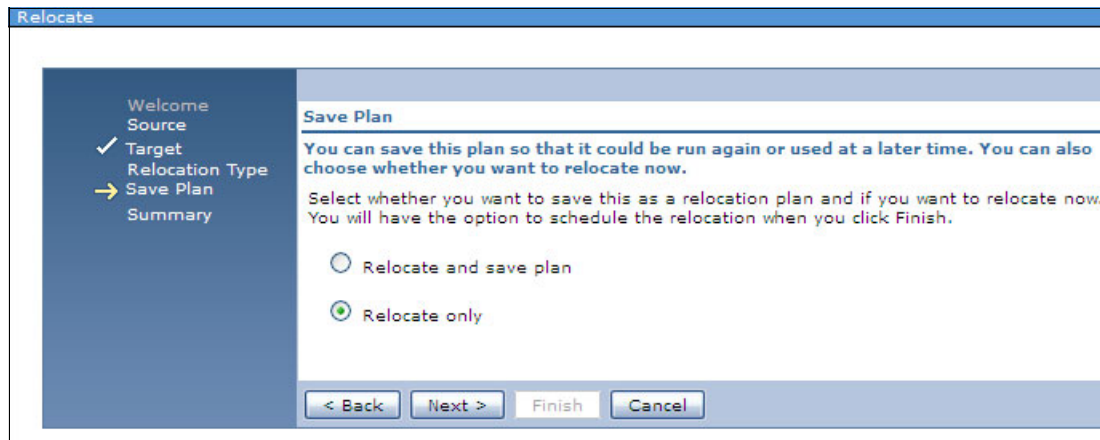


Figure 11-27 Relocate Save Plan window

5. Verify the relocation Summary and click **Finish** (see Figure 11-28).

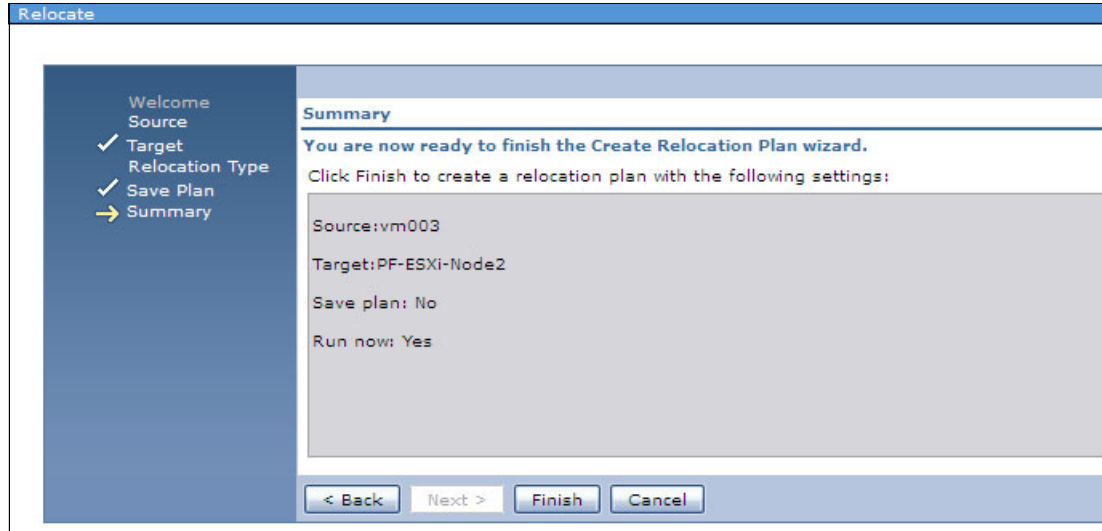


Figure 11-28 Relocate Summary window

6. The virtual server Status in the Virtual Servers and Hosts window changes to Relocating during the relocation from PF-ESXi01 to PF-ESXi02, as shown in Figure 11-29.

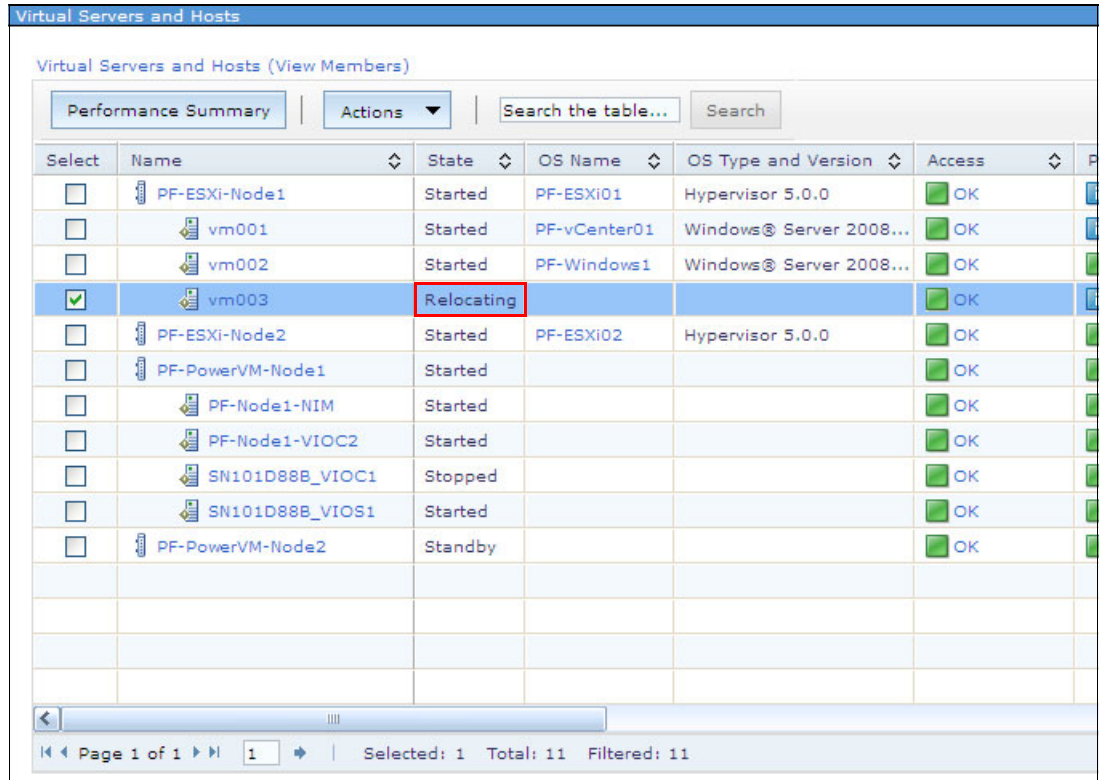


Figure 11-29 Virtual Servers and Hosts window that shows a virtual server in the Relocating state

FSM triggers a command for vCenter to run a vMotion migration of vm003 from FP-ESXi01 to FP-ESXi02. The migration completes successfully and vm003 is now running on FP-ESXi02, as shown in Figure 11-30.




Figure 11-30 vSphere Client window that shows the migrated virtual machine

- Return to the Virtual Servers and Hosts window in FSM. The virtual machine vm003 is now listed under FP-ESXI02 and it is in the Started state, as shown in Figure 11-31.

Virtual Servers and Hosts (View Members)

Performance Summary | Actions | Search the table... Search

Select	Name	State	OS Name	OS Type and Version	Access	Problems
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008...	OK	Information
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008...	OK	OK
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm003	Started			OK	Information
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK	OK
<input type="checkbox"/>	PF-Node1-NIM	Stopped			OK	Information
<input type="checkbox"/>	PF-Node1-VIOC2	Started			OK	OK
<input type="checkbox"/>	SN101D88B_VIOC1	Stopped			OK	OK
<input type="checkbox"/>	SN101D88B_VIOS1	Started			OK	OK
<input type="checkbox"/>	PF-PowerVM-Node2	Standby			OK	OK

Page 1 of 1 | 1 | Selected: 0 Total: 11 Filtered: 11

Figure 11-31 Virtual Servers and Hosts window that shows the migrated virtual server

11.4 Relocating all VMs from a host and saving a relocation plan

In certain cases, you might need to relocate all virtual servers away from a specific host to perform service tasks. For this example, move vm003 back to PF-ESXi-Node1. To relocate all VMs from a host and save a relocation plan, which can be run later or used in an automation plan, perform these steps:

1. Right-click the host (**PF-ESXi-Node1**) and select **Availability** → **Relocate Virtual Servers**, as shown in Figure 11-32.

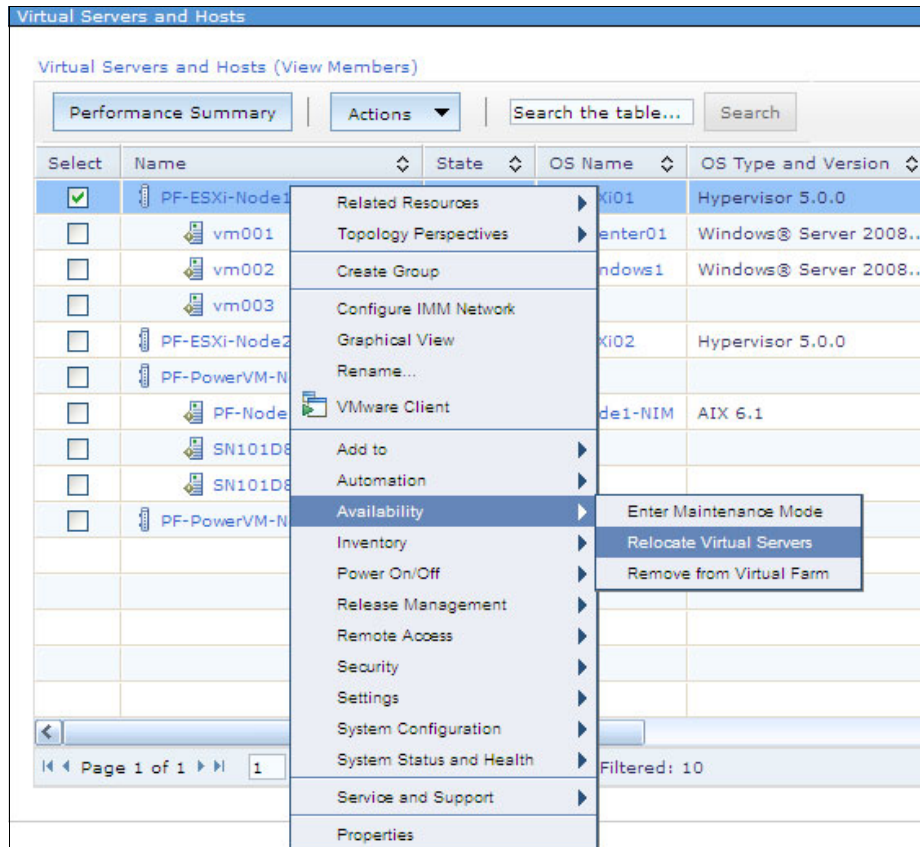


Figure 11-32 Relocate Virtual Servers menu item for the selected host

2. Verify the source host and click **Next**, as shown in Figure 11-33. If necessary, you can select “Put host in maintenance mode after all virtual servers are relocated”.




Figure 11-33 Relocate Source window

3. You can select a specific target host or choose “Relocate by CPU utilization”. Select **PF-ESXi-Node2** and click **Next**, as shown in Figure 11-34.




Figure 11-34 Relocate Target window

4. Select **Relocate and save plan** and provide a descriptive relocation plan name, as shown in Figure 11-35.

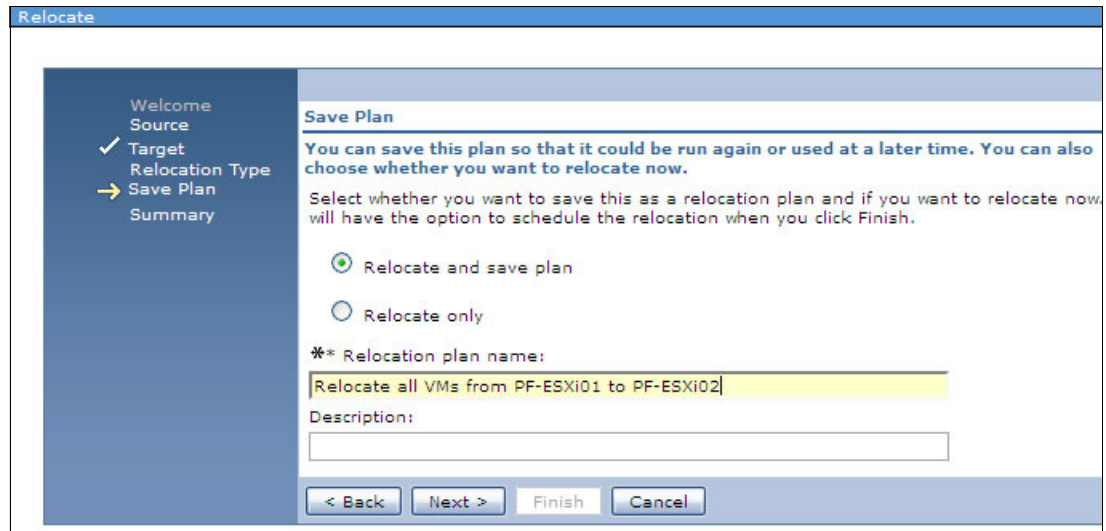


Figure 11-35 Relocate Save Plan window

5. Review the Summary window and click **Finish**, see Figure 11-36.

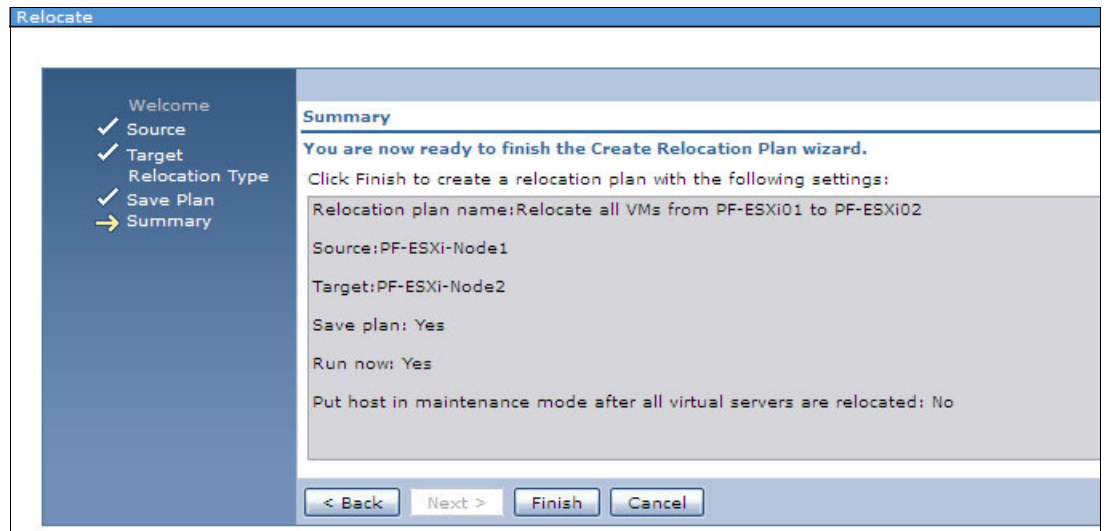


Figure 11-36 Relocate Summary window

- Click **OK** to start the relocation job immediately and observe the Virtual Servers and Hosts window to ensure that all virtual machines from PF-ESXi01 relocate to PF-ESXi02, as shown in Figure 11-37.




Figure 11-37 Virtual Servers and Hosts window

- In the IBM Flex System Manager web interface navigation area, expand **Availability** and click **Relocation Plans for Farms**, as shown in Figure 11-38. In this window, you can view, run, and manage all relocation plans for farms.




Figure 11-38 Relocation Plans for Farms window

11.5 Modifying the Virtual Server resource allocation

Changing the resource allocation for a virtual server is another task that is often the responsibility of a full administrator. To modify the memory allocation of a virtual server, perform these steps:

1. Right-click the virtual server **vm002** and select **Power On/Off** → **Shut down and power off**, as shown in Figure 11-39. This process starts a graceful OS shutdown before powering off the virtual server.

Important: Ensure that you always have an up-to-date version of VMware tools installed in the guest OS of your VMware virtual machines. Graceful OS shutdown is just one of the many features that make VMware tools extremely useful.

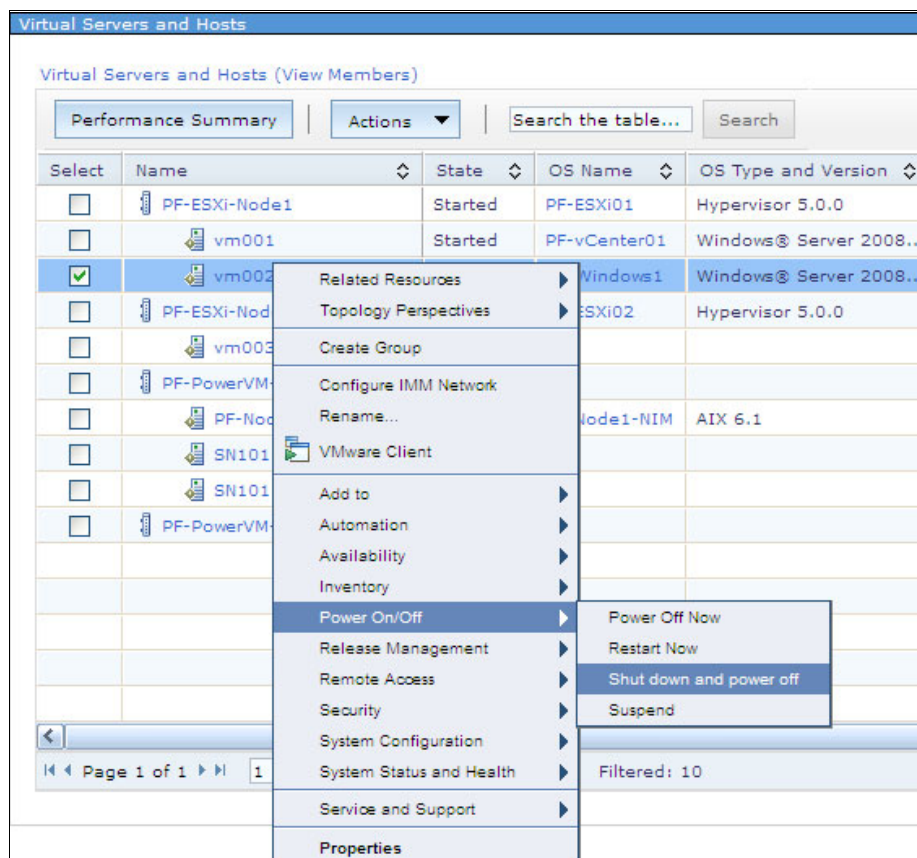


Figure 11-39 Shut down and power off menu item in the Virtual Servers and Hosts window

2. Click **OK** in the Launch Job window to run the job immediately as shown in Figure 11-40.

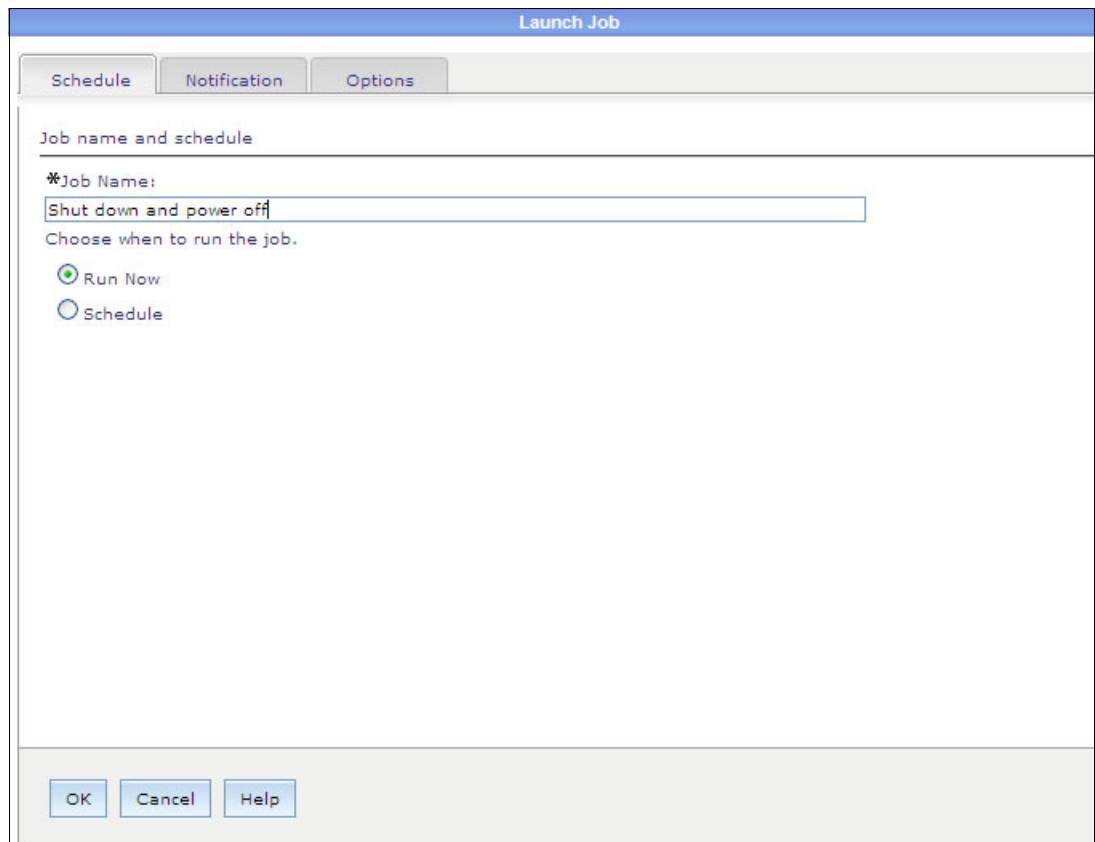


Figure 11-40 Shut down and power off Launch Job window

3. Check the state in vCenter, as shown in Figure 11-41. The “Initiate guest OS shutdown” task that was triggered by Administrator completed successfully, and the virtual machine vm002 is in the Powered Off state.




Figure 11-41 vSphere Client window that shows the Powered Off virtual machine

- Return to the Virtual Servers and Hosts window, right-click the powered-off **vm002**, and select **System Configuration** → **Edit Virtual Server**, as shown in Figure 11-42.

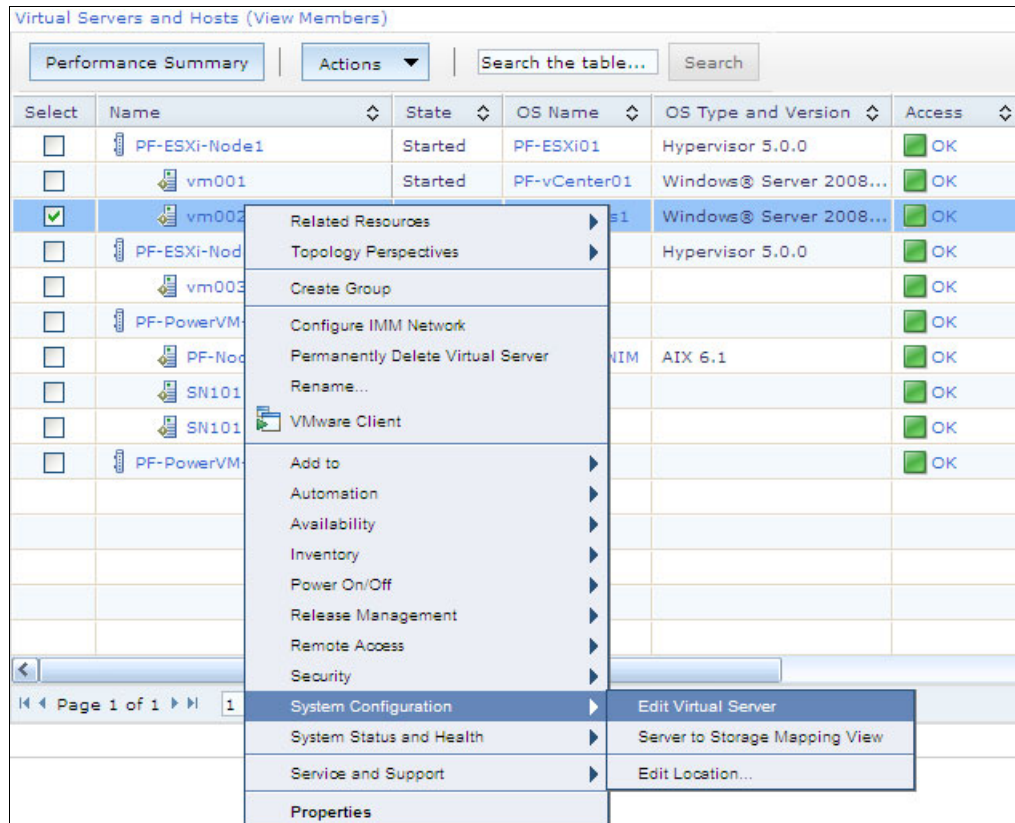


Figure 11-42 Edit Virtual Server menu item in the Virtual Servers and Hosts window

- The Edit Virtual Server window opens. Click the **Memory** tab, as shown in Figure 11-43.




Figure 11-43 Edit Virtual Server window

6. Observe the current memory that is assigned to the virtual server, as shown in Figure 11-44.

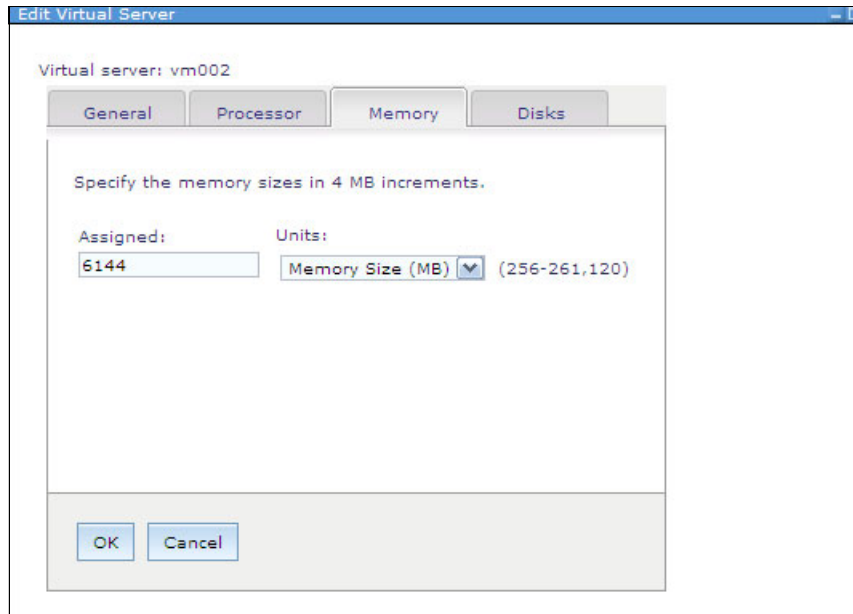


Figure 11-44 Memory tab of the Edit Virtual Server window

7. Change the assigned memory value from 6144 to 8192 to increase the virtual server memory to 8 GB (see Figure 11-45). Click **OK** to apply the new configuration.




Figure 11-45 Updated memory value in the Edit Virtual Server window

8. Right-click **vm002** and power it on by selecting **Power On/Off** → **Power On**. See Figure 11-46.

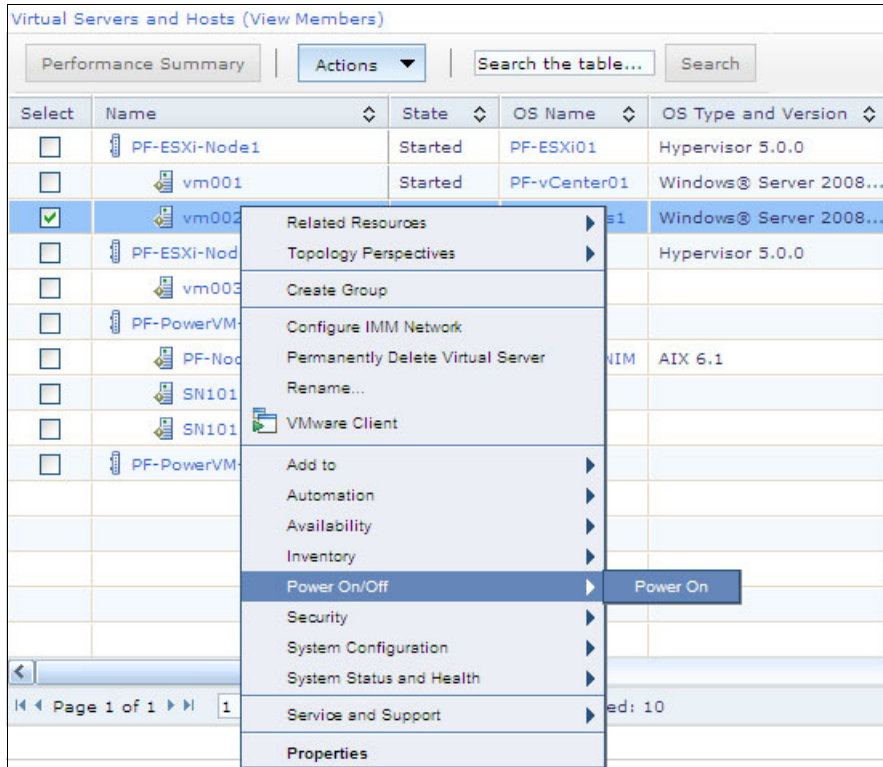


Figure 11-46 Power On menu item in the Virtual Servers and Hosts window

- The virtual server is now running with 8 GB of allocated memory. Observe the tasks that FSM sent to vCenter on Figure 11-47.

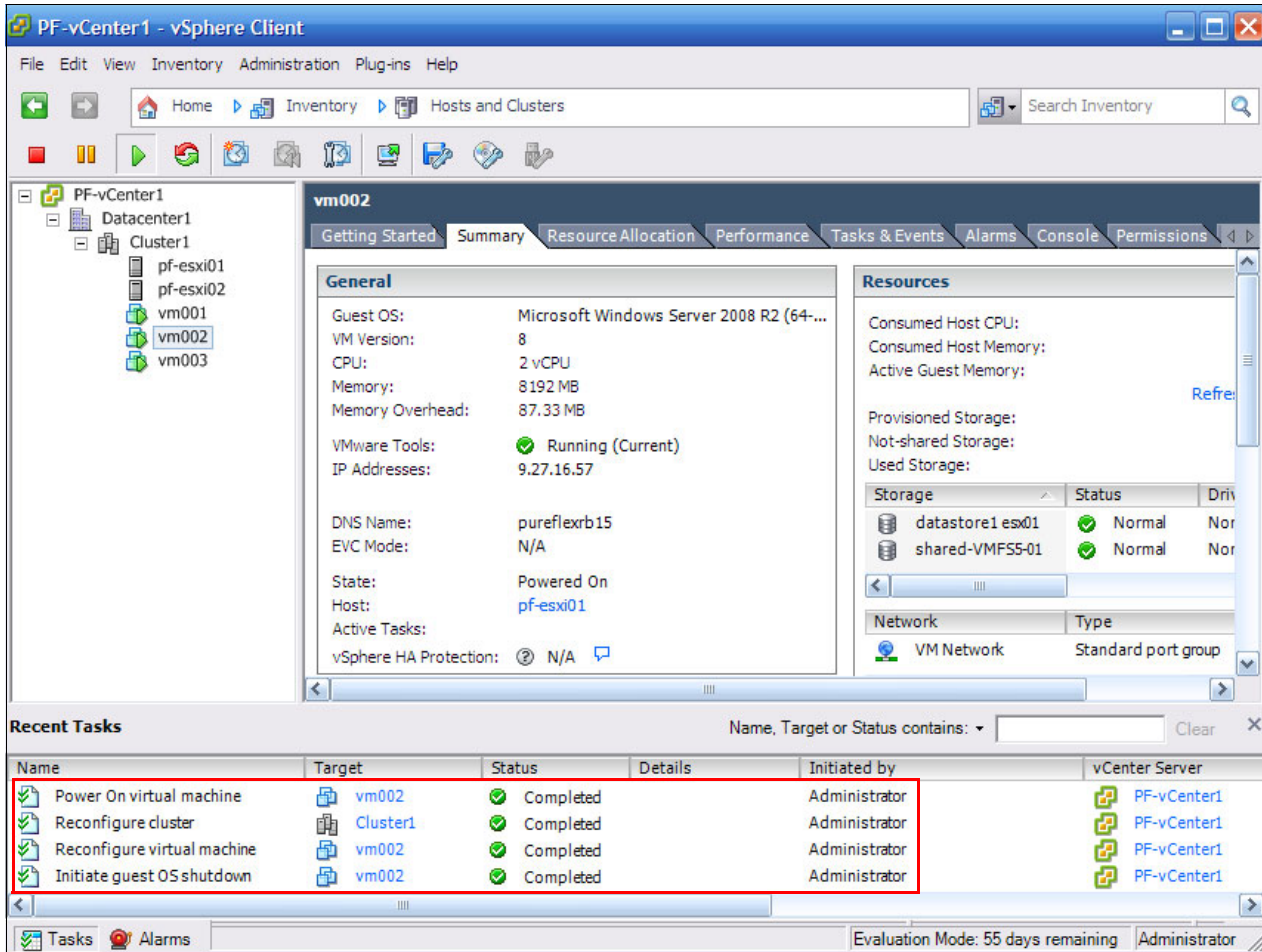


Figure 11-47 vSphere Client window that shows tasks for edited virtual machine

11.6 Enabling VMware Distributed Resource Scheduler (DRS)

VMware Distributed Resource Scheduler (DRS) is a cluster feature that can perform dynamic load balancing of compute resources (processor and memory) across physical hosts that are members of the cluster. When configured in *Fully Automated mode*, DRS uses VMware vMotion to run live migration of virtual machines (VMs) whenever needed. DRS continuously monitors the processor and memory resource usage for all cluster physical hosts and their VMs. DRS evaluates these metrics and ensures an optimal VM placement to achieve a relatively even load on all cluster physical hosts.

Explanation: For most vSphere environments, configure DRS in Fully Automated mode. To use DRS for load balancing, you must have a vMotion network that is configured in your cluster, and your virtual machines must meet vMotion requirements. Ensure that you have the correct vSphere license to use DRS. For more information about DRS, see the *vSphere Resource Management Guide* at this website:

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcen-ter-server-501-resource-management-guide.pdf>

2. Right-click **Cluster1** and select **Availability** → **Edit Virtual Farm**, as shown in Figure 11-49.

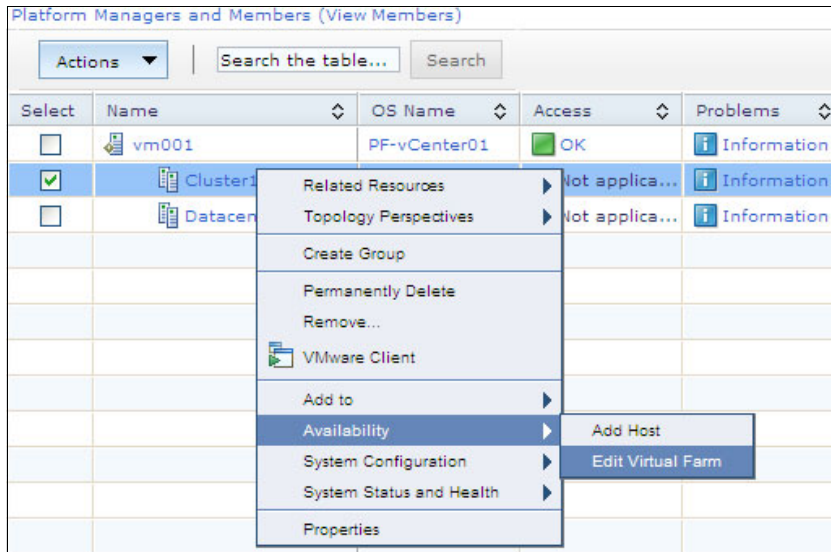


Figure 11-49 Edit Virtual Farm menu item in the Platform Managers and Members window

3. Click **Next** in the Welcome window, as shown in Figure 11-50.



Figure 11-50 Edit Virtual Farm Welcome window

4. Verify the cluster name and click **Next**, as shown in Figure 11-51.

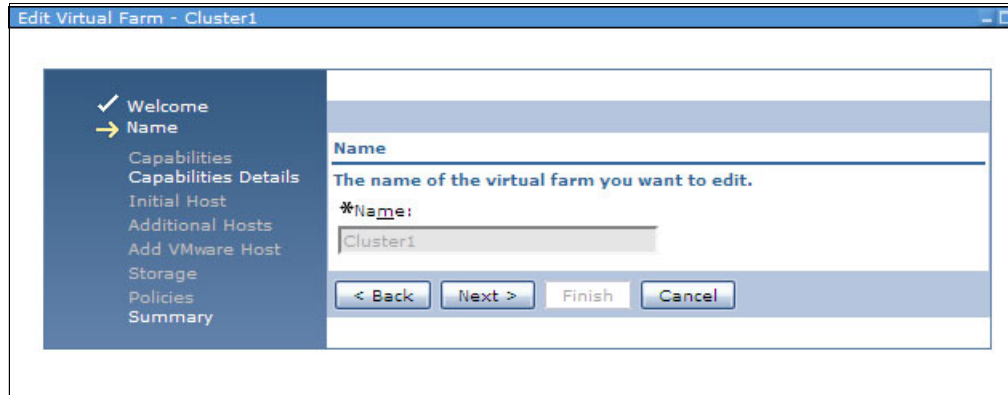


Figure 11-51 Edit Virtual Farm Name window

5. Select **VMware Distributed Resource Scheduler (DRS)** and leave VMotion rate as Normal, as shown in Figure 11-52. Click **Next**.

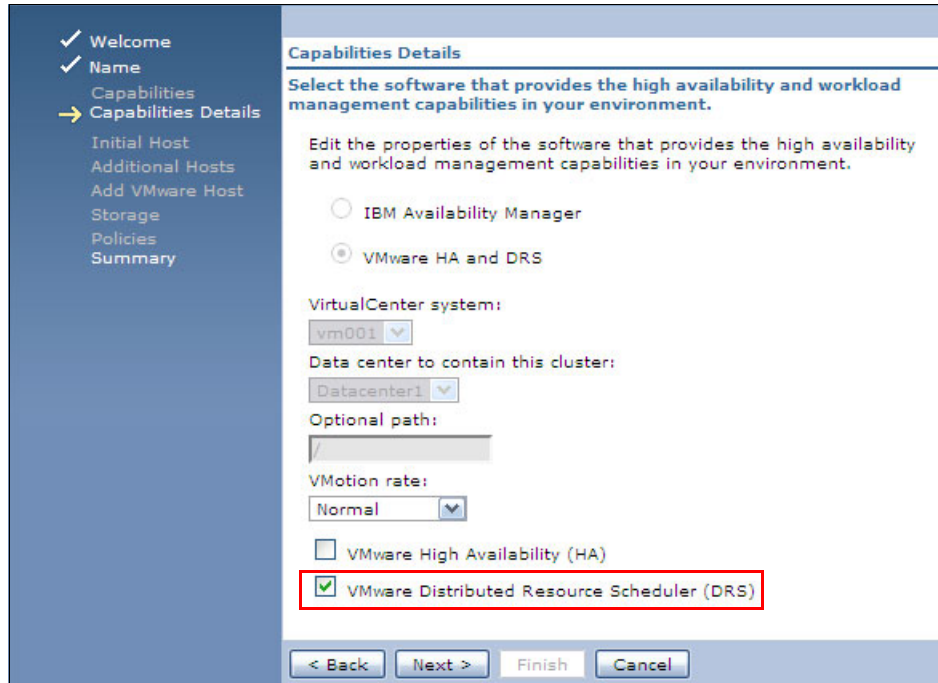


Figure 11-52 Edit Virtual Farm Capabilities Details window

6. Review the Summary window and click **Finish**, as shown in Figure 11-53.




Figure 11-53 Edit Virtual Farm Summary window

FSM sends a command to vCenter to enable DRS on Cluster1 with a normal migration threshold and in Fully Automated mode. The status of vCenter is shown in Figure 11-54.

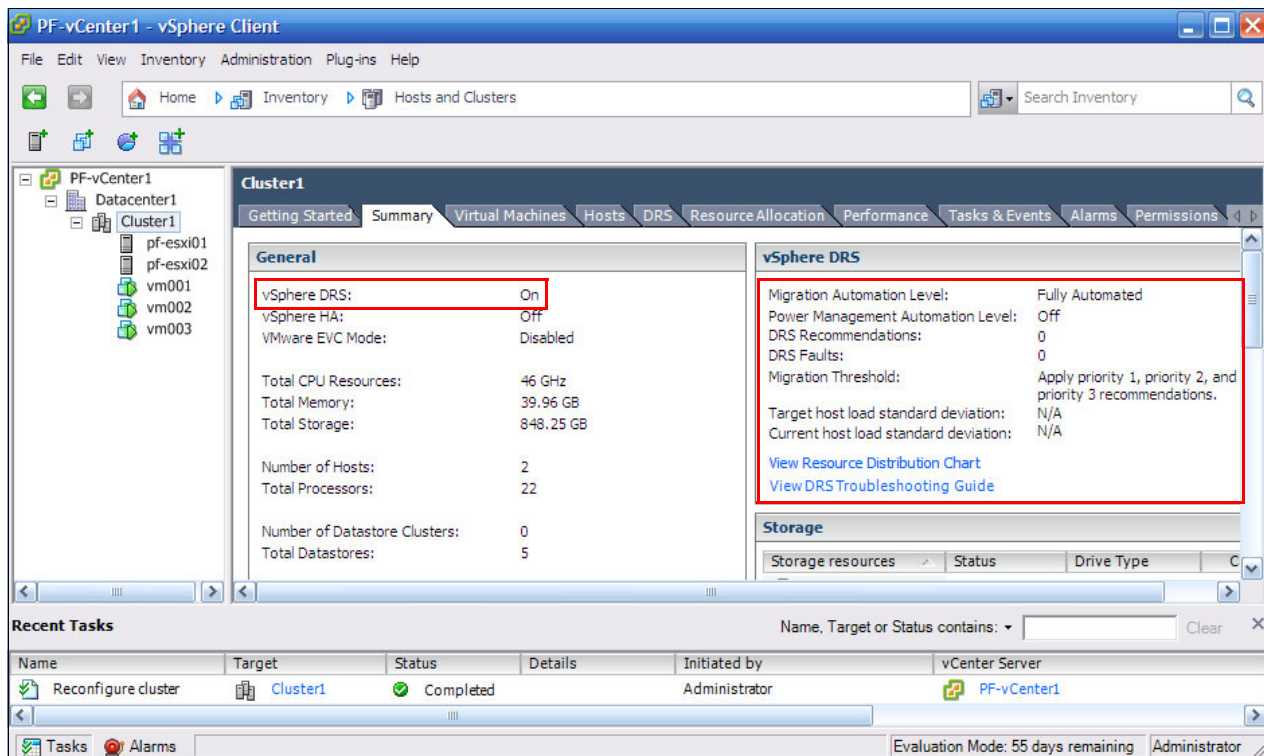


Figure 11-54 vSphere Client window that shows a DRS-enabled cluster

11.7 Putting a host in maintenance mode

You can place a host in maintenance mode to perform service tasks on it. A vSphere host in maintenance mode cannot have any virtual machines in the powered-on state. If a host entering maintenance mode has powered on virtual machines and is a member of a fully automated DRS cluster, DRS automatically migrates all running virtual machines to other hosts in the cluster. DRS then places the host in maintenance mode.

In the previous section, you enabled DRS in Fully Automated mode for Cluster1. Now, place PF-ESXi-Node2 in maintenance mode. PF-ESXi-Node2 is running one virtual server: vm003. To place a VMware host in maintenance mode, perform these steps:

1. Open the Virtual Servers and Hosts window and select **PF-ESXi-Node2**. Select **Actions** → **Availability** → **Enter Maintenance Mode**, as shown in Figure 11-55.

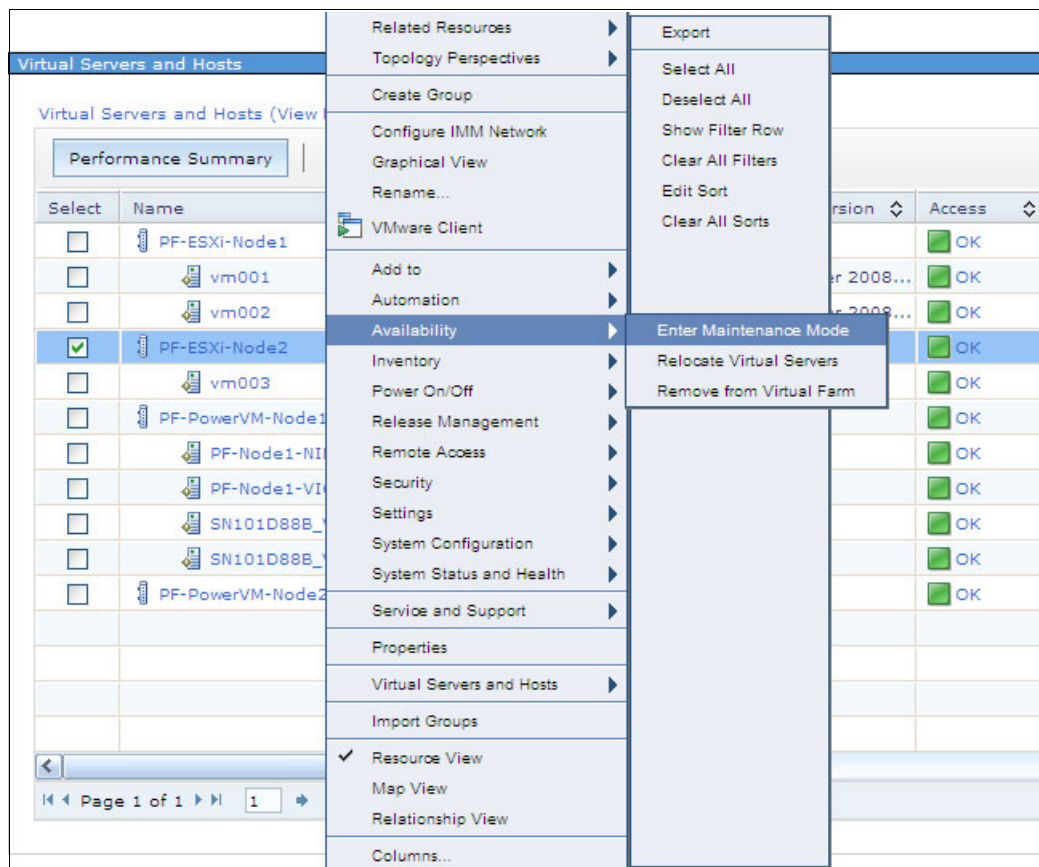


Figure 11-55 Enter Maintenance Mode menu item in Virtual Servers and Hosts window

FSM places the ESXi host in maintenance mode by sending a command to vCenter. This mode is also evident in vCenter (Figure 11-56). The virtual machine vm003 originally was on PF-ESXi02, but DRS migrated it to PF-ESXi01 before placing PF-ESXi02 in maintenance mode.

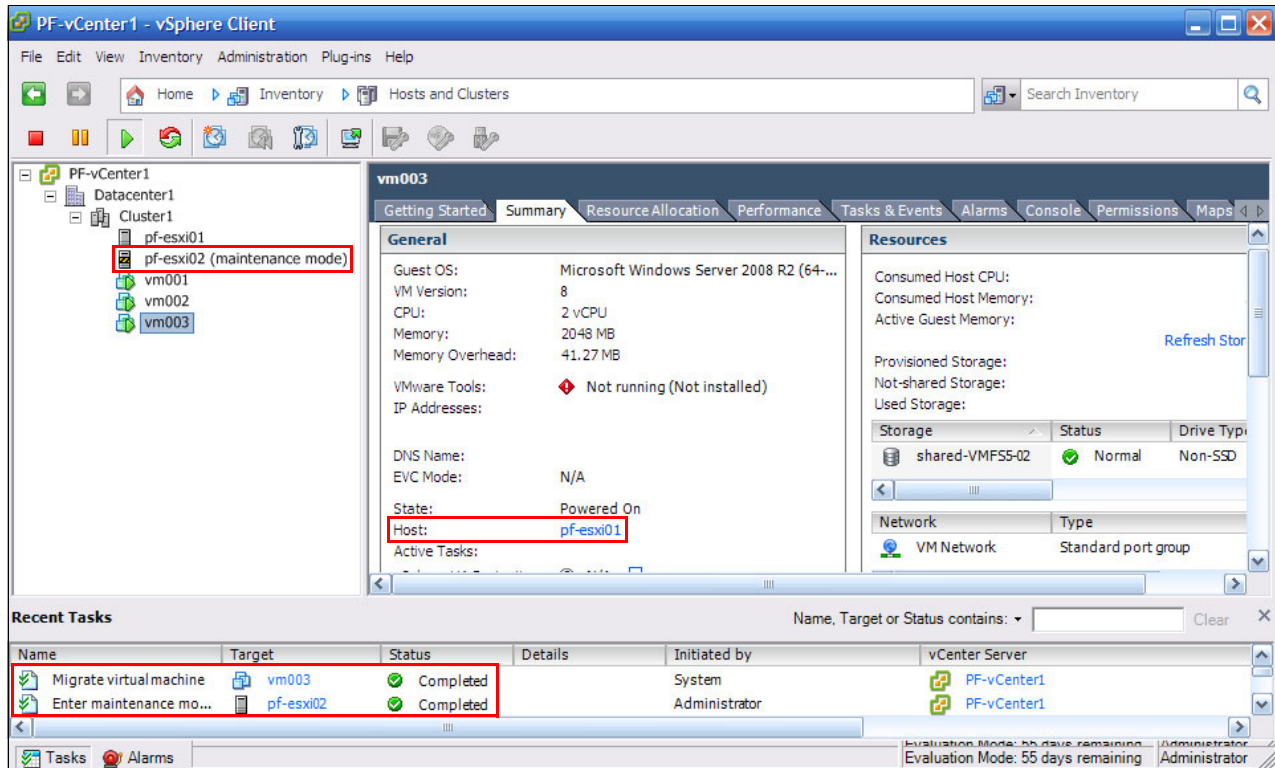


Figure 11-56 vSphere Client window that shows the host in maintenance mode

- Return to the Virtual Servers and Hosts window and observe that all virtual servers are now running on PF-ESXi01 (Figure 11-57).

Select	Name	State	OS Name	OS Type and Version	Access	Problems
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	Information
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008...	OK	Information
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008...	OK	OK
<input type="checkbox"/>	vm003	Started			OK	OK
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			Partial a...	OK
<input type="checkbox"/>	PF-Node1-NIM	Started			OK	OK
<input type="checkbox"/>	SN101D88B_VIOC1	Started			OK	OK
<input type="checkbox"/>	SN101D88B_VIOS1	Started			OK	OK
<input type="checkbox"/>	PF-PowerVM-Node2	Standby			Partial a...	OK

Page 1 of 1 | 1 | Selected: 0 Total: 10 Filtered: 10

Figure 11-57 Virtual Servers and Hosts window

3. Select **PF-ESXi-Node2** and click **Actions** → **Availability** → **Exit Maintenance Mode**, as shown in Figure 11-58.




Figure 11-58 Exit Maintenance Mode menu item in Virtual Servers and Hosts window

4. Click **OK** to start the job immediately, as shown in Figure 11-59.




Figure 11-59 Exit Maintenance Mode Launch Job window

As shown in Figure 11-60, the Exit maintenance mode task that was triggered by FSM completed successfully.

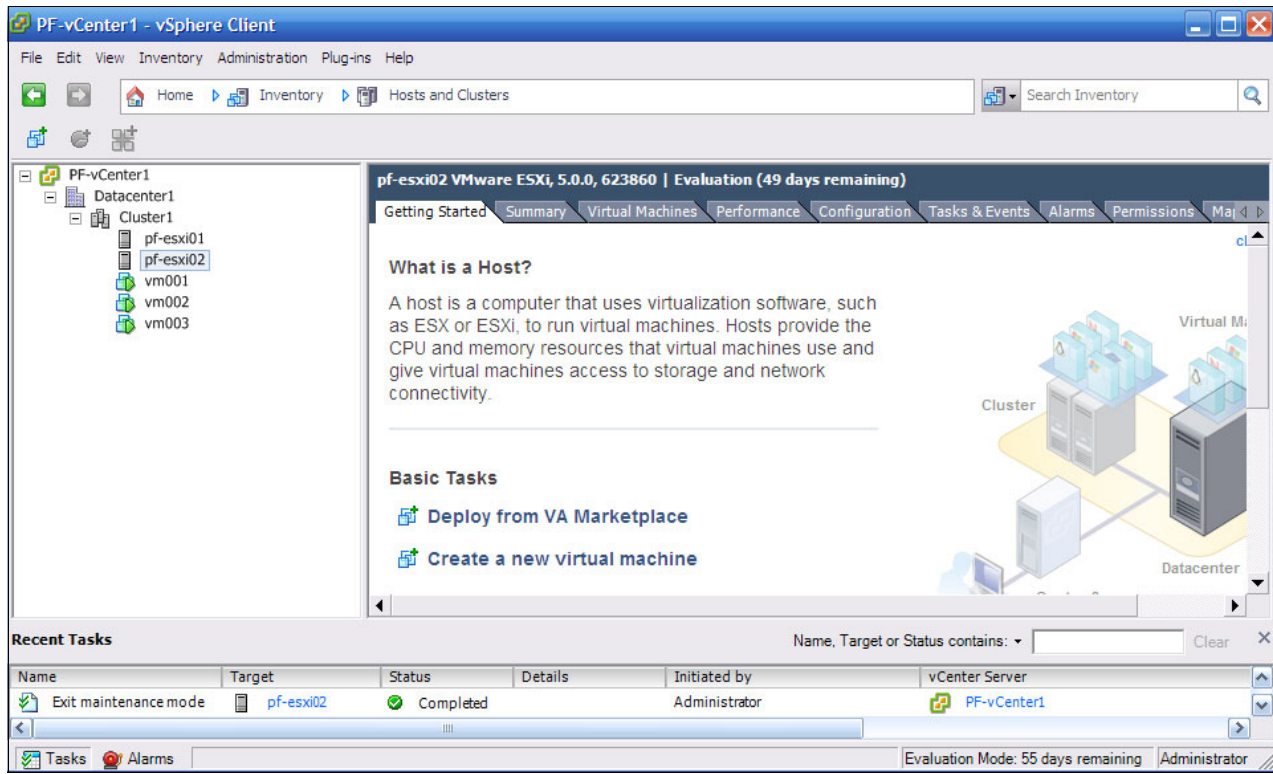


Figure 11-60 vSphere Client window that shows the Exit maintenance mode task

11.8 Topology view

Using the topology view in FSM, you can to view and manage your virtual infrastructure. You can use the Virtualization Basic Topology perspective to view and manage your vCenter, data centers, clusters, hosts, virtual servers, operating systems, and physical compute nodes. The Virtualization Basic Topology interactive map shows you the logical relationships between the components of your virtual infrastructure. This view can also be useful to troubleshoot a problem by determining problematic components and their logical connections to the rest of the infrastructure.

To enable the topology view, perform these steps:

1. In the IBM Flex System Manager web interface navigation area, expand **Inventory** and **Views** and click **Platform Managers and Members**, as shown in Figure 11-61. For this example, select the vCenter server PF-vCenter01 that is installed in virtual machine **vm001**.

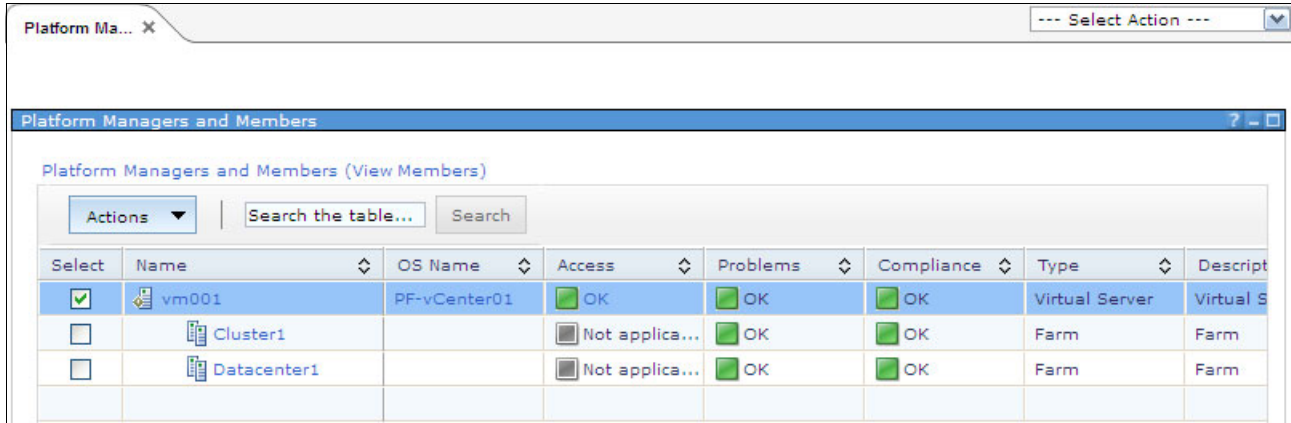


Figure 11-61 Platform Managers and Members window

2. Click **Actions** → **Topology Perspectives** → **Virtualization Basic**, as shown in Figure 11-62.




Figure 11-62 Virtualization Basic menu item in Platform Managers and Members window

3. Click the **Cluster1** icon and click the **Maximize** icon in the Details window, as shown in Figure 11-63.




Figure 11-63 Virtualization Basic Topology view

The details for the selected object are displayed, as shown in Figure 11-64.

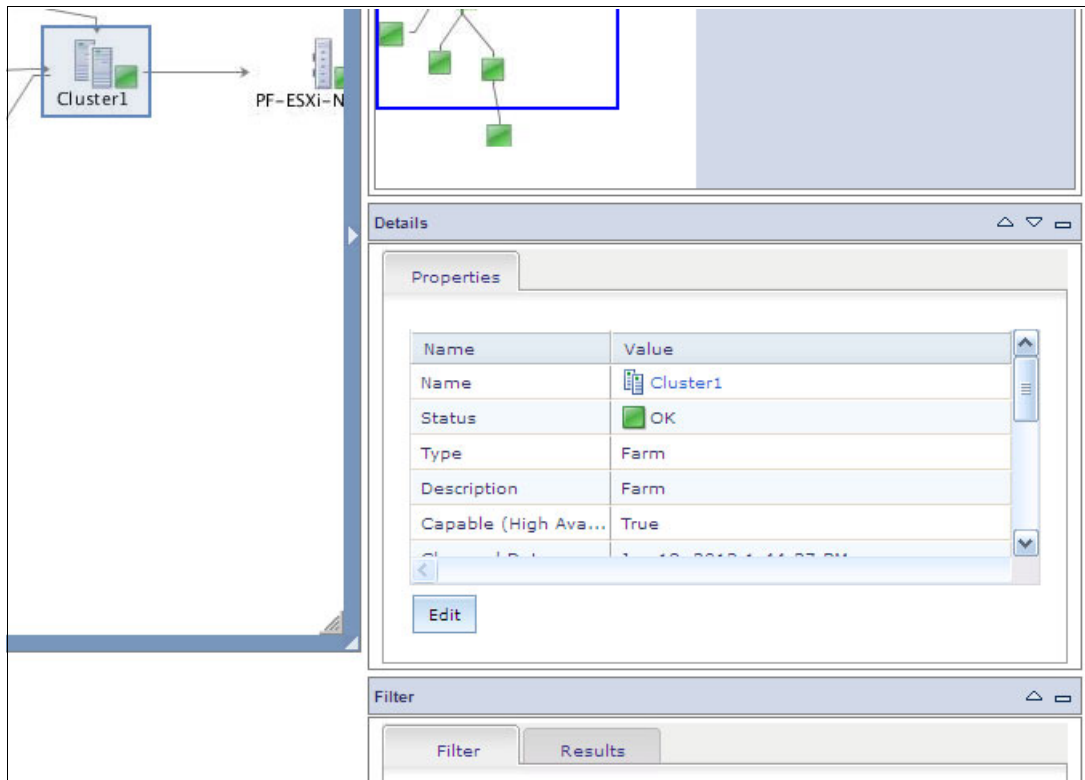


Figure 11-64 Details window in the Virtualization Basic Topology view

- Click the **Hide Palette View** icon and **Zoom To Fit** icon to get a full diagram of the base components of your virtual infrastructure. Right-click any component to get the Actions menu that is relevant to the selected component, as shown in Figure 11-65.

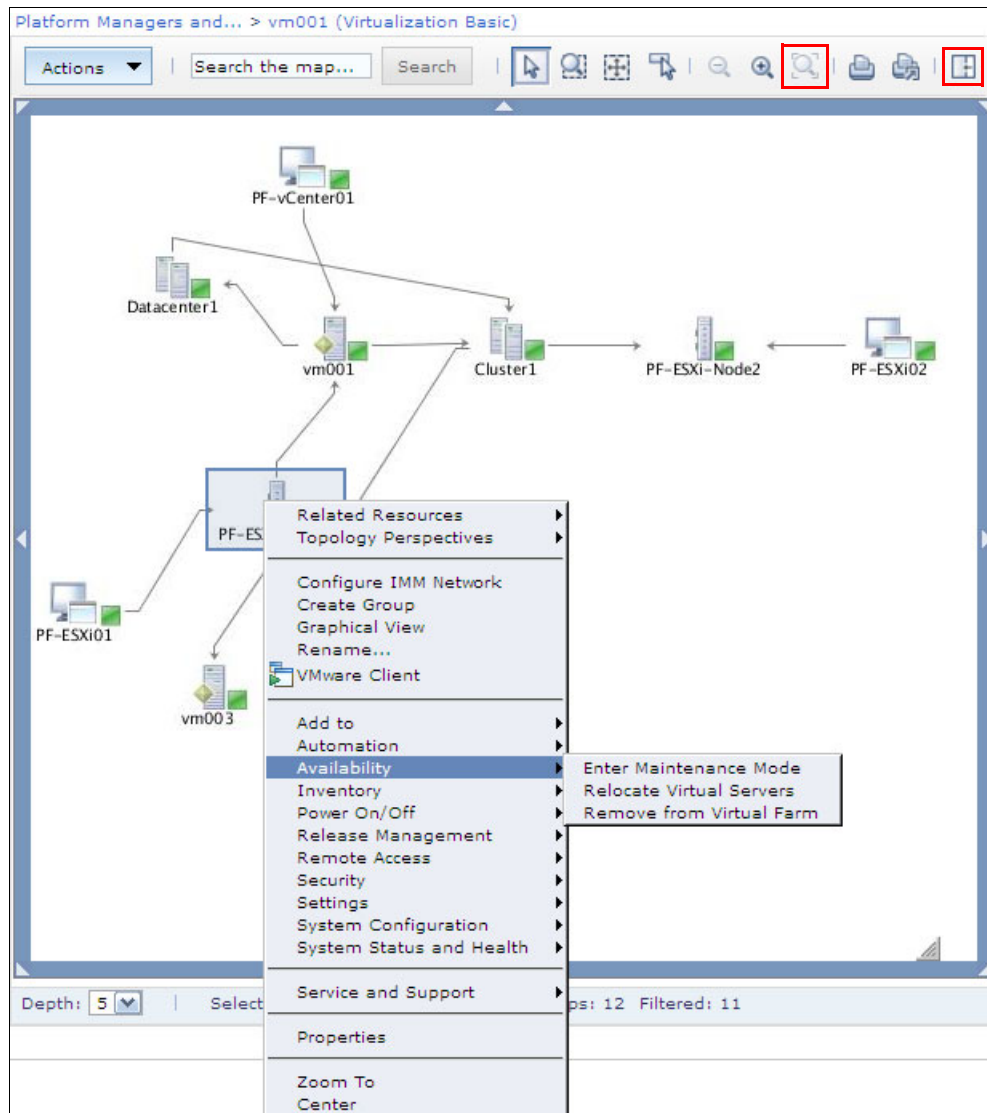


Figure 11-65 Menu items in the Virtualization Basic Topology view

Tip: Click **Actions** → **Layout** → **Tree** to change the default Radial layout if it does not suit your purposes.

11.9 Automating preventive actions in response to hardware alerts

This section addresses how to automate tasks that can prevent service outages. This example involves configuring automation that is based on hardware alerts. To automate preventive actions, perform these steps:

1. In the IBM Flex System Manager web interface navigation area, expand **Automation** and click **Event Automation Plans**, as shown in Figure 11-66. Click **Create** for a new automation plan.

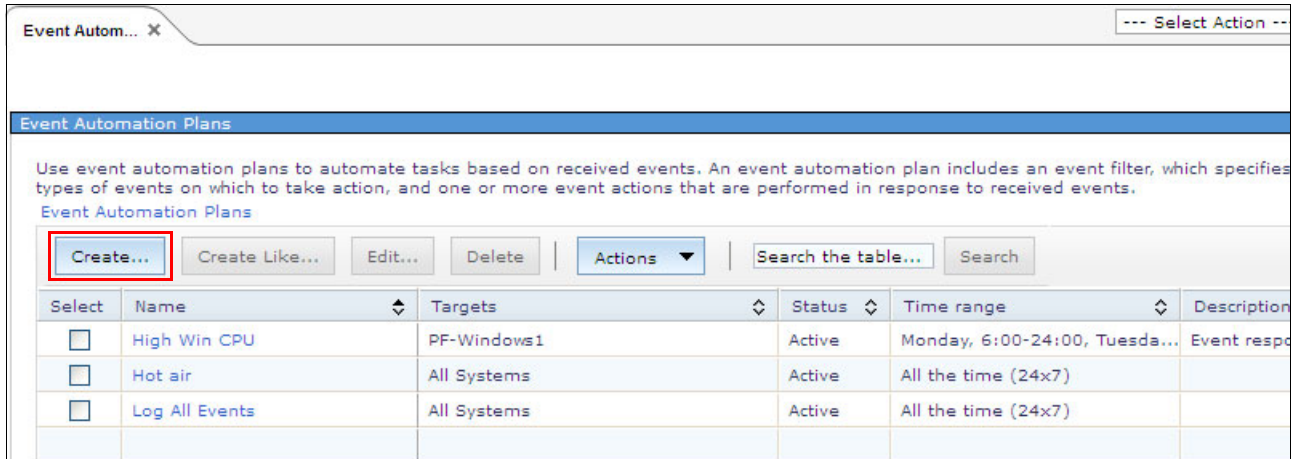


Figure 11-66 Event Automation Plans window

2. Click **Next** in the Welcome window, as shown in Figure 11-67.




Figure 11-67 Create Event Automation Plan Wizard Welcome window

3. Enter a name and description for the automation plan, as shown in Figure 11-68.




Figure 11-68 Create Event Automation Plan Wizard Name and Description window

4. In the Targets window, select the systems to be affected by the event automation plan. For this example, select the two ESXi servers, **PF-ESXi-Node1** and **PF-ESXi-Node2**, as shown in Figure 11-69, and click **Add**. Click **Next** to proceed.

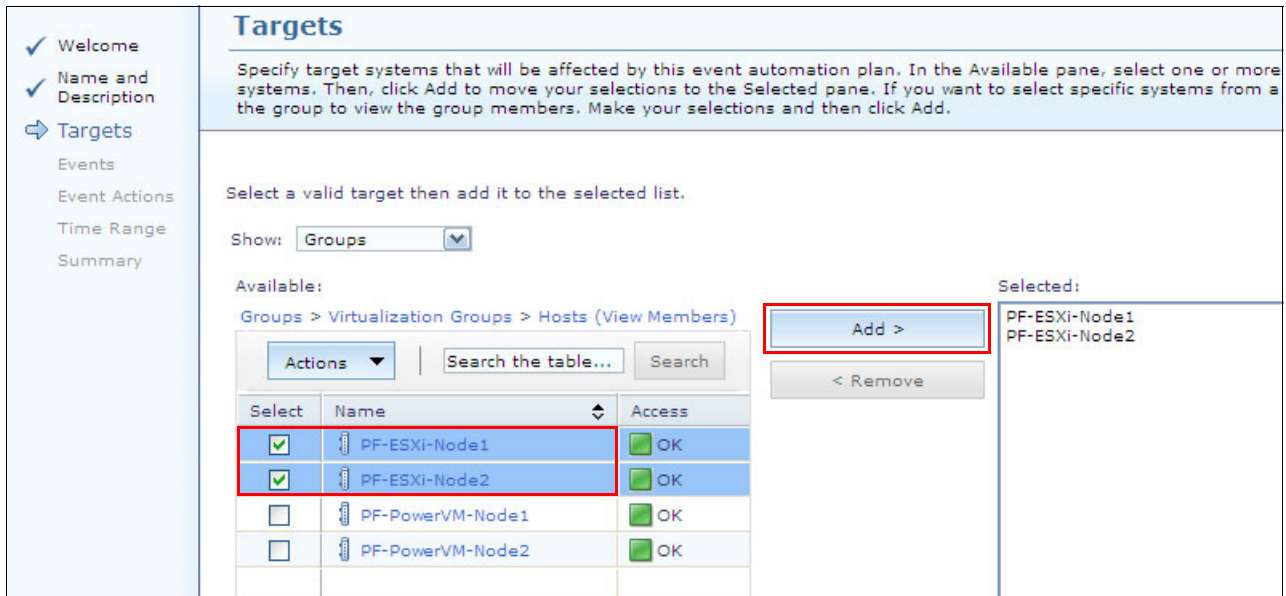


Figure 11-69 Create Event Automation Plan Wizard Targets window

5. Select **Advanced Event Filters** from the Events menu. Select **Critical Events** from the Event Filters list to process all events that have a Critical severity (Figure 11-70).

Tip: If needed, you can also select the Hardware Predictive Failure Alerts event filter. For this exercise, filter all Critical events.



Figure 11-70 Create Event Automation Plan Wizard Events window

6. Click **Create** for a new event action, as shown in Figure 11-71.

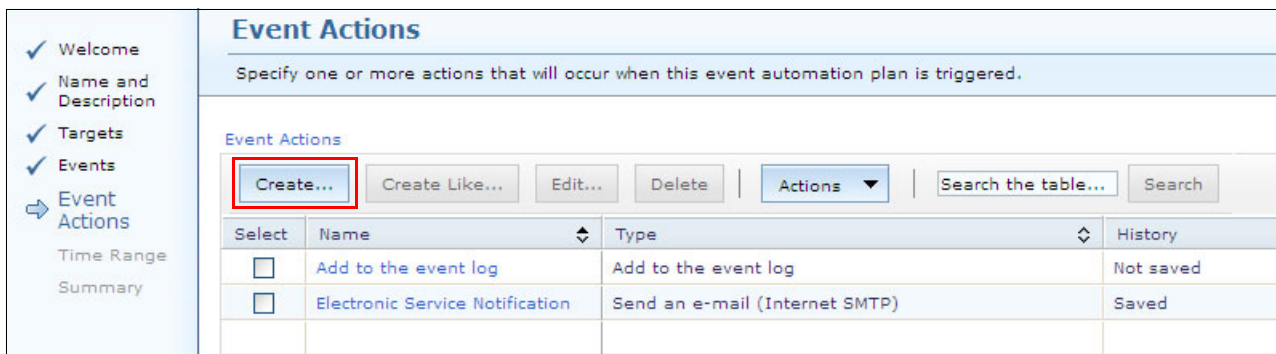


Figure 11-71 Create Event Automation Plan Wizard Event Actions window

- Enter an action name and description for the event action and select **Enter Maintenance Mode** from the Select a task to run menu, as shown in Figure 11-73. Get familiar with the broad choice of tasks that you can run as an action. Click **OK**.

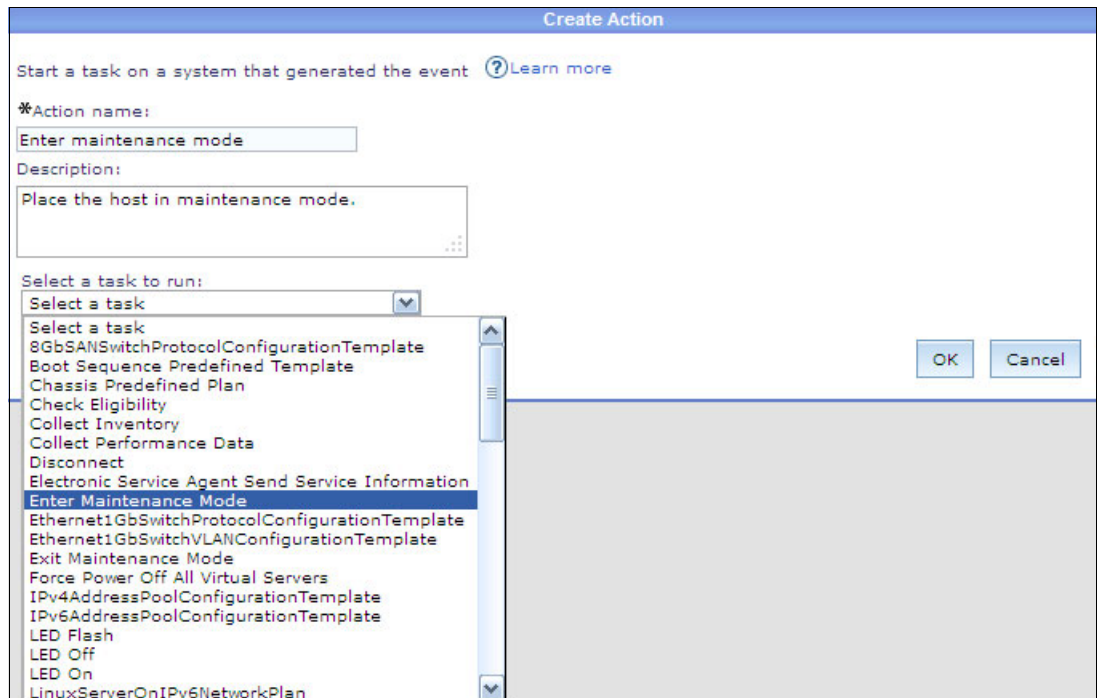


Figure 11-73 Create Action window action properties

- Select the newly created **Enter maintenance mode** event action, as shown in Figure 11-74. Click **Next**.

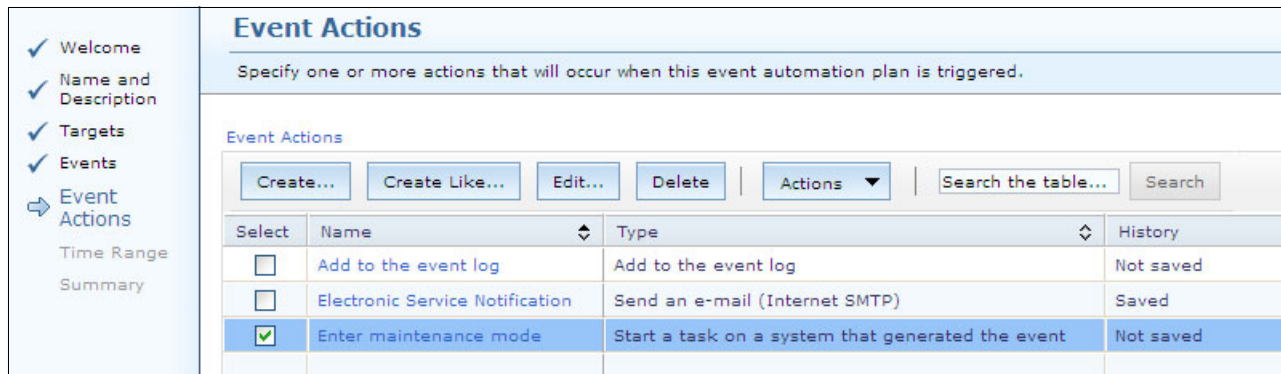


Figure 11-74 Create Event Automation Plan Wizard Event Actions window

10. Click **Next** in the Time Range window, as shown in Figure 11-75.

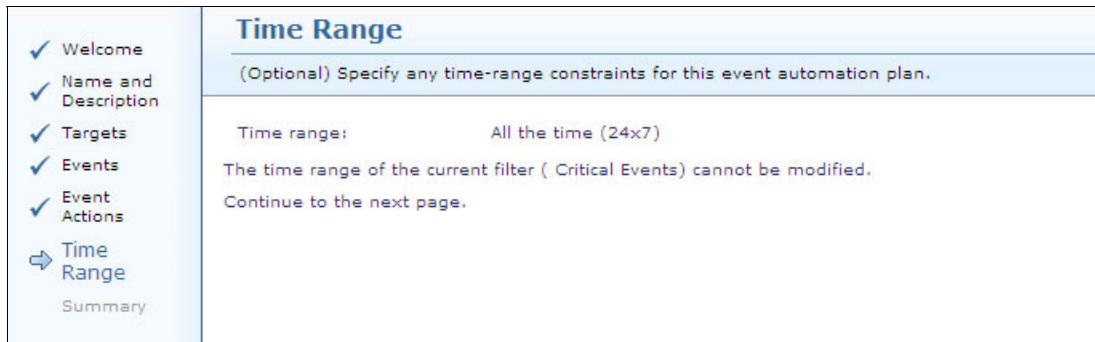


Figure 11-75 Create Event Automation Plan Wizard Time Range window

11. Review the Summary window and click **Finish** to create and apply the event automation plan, as shown in Figure 11-76.

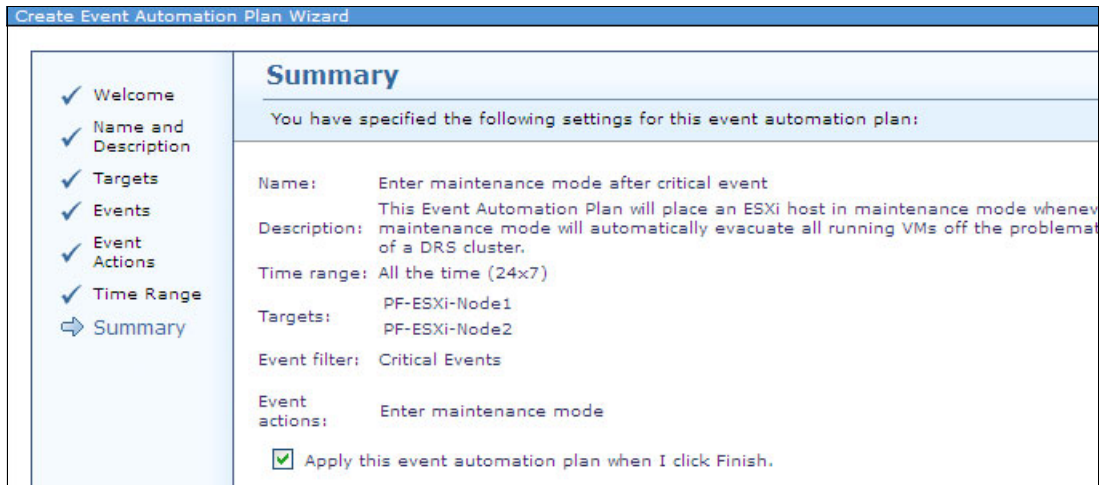


Figure 11-76 Create Event Automation Plan Wizard Summary window

The newly created event automation plan is displayed in the Event Automation Plans window, as shown in Figure 11-77.

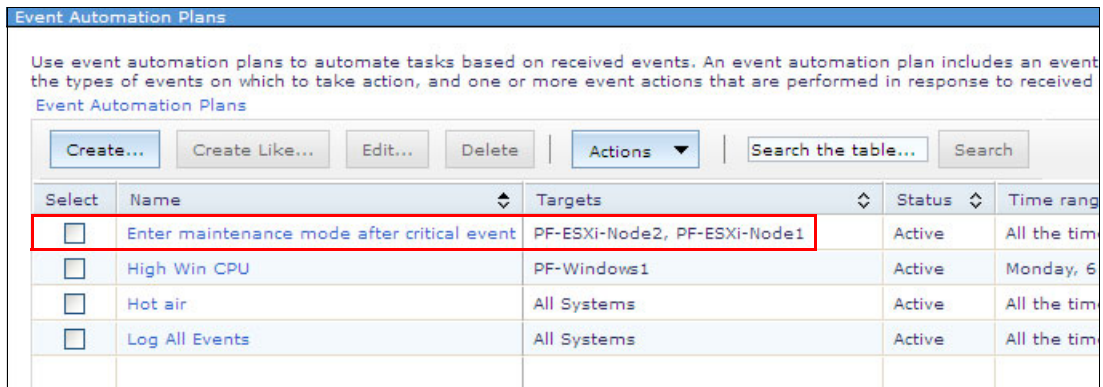


Figure 11-77 Event Automation Plans window

12. For testing, generate a Critical System error with source PF-ESXi-Node2. The status of the PF-ESXi-Node2 compute node shows Critical on the Chassis Map, as shown in Figure 11-78.




Figure 11-78 Chassis Map showing PF-ESXi-Node2 with a critical error

The generated error is also listed in PF-ESXi-Node2 Event Log, as shown in Figure 11-79.

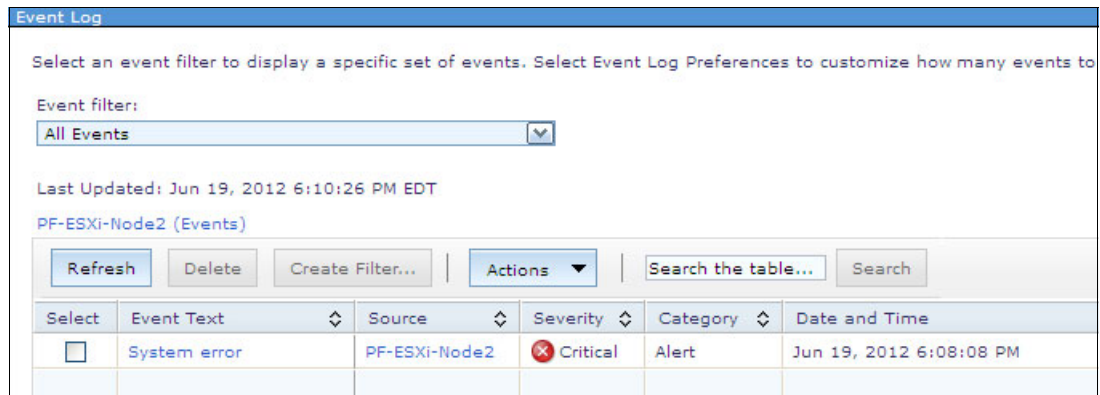


Figure 11-79 Event Log window for selected node

13. Return to the Virtual Servers and Hosts window and make sure that all virtual servers were automatically migrated away from the host with the Critical error. See Figure 11-80.

Select	Name	State	OS Name	OS Type and Version	Access	Problems	Compliance
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	OK	OK
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008...	OK	Information	OK
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008...	OK	Information	OK
<input type="checkbox"/>	vm003	Started			OK	Information	OK
<input type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	Critical	OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK	OK	OK
<input type="checkbox"/>	PF-Node1-NIM	Started	PF-Node1-NIM	AIX 6.1	OK	OK	OK
<input type="checkbox"/>	SN101D88B_VIOC1	Started			OK	OK	OK
<input type="checkbox"/>	SN101D88B_VIOS1	Started			OK	OK	OK
<input type="checkbox"/>	PF-PowerVM-Node2	Standby			OK	OK	OK

Figure 11-80 Virtual Servers and Hosts window

14. Select PF-ESXi-Node2 and click **Actions** → **System Status and Health** → **Active Status**, as shown in Figure 11-81.

Select	Name	State	OS Name	OS Type and Version	Access	Problems	Compliance
<input type="checkbox"/>	PF-ESXi-Node1	Started	PF-ESXi01	Hypervisor 5.0.0	OK	OK	OK
<input type="checkbox"/>	vm001	Started	PF-vCenter01	Windows® Server 2008...	OK	Information	OK
<input type="checkbox"/>	vm002	Started	PF-Windows1	Windows® Server 2008...	OK	Information	OK
<input type="checkbox"/>	vm003	Started			OK	Information	OK
<input checked="" type="checkbox"/>	PF-ESXi-Node2	Started	PF-ESXi02	Hypervisor 5.0.0	OK	Critical	OK
<input type="checkbox"/>	PF-PowerVM-Node1	Started			OK	OK	OK
<input type="checkbox"/>	PF-Node1-NIM	Started	PF-Node1-NIM	AIX 6.1	OK	OK	OK
<input type="checkbox"/>	SN101D88B_VIOC1	Started			OK	OK	OK
<input type="checkbox"/>	SN101D88B_VIOS1	Started			OK	OK	OK
<input type="checkbox"/>	PF-PowerVM-Node2	Standby			OK	OK	OK

Figure 11-81 Selecting Active Status

15. Observe the Information events that describe the actions that were performed by the event automation plan configured earlier. See Figure 11-82. PF-ESXi-Node2 is now in maintenance mode.

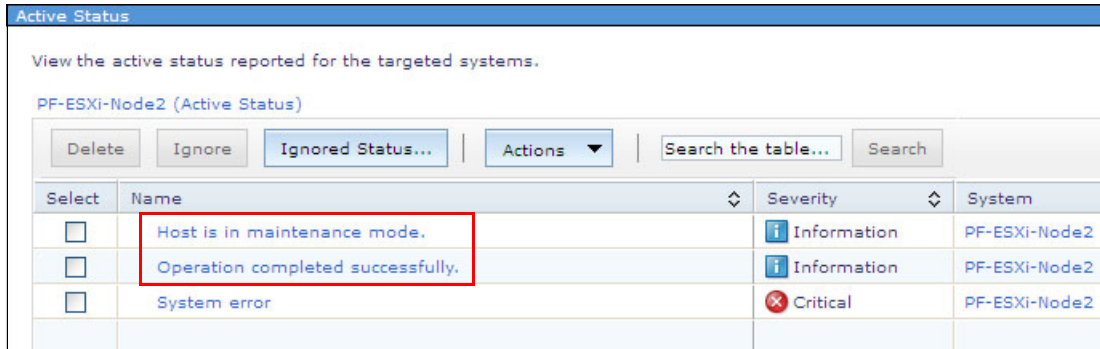


Figure 11-82 Active Status window for selected node

FSM triggered the “Enter maintenance mode” command on vCenter and DRS migrated all virtual machines away from pf-esxi02 before placing it in maintenance mode, as shown in Figure 11-83.

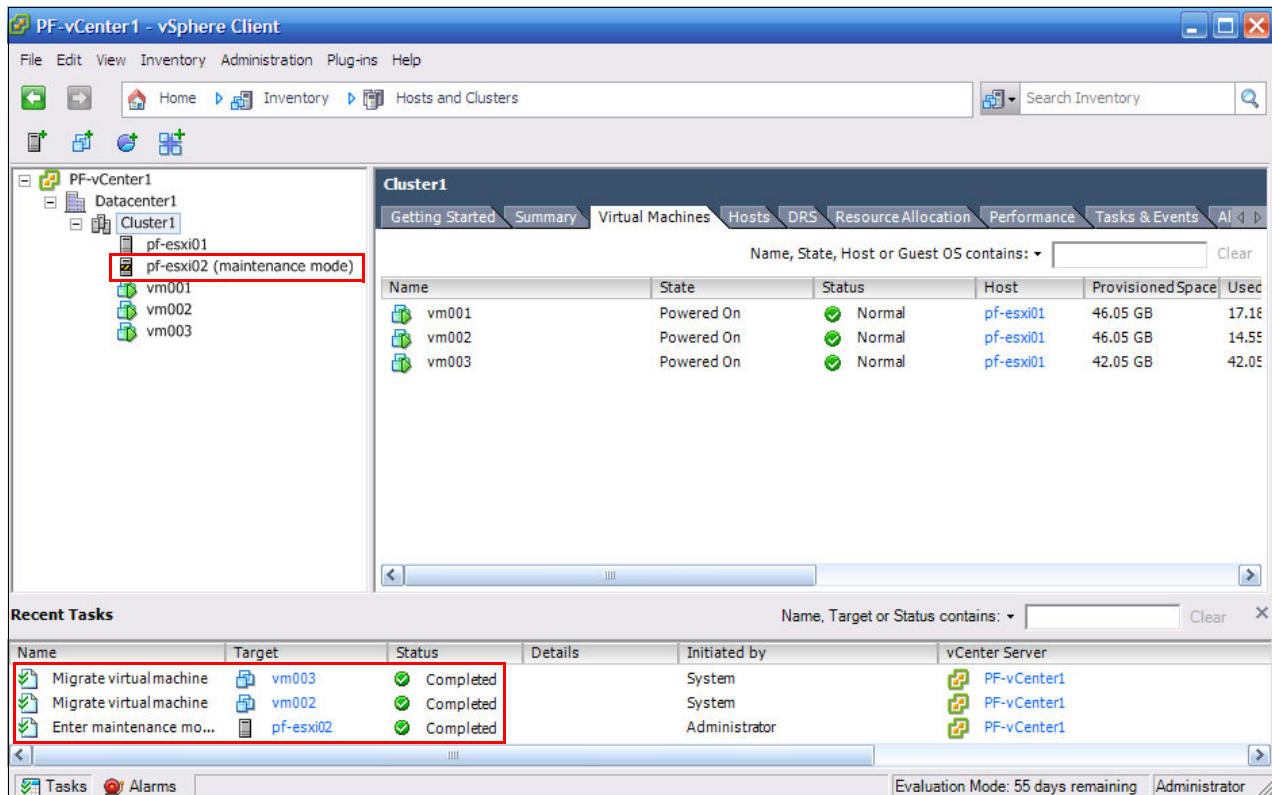


Figure 11-83 vSphere Client window that shows a host in maintenance mode

16. Return to the Virtual Servers and Hosts window, right-click **PF-ESXi-Node2** and select **Availability** → **Exit Maintenance Mode**, as shown in Figure 11-84.

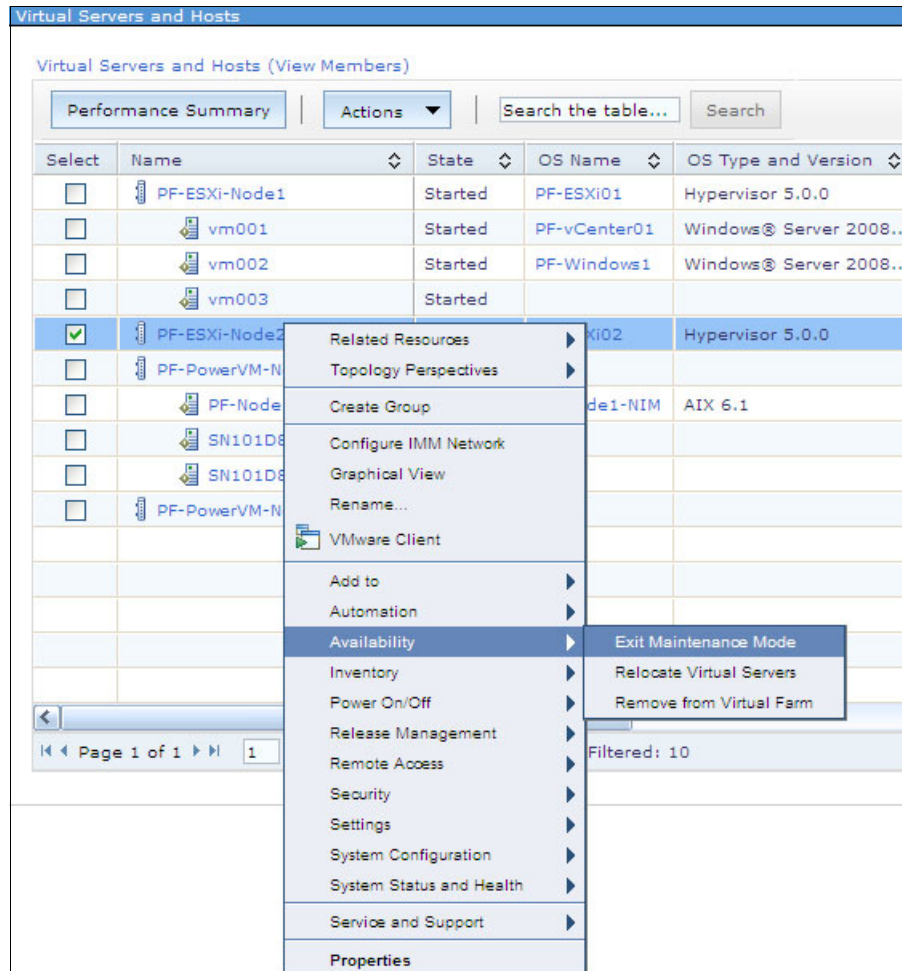


Figure 11-84 Exiting Maintenance Mode



Managing the Hyper-V environment with IBM Flex System Manager

This chapter is focused on managing the Hyper-V-based virtualization environment with IBM Flex System Manager (FSM). It addresses how to enable Hyper-V to be managed by FSM, and how to perform typical virtualization management tasks. These tasks include virtual machine lifecycle management, automation capabilities, and maintenance.

The following topics are covered:

- ▶ 12.1, “Initial setup tasks for a Hyper-V node” on page 556
- ▶ 12.2, “Managing Hyper-V with IBM Flex System Manager” on page 562

12.1 Initial setup tasks for a Hyper-V node

IBM Flex System Manager can run basic tasks for the Microsoft Hyper-V hypervisor. You can start, stop, restart, suspend, create, and delete your virtual servers that run on Microsoft hypervisors with the same tool that manages other hypervisors in the market.

Before you can manage a Hyper-V virtual environment, you must perform these tasks:

- ▶ Prepare your Hyper-V system as addressed in 5.2.5, “Planning for Hyper-V virtualization” on page 118.
- ▶ Download the Windows Common Agent for Remote Installation from the URL:
https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=dmp&lang=en_US&S_PKG=dir_63_x86_RDagents
- ▶ Import the Common Agent for Windows x64 operating systems as described in “Importing the Common Agent for Windows” on page 557.
- ▶ Discover your system without any agent, grant access to it, and run the Collect Inventory task before you install the agent with FSM.
- ▶ Install a Platform or a Common Agent on your hypervisor to allow your FSM management appliance to manage your Microsoft hypervisor.

This section describes the following tasks:

- ▶ 12.1.1, “Discovering your Hyper-V server” on page 556
- ▶ 12.1.2, “Importing the Common Agent for Windows” on page 557
- ▶ 12.1.3, “Granting access and collecting inventory on a Hyper-V node” on page 558
- ▶ 12.1.4, “Installing the Common Agent on a Hyper-V host” on page 559

12.1.1 Discovering your Hyper-V server

To discover a Hyper-V server, use System Discovery as described in 6.8, “System discovery, access, and inventory collection” on page 176.

Specify the IP address of your operating system and click **Discover Now**, as shown in Figure 12-1. Then, wait until the job is complete.

System Discovery

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a resource for a single discover resources of the same type for a range of IP addresses, or use a discovery profile. Discovery profiles enable you to customize discovery IP addresses, and requesting access to and collecting inventory for the discovered resources.

[Learn more about using discovery](#)

Select a discovery option:
Single IPv4 address

IP address:
9 . 27 . 16 . 125

Select the resource type to discover:
All

Discover Now

Schedule...

Advanced Tasks

- Create new profile
- Manage discovery profiles
- Discovery jobs

Figure 12-1 Discovering a Hyper-V node

12.1.2 Importing the Common Agent for Windows

To import the Common Agent for Windows, perform these steps:

1. Select **Release Management**. Then, click **Agents** as shown in Figure 12-2.



Figure 12-2 Release Management

2. Select an agent and click **Import Agent** (Figure 12-3).




Figure 12-3 Importing an agent

3. Specify the path that contains the agent that you want to import and click **OK**, as shown in Figure 12-4.

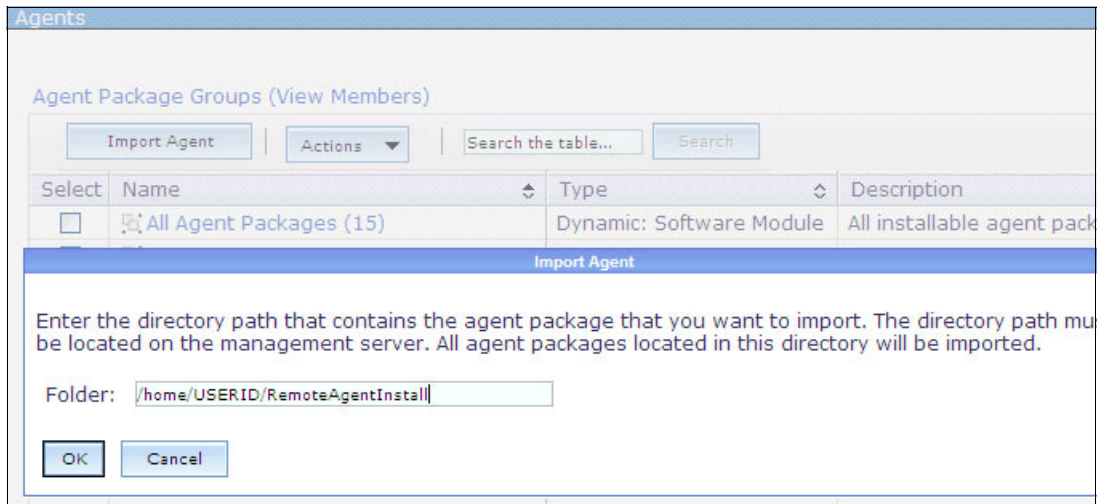


Figure 12-4 Path to import agent

After a few minutes, you get a blue information window as shown in Figure 12-5.




Figure 12-5 Agent imported

Tip: In this example, six agents are imported because six agents were present in the local FSM directory /home/USERID/RemoteAgentInstall.

12.1.3 Granting access and collecting inventory on a Hyper-V node

To grant access and collect inventory on a Hyper-V node, perform these steps:

1. Find your server in the Resource Explorer, and click **No Access** (Figure 12-6).




Figure 12-6 Request access to Hyper-V node

2. When the job is complete, your access is OK as shown in Figure 12-7. You can install your Common Agent that you imported previously.


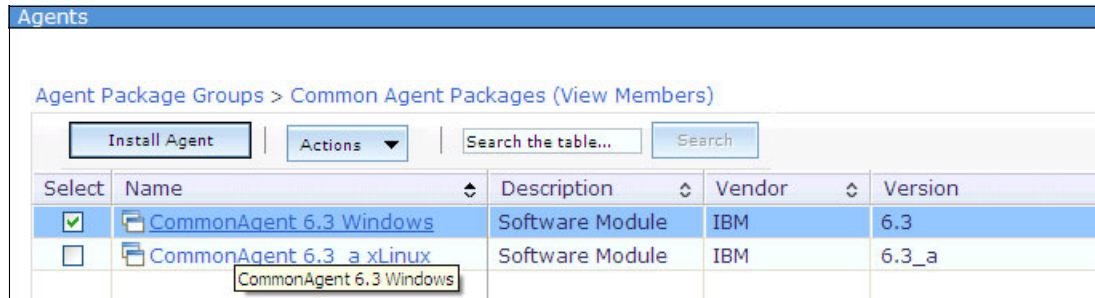


Figure 12-7 Hyper-V access is granted

12.1.4 Installing the Common Agent on a Hyper-V host

To install the Common Agent on a Hyper-V host, perform these steps:

1. Go to **Release Management** as shown in Figure 12-8.



The screenshot shows the 'Agents' console window. At the top, it displays 'Agent Package Groups > Common Agent Packages (View Members)'. Below this, there is a toolbar with an 'Install Agent' button, an 'Actions' dropdown menu, a search box labeled 'Search the table...', and a 'Search' button. A table lists the available agent packages:

Select	Name	Description	Vendor	Version
<input checked="" type="checkbox"/>	CommonAgent 6.3 Windows	Software Module	IBM	6.3
<input type="checkbox"/>	CommonAgent 6.3 a xLinux	Software Module	IBM	6.3_a

Figure 12-8 Select CommonAgent 6.3 Windows for an installation on Hyper-V

An Agent Installation Welcome window opens as shown in Figure 12-9.

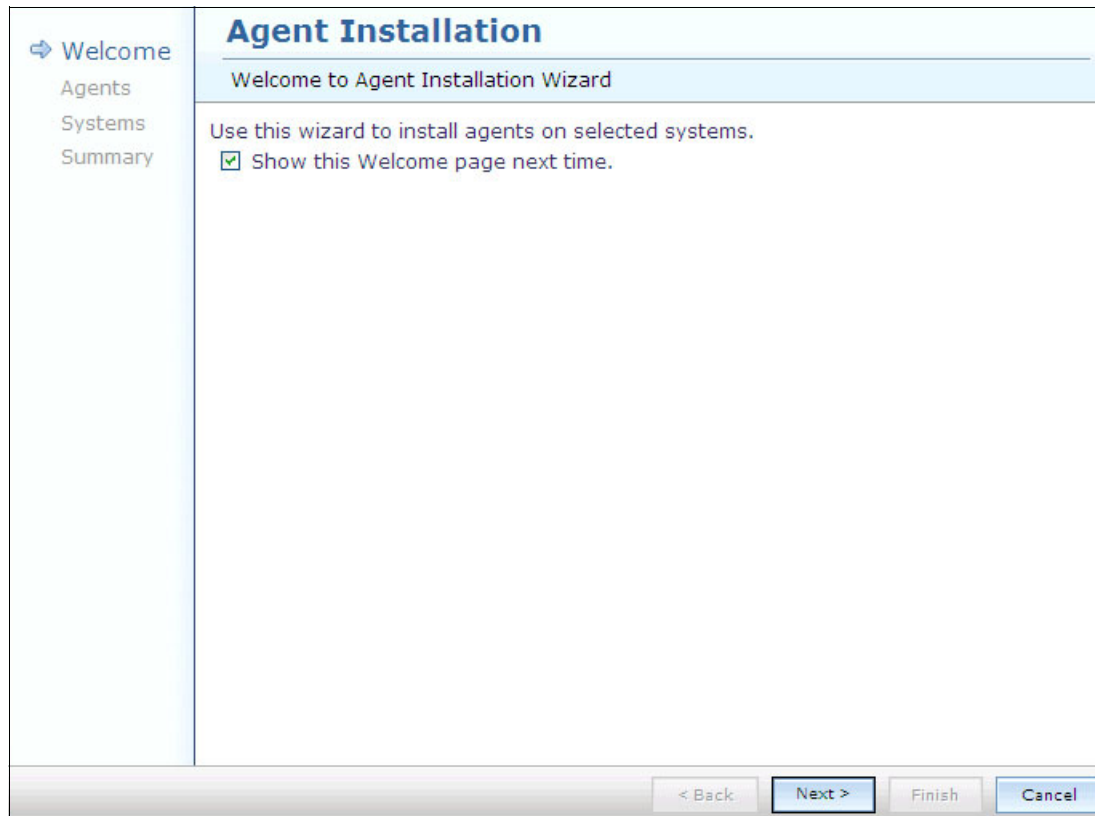


Figure 12-9 Agent Installation: Welcome page

- You can see a list of agents. Click **Next** because the Common Agent for Windows is already in the selected list (Figure 12-10). Otherwise, add the Common Agent for Windows from the Common Agent Packages by clicking **Common Agent Packages**.

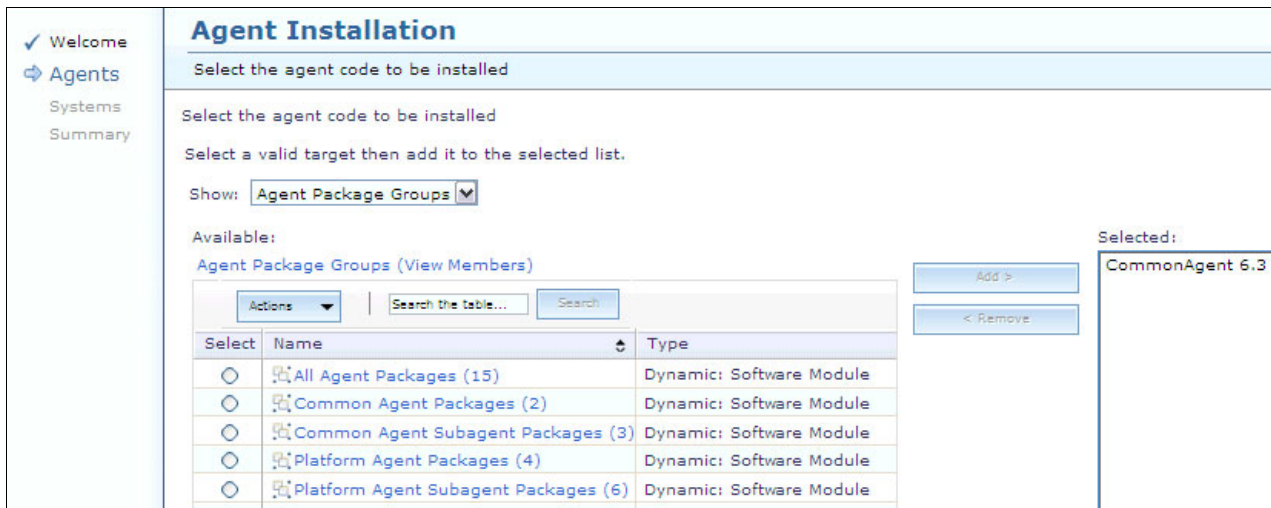


Figure 12-10 Agent Installation: Package selection

- Select your Hyper-V host as shown in Figure 12-11. Click **Add**, and then click **Next**.




Figure 12-11 Agent Installation: Target selection

4. Check the summary of the Common Agent installation as shown in Figure 12-12.




Figure 12-12 Agent Installation: Summary

5. In the Launch Job window, click **OK** to run the installation as shown in Figure 12-13.

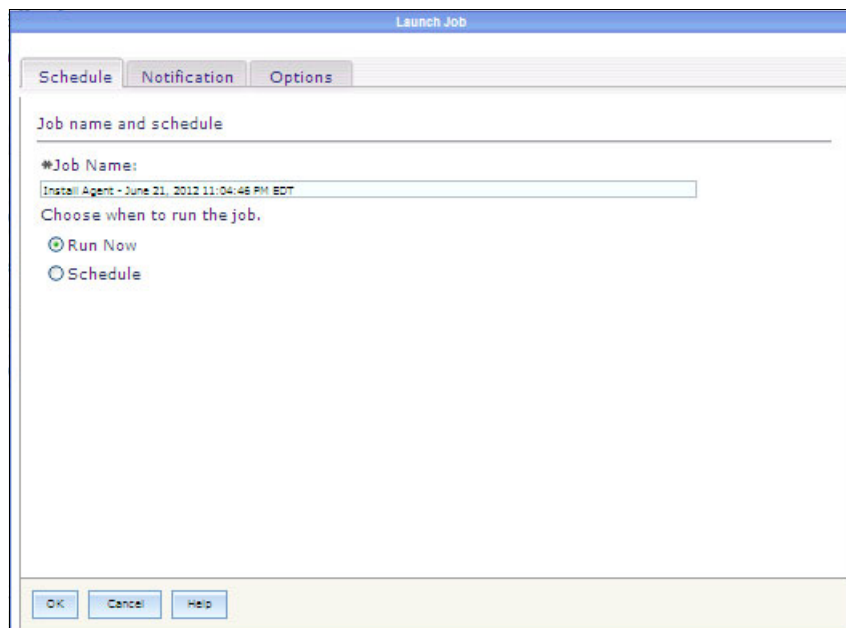


Figure 12-13 Agent Installation: Launch Job

6. Click **Display Properties** as shown in Figure 12-14.




Figure 12-14 Display Properties

- Go to your Windows 2008 R2 Hyper-V server and make sure that the Common Agent is present in the list of installed programs.

12.2 Managing Hyper-V with IBM Flex System Manager

The following tasks can be performed in a Hyper-V environment from the FSM appliance:

- ▶ Deploying virtual servers
- ▶ Editing a virtual server
- ▶ Deleting a virtual server
- ▶ Viewing the virtual server network topology

12.2.1 Deploying virtual servers

To deploy virtual servers, perform these steps:

- Go to VMControl and click the **Virtual Servers/Hosts** tab as shown in Figure 12-15.

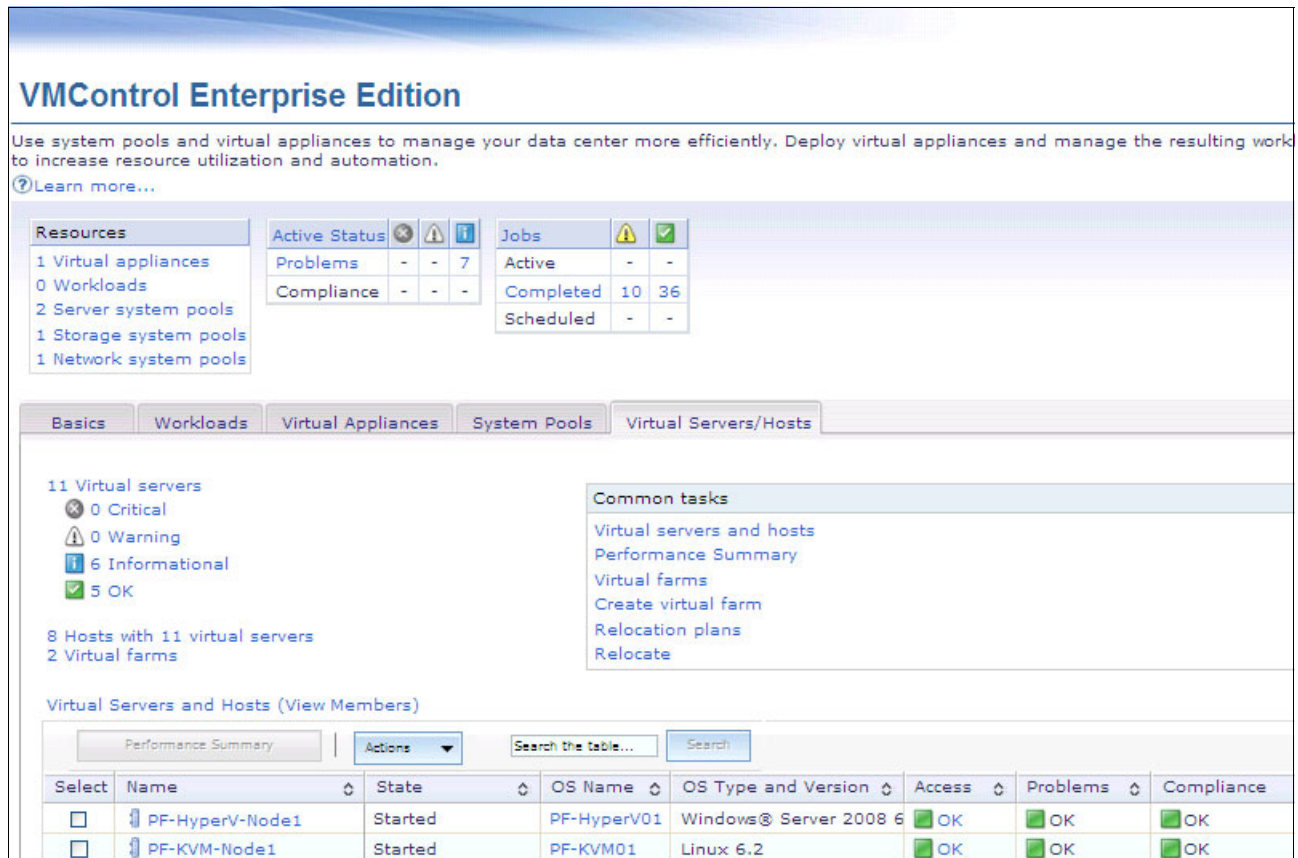


Figure 12-15 VMControl Enterprise Edition

- Right-click your Hyper-V host and select **System Configuration** → **Create Virtual Server**, as shown in Figure 12-16.




Figure 12-16 VMControl: Create Virtual Server

- A Welcome window opens as shown in Figure 12-17. Click **Next**.

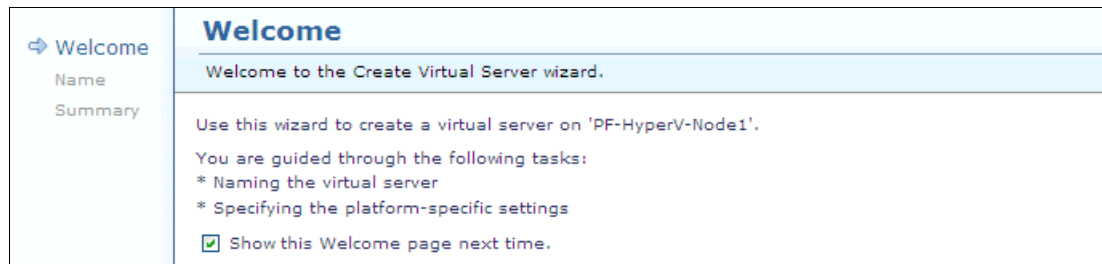


Figure 12-17 Create Virtual Server: Welcome page

- Enter a name for your virtual server and click **Next**, as shown in Figure 12-18.



Figure 12-18 Naming a virtual server

5. Enter the number of processors and click **Next**, as shown in Figure 12-19.

The screenshot shows a configuration wizard with a sidebar on the left containing a list of steps: Welcome, Name, Processor (highlighted with a blue arrow), Memory, Disks, Network, and Summary. The main panel is titled "Processor" and contains the instruction "Specify the number of processors to assign to this virtual server." Below this, there is a label "Number Of Processors:" followed by a text input field containing the number "1" and a range indicator "(1-4)".

Figure 12-19 Number of virtual server processors

6. Enter your virtual server memory size and click **Next**, as shown in Figure 12-20.

The screenshot shows a configuration wizard with a sidebar on the left containing a list of steps: Welcome, Name, Processor, Memory (highlighted with a blue arrow), Disks, Network, and Summary. The main panel is titled "Memory" and contains the instruction "Specify the amount of memory to assign to this virtual server." Below this, there are two fields: "Memory Size:" with a text input field containing "8", and "Units:" with a dropdown menu set to "MB" and a range indicator "(8-65,536)".

Figure 12-20 Virtual server memory size

7. Enter your disk size and click **Next**, as shown in Figure 12-21.

The screenshot shows a configuration wizard with a sidebar on the left containing a list of steps: Welcome, Name, Processor, Memory, Disks (highlighted with a blue arrow), Network, and Summary. The main panel is titled "Disks" and contains the instruction "Specify the disk settings to use for this virtual server." Below this, there is a label "Select the amount of disk space to assign to this virtual server." followed by "Size:" and a text input field containing "20" and a range indicator "(3-2,040) GB".

Figure 12-21 Virtual server disk size

8. Select the virtual switch to which to connect your virtual server and click **Next**, as shown in Figure 12-22.

The screenshot shows a configuration wizard with a sidebar on the left containing a list of steps: Welcome, Name, Processor, Memory, Disks, Network (highlighted with a blue arrow), and Summary. The main panel is titled "Network" and contains the instruction "Select the network label for this virtual server." Below this, there is a label "Network Label:" followed by a dropdown menu showing "Local Area Connection 2 - Virtual Network".

Figure 12-22 Configuring a virtual server network

9. Review your virtual server summary configuration and click **Finish**, as shown in Figure 12-23.




Figure 12-23 Virtual server configuration summary

10. Click **OK** to run your virtual server creation task as shown in Figure 12-24.

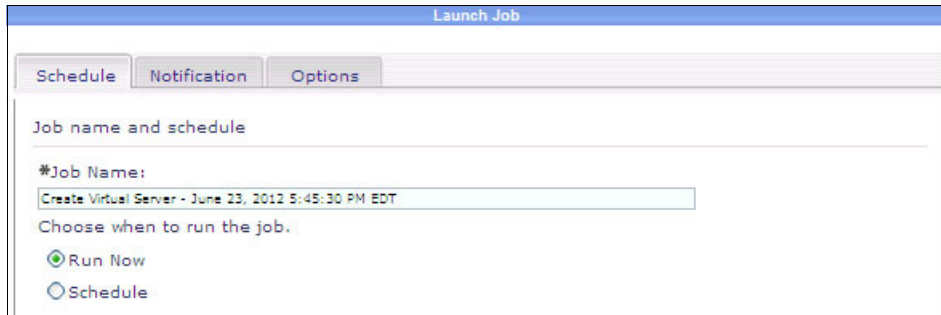


Figure 12-24 Run now

11. A blue box information message is displayed as shown in Figure 12-25. Click **Display Properties**.

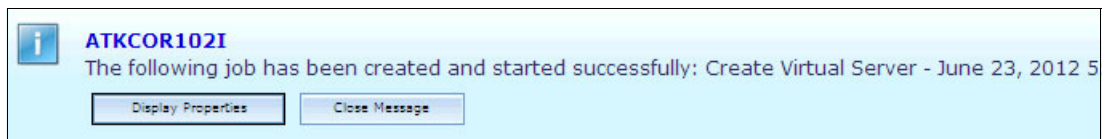


Figure 12-25 Information blue box

Wait until the job is complete as shown in Figure 12-26.

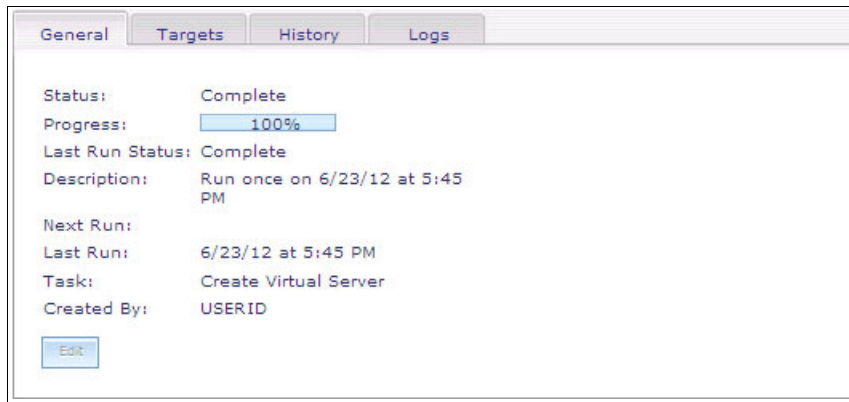


Figure 12-26 Job is complete

12. Go back to VMControl and select **Virtual Servers/Hosts** to make sure that your virtual server was created as shown in Figure 12-27.

The screenshot shows a table titled 'Virtual Servers and Hosts (View Members)'. The table has columns for 'Select', 'Name', 'State', 'OS Name', 'OS Type and Version', and 'Access'. There are two rows of data:

Select	Name	State	OS Name	OS Type and Version	Access
<input type="checkbox"/>	PF-HyperV-Node1	Started	PF-HyperV01	Windows® Server 2008 6	OK
<input type="checkbox"/>	VMWindowsHyperV	Stopped			OK

Figure 12-27 Virtual server was created on Hyper-V

12.2.2 Editing a virtual server

To edit a virtual server, right-click your virtual server and select **System Configuration** → **Edit Virtual Server**, as shown in Figure 12-28.

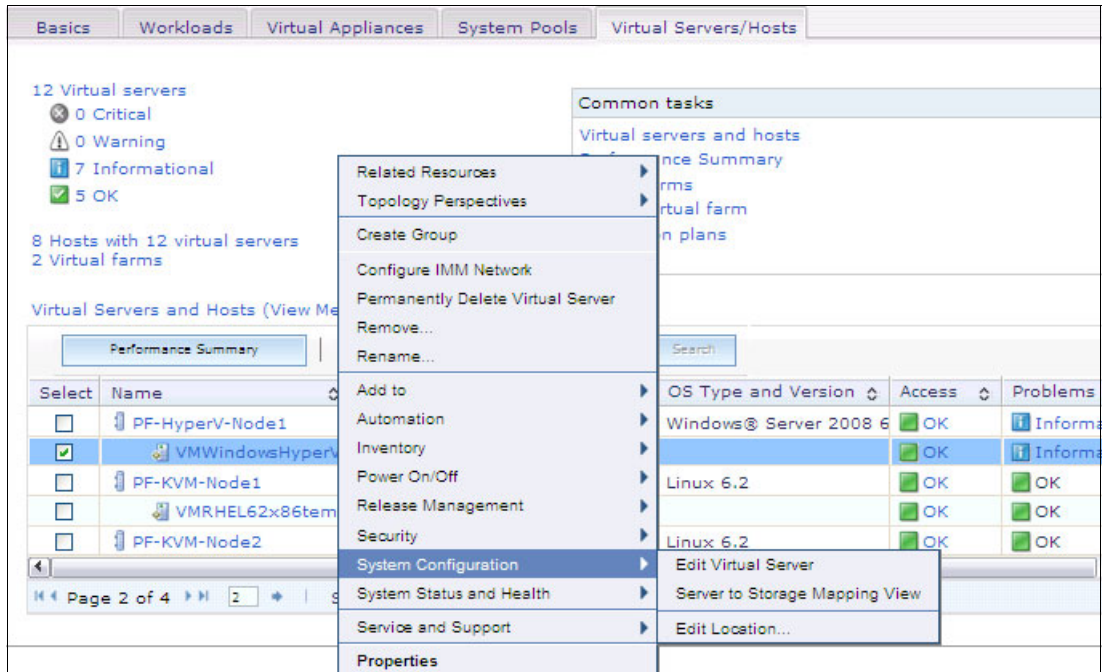


Figure 12-28 Editing a virtual server

You get information about processor and memory size, as shown in Figure 12-29.

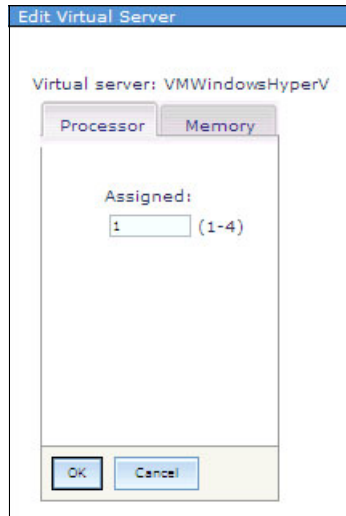


Figure 12-29 Virtual server details

12.2.3 Deleting a virtual server

To delete a virtual server, right-click your virtual server and select **Permanently Delete Virtual Server**, as shown in Figure 12-30.

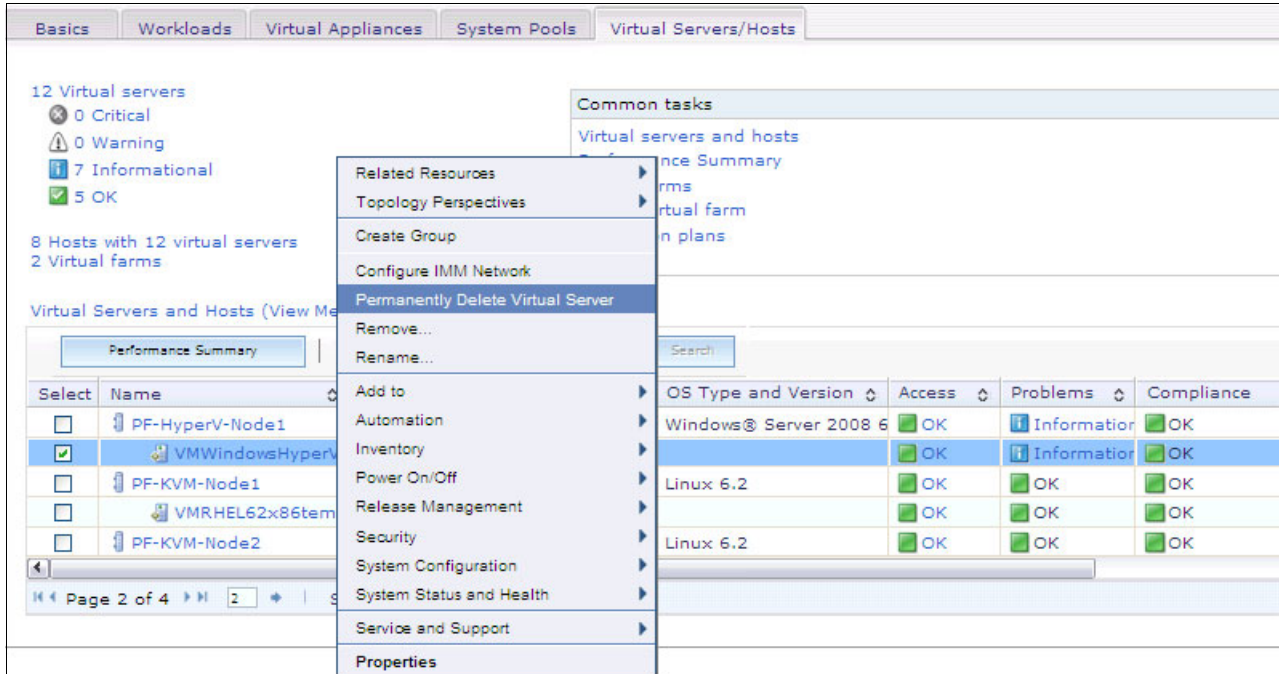


Figure 12-30 Deleting a virtual server

12.2.4 Viewing the virtual server network topology

To view the virtual network topology, perform these steps:

1. Select your virtual server as shown in Figure 12-31.




Figure 12-31 Selecting a virtual server

2. Right-click the server and select **Topology Perspective** → **Network** → **Basic**, as shown in Figure 12-32.




Figure 12-32 Selecting a basic network topology

A network topology view of your virtual server within the infrastructure is displayed as shown in Figure 12-33.

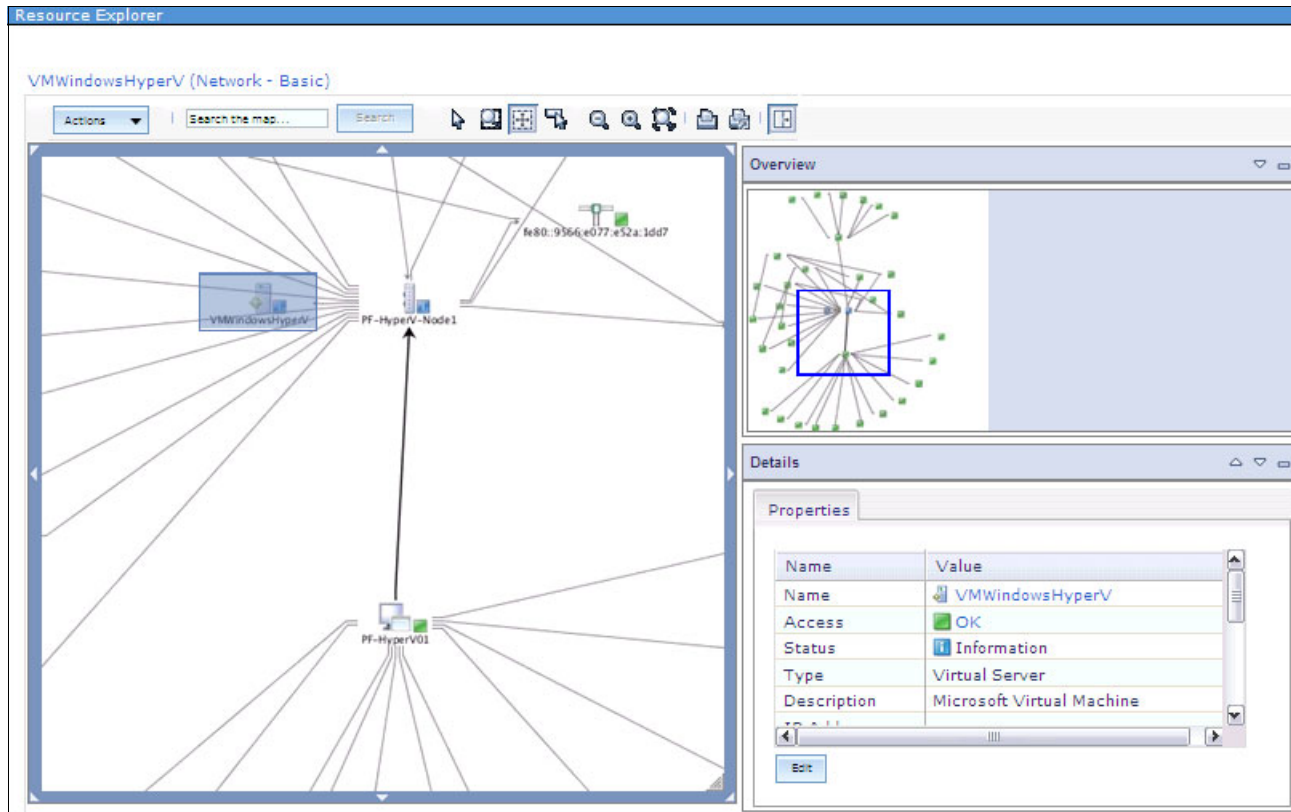


Figure 12-33 Topology view



Mobile management

This chapter describes the features and advantages of the IBM Flex System Manager (FSM) application, which is now available for mobile devices.

In this chapter, we provide information about the mobile application, which enables you to manage your IBM Flex System and PureFlex System hardware remotely through the Flex System Manager.

The following topics are covered:

- ▶ 13.1, “Obtaining the mobile application” on page 572
- ▶ 13.2, “Configuring secure communications to the FSM” on page 572
- ▶ 13.3, “Using the Flex System Manager mobile application” on page 576

13.1 Obtaining the mobile application

The Flex System Manager mobile application can be downloaded from the following application stores (app stores):

- ▶ For the Android operating system, see the IBM Flex System Manager for mobile devices Google Play page:

<http://play.google.com/store/apps/details?id=com.ibm.msm.android>

Note: IBM Flex System Manager for mobile devices is not supported for Android Version 3.0.

- ▶ For Apple iOS, see the IBM Flex System Manager for mobile devices iTunes page:

<http://bit.ly/1jaDIFF>

- ▶ For BlackBerry OS, see the IBM Flex System Manager for mobile devices BlackBerry App World page:

<http://appworld.blackberry.com/webstore/content/20199697/?lang=en>

Note: The IBM Flex System Manager for mobile devices application does not support BlackBerry OS Version 7.0 or earlier.

13.2 Configuring secure communications to the FSM

IBM Flex System Manager for mobile devices can use Secure Sockets Layer (SSL) certificates to create a secure connection to the IBM Flex System Manager management software.

Important: If you have not configured SSL certificates on the IBM Flex System Manager, the app will prompt you with the option to connect in an insecure manner.

To ensure that the connection is secure, the certificate that is installed on the IBM Flex System Manager must be a valid certificate for the URI that will be used to access the IBM Flex System Manager and be signed by a separate certificate authority (CA). Self-signed certificates are not accepted by IBM Flex System Manager for Android, BlackBerry, and iOS.

In addition, there is a current issue with importing certificates that were signed by an intermediate CA instead of a root CA into the default keystore on the IBM Flex System Manager. This issue requires you to create a new keystore on a separate system (with Java installed) from the Flex System Manager and replace the existing keystore with the new keystore. To replace the existing keystore, the IBM Flex System Manager must be at level 1.2.0 or higher.

13.2.1 Generating a Java keystore and Certificate Signing Request

Java ships with a utility named *Keytool* in its bin directory that can be used to create and edit keystore files. The first step is to create a keystore by using the following command:

```
keytool -genkey -alias <keystore_alias> -keyalg <encryption_algorithm> -keystore <path_to_the_keystore_being_created> -keysize <size_of_encryption_key>
```

Replace the specific options with the ones for your keystore, as shown in Example 13-1.

Example 13-1 Creating a keystore

```
keytool -genkey -alias Flex_Manager -keyalg RSA -keystore flexStore.jks -keysize 2048
```

This command will prompt you to create a password for the keystore being generated. Remember this password because it will be required later when you replace the keystore on the IBM Flex System Manager. This password will be requested for each subsequent **keytool** command run against the created keystore. In addition, it will prompt for organization and location information to create the keystore. Also, it will prompt for a password for the alias specified, which can be the same or different from the previous password. In this example, the alias was `Flex_Manager` and the keystore file was `flexStore.jks`. Using a strong keysize is advised because some mobile operating systems have restrictions on the keysize that they will accept.

After a keystore is generated, a Certificate Signing Request (CSR) can be created from the keystore by using the following command:

```
keytool -certreq -alias <keystore_alias> -keystore <path_to_the_keystore> -file <path_to_the_csr_file_being_created>
```

Replace the specific options with the options for your environment, as shown in Example 13-2.

Example 13-2 Creating the Certificate Signing Request

```
keytool -certreq -alias Flex_Manager -keystore flexStore.jks -file mydomain.csr
```

The Certificate Signing Request that is generated can be submitted to a CA to create a certificate signed by the CA. Send the certificate-signing request file to the CA. See the CA website for specific instructions about requesting a new certificate.

You can request either a test certificate or a production certificate from the CA. However, in a production environment, you must request a production certificate.

The next steps involve installing the CA root and any intermediate certificates into the keystore, and then installing the generated server certificate into the keystore. These certificates can be acquired from the CA that is used to generate the server certificate.

To install the root and intermediate certificates (start with the root certificate first), run the following command:

```
keytool -import -trustcacerts -alias <root_certificate_alias> -file <path_to_the_root_certificate> -keystore <path_to_the_keystore>
```

Replace the specific options with the options for your environment, as shown in Example 13-3.

Example 13-3 Installing root certificates

```
keytool -import -trustcacerts -alias root -file root.crt -keystore flexStore.jks
```

In Example 13-3, `root.crt` is the CA root or intermediate certificate and `flexStore.jks` is the name of the previously generated keystore. When prompted, select to trust the certificate that is being installed. Run this command for each certificate in the certificate chain.

Import the server certificate that is returned from the CA by running the following command:

```
keytool -import -trustcacerts -alias <server_certificate_alias> -file  
<path_to_server_certificate> -keystore <path_to_the_keystore>
```

Replace the specific options with the options for your environment, as shown in Example 13-4.

Example 13-4 Importing server certificates

```
keytool -import -trustcacerts -alias Flex_Manager_Server -file mydomain.crt  
-keystore flexStore.jks
```

In Example 13-4, the alias that is used is the alias for the server certificate, and the file that is provided is the server certificate file.

13.2.2 Installing the keystore into the IBM Flex System Manager

To install the keystore into the IBM Flex System Manager, follow these steps:

1. Copy the keystore to the IBM Flex System Manager by using Secure Copy Protocol (SCP).
2. Log in to the Flex System Manager by using Secure Shell (SSH).
3. Run the **smstop** command on the IBM Flex System Manager to stop the web server, as shown in Example 13-5.

Example 13-5 Stopping the web server

```
USERID@fsm1:~> smstop  
Shutting down IBM Director... done
```

4. Run the following command:

```
updcert -I -n <password of the keystore> -f <path to the keystore file>
```

The **-f** parameter is the location of the keystore file that is copied over SCP, as shown in Example 13-6.

Example 13-6 updcert command to install keystore onto Flex System Manager node

```
USERID@fsm1:~>updcert -I -n password -f /home/USERID/flexStore.jks
```

5. The web server restarts automatically, but its progress can be checked by using the **smstatus** command, as shown in Example 13-7.

Example 13-7 Check status of the web server

```
USERID@fsm1:~> smstatus  
Starting IBM Director...The starting process may take a while. Please use  
smstatus to check if the server is active  
USERID@fsm1:~> smstatus  
Active
```

13.2.3 Installation on Android

The installation of custom CA certificates on Android 2.3 is only supported on Motorola devices.

Motorola allows you to install custom CA certificates through its custom certificate manager. For more information, see this website:

[https://motorola-enterprise.custhelp.com/app/answers/detail/a_id/57093/~android--root-certificate-management](https://motorola-enterprise.custhelp.com/app/answers/detail/a_id/57093/~/android--root-certificate-management)

To use IBM Flex System Manager for Android with Android 2.3 devices by manufacturers other than Motorola, the certificate installed on IBM Flex System Manager must be recognized by one of the preinstalled certificate authorities on the Android device. Installing a certificate that is trusted by one of these preinstalled certificates allows a successful connection to an IBM Flex System Manager with the Android device.

Starting with Android 4.0, installing CA certificates is supported by Android natively. See this website:

<http://support.google.com/android/bin/answer.py?hl=en&answer=1649774>

On Android 4.0, the installed certificates can be seen inside “Trusted Credentials” in the “Security” section of “Settings”.

It is now possible to connect to IBM Flex System Manager systems that have a server certificate signed by the CA certificate that is installed by using the IBM Flex System Manager for Android application.

If the Android device does not connect successfully after the installation of a CA certificate, try restarting the Android device.

13.2.4 Installation on BlackBerry

There are two ways to get a CA certificate onto a BlackBerry device. A CA certificate can be installed by using BlackBerry Desktop Software or by importing it directly to the device.

To import the certificate by using the BlackBerry Desktop Software, follow these steps:

1. Download the certificate onto a device management system.
2. Import the certificate onto the management system through the web browser.
3. Connect to the device by using BlackBerry Desktop Software.
4. Select **Tools** → **Desktop Options**.
5. In the dialog that displays, select the **General** tab.
6. Select **Use certificate synchronization** and select **OK**.
7. In the left pane, select **Certificates**.
8. Select the store into which the CA certificate was imported.
9. Select the certificate and select **Sync Certificates**.

Or, to install the certificate directly, follow these steps:

1. Download the CA certificate to the device.
2. Opening the file prompts you to import the certificate.
3. Click **Import**, and then create a password for the keystore. This password can be any password that you want. It is used if you want to uninstall the certificate later.
4. Click **OK** after setting the password. The BlackBerry shows the certificate details and a green check mark indicating that the certificate is successfully installed.

To verify that the certificate was installed, go to **Home** → **Options** → **Security** → **Advanced Security** → **Certificates** → **<CA Certificate>**. It is now possible to connect to IBM Flex System Manager systems with a server certificate that is signed by the CA certificate that was installed using the IBM Flex System Manager for BlackBerry application. No warning message appears.

13.2.5 Installation on iOS

Acquire the CA root certificate on the iOS device through email, a website link, or another method.

After clicking the link or file, iOS automatically brings you to another window that is labeled Install Profile. In this window, press **Install**, then press **Install Now**. To verify the certificate was installed, open iOS Settings and go to **General** → **Profiles**. The imported CA certificate is listed. It is now possible to connect to IBM Flex System Manager systems with a server certificate that is signed by the CA certificate that was installed by using the IBM Flex System Manager for iOS application.

13.3 Using the Flex System Manager mobile application

The Flex System Manager mobile application enables you to view the following types of IBM Flex System information:

- ▶ Managed resource health problems and status
- ▶ Event history for chassis, compute nodes, and network devices
- ▶ Front and rear graphical views of a chassis
- ▶ Hardware components installed in a chassis
- ▶ Manage resource Vital Product Data (VPD) and firmware levels
- ▶ Recent scheduled jobs

The Flex System Manager for mobile devices enables you to manage your IBM Flex System and PureFlex System hardware remotely with the following hardware-management actions:

- ▶ Manage multiple chassis and multiple management nodes from a single application
- ▶ Perform actions on compute nodes, such as Power On, Power Off, Restart, and Shut Down and Power Off
- ▶ Perform actions on the Chassis Management Module (CMM), such as Virtual Reseat and Restart Primary CMM

Figure 13-1 shows the initial setup window where you can enter a passcode, which is used when executing important commands, such as node power on and off.




Figure 13-1 Entering your application passcode

Choose **Add Connection** as shown on Figure 13-2.




Figure 13-2 Setting up a new chassis connection

Enter Flex System Manager login information as shown in Figure 13-3 on page 578.

Figure 13-3 Enter your FSM login credentials

The successful connection and addition are shown in Figure 13-4.



Figure 13-4 Successful addition of a new chassis

Choosing the connection displays the chassis view with status indicators, as shown in Figure 13-5.

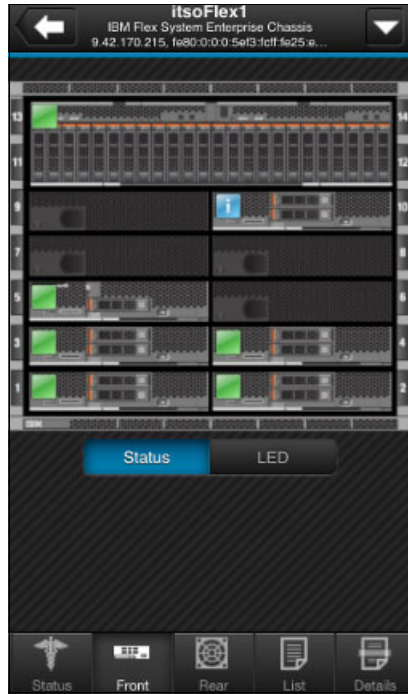


Figure 13-5 Front view of the chassis

Choose the LED button to see the front panel indicators for each node as shown in Figure 13-6.

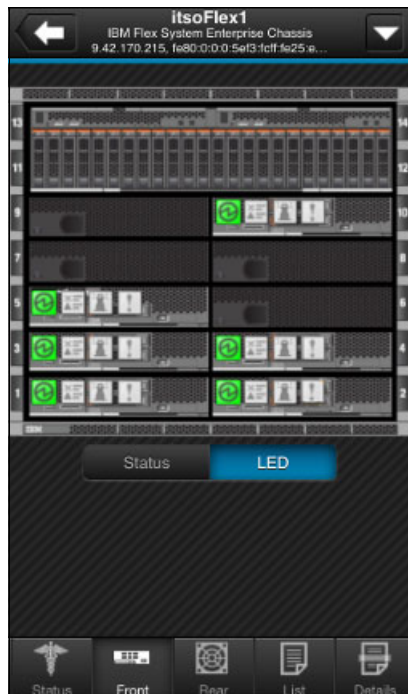


Figure 13-6 Showing front panel indicators

From the chassis pull-down menu, multiple options are available to manipulate the CMM as shown in Figure 13-7.

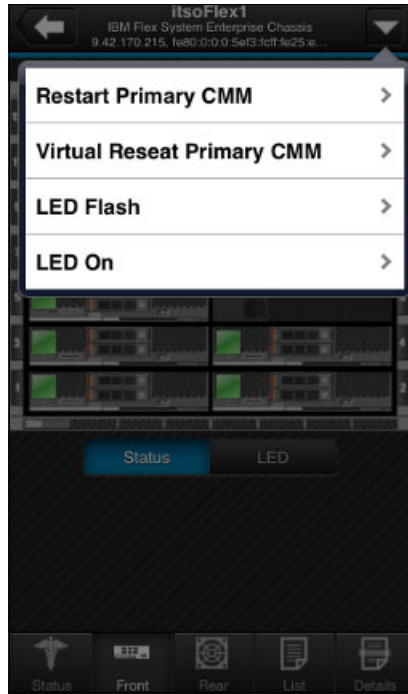


Figure 13-7 CMM functions in the chassis pull-down menu

You can view and perform actions on a specific node by choosing it as shown in Figure 13-8.




Figure 13-8 Node status

The event log for a specific node can be chosen by pressing the Event log button as shown in Figure 13-9.

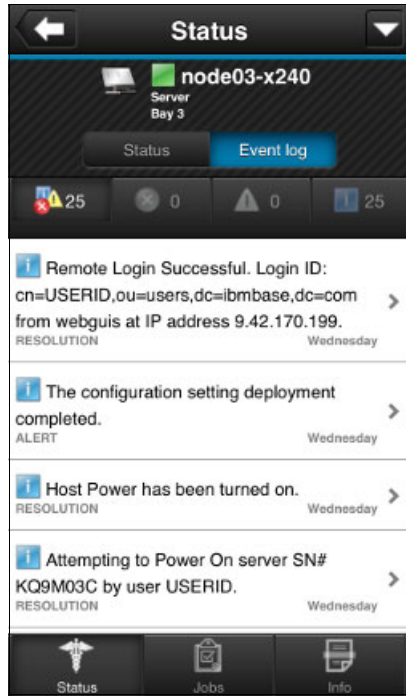


Figure 13-9 Viewing the event log for a specific node

You can press the Status button to perform actions on the node as shown in Figure 13-10.

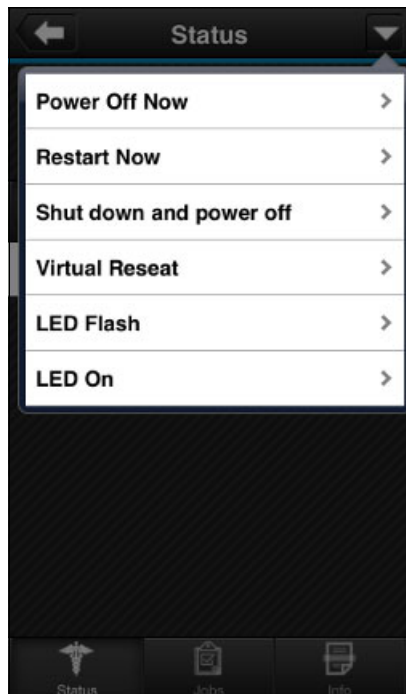


Figure 13-10 Status menu for a specific node

If you choose to perform an action, it prompts you for your app passcode as shown in Figure 13-11.




Figure 13-11 Prompt for your passcode to perform an action, such as LED flash

The LED flash job shows in the Recent Jobs list by pressing the **Jobs** button on the bottom of the display, as shown in Figure 13-12.




Figure 13-12 Recent jobs showing LED flash job

You can see information for a certain node by pressing the **Info** button as shown in Figure 13-13.




Figure 13-13 Node information

You can see the LED status for the node by scrolling to the bottom of the System Information window as shown in Figure 13-14.



Figure 13-14 LED information as shown on the System Information window

You can view status indicators from the rear of the chassis as shown in Figure 13-15.

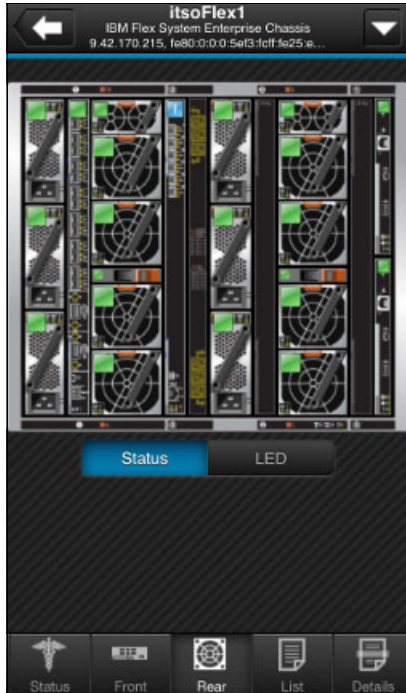


Figure 13-15 Viewing status indicators for the rear of the chassis

You can view a list of chassis components by using the List button as shown in Figure 13-16.



Figure 13-16 The list option to view chassis components

The Chassis button allows you to choose the chassis to work with as shown in Figure 13-17.

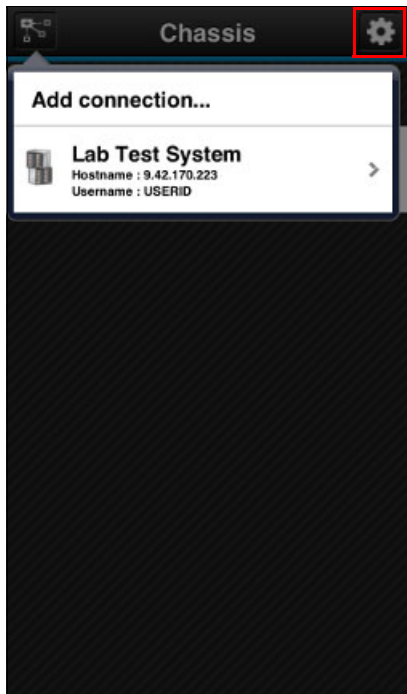


Figure 13-17 Choosing the chassis to work with

Press the gear icon (upper-right corner in Figure 13-17) to update settings.

Network information, login credentials, and the passcode can be updated, as shown in Figure 13-18.




Figure 13-18 Updating login or passcode information

Abbreviations and acronyms

AMM	advanced management module	HTTP	Hypertext Transfer Protocol
ASU	Advanced Settings Utility	HTTPS	HTTP over SSL
ATS	Advanced Technical Skills	I/O	input/output
BE3	BladeEngine 3	IBM	International Business Machines
CA	certificate authority	ID	identifier
CAS	Common Agent Services	IFM	IBM Fabric Manager
CD	compact disc	IMM	integrated management module
CD-ROM	compact-disc read-only memory	IMM2	Integrated management module II
CIM	Common Information Model	IP	Internet Protocol
CIMOM	CIM object manager	IPC	interprocess communication
CLI	command-line interface	ISV	independent software vendor
CMM	Chassis Management Module	IT	information technology
COM	Component Object Model	ITIL	Information Technology Infrastructure Library
CPU	central processing unit	ITSO	International Technical Support Organization
CRTM	Core Root of Trust Measurement	JRE	Java Runtime Environment
CSR	certificate signing request	KMS	Key Management System
DCOM	distributed component object model	KVM	kernel-based virtual machine
DHCP	Dynamic Host Configuration Protocol	LAN	local area network
DNS	Domain Name System	LDAP	Lightweight Directory Access Protocol
DRS	Distributed Resource Scheduler	LED	light emitting diode
DRTM	Dynamic Root of Trust Measurement	LLDP	Link Layer Discovery Protocol
DVD	digital versatile disc	LOM	LAN on motherboard
ECC	error checking and correcting	LP	low profile
ESP	Early Shipment Program	LUN	logical unit number
FC	Fibre Channel	MAC	Media Access Control
FCoE	Fibre Channel over Ethernet	MB	megabyte
FDR	fourteen data rate	MIB	Management Information Base
FFDC	first-failure data capture	MLC	multi-level cell
FSM	Flex System Manager	MPIO	multi-path I/O
FSP	flexible service processor	NFS	Network File System
FTP	File Transfer Protocol	NIM	Network Installation Manager
FTSS	Field Technical Sales Support	NL	nearline
FoD	Features on Demand	NTP	Network Time Protocol
GB	gigabyte	OID	object identifier
GUI	graphical user interface	OS	operating system
HA	high availability	OVF	Open Virtualization Format
HBA	host bus adapter	PEP	policy enforcement point
HDD	hard disk drive	QoS	quality of service
HS	hot swap		

RAID	redundant array of independent disks	USB	Universal Serial Bus
RAM	random access memory	VEB	Virtual Ethernet Bridging
RBAC	role-based access control	VEPA	Virtual Ethernet Port Aggregator
RSAP	Remote Service Access Point	VIOS	Virtual I/O Server
RDIMM	registered DIMM	VLAN	virtual local area network
RHEL	Red Hat Enterprise Linux	VM	virtual machine
RPM	Red Hat Package Manager	VMs	virtual machine
RSS	Receive-Side Scaling	VPD	Vital Product Data
SAN	storage area network	WWN	worldwide name
SAS	Serial Attached SCSI	WWPN	worldwide port name
SATA	Serial Advanced Technology Attachment	XML	Extensible Markup Language
SCP	Secure Copy Protocol		
SCS	Storage Copy Services		
SCSI	Small Computer System Interface		
SFF	small form factor		
SLP	Service Location Protocol		
SMI-S	Storage Management Initiative Specification		
SMS	Software Management Services		
SNIA	Storage Networking Industry Association		
SNMP	Simple Network Management Protocol		
SNMPv3	Simple Network Management Protocol v3		
SOL	Serial over LAN		
SSD	solid-state drive		
SSH	Secure Shell		
SSL	Secure Sockets Layer		
SW	special weight		
TB	terabyte		
TCG	Trusted Computing Group		
TCP	Transmission Control Protocol		
TCP/IP	Transmission Control Protocol/Internet Protocol		
TFTP	Trivial File Transfer Protocol		
TPM	Trusted Platform Module		
TXT	text		
UEFI	Unified Extensible Firmware Interface		
UI	user interface		
ULA	Unique Local Address		
ULAs	Unique Local Addresses		
URL	Uniform Resource Locator		

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- ▶ *IBM Flex System p260 and p460 Planning and Implementation Guide*, SG24-7989
- ▶ *IBM Flex System p270 Compute Node Planning and Implementation Guide*, SG24-8166
- ▶ *IBM PowerVM Live Partition Mobility*, SG24-7460
- ▶ *Implementing the IBM Storwize V7000 V6.3*, SG24-7938
- ▶ *Introduction to Storage Area Networks*, SG24-5470
- ▶ *Moving to IBM PureFlex System x86-to-x86 Migration*, REDP-4887
- ▶ *IBM Flex System Networking in an Enterprise Data Center*, REDP-4834

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ IBM Flex System Information Center
<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>
- ▶ IBM PureSystems offerings
<http://www.ibm.com/ibm/puresystems>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Redbooks

Implementing Systems Management of IBM PureFlex System

(1.0" spine)

0.875" x 1.498"

460 <-> 788 pages



Implementing Systems Management of IBM PureFlex System

Explores IBM PureFlex System and its systems management capabilities

Provides planning and deployment considerations

Gives step-by-step implementation instructions

To meet today's complex and ever-changing business demands, you need a solid foundation of compute, storage, networking, and software resources. This system must be simple to deploy and be able to quickly and automatically adapt to changing conditions. You also need to be able to take advantage of broad expertise and proven guidelines in systems management, applications, industry solutions, and more.

IBM PureFlex System combines no-compromise system designs along with built-in expertise and integrates them into complete, optimized scalable solutions. With IBM Flex System Manager, multiple solution components that include compute nodes, network and storage infrastructures, storage systems, and heterogeneous virtualization environments can be managed from a single panel.

This IBM Redbooks publication introduces IBM PureFlex System and IBM Flex System and their management devices and appliances. It provides implementation guidelines for managing Linux kernel-based virtual machine (KVM), IBM PowerVM, VMware vSphere, and Microsoft Hyper-V virtualization environments.

This book is intended for the IT community of clients, IBM Business Partners, and IBM employees who are interested in planning and implementing systems management of the IBM PureFlex System.

**INTERNATIONAL
TECHNICAL
SUPPORT
ORGANIZATION**

**BUILDING TECHNICAL
INFORMATION BASED ON
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**

SG24-8060-01

ISBN 0738439584